
Akka Java Documentation

Release 2.4.0

Typesafe Inc

September 30, 2015

CONTENTS

1	Introduction	1
1.1	What is Akka?	1
1.2	Why Akka?	2
1.3	Getting Started	3
1.4	The Obligatory Hello World	7
1.5	Use-case and Deployment Scenarios	7
1.6	Examples of use-cases for Akka	8
2	General	10
2.1	Terminology, Concepts	10
2.2	Actor Systems	12
2.3	What is an Actor?	14
2.4	Supervision and Monitoring	16
2.5	Actor References, Paths and Addresses	20
2.6	Location Transparency	25
2.7	Akka and the Java Memory Model	26
2.8	Message Delivery Reliability	28
2.9	Configuration	33
3	Actors	82
3.1	Actors	82
3.2	Typed Actors	102
3.3	Fault Tolerance	111
3.4	Dispatchers	126
3.5	Mailboxes	129
3.6	Routing	136
3.7	Building Finite State Machine Actors	153
3.8	Persistence	157
3.9	Persistence - Schema Evolution	184
3.10	Persistence Query	199
3.11	Persistence Query for LevelDB	209
3.12	Testing Actor Systems	212
4	Actors (Java with Lambda Support)	228
4.1	Actors (Java with Lambda Support)	228
4.2	Fault Tolerance (Java with Lambda Support)	247
4.3	FSM (Java with Lambda Support)	261
4.4	Persistence (Java with Lambda Support)	270
5	Futures and Agents	300
5.1	Futures	300
5.2	Agents	308
6	Networking	311

6.1	Cluster Specification	311
6.2	Cluster Usage	317
6.3	Cluster Singleton	336
6.4	Distributed Publish Subscribe in Cluster	339
6.5	Cluster Client	343
6.6	Cluster Sharding	346
6.7	Cluster Metrics Extension	356
6.8	Distributed Data	363
6.9	Remoting	381
6.10	Serialization	390
6.11	I/O	397
6.12	Using TCP	398
6.13	Using UDP	409
6.14	Camel	413
7	Utilities	426
7.1	Event Bus	426
7.2	Logging	433
7.3	Scheduler	439
7.4	Duration	442
7.5	Circuit Breaker	443
7.6	Akka Extensions	447
7.7	Use-case and Deployment Scenarios	449
8	HowTo: Common Patterns	451
8.1	Scheduling Periodic Messages	451
8.2	Single-Use Actor Trees with High-Level Error Reporting	452
8.3	Template Pattern	455
9	Experimental Modules	456
9.1	Multi Node Testing	456
9.2	Actors (Java with Lambda Support)	461
9.3	FSM (Java with Lambda Support)	480
9.4	Persistence Query	490
9.5	External Contributions	499
10	Information for Akka Developers	516
10.1	Building Akka	516
10.2	Multi JVM Testing	518
10.3	I/O Layer Design	521
10.4	Developer Guidelines	523
10.5	Documentation Guidelines	524
10.6	Team	526
11	Project Information	527
11.1	Migration Guides	527
11.2	Issue Tracking	544
11.3	Licenses	545
11.4	Sponsors	545
11.5	Project	545
12	Additional Information	548
12.1	Frequently Asked Questions	548
12.2	Books	551
12.3	Other Language Bindings	551
12.4	Akka in OSGi	551
12.5	Incomplete List of HTTP Frameworks	552

INTRODUCTION

1.1 What is Akka?

Scalable real-time transaction processing

We believe that writing correct concurrent, fault-tolerant and scalable applications is too hard. Most of the time it's because we are using the wrong tools and the wrong level of abstraction. Akka is here to change that. Using the Actor Model we raise the abstraction level and provide a better platform to build scalable, resilient and responsive applications—see the [Reactive Manifesto](#) for more details. For fault-tolerance we adopt the “let it crash” model which the telecom industry has used with great success to build applications that self-heal and systems that never stop. Actors also provide the abstraction for transparent distribution and the basis for truly scalable and fault-tolerant applications.

Akka is Open Source and available under the Apache 2 License.

Download from <http://akka.io/downloads>.

Please note that all code samples compile, so if you want direct access to the sources, have a look over at the Akka Docs subproject on github: for [Java](#) and [Scala](#).

1.1.1 Akka implements a unique hybrid

Actors

Actors give you:

- Simple and high-level abstractions for concurrency and parallelism.
- Asynchronous, non-blocking and highly performant event-driven programming model.
- Very lightweight event-driven processes (several million actors per GB of heap memory).

See the chapter for [Scala](#) or [Java](#).

Fault Tolerance

- Supervisor hierarchies with “let-it-crash” semantics.
- Supervisor hierarchies can span over multiple JVMs to provide truly fault-tolerant systems.
- Excellent for writing highly fault-tolerant systems that self-heal and never stop.

See [Fault Tolerance \(Scala\)](#) and [Fault Tolerance \(Java\)](#).

Location Transparency

Everything in Akka is designed to work in a distributed environment: all interactions of actors use pure message passing and everything is asynchronous.

For an overview of the cluster support see the *Java* and *Scala* documentation chapters.

Persistence

Messages received by an actor can optionally be persisted and replayed when the actor is started or restarted. This allows actors to recover their state, even after JVM crashes or when being migrated to another node.

You can find more details in the respective chapter for *Java* or *Scala*.

1.1.2 Scala and Java APIs

Akka has both a *scala-api* and a *Java Documentation*.

1.1.3 Akka can be used in two different ways

Akka can be used and deployed in different ways:

- As a library: used as a regular JAR on the classpath and/or in a web app, to be put into `WEB-INF/lib`
- Package with `sbt-native-packager`
- Package and deploy using `Typesafe ConductR`.

1.1.4 Commercial Support

Akka is available from Typesafe Inc. under a commercial license which includes development or production support, read more [here](#).

1.2 Why Akka?

1.2.1 What features can the Akka platform offer, over the competition?

Akka provides scalable real-time transaction processing.

Akka is an unified runtime and programming model for:

- Scale up (Concurrency)
- Scale out (Remoting)
- Fault tolerance

One thing to learn and admin, with high cohesion and coherent semantics.

Akka is a very scalable piece of software, not only in the context of performance but also in the size of applications it is useful for. The core of Akka, `akka-actor`, is very small and easily dropped into an existing project where you need asynchronicity and lockless concurrency without hassle.

You can choose to include only the parts of akka you need in your application. With CPUs growing more and more cores every cycle, Akka is the alternative that provides outstanding performance even if you're only running it on one machine. Akka also supplies a wide array of concurrency-paradigms, allowing users to choose the right tool for the job.

1.2.2 What's a good use-case for Akka?

We see Akka being adopted by many large organizations in a big range of industries:

- Investment and Merchant Banking
- Retail
- Social Media
- Simulation
- Gaming and Betting
- Automobile and Traffic Systems
- Health Care
- Data Analytics

and much more. Any system with the need for high-throughput and low latency is a good candidate for using Akka.

Actors let you manage service failures (Supervisors), load management (back-off strategies, timeouts and processing-isolation), as well as both horizontal and vertical scalability (add more cores and/or add more machines).

Here's what some of the Akka users have to say about how they are using Akka: <http://stackoverflow.com/questions/4493001/good-use-case-for-akka>

All this in the ApacheV2-licensed open source project.

1.3 Getting Started

1.3.1 Prerequisites

Akka requires that you have [Java 8](#) or later installed on your machine.

[Typesafe](#) provides a commercial build of Akka and related projects such as Scala or Play as part of the [Reactive Platform](#) which is made available for Java 6 in case your project can not upgrade to Java 8 just yet. It also includes additional commercial features or libraries.

1.3.2 Getting Started Guides and Template Projects

The best way to start learning Akka is to download [Typesafe Activator](#) and try out one of Akka Template Projects.

1.3.3 Download

There are several ways to download Akka. You can download it as part of the Typesafe Platform (as described above). You can download the full distribution, which includes all modules. Or you can use a build tool like Maven or SBT to download dependencies from the Akka Maven repository.

1.3.4 Modules

Akka is very modular and consists of several JARs containing different features.

- `akka-actor` – Classic Actors, Typed Actors, IO Actor etc.
- `akka-agent` – Agents, integrated with Scala STM
- `akka-camel` – Apache Camel integration

- `akka-cluster` – Cluster membership management, elastic routers.
- `akka-osgi` – utilities for using Akka in OSGi containers
- `akka-osgi-aries` – Aries blueprint for provisioning actor systems
- `akka-remote` – Remote Actors
- `akka-slf4j` – SLF4J Logger (event bus listener)
- `akka-testkit` – Toolkit for testing Actor systems

In addition to these stable modules there are several which are on their way into the stable core but are still marked “experimental” at this point. This does not mean that they do not function as intended, it primarily means that their API has not yet solidified enough in order to be considered frozen. You can help accelerating this process by giving feedback on these modules on our mailing list.

- `akka-contrib` – an assortment of contributions which may or may not be moved into core modules, see *External Contributions* for more details.

The filename of the actual JAR is for example `akka-actor_2.11-2.4.0.jar` (and analog for the other modules).

How to see the JARs dependencies of each Akka module is described in the *Dependencies* section.

1.3.5 Using a release distribution

Download the release you need from <http://akka.io/downloads> and unzip it.

1.3.6 Using a snapshot version

The Akka nightly snapshots are published to <http://repo.akka.io/snapshots/> and are versioned with both SNAPSHOT and timestamps. You can choose a timestamped version to work with and can decide when to update to a newer version.

Warning: The use of Akka SNAPSHOTs, nightlies and milestone releases is discouraged unless you know what you are doing.

1.3.7 Using a build tool

Akka can be used with build tools that support Maven repositories.

1.3.8 Maven repositories

For Akka version 2.1-M2 and onwards:

[Maven Central](#)

For previous Akka versions:

[Akka Repo](#)

1.3.9 Using Akka with Maven

The simplest way to get started with Akka and Maven is to check out the [Typesafe Activator](#) tutorial named [Akka Main in Java](#).

Since Akka is published to Maven Central (for versions since 2.1-M2), it is enough to add the Akka dependencies to the POM. For example, here is the dependency for `akka-actor`:

```
<dependency>
  <groupId>com.typesafe.akka</groupId>
  <artifactId>akka-actor_2.11</artifactId>
  <version>2.4.0</version>
</dependency>
```

For snapshot versions, the snapshot repository needs to be added as well:

```
<repositories>
  <repository>
    <id>akka-snapshots</id>
    <snapshots>
      <enabled>>true</enabled>
    </snapshots>
    <url>http://repo.akka.io/snapshots/</url>
  </repository>
</repositories>
```

Note: for snapshot versions both SNAPSHOT and timestamped versions are published.

1.3.10 Using Akka with SBT

The simplest way to get started with Akka and SBT is to check out the [Akka/SBT template project](#).

Summary of the essential parts for using Akka with SBT:

SBT installation instructions on <https://github.com/harrah/xsbt/wiki/Setup>

build.sbt file:

```
name := "My Project"

version := "1.0"

scalaVersion := "2.11.7"

libraryDependencies +=
  "com.typesafe.akka" %% "akka-actor" % "2.4.0"
```

Note: the libraryDependencies setting above is specific to SBT v0.12.x and higher. If you are using an older version of SBT, the libraryDependencies should look like this:

```
libraryDependencies +=
  "com.typesafe.akka" % "akka-actor_2.11" % "2.4.0"
```

For snapshot versions, the snapshot repository needs to be added as well:

```
resolvers += "Akka Snapshot Repository" at "http://repo.akka.io/snapshots/"
```

1.3.11 Using Akka with Gradle

Requires at least Gradle 1.4 Uses the Scala plugin

```
apply plugin: 'scala'

repositories {
  mavenCentral()
}

dependencies {
  compile 'org.scala-lang:scala-library:2.11.7'
}
```

```
tasks.withType(ScalaCompile) {
  scalaCompileOptions.useAnt = false
}

dependencies {
  compile group: 'com.typesafe.akka', name: 'akka-actor_2.11', version: '2.4.0'
  compile group: 'org.scala-lang', name: 'scala-library', version: '2.11.7'
}
```

For snapshot versions, the snapshot repository needs to be added as well:

```
repositories {
  mavenCentral()
  maven {
    url "http://repo.akka.io/snapshots/"
  }
}
```

1.3.12 Using Akka with Eclipse

Setup SBT project and then use `sbteclipse` to generate a Eclipse project.

1.3.13 Using Akka with IntelliJ IDEA

Setup SBT project and then use `sbt-idea` to generate a IntelliJ IDEA project.

1.3.14 Using Akka with NetBeans

Setup SBT project and then use `nbsbt` to generate a NetBeans project.

You should also use `nbscala` for general scala support in the IDE.

1.3.15 Do not use -optimize Scala compiler flag

Warning: Akka has not been compiled or tested with -optimize Scala compiler flag. Strange behavior has been reported by users that have tried it.

1.3.16 Build from sources

Akka uses Git and is hosted at [Github](#).

- Akka: clone the Akka repository from <https://github.com/akka/akka>

Continue reading the page on *Building Akka*

1.3.17 Need help?

If you have questions you can get help on the [Akka Mailing List](#).

You can also ask for [commercial support](#).

Thanks for being a part of the Akka community.

1.4 The Obligatory Hello World

The actor based version of the tough problem of printing a well-known greeting to the console is introduced in a [Typesafe Activator](#) tutorial named [Akka Main in Java](#).

The tutorial illustrates the generic launcher class `akka.Main` which expects only one command line argument: the class name of the application's main actor. This main method will then create the infrastructure needed for running the actors, start the given main actor and arrange for the whole application to shut down once the main actor terminates.

There is also another [Typesafe Activator](#) tutorial in the same problem domain that is named [Hello Akka!](#). It describes the basics of Akka in more depth.

1.5 Use-case and Deployment Scenarios

1.5.1 How can I use and deploy Akka?

Akka can be used in different ways:

- As a library: used as a regular JAR on the classpath and/or in a web app, to be put into `WEB-INF/lib`
- Package with [sbt-native-packager](#)
- Package and deploy using [Typesafe ConductR](#).

1.5.2 Native Packager

[sbt-native-packager](#) is a tool for creating distributions of any type of application, including an Akka applications.

Define sbt version in `project/build.properties` file:

```
sbt.version=0.13.7
```

Add [sbt-native-packager](#) in `project/plugins.sbt` file:

```
addSbtPlugin("com.typesafe.sbt" % "sbt-native-packager" % "1.0.0-RC1")
```

Use the package settings and optionally specify the `mainClass` in `build.sbt` file:

```
import NativePackagerHelper._

name := "akka-sample-main-scala"

version := "2.4.0"

scalaVersion := "2.11.7"

libraryDependencies += Seq(
  "com.typesafe.akka" %% "akka-actor" % "2.4.0"
)

enablePlugins(JavaServerAppPackaging)

mainClass in Compile := Some("sample.hello.Main")

mappings in Universal += {
  // optional example illustrating how to copy additional directory
  directory("scripts") ++
  // copy configuration files to config directory
  contentOf("src/main/resources").toMap.mapValues("config/" + _)
}
```

```
// add 'config' directory first in the classpath of the start script,
// an alternative is to set the config file locations via CLI parameters
// when starting the application
scriptClasspath := Seq("../config/") ++ scriptClasspath.value
```

Note: Use the `JavaServerAppPackaging`. Don't use `AkkaAppPackaging` (previously named `packageArchetype.akka_application`, since it doesn't have the same flexibility and quality as the `JavaServerAppPackaging`.

Use `sbt task dist` package the application.

To start the application (on a unix-based system):

```
cd target/universal/
unzip akka-sample-main-scala-2.4.0.zip
chmod u+x akka-sample-main-scala-2.4.0/bin/akka-sample-main-scala
akka-sample-main-scala-2.4.0/bin/akka-sample-main-scala sample.hello.Main
```

Use `Ctrl-C` to interrupt and exit the application.

On a Windows machine you can also use the `bin\akka-sample-main-scala.bat` script.

1.6 Examples of use-cases for Akka

We see Akka being adopted by many large organizations in a big range of industries all from investment and merchant banking, retail and social media, simulation, gaming and betting, automobile and traffic systems, health care, data analytics and much more. Any system that have the need for high-throughput and low latency is a good candidate for using Akka.

There is a great discussion on use-cases for Akka with some good write-ups by production users [here](#)

1.6.1 Here are some of the areas where Akka is being deployed into production

Transaction processing (Online Gaming, Finance/Banking, Trading, Statistics, Betting, Social Media, Telecom)

Scale up, scale out, fault-tolerance / HA

Service backend (any industry, any app)

Service REST, SOAP, Cometd, WebSockets etc Act as message hub / integration layer Scale up, scale out, fault-tolerance / HA

Concurrency/parallelism (any app)

Correct Simple to work with and understand Just add the jars to your existing JVM project (use Scala, Java, Groovy or JRuby)

Simulation

Master/Worker, Compute Grid, MapReduce etc.

Batch processing (any industry)

Camel integration to hook up with batch data sources Actors divide and conquer the batch workloads

Communications Hub (Telecom, Web media, Mobile media)

Scale up, scale out, fault-tolerance / HA

Gaming and Betting (MOM, online gaming, betting)

Scale up, scale out, fault-tolerance / HA

Business Intelligence/Data Mining/general purpose crunching

Scale up, scale out, fault-tolerance / HA

Complex Event Stream Processing

Scale up, scale out, fault-tolerance / HA

GENERAL

2.1 Terminology, Concepts

In this chapter we attempt to establish a common terminology to define a solid ground for communicating about concurrent, distributed systems which Akka targets. Please note that, for many of these terms, there is no single agreed definition. We simply seek to give working definitions that will be used in the scope of the Akka documentation.

2.1.1 Concurrency vs. Parallelism

Concurrency and parallelism are related concepts, but there are small differences. *Concurrency* means that two or more tasks are making progress even though they might not be executing simultaneously. This can for example be realized with time slicing where parts of tasks are executed sequentially and mixed with parts of other tasks. *Parallelism* on the other hand arise when the execution can be truly simultaneous.

2.1.2 Asynchronous vs. Synchronous

A method call is considered *synchronous* if the caller cannot make progress until the method returns a value or throws an exception. On the other hand, an *asynchronous* call allows the caller to progress after a finite number of steps, and the completion of the method may be signalled via some additional mechanism (it might be a registered callback, a Future, or a message).

A synchronous API may use blocking to implement synchrony, but this is not a necessity. A very CPU intensive task might give a similar behavior as blocking. In general, it is preferred to use asynchronous APIs, as they guarantee that the system is able to progress. Actors are asynchronous by nature: an actor can progress after a message send without waiting for the actual delivery to happen.

2.1.3 Non-blocking vs. Blocking

We talk about *blocking* if the delay of one thread can indefinitely delay some of the other threads. A good example is a resource which can be used exclusively by one thread using mutual exclusion. If a thread holds on to the resource indefinitely (for example accidentally running an infinite loop) other threads waiting on the resource can not progress. In contrast, *non-blocking* means that no thread is able to indefinitely delay others.

Non-blocking operations are preferred to blocking ones, as the overall progress of the system is not trivially guaranteed when it contains blocking operations.

2.1.4 Deadlock vs. Starvation vs. Live-lock

Deadlock arises when several participants are waiting on each other to reach a specific state to be able to progress. As none of them can progress without some other participant to reach a certain state (a “Catch-22” problem) all

affected subsystems stall. Deadlock is closely related to *blocking*, as it is necessary that a participant thread be able to delay the progression of other threads indefinitely.

In the case of *deadlock*, no participants can make progress, while in contrast *Starvation* happens, when there are participants that can make progress, but there might be one or more that cannot. Typical scenario is the case of a naive scheduling algorithm that always selects high-priority tasks over low-priority ones. If the number of incoming high-priority tasks is constantly high enough, no low-priority ones will be ever finished.

Livelock is similar to *deadlock* as none of the participants make progress. The difference though is that instead of being frozen in a state of waiting for others to progress, the participants continuously change their state. An example scenario when two participants have two identical resources available. They each try to get the resource, but they also check if the other needs the resource, too. If the resource is requested by the other participant, they try to get the other instance of the resource. In the unfortunate case it might happen that the two participants “bounce” between the two resources, never acquiring it, but always yielding to the other.

2.1.5 Race Condition

We call it a *Race condition* when an assumption about the ordering of a set of events might be violated by external non-deterministic effects. Race conditions often arise when multiple threads have a shared mutable state, and the operations of thread on the state might be interleaved causing unexpected behavior. While this is a common case, shared state is not necessary to have race conditions. One example could be a client sending unordered packets (e.g UDP datagrams) P1, P2 to a server. As the packets might potentially travel via different network routes, it is possible that the server receives P2 first and P1 afterwards. If the messages contain no information about their sending order it is impossible to determine by the server that they were sent in a different order. Depending on the meaning of the packets this can cause race conditions.

Note: The only guarantee that Akka provides about messages sent between a given pair of actors is that their order is always preserved. see [Message Delivery Reliability](#)

2.1.6 Non-blocking Guarantees (Progress Conditions)

As discussed in the previous sections blocking is undesirable for several reasons, including the dangers of deadlocks and reduced throughput in the system. In the following sections we discuss various non-blocking properties with different strength.

Wait-freedom

A method is *wait-free* if every call is guaranteed to finish in a finite number of steps. If a method is *bounded wait-free* then the number of steps has a finite upper bound.

From this definition it follows that wait-free methods are never blocking, therefore deadlock can not happen. Additionally, as each participant can progress after a finite number of steps (when the call finishes), wait-free methods are free of starvation.

Lock-freedom

Lock-freedom is a weaker property than *wait-freedom*. In the case of lock-free calls, infinitely often some method finishes in a finite number of steps. This definition implies that no deadlock is possible for lock-free calls. On the other hand, the guarantee that *some call finishes* in a finite number of steps is not enough to guarantee that *all of them eventually finish*. In other words, lock-freedom is not enough to guarantee the lack of starvation.

Obstruction-freedom

Obstruction-freedom is the weakest non-blocking guarantee discussed here. A method is called *obstruction-free* if there is a point in time after which it executes in isolation (other threads make no steps, e.g.: become suspended),

it finishes in a bounded number of steps. All lock-free objects are obstruction-free, but the opposite is generally not true.

Optimistic concurrency control (OCC) methods are usually obstruction-free. The OCC approach is that every participant tries to execute its operation on the shared object, but if a participant detects conflicts from others, it rolls back the modifications, and tries again according to some schedule. If there is a point in time, where one of the participants is the only one trying, the operation will succeed.

2.1.7 Recommended literature

- The Art of Multiprocessor Programming, M. Herlihy and N Shavit, 2008. ISBN 978-0123705914
- Java Concurrency in Practice, B. Goetz, T. Peierls, J. Bloch, J. Bowbeer, D. Holmes and D. Lea, 2006. ISBN 978-0321349606

2.2 Actor Systems

Actors are objects which encapsulate state and behavior, they communicate exclusively by exchanging messages which are placed into the recipient's mailbox. In a sense, actors are the most stringent form of object-oriented programming, but it serves better to view them as persons: while modeling a solution with actors, envision a group of people and assign sub-tasks to them, arrange their functions into an organizational structure and think about how to escalate failure (all with the benefit of not actually dealing with people, which means that we need not concern ourselves with their emotional state or moral issues). The result can then serve as a mental scaffolding for building the software implementation.

Note: An ActorSystem is a heavyweight structure that will allocate 1...N Threads, so create one per logical application.

2.2.1 Hierarchical Structure

Like in an economic organization, actors naturally form hierarchies. One actor, which is to oversee a certain function in the program might want to split up its task into smaller, more manageable pieces. For this purpose it starts child actors which it supervises. While the details of supervision are explained [here](#), we shall concentrate on the underlying concepts in this section. The only prerequisite is to know that each actor has exactly one supervisor, which is the actor that created it.

The quintessential feature of actor systems is that tasks are split up and delegated until they become small enough to be handled in one piece. In doing so, not only is the task itself clearly structured, but the resulting actors can be reasoned about in terms of which messages they should process, how they should react normally and how failure should be handled. If one actor does not have the means for dealing with a certain situation, it sends a corresponding failure message to its supervisor, asking for help. The recursive structure then allows to handle failure at the right level.

Compare this to layered software design which easily devolves into defensive programming with the aim of not leaking any failure out: if the problem is communicated to the right person, a better solution can be found than if trying to keep everything "under the carpet".

Now, the difficulty in designing such a system is how to decide who should supervise what. There is of course no single best solution, but there are a few guidelines which might be helpful:

- If one actor manages the work another actor is doing, e.g. by passing on sub-tasks, then the manager should supervise the child. The reason is that the manager knows which kind of failures are expected and how to handle them.
- If one actor carries very important data (i.e. its state shall not be lost if avoidable), this actor should source out any possibly dangerous sub-tasks to children it supervises and handle failures of these children as appropriate. Depending on the nature of the requests, it may be best to create a new child for each request,

which simplifies state management for collecting the replies. This is known as the “Error Kernel Pattern” from Erlang.

- If one actor depends on another actor for carrying out its duty, it should watch that other actor’s liveness and act upon receiving a termination notice. This is different from supervision, as the watching party has no influence on the supervisor strategy, and it should be noted that a functional dependency alone is not a criterion for deciding where to place a certain child actor in the hierarchy.

There are of course always exceptions to these rules, but no matter whether you follow the rules or break them, you should always have a reason.

2.2.2 Configuration Container

The actor system as a collaborating ensemble of actors is the natural unit for managing shared facilities like scheduling services, configuration, logging, etc. Several actor systems with different configuration may co-exist within the same JVM without problems, there is no global shared state within Akka itself. Couple this with the transparent communication between actor systems—within one node or across a network connection—to see that actor systems themselves can be used as building blocks in a functional hierarchy.

2.2.3 Actor Best Practices

1. Actors should be like nice co-workers: do their job efficiently without bothering everyone else needlessly and avoid hogging resources. Translated to programming this means to process events and generate responses (or more requests) in an event-driven manner. Actors should not block (i.e. passively wait while occupying a Thread) on some external entity—which might be a lock, a network socket, etc.—unless it is unavoidable; in the latter case see below.
2. Do not pass mutable objects between actors. In order to ensure that, prefer immutable messages. If the encapsulation of actors is broken by exposing their mutable state to the outside, you are back in normal Java concurrency land with all the drawbacks.
3. Actors are made to be containers for behavior and state, embracing this means to not routinely send behavior within messages (which may be tempting using Scala closures). One of the risks is to accidentally share mutable state between actors, and this violation of the actor model unfortunately breaks all the properties which make programming in actors such a nice experience.
4. Top-level actors are the innermost part of your Error Kernel, so create them sparingly and prefer truly hierarchical systems. This has benefits with respect to fault-handling (both considering the granularity of configuration and the performance) and it also reduces the strain on the guardian actor, which is a single point of contention if over-used.

2.2.4 Blocking Needs Careful Management

In some cases it is unavoidable to do blocking operations, i.e. to put a thread to sleep for an indeterminate time, waiting for an external event to occur. Examples are legacy RDBMS drivers or messaging APIs, and the underlying reason is typically that (network) I/O occurs under the covers. When facing this, you may be tempted to just wrap the blocking call inside a `Future` and work with that instead, but this strategy is too simple: you are quite likely to find bottlenecks or run out of memory or threads when the application runs under increased load.

The non-exhaustive list of adequate solutions to the “blocking problem” includes the following suggestions:

- Do the blocking call within an actor (or a set of actors managed by a router [*Java*, *Scala*]), making sure to configure a thread pool which is either dedicated for this purpose or sufficiently sized.
- Do the blocking call within a `Future`, ensuring an upper bound on the number of such calls at any point in time (submitting an unbounded number of tasks of this nature will exhaust your memory or thread limits).
- Do the blocking call within a `Future`, providing a thread pool with an upper limit on the number of threads which is appropriate for the hardware on which the application runs.

- Dedicate a single thread to manage a set of blocking resources (e.g. a NIO selector driving multiple channels) and dispatch events as they occur as actor messages.

The first possibility is especially well-suited for resources which are single-threaded in nature, like database handles which traditionally can only execute one outstanding query at a time and use internal synchronization to ensure this. A common pattern is to create a router for N actors, each of which wraps a single DB connection and handles queries as sent to the router. The number N must then be tuned for maximum throughput, which will vary depending on which DBMS is deployed on what hardware.

Note: Configuring thread pools is a task best delegated to Akka, simply configure in the `application.conf` and instantiate through an `ActorSystem` [[Java](#), [Scala](#)]

2.2.5 What you should not concern yourself with

An actor system manages the resources it is configured to use in order to run the actors which it contains. There may be millions of actors within one such system, after all the mantra is to view them as abundant and they weigh in at an overhead of only roughly 300 bytes per instance. Naturally, the exact order in which messages are processed in large systems is not controllable by the application author, but this is also not intended. Take a step back and relax while Akka does the heavy lifting under the hood.

2.3 What is an Actor?

The previous section about [Actor Systems](#) explained how actors form hierarchies and are the smallest unit when building an application. This section looks at one such actor in isolation, explaining the concepts you encounter while implementing it. For a more in depth reference with all the details please refer to [Actors \(Scala\)](#) and [Untyped Actors \(Java\)](#).

An actor is a container for [State](#), [Behavior](#), a [Mailbox](#), [Children](#) and a [Supervisor Strategy](#). All of this is encapsulated behind an [Actor Reference](#). Finally, this happens [When an Actor Terminates](#).

2.3.1 Actor Reference

As detailed below, an actor object needs to be shielded from the outside in order to benefit from the actor model. Therefore, actors are represented to the outside using actor references, which are objects that can be passed around freely and without restriction. This split into inner and outer object enables transparency for all the desired operations: restarting an actor without needing to update references elsewhere, placing the actual actor object on remote hosts, sending messages to actors in completely different applications. But the most important aspect is that it is not possible to look inside an actor and get hold of its state from the outside, unless the actor unwisely publishes this information itself.

2.3.2 State

Actor objects will typically contain some variables which reflect possible states the actor may be in. This can be an explicit state machine (e.g. using the `fsm-scala` module), or it could be a counter, set of listeners, pending requests, etc. These data are what make an actor valuable, and they must be protected from corruption by other actors. The good news is that Akka actors conceptually each have their own light-weight thread, which is completely shielded from the rest of the system. This means that instead of having to synchronize access using locks you can just write your actor code without worrying about concurrency at all.

Behind the scenes Akka will run sets of actors on sets of real threads, where typically many actors share one thread, and subsequent invocations of one actor may end up being processed on different threads. Akka ensures that this implementation detail does not affect the single-threadedness of handling the actor's state.

Because the internal state is vital to an actor's operations, having inconsistent state is fatal. Thus, when the actor fails and is restarted by its supervisor, the state will be created from scratch, like upon first creating the actor. This is to enable the ability of self-healing of the system.

Optionally, an actor's state can be automatically recovered to the state before a restart by persisting received messages and replaying them after restart (see *persistence-scala*).

2.3.3 Behavior

Every time a message is processed, it is matched against the current behavior of the actor. Behavior means a function which defines the actions to be taken in reaction to the message at that point in time, say forward a request if the client is authorized, deny it otherwise. This behavior may change over time, e.g. because different clients obtain authorization over time, or because the actor may go into an "out-of-service" mode and later come back. These changes are achieved by either encoding them in state variables which are read from the behavior logic, or the function itself may be swapped out at runtime, see the `become` and `unbecome` operations. However, the initial behavior defined during construction of the actor object is special in the sense that a restart of the actor will reset its behavior to this initial one.

2.3.4 Mailbox

An actor's purpose is the processing of messages, and these messages were sent to the actor from other actors (or from outside the actor system). The piece which connects sender and receiver is the actor's mailbox: each actor has exactly one mailbox to which all senders enqueue their messages. Enqueuing happens in the time-order of send operations, which means that messages sent from different actors may not have a defined order at runtime due to the apparent randomness of distributing actors across threads. Sending multiple messages to the same target from the same actor, on the other hand, will enqueue them in the same order.

There are different mailbox implementations to choose from, the default being a FIFO: the order of the messages processed by the actor matches the order in which they were enqueued. This is usually a good default, but applications may need to prioritize some messages over others. In this case, a priority mailbox will enqueue not always at the end but at a position as given by the message priority, which might even be at the front. While using such a queue, the order of messages processed will naturally be defined by the queue's algorithm and in general not be FIFO.

An important feature in which Akka differs from some other actor model implementations is that the current behavior must always handle the next dequeued message, there is no scanning the mailbox for the next matching one. Failure to handle a message will typically be treated as a failure, unless this behavior is overridden.

2.3.5 Children

Each actor is potentially a supervisor: if it creates children for delegating sub-tasks, it will automatically supervise them. The list of children is maintained within the actor's context and the actor has access to it. Modifications to the list are done by creating (`context.actorOf(...)`) or stopping (`context.stop(child)`) children and these actions are reflected immediately. The actual creation and termination actions happen behind the scenes in an asynchronous way, so they do not "block" their supervisor.

2.3.6 Supervisor Strategy

The final piece of an actor is its strategy for handling faults of its children. Fault handling is then done transparently by Akka, applying one of the strategies described in *Supervision and Monitoring* for each incoming failure. As this strategy is fundamental to how an actor system is structured, it cannot be changed once an actor has been created.

Considering that there is only one such strategy for each actor, this means that if different strategies apply to the various children of an actor, the children should be grouped beneath intermediate supervisors with matching strategies, preferring once more the structuring of actor systems according to the splitting of tasks into sub-tasks.

2.3.7 When an Actor Terminates

Once an actor terminates, i.e. fails in a way which is not handled by a restart, stops itself or is stopped by its supervisor, it will free up its resources, draining all remaining messages from its mailbox into the system's "dead letter mailbox" which will forward them to the `EventStream` as `DeadLetters`. The mailbox is then replaced within the actor reference with a system mailbox, redirecting all new messages to the `EventStream` as `DeadLetters`. This is done on a best effort basis, though, so do not rely on it in order to construct "guaranteed delivery".

The reason for not just silently dumping the messages was inspired by our tests: we register the `TestEventListener` on the event bus to which the dead letters are forwarded, and that will log a warning for every dead letter received—this has been very helpful for deciphering test failures more quickly. It is conceivable that this feature may also be of use for other purposes.

2.4 Supervision and Monitoring

This chapter outlines the concept behind supervision, the primitives offered and their semantics. For details on how that translates into real code, please refer to the corresponding chapters for Scala and Java APIs.

2.4.1 What Supervision Means

As described in *Actor Systems* supervision describes a dependency relationship between actors: the supervisor delegates tasks to subordinates and therefore must respond to their failures. When a subordinate detects a failure (i.e. throws an exception), it suspends itself and all its subordinates and sends a message to its supervisor, signaling failure. Depending on the nature of the work to be supervised and the nature of the failure, the supervisor has a choice of the following four options:

1. Resume the subordinate, keeping its accumulated internal state
2. Restart the subordinate, clearing out its accumulated internal state
3. Stop the subordinate permanently
4. Escalate the failure, thereby failing itself


It is important to always view an actor as part of a supervision hierarchy, which explains the existence of the fourth choice (as a supervisor also is subordinate to another supervisor higher up) and has implications on the first three: resuming an actor resumes all its subordinates, restarting an actor entails restarting all its subordinates (but see below for more details), similarly terminating an actor will also terminate all its subordinates. It should be noted that the default behavior of the `preRestart` hook of the `Actor` class is to terminate all its children before restarting, but this hook can be overridden; the recursive restart applies to all children left after this hook has been executed.

Each supervisor is configured with a function translating all possible failure causes (i.e. exceptions) into one of the four choices given above; notably, this function does not take the failed actor's identity as an input. It is quite easy to come up with examples of structures where this might not seem flexible enough, e.g. wishing for different strategies to be applied to different subordinates. At this point it is vital to understand that supervision is about forming a recursive fault handling structure. If you try to do too much at one level, it will become hard to reason about, hence the recommended way in this case is to add a level of supervision.

Akka implements a specific form called "parental supervision". Actors can only be created by other actors—where the top-level actor is provided by the library—and each created actor is supervised by its parent. This restriction makes the formation of actor supervision hierarchies implicit and encourages sound design decisions. It should be noted that this also guarantees that actors cannot be orphaned or attached to supervisors from the outside, which might otherwise catch them unawares. In addition, this yields a natural and clean shutdown procedure for (sub-trees of) actor applications.

Warning: Supervision related parent-child communication happens by special system messages that have their own mailboxes separate from user messages. This implies that supervision related events are not deterministically ordered relative to ordinary messages. In general, the user cannot influence the order of normal messages and failure notifications. For details and example see the *Discussion: Message Ordering* section.

2.4.2 The Top-Level Supervisors



An actor system will during its creation start at least three actors, shown in the image above. For more information about the consequences for actor paths see *Top-Level Scopes for Actor Paths*.

/user: The Guardian Actor

The actor which is probably most interacted with is the parent of all user-created actors, the guardian named `"/user"`. Actors created using `system.actorOf()` are children of this actor. This means that when this guardian terminates, all normal actors in the system will be shutdown, too. It also means that this guardian's supervisor strategy determines how the top-level normal actors are supervised. Since Akka 2.1 it is possible to configure this using the setting `akka.actor.guardian-supervisor-strategy`, which takes the fully-qualified class-name of a `SupervisorStrategyConfigurator`. When the guardian escalates a failure, the root guardian's response will be to terminate the guardian, which in effect will shut down the whole actor system.

/system: The System Guardian

This special guardian has been introduced in order to achieve an orderly shut-down sequence where logging remains active while all normal actors terminate, even though logging itself is implemented using actors. This is realized by having the system guardian watch the user guardian and initiate its own shut-down upon reception of the `Terminated` message. The top-level system actors are supervised using a strategy which will restart indefinitely upon all types of `Exception` except for `ActorInitializationException` and `ActorKilledException`, which will terminate the child in question. All other throwables are escalated, which will shut down the whole actor system.

/: The Root Guardian

The root guardian is the grand-parent of all so-called "top-level" actors and supervises all the special actors mentioned in *Top-Level Scopes for Actor Paths* using the `SupervisorStrategy.stoppingStrategy`, whose purpose is to terminate the child upon any type of `Exception`. All other throwables will be escalated ... but to whom? Since every real actor has a supervisor, the supervisor of the root guardian cannot be a real

actor. And because this means that it is “outside of the bubble”, it is called the “bubble-walker”. This is a synthetic `ActorRef` which in effect stops its child upon the first sign of trouble and sets the actor system’s `isTerminated` status to `true` as soon as the root guardian is fully terminated (all children recursively stopped).

2.4.3 What Restarting Means

When presented with an actor which failed while processing a certain message, causes for the failure fall into three categories:

- Systematic (i.e. programming) error for the specific message received
- (Transient) failure of some external resource used during processing the message
- Corrupt internal state of the actor

Unless the failure is specifically recognizable, the third cause cannot be ruled out, which leads to the conclusion that the internal state needs to be cleared out. If the supervisor decides that its other children or itself is not affected by the corruption—e.g. because of conscious application of the error kernel pattern—it is therefore best to restart the child. This is carried out by creating a new instance of the underlying `Actor` class and replacing the failed instance with the fresh one inside the child’s `ActorRef`; the ability to do this is one of the reasons for encapsulating actors within special references. The new actor then resumes processing its mailbox, meaning that the restart is not visible outside of the actor itself with the notable exception that the message during which the failure occurred is not re-processed.

The precise sequence of events during a restart is the following:

1. suspend the actor (which means that it will not process normal messages until resumed), and recursively suspend all children
2. call the old instance’s `preRestart` hook (defaults to sending termination requests to all children and calling `postStop`)
3. wait for all children which were requested to terminate (using `context.stop()`) during `preRestart` to actually terminate; this—like all actor operations—is non-blocking, the termination notice from the last killed child will effect the progression to the next step
4. create new actor instance by invoking the originally provided factory again
5. invoke `postRestart` on the new instance (which by default also calls `preStart`)
6. send restart request to all children which were not killed in step 3; restarted children will follow the same process recursively, from step 2
7. resume the actor

2.4.4 What Lifecycle Monitoring Means

Note: Lifecycle Monitoring in Akka is usually referred to as `DeathWatch`

In contrast to the special relationship between parent and child described above, each actor may monitor any other actor. Since actors emerge from creation fully alive and restarts are not visible outside of the affected supervisors, the only state change available for monitoring is the transition from alive to dead. Monitoring is thus used to tie one actor to another so that it may react to the other actor’s termination, in contrast to supervision which reacts to failure.

Lifecycle monitoring is implemented using a `Terminated` message to be received by the monitoring actor, where the default behavior is to throw a special `DeathPactException` if not otherwise handled. In order to start listening for `Terminated` messages, invoke `ActorContext.watch(targetActorRef)`. To stop listening, invoke `ActorContext.unwatch(targetActorRef)`. One important property is that the message will be delivered irrespective of the order in which the monitoring request and target’s termination occur, i.e. you still get the message even if at the time of registration the target is already dead.

Monitoring is particularly useful if a supervisor cannot simply restart its children and has to terminate them, e.g. in case of errors during actor initialization. In that case it should monitor those children and re-create them or schedule itself to retry this at a later time.

Another common use case is that an actor needs to fail in the absence of an external resource, which may also be one of its own children. If a third party terminates a child by way of the `system.stop(child)` method or sending a `PoisonPill`, the supervisor might well be affected.

Delayed restarts with the `BackoffSupervisor` pattern

Provided as a build-in pattern the `akka.pattern.BackoffSupervisor` actor implements the so-called *exponential backoff supervision strategy*, which can be used to death-watch an actor, and when it terminates try to start it again, each time with a growing time delay between those restarts.

This pattern is useful when the started actor fails because some external resource is not available, and we need to give it some time to start-up again. One of the prime examples when this is useful is when a *PersistentActor* fails with an persistence failure - which indicates that the database may be down or overloaded, in such situations it makes most sense to give it a little bit of time to recover before the persistent actor is restarted.

The following Scala snippet shows how to create a backoff supervisor which will start the given echo actor in increasing intervals of 3, 6, 12, 24 and finally 30 seconds:

```
val childProps = Props(classOf[EchoActor])

val supervisor = BackoffSupervisor.props(
  childProps,
  childName = "myEcho",
  minBackoff = 3.seconds,
  maxBackoff = 30.seconds,
  randomFactor = 0.2) // adds 20% "noise" to vary the intervals slightly

system.actorOf(supervisor, name = "echoSupervisor")
```

The above is equivalent to this Java code:

```
import scala.concurrent.duration.Duration;

final Props childProps = Props.create(EchoActor.class);

final Props supervisorProps = BackoffSupervisor.props(
  childProps,
  "myEcho",
  Duration.create(3, TimeUnit.SECONDS),
  Duration.create(30, TimeUnit.SECONDS),
  0.2); // adds 20% "noise" to vary the intervals slightly

system.actorOf(supervisorProps, "echoSupervisor");
```

Using a `randomFactor` to add a little bit of additional variance to the backoff intervals is highly recommended, in order to avoid multiple actors re-start at the exact same point in time, for example because they were stopped due to a shared resource such as a database going down and re-starting after the same configured interval. By adding additional randomness to the re-start intervals the actors will start in slightly different points in time, thus avoiding large spikes of traffic hitting the recovering shared database or other resource that they all need to contact.

2.4.5 One-For-One Strategy vs. All-For-One Strategy

There are two classes of supervision strategies which come with Akka: `OneForOneStrategy` and `AllForOneStrategy`. Both are configured with a mapping from exception type to supervision directive (see [above](#)) and limits on how often a child is allowed to fail before terminating it. The difference between them is that the former applies the obtained directive only to the failed child, whereas the latter applies it to all siblings as well. Normally, you should use the `OneForOneStrategy`, which also is the default if none is specified explicitly.


The `AllForOneStrategy` is applicable in cases where the ensemble of children has such tight dependencies among them, that a failure of one child affects the function of the others, i.e. they are inextricably linked. Since a restart does not clear out the mailbox, it often is best to terminate the children upon failure and re-create them explicitly from the supervisor (by watching the children's lifecycle); otherwise you have to make sure that it is no problem for any of the actors to receive a message which was queued before the restart but processed afterwards.

Normally stopping a child (i.e. not in response to a failure) will not automatically terminate the other children in an all-for-one strategy; this can easily be done by watching their lifecycle: if the `Terminated` message is not handled by the supervisor, it will throw a `DeathPactException` which (depending on its supervisor) will restart it, and the default `preRestart` action will terminate all children. Of course this can be handled explicitly as well.

Please note that creating one-off actors from an all-for-one supervisor entails that failures escalated by the temporary actor will affect all the permanent ones. If this is not desired, install an intermediate supervisor; this can very easily be done by declaring a router of size 1 for the worker, see *routing-scala* or *Routing*.

2.5 Actor References, Paths and Addresses

This chapter describes how actors are identified and located within a possibly distributed actor system. It ties into the central idea that *Actor Systems* form intrinsic supervision hierarchies as well as that communication between actors is transparent with respect to their placement across multiple network nodes.



The above image displays the relationship between the most important entities within an actor system, please read on for the details.

2.5.1 What is an Actor Reference?

An actor reference is a subtype of `ActorRef`, whose foremost purpose is to support sending messages to the actor it represents. Each actor has access to its canonical (local) reference through the `self` field; this reference is also included as sender reference by default for all messages sent to other actors. Conversely, during message processing the actor has access to a reference representing the sender of the current message through the `sender` method.

There are several different types of actor references that are supported depending on the configuration of the actor system:

- Purely local actor references are used by actor systems which are not configured to support networking functions. These actor references will not function if sent across a network connection to a remote JVM.

- Local actor references when remoting is enabled are used by actor systems which support networking functions for those references which represent actors within the same JVM. In order to also be reachable when sent to other network nodes, these references include protocol and remote addressing information.
- There is a subtype of local actor references which is used for routers (i.e. actors mixing in the `Router` trait). Its logical structure is the same as for the aforementioned local references, but sending a message to them dispatches to one of their children directly instead.
- Remote actor references represent actors which are reachable using remote communication, i.e. sending messages to them will serialize the messages transparently and send them to the remote JVM.
- There are several special types of actor references which behave like local actor references for all practical purposes:
 - `PromiseActorRef` is the special representation of a `Promise` for the purpose of being completed by the response from an actor. `akka.pattern.ask` creates this actor reference.
 - `DeadLetterActorRef` is the default implementation of the dead letters service to which Akka routes all messages whose destinations are shut down or non-existent.
 - `EmptyLocalActorRef` is what Akka returns when looking up a non-existent local actor path: it is equivalent to a `DeadLetterActorRef`, but it retains its path so that Akka can send it over the network and compare it to other existing actor references for that path, some of which might have been obtained before the actor died.
- And then there are some one-off internal implementations which you should never really see:
 - There is an actor reference which does not represent an actor but acts only as a pseudo-supervisor for the root guardian, we call it “the one who walks the bubbles of space-time”.
 - The first logging service started before actually firing up actor creation facilities is a fake actor reference which accepts log events and prints them directly to standard output; it is `Logging.StandardOutLogger`.

2.5.2 What is an Actor Path?

Since actors are created in a strictly hierarchical fashion, there exists a unique sequence of actor names given by recursively following the supervision links between child and parent down towards the root of the actor system. This sequence can be seen as enclosing folders in a file system, hence we adopted the name “path” to refer to it. As in some real file-systems there also are “symbolic links”, i.e. one actor may be reachable using more than one path, where all but one involve some translation which decouples part of the path from the actor’s actual supervision ancestor line; these specialities are described in the sub-sections to follow.

An actor path consists of an anchor, which identifies the actor system, followed by the concatenation of the path elements, from root guardian to the designated actor; the path elements are the names of the traversed actors and are separated by slashes.

What is the Difference Between Actor Reference and Path?

An actor reference designates a single actor and the life-cycle of the reference matches that actor’s life-cycle; an actor path represents a name which may or may not be inhabited by an actor and the path itself does not have a life-cycle, it never becomes invalid. You can create an actor path without creating an actor, but you cannot create an actor reference without creating corresponding actor.

You can create an actor, terminate it, and then create a new actor with the same actor path. The newly created actor is a new incarnation of the actor. It is not the same actor. An actor reference to the old incarnation is not valid for the new incarnation. Messages sent to the old actor reference will not be delivered to the new incarnation even though they have the same path.

Actor Path Anchors

Each actor path has an address component, describing the protocol and location by which the corresponding actor is reachable, followed by the names of the actors in the hierarchy from the root up. Examples are:

```
"akka://my-sys/user/service-a/worker1" // purely local
"akka.tcp://my-sys@host.example.com:5678/user/service-b" // remote
```

Here, `akka.tcp` is the default remote transport for the 2.2 release; other transports are pluggable. A remote host using UDP would be accessible by using `akka.udp`. The interpretation of the host and port part (i.e. `host.example.com:5678` in the example) depends on the transport mechanism used, but it must abide by the URI structural rules.

Logical Actor Paths

The unique path obtained by following the parental supervision links towards the root guardian is called the logical actor path. This path matches exactly the creation ancestry of an actor, so it is completely deterministic as soon as the actor system's remoting configuration (and with it the address component of the path) is set.

Physical Actor Paths

While the logical actor path describes the functional location within one actor system, configuration-based remote deployment means that an actor may be created on a different network host than its parent, i.e. within a different actor system. In this case, following the actor path from the root guardian up entails traversing the network, which is a costly operation. Therefore, each actor also has a physical path, starting at the root guardian of the actor system where the actual actor object resides. Using this path as sender reference when querying other actors will let them reply directly to this actor, minimizing delays incurred by routing.

One important aspect is that a physical actor path never spans multiple actor systems or JVMs. This means that the logical path (supervision hierarchy) and the physical path (actor deployment) of an actor may diverge if one of its ancestors is remotely supervised.

2.5.3 How are Actor References obtained?

There are two general categories to how actor references may be obtained: by creating actors or by looking them up, where the latter functionality comes in the two flavours of creating actor references from concrete actor paths and querying the logical actor hierarchy.

Creating Actors

An actor system is typically started by creating actors beneath the guardian actor using the `ActorSystem.actorOf` method and then using `ActorContext.actorOf` from within the created actors to spawn the actor tree. These methods return a reference to the newly created actor. Each actor has direct access (through its `ActorContext`) to references for its parent, itself and its children. These references may be sent within messages to other actors, enabling those to reply directly.

Looking up Actors by Concrete Path

In addition, actor references may be looked up using the `ActorSystem.actorSelection` method. The selection can be used for communicating with said actor and the actor corresponding to the selection is looked up when delivering each message.

To acquire an `ActorRef` that is bound to the life-cycle of a specific actor you need to send a message, such as the built-in `Identify` message, to the actor and use the `sender()` reference of a reply from the actor.

Absolute vs. Relative Paths

In addition to `ActorSystem.actorSelection` there is also `ActorContext.actorSelection`, which is available inside any actor as `context.actorSelection`. This yields an actor selection much like its twin on `ActorSystem`, but instead of looking up the path starting from the root of the actor tree it starts out on the current actor. Path elements consisting of two dots ("`..`") may be used to access the parent actor. You can for example send a message to a specific sibling:

```
context.actorSelection("../brother") ! msg
```

Absolute paths may of course also be looked up on `context` in the usual way, i.e.

```
context.actorSelection("/user/serviceA") ! msg
```

will work as expected.

Querying the Logical Actor Hierarchy

Since the actor system forms a file-system like hierarchy, matching on paths is possible in the same way as supported by Unix shells: you may replace (parts of) path element names with wildcards (`«*»` and `«?»`) to formulate a selection which may match zero or more actual actors. Because the result is not a single actor reference, it has a different type `ActorSelection` and does not support the full set of operations an `ActorRef` does. Selections may be formulated using the `ActorSystem.actorSelection` and `ActorContext.actorSelection` methods and do support sending messages:

```
context.actorSelection("../*") ! msg
```

will send `msg` to all siblings including the current actor. As for references obtained using `actorSelection`, a traversal of the supervision hierarchy is done in order to perform the message send. As the exact set of actors which match a selection may change even while a message is making its way to the recipients, it is not possible to watch a selection for liveness changes. In order to do that, resolve the uncertainty by sending a request and gathering all answers, extracting the sender references, and then watch all discovered concrete actors. This scheme of resolving a selection may be improved upon in a future release.

Summary: `actorOf` vs. `actorSelection`

Note: What the above sections described in some detail can be summarized and memorized easily as follows:

- `actorOf` only ever creates a new actor, and it creates it as a direct child of the context on which this method is invoked (which may be any actor or actor system).
- `actorSelection` only ever looks up existing actors when messages are delivered, i.e. does not create actors, or verify existence of actors when the selection is created.

2.5.4 Actor Reference and Path Equality

Equality of `ActorRef` match the intention that an `ActorRef` corresponds to the target actor incarnation. Two actor references are compared equal when they have the same path and point to the same actor incarnation. A reference pointing to a terminated actor does not compare equal to a reference pointing to another (re-created) actor with the same path. Note that a restart of an actor caused by a failure still means that it is the same actor incarnation, i.e. a restart is not visible for the consumer of the `ActorRef`.

If you need to keep track of actor references in a collection and do not care about the exact actor incarnation you can use the `ActorPath` as key, because the identifier of the target actor is not taken into account when comparing actor paths.

2.5.5 Reusing Actor Paths


When an actor is terminated, its reference will point to the dead letter mailbox, DeathWatch will publish its final transition and in general it is not expected to come back to life again (since the actor life cycle does not allow this). While it is possible to create an actor at a later time with an identical path—simply due to it being impossible to enforce the opposite without keeping the set of all actors ever created available—this is not good practice: messages sent with `actorSelection` to an actor which “died” suddenly start to work again, but without any guarantee of ordering between this transition and any other event, hence the new inhabitant of the path may receive messages which were destined for the previous tenant.

It may be the right thing to do in very specific circumstances, but make sure to confine the handling of this precisely to the actor’s supervisor, because that is the only actor which can reliably detect proper deregistration of the name, before which creation of the new child will fail.

It may also be required during testing, when the test subject depends on being instantiated at a specific path. In that case it is best to mock its supervisor so that it will forward the `Terminated` message to the appropriate point in the test procedure, enabling the latter to await proper deregistration of the name.

2.5.6 The Interplay with Remote Deployment

When an actor creates a child, the actor system’s deployer will decide whether the new actor resides in the same JVM or on another node. In the second case, creation of the actor will be triggered via a network connection to happen in a different JVM and consequently within a different actor system. The remote system will place the new actor below a special path reserved for this purpose and the supervisor of the new actor will be a remote actor reference (representing that actor which triggered its creation). In this case, `context.parent` (the supervisor reference) and `context.path.parent` (the parent node in the actor’s path) do not represent the same actor. However, looking up the child’s name within the supervisor will find it on the remote node, preserving logical structure e.g. when sending to an unresolved actor reference.



logical actor path: akka.tcp://sys@A:2552/user/parent/child

physical actor path: akka.tcp://sys@B:2552/remote/sys@A:2552/user/parent/child

2.5.7 What is the Address part used for?

When sending an actor reference across the network, it is represented by its path. Hence, the path must fully encode all information necessary to send messages to the underlying actor. This is achieved by encoding protocol, host and port in the address part of the path string. When an actor system receives an actor path from a remote node, it checks whether that path's address matches the address of this actor system, in which case it will be resolved to the actor's local reference. Otherwise, it will be represented by a remote actor reference.

2.5.8 Top-Level Scopes for Actor Paths

At the root of the path hierarchy resides the root guardian above which all other actors are found; its name is `"/`. The next level consists of the following:

- `"/user"` is the guardian actor for all user-created top-level actors; actors created using `ActorSystem.actorOf` are found below this one.
- `"/system"` is the guardian actor for all system-created top-level actors, e.g. logging listeners or actors automatically deployed by configuration at the start of the actor system.
- `"/deadLetters"` is the dead letter actor, which is where all messages sent to stopped or non-existing actors are re-routed (on a best-effort basis: messages may be lost even within the local JVM).
- `"/temp"` is the guardian for all short-lived system-created actors, e.g. those which are used in the implementation of `ActorRef.ask`.
- `"/remote"` is an artificial path below which all actors reside whose supervisors are remote actor references

The need to structure the name space for actors like this arises from a central and very simple design goal: everything in the hierarchy is an actor, and all actors function in the same way. Hence you can not only look up the actors you created, you can also look up the system guardian and send it a message (which it will dutifully discard in this case). This powerful principle means that there are no quirks to remember, it makes the whole system more uniform and consistent.

If you want to read more about the top-level structure of an actor system, have a look at *The Top-Level Supervisors*.

2.6 Location Transparency

The previous section describes how actor paths are used to enable location transparency. This special feature deserves some extra explanation, because the related term “transparent remoting” was used quite differently in the context of programming languages, platforms and technologies.

2.6.1 Distributed by Default

Everything in Akka is designed to work in a distributed setting: all interactions of actors use purely message passing and everything is asynchronous. This effort has been undertaken to ensure that all functions are available equally when running within a single JVM or on a cluster of hundreds of machines. The key for enabling this is to go from remote to local by way of optimization instead of trying to go from local to remote by way of generalization. See [this classic paper](#) for a detailed discussion on why the second approach is bound to fail.

2.6.2 Ways in which Transparency is Broken

What is true of Akka need not be true of the application which uses it, since designing for distributed execution poses some restrictions on what is possible. The most obvious one is that all messages sent over the wire must be serializable. While being a little less obvious this includes closures which are used as actor factories (i.e. within `Props`) if the actor is to be created on a remote node.

Another consequence is that everything needs to be aware of all interactions being fully asynchronous, which in a computer network might mean that it may take several minutes for a message to reach its recipient (depending

on configuration). It also means that the probability for a message to be lost is much higher than within one JVM, where it is close to zero (still: no hard guarantee!).

2.6.3 How is Remoting Used?

We took the idea of transparency to the limit in that there is nearly no API for the remoting layer of Akka: it is purely driven by configuration. Just write your application according to the principles outlined in the previous sections, then specify remote deployment of actor sub-trees in the configuration file. This way, your application can be scaled out without having to touch the code. The only piece of the API which allows programmatic influence on remote deployment is that `Props` contain a field which may be set to a specific `Deploy` instance; this has the same effect as putting an equivalent deployment into the configuration file (if both are given, configuration file wins).

2.6.4 Peer-to-Peer vs. Client-Server

Akka Remoting is a communication module for connecting actor systems in a peer-to-peer fashion, and it is the foundation for Akka Clustering. The design of remoting is driven by two (related) design decisions:

1. Communication between involved systems is symmetric: if a system A can connect to a system B then system B must also be able to connect to system A independently.
2. The role of the communicating systems are symmetric in regards to connection patterns: there is no system that only accepts connections, and there is no system that only initiates connections.

The consequence of these decisions is that it is not possible to safely create pure client-server setups with predefined roles (violates assumption 2). For client-server setups it is better to use HTTP or Akka I/O.

Important: Using setups involving Network Address Translation, Load Balancers or Docker containers violates assumption 1, unless additional steps are taken in the network configuration to allow symmetric communication between involved systems. In such situations Akka can be configured to bind to a different network address than the one used for establishing connections between Akka nodes. See *remote-configuration-nat*.

2.6.5 Marking Points for Scaling Up with Routers

In addition to being able to run different parts of an actor system on different nodes of a cluster, it is also possible to scale up onto more cores by multiplying actor sub-trees which support parallelization (think for example a search engine processing different queries in parallel). The clones can then be routed to in different fashions, e.g. round-robin. The only thing necessary to achieve this is that the developer needs to declare a certain actor as “withRouter”, then—in its stead—a router actor will be created which will spawn up a configurable number of children of the desired type and route to them in the configured fashion. Once such a router has been declared, its configuration can be freely overridden from the configuration file, including mixing it with the remote deployment of (some of) the children. Read more about this in *Routing (Scala)* and *Routing (Java)*.

2.7 Akka and the Java Memory Model

A major benefit of using the Typesafe Platform, including Scala and Akka, is that it simplifies the process of writing concurrent software. This article discusses how the Typesafe Platform, and Akka in particular, approaches shared memory in concurrent applications.

2.7.1 The Java Memory Model

Prior to Java 5, the Java Memory Model (JMM) was ill defined. It was possible to get all kinds of strange results when shared memory was accessed by multiple threads, such as:

- a thread not seeing values written by other threads: a visibility problem

- a thread observing ‘impossible’ behavior of other threads, caused by instructions not being executed in the order expected: an instruction reordering problem.

With the implementation of JSR 133 in Java 5, a lot of these issues have been resolved. The JMM is a set of rules based on the “happens-before” relation, which constrain when one memory access must happen before another, and conversely, when they are allowed to happen out of order. Two examples of these rules are:

- **The monitor lock rule:** a release of a lock happens before every subsequent acquire of the same lock.
- **The volatile variable rule:** a write of a volatile variable happens before every subsequent read of the same volatile variable

Although the JMM can seem complicated, the specification tries to find a balance between ease of use and the ability to write performant and scalable concurrent data structures.

2.7.2 Actors and the Java Memory Model

With the Actors implementation in Akka, there are two ways multiple threads can execute actions on shared memory:

- if a message is sent to an actor (e.g. by another actor). In most cases messages are immutable, but if that message is not a properly constructed immutable object, without a “happens before” rule, it would be possible for the receiver to see partially initialized data structures and possibly even values out of thin air (longs/doubles).
- if an actor makes changes to its internal state while processing a message, and accesses that state while processing another message moments later. It is important to realize that with the actor model you don’t get any guarantee that the same thread will be executing the same actor for different messages.

To prevent visibility and reordering problems on actors, Akka guarantees the following two “happens before” rules:

- **The actor send rule:** the send of the message to an actor happens before the receive of that message by the same actor.
- **The actor subsequent processing rule:** processing of one message happens before processing of the next message by the same actor.

Note: In layman’s terms this means that changes to internal fields of the actor are visible when the next message is processed by that actor. So fields in your actor need not be volatile or equivalent.

Both rules only apply for the same actor instance and are not valid if different actors are used.

2.7.3 Futures and the Java Memory Model

The completion of a Future “happens before” the invocation of any callbacks registered to it are executed.

We recommend not to close over non-final fields (final in Java and val in Scala), and if you *do* choose to close over non-final fields, they must be marked *volatile* in order for the current value of the field to be visible to the callback.

If you close over a reference, you must also ensure that the instance that is referred to is thread safe. We highly recommend staying away from objects that use locking, since it can introduce performance problems and in the worst case, deadlocks. Such are the perils of synchronized.

2.7.4 STM and the Java Memory Model

Akka’s Software Transactional Memory (STM) also provides a “happens before” rule:

- **The transactional reference rule:** a successful write during commit, on a transactional reference, happens before every subsequent read of the same transactional reference.

This rule looks a lot like the ‘volatile variable’ rule from the JMM. Currently the Akka STM only supports deferred writes, so the actual writing to shared memory is deferred until the transaction commits. Writes during the transaction are placed in a local buffer (the writeset of the transaction) and are not visible to other transactions. That is why dirty reads are not possible.

How these rules are realized in Akka is an implementation detail and can change over time, and the exact details could even depend on the used configuration. But they will build on the other JMM rules like the monitor lock rule or the volatile variable rule. This means that you, the Akka user, do not need to worry about adding synchronization to provide such a “happens before” relation, because it is the responsibility of Akka. So you have your hands free to deal with your business logic, and the Akka framework makes sure that those rules are guaranteed on your behalf.

2.7.5 Actors and shared mutable state

Since Akka runs on the JVM there are still some rules to be followed.

- Closing over internal Actor state and exposing it to other threads

```
class MyActor extends Actor {
  var state = ...
  def receive = {
    case _ =>
      //Wrongs

    // Very bad, shared mutable state,
    // will break your application in weird ways
    Future { state = NewState }
    anotherActor ? message onSuccess { r => state = r }

    // Very bad, "sender" changes for every message,
    // shared mutable state bug
    Future { expensiveCalculation(sender()) }

    //Rights

    // Completely safe, "self" is OK to close over
    // and it's an ActorRef, which is thread-safe
    Future { expensiveCalculation() } onComplete { f => self ! f.value.get }

    // Completely safe, we close over a fixed value
    // and it's an ActorRef, which is thread-safe
    val currentSender = sender()
    Future { expensiveCalculation(currentSender) }
  }
}
```

- Messages **should** be immutable, this is to avoid the shared mutable state trap.

2.8 Message Delivery Reliability

Akka helps you build reliable applications which make use of multiple processor cores in one machine (“scaling up”) or distributed across a computer network (“scaling out”). The key abstraction to make this work is that all interactions between your code units—actors—happen via message passing, which is why the precise semantics of how messages are passed between actors deserve their own chapter.

In order to give some context to the discussion below, consider an application which spans multiple network hosts. The basic mechanism for communication is the same whether sending to an actor on the local JVM or to a remote actor, but of course there will be observable differences in the latency of delivery (possibly also depending on the bandwidth of the network link and the message size) and the reliability. In case of a remote message send there are obviously more steps involved which means that more can go wrong. Another aspect is that local sending will

just pass a reference to the message inside the same JVM, without any restrictions on the underlying object which is sent, whereas a remote transport will place a limit on the message size.

Writing your actors such that every interaction could possibly be remote is the safe, pessimistic bet. It means to only rely on those properties which are always guaranteed and which are discussed in detail below. This has of course some overhead in the actor's implementation. If you are willing to sacrifice full location transparency—for example in case of a group of closely collaborating actors—you can place them always on the same JVM and enjoy stricter guarantees on message delivery. The details of this trade-off are discussed further below.

As a supplementary part we give a few pointers at how to build stronger reliability on top of the built-in ones. The chapter closes by discussing the role of the “Dead Letter Office”.

2.8.1 The General Rules

These are the rules for message sends (i.e. the `tell` or `!` method, which also underlies the `ask` pattern):

- **at-most-once delivery**, i.e. no guaranteed delivery
- **message ordering per sender–receiver pair**

The first rule is typically found also in other actor implementations while the second is specific to Akka.

Discussion: What does “at-most-once” mean?

When it comes to describing the semantics of a delivery mechanism, there are three basic categories:

- **at-most-once** delivery means that for each message handed to the mechanism, that message is delivered zero or one times; in more casual terms it means that messages may be lost.
- **at-least-once** delivery means that for each message handed to the mechanism potentially multiple attempts are made at delivering it, such that at least one succeeds; again, in more casual terms this means that messages may be duplicated but not lost.
- **exactly-once** delivery means that for each message handed to the mechanism exactly one delivery is made to the recipient; the message can neither be lost nor duplicated.

The first one is the cheapest—highest performance, least implementation overhead—because it can be done in a fire-and-forget fashion without keeping state at the sending end or in the transport mechanism. The second one requires retries to counter transport losses, which means keeping state at the sending end and having an acknowledgement mechanism at the receiving end. The third is most expensive—and has consequently worst performance—because in addition to the second it requires state to be kept at the receiving end in order to filter out duplicate deliveries.

Discussion: Why No Guaranteed Delivery?

At the core of the problem lies the question what exactly this guarantee shall mean:

1. The message is sent out on the network?
2. The message is received by the other host?
3. The message is put into the target actor's mailbox?
4. The message is starting to be processed by the target actor?
5. The message is processed successfully by the target actor?

Each one of these have different challenges and costs, and it is obvious that there are conditions under which any message passing library would be unable to comply; think for example about configurable mailbox types and how a bounded mailbox would interact with the third point, or even what it would mean to decide upon the “successfully” part of point five.

Along those same lines goes the reasoning in [Nobody Needs Reliable Messaging](#). The only meaningful way for a sender to know whether an interaction was successful is by receiving a business-level acknowledgement message,

which is not something Akka could make up on its own (neither are we writing a “do what I mean” framework nor would you want us to).

Akka embraces distributed computing and makes the fallibility of communication explicit through message passing, therefore it does not try to lie and emulate a leaky abstraction. This is a model that has been used with great success in Erlang and requires the users to design their applications around it. You can read more about this approach in the [Erlang documentation](#) (section 10.9 and 10.10), Akka follows it closely.

Another angle on this issue is that by providing only basic guarantees those use cases which do not need stronger reliability do not pay the cost of their implementation; it is always possible to add stronger reliability on top of basic ones, but it is not possible to retro-actively remove reliability in order to gain more performance.

Discussion: Message Ordering

The rule more specifically is that *for a given pair of actors, messages sent directly from the first to the second will not be received out-of-order*. The word *directly* emphasizes that this guarantee only applies when sending with the *tell* operator to the final destination, not when employing mediators or other message dissemination features (unless stated otherwise).

The guarantee is illustrated in the following:

Actor A1 sends messages M1, M2, M3 to A2

Actor A3 sends messages M4, M5, M6 to A2

This means that:

1. If M1 is delivered it must be delivered before M2 and M3
2. If M2 is delivered it must be delivered before M3
3. If M4 is delivered it must be delivered before M5 and M6
4. If M5 is delivered it must be delivered before M6
5. A2 can see messages from A1 interleaved with messages from A3
6. Since there is no guaranteed delivery, any of the messages may be dropped, i.e. not arrive at A2

Note: It is important to note that Akka’s guarantee applies to the order in which messages are enqueued into the recipient’s mailbox. If the mailbox implementation does not respect FIFO order (e.g. a `PriorityMailbox`), then the order of processing by the actor can deviate from the enqueueing order.

Please note that this rule is **not transitive**:

Actor A sends message M1 to actor C

Actor A then sends message M2 to actor B

Actor B forwards message M2 to actor C

Actor C may receive M1 and M2 in any order

Causal transitive ordering would imply that M2 is never received before M1 at actor C (though any of them might be lost). This ordering can be violated due to different message delivery latencies when A, B and C reside on different network hosts, see more below.

Note: Actor creation is treated as a message sent from the parent to the child, with the same semantics as discussed above. Sending a message to an actor in a way which could be reordered with this initial creation message means that the message might not arrive because the actor does not exist yet. An example where the message might arrive too early would be to create a remote-deployed actor R1, send its reference to another remote actor R2 and have R2 send a message to R1. An example of well-defined ordering is a parent which creates an actor and immediately sends a message to it.

Communication of failure

Please note, that the ordering guarantees discussed above only hold for user messages between actors. Failure of a child of an actor is communicated by special system messages that are not ordered relative to ordinary user messages. In particular:

Child actor C sends message M to its parent P

Child actor fails with failure F

Parent actor P might receive the two events either in order M, F or F, M

The reason for this is that internal system messages has their own mailboxes therefore the ordering of enqueue calls of a user and system message cannot guarantee the ordering of their dequeue times.

2.8.2 The Rules for In-JVM (Local) Message Sends

Be careful what you do with this section!

Relying on the stronger reliability in this section is not recommended since it will bind your application to local-only deployment: an application may have to be designed differently (as opposed to just employing some message exchange patterns local to some actors) in order to be fit for running on a cluster of machines. Our credo is “design once, deploy any way you wish”, and to achieve this you should only rely on [The General Rules](#).

Reliability of Local Message Sends

The Akka test suite relies on not losing messages in the local context (and for non-error condition tests also for remote deployment), meaning that we actually do apply the best effort to keep our tests stable. A local `tell` operation can however fail for the same reasons as a normal method call can on the JVM:

- `StackOverflowError`
- `OutOfMemoryError`
- other `VirtualMachineError`

In addition, local sends can fail in Akka-specific ways:

- if the mailbox does not accept the message (e.g. full `BoundedMailbox`)
- if the receiving actor fails while processing the message or is already terminated

While the first is clearly a matter of configuration the second deserves some thought: the sender of a message does not get feedback if there was an exception while processing, that notification goes to the supervisor instead. This is in general not distinguishable from a lost message for an outside observer.

Ordering of Local Message Sends

Assuming strict FIFO mailboxes the aforementioned caveat of non-transitivity of the message ordering guarantee is eliminated under certain conditions. As you will note, these are quite subtle as it stands, and it is even possible that future performance optimizations will invalidate this whole paragraph. The possibly non-exhaustive list of counter-indications is:

- Before receiving the first reply from a top-level actor, there is a lock which protects an internal interim queue, and this lock is not fair; the implication is that enqueue requests from different senders which arrive during the actor’s construction (figuratively, the details are more involved) may be reordered depending on low-level thread scheduling. Since completely fair locks do not exist on the JVM this is unfixable.
- The same mechanism is used during the construction of a `Router`, more precisely the routed `ActorRef`, hence the same problem exists for actors deployed with `Routers`.

- As mentioned above, the problem occurs anywhere a lock is involved during enqueueing, which may also apply to custom mailboxes.

This list has been compiled carefully, but other problematic scenarios may have escaped our analysis.

How does Local Ordering relate to Network Ordering

The rule that *for a given pair of actors, messages sent directly from the first to the second will not be received out-of-order* holds for messages sent over the network with the TCP based Akka remote transport protocol.

As explained in the previous section local message sends obey transitive causal ordering under certain conditions. This ordering can be violated due to different message delivery latencies. For example:

Actor A on node-1 sends message M1 to actor C on node-3

Actor A on node-1 then sends message M2 to actor B on node-2

Actor B on node-2 forwards message M2 to actor C on node-3

Actor C may receive M1 and M2 in any order

It might take longer time for M1 to “travel” to node-3 than it takes for M2 to “travel” to node-3 via node-2.

2.8.3 Higher-level abstractions

Based on a small and consistent tool set in Akka’s core, Akka also provides powerful, higher-level abstractions on top it.

Messaging Patterns

As discussed above a straight-forward answer to the requirement of reliable delivery is an explicit ACK-RETRY protocol. In its simplest form this requires

- a way to identify individual messages to correlate message with acknowledgement
- a retry mechanism which will resend messages if not acknowledged in time
- a way for the receiver to detect and discard duplicates

The third becomes necessary by virtue of the acknowledgements not being guaranteed to arrive either. An ACK-RETRY protocol with business-level acknowledgements is supported by *at-least-once-delivery-scala* of the Akka Persistence module. Duplicates can be detected by tracking the identifiers of messages sent via *at-least-once-delivery-scala*. Another way of implementing the third part would be to make processing the messages idempotent on the level of the business logic.

Another example of implementing all three requirements is shown at *Reliable Proxy Pattern* (which is now superseded by *at-least-once-delivery-scala*).

Event Sourcing

Event sourcing (and sharding) is what makes large websites scale to billions of users, and the idea is quite simple: when a component (think actor) processes a command it will generate a list of events representing the effect of the command. These events are stored in addition to being applied to the component’s state. The nice thing about this scheme is that events only ever are appended to the storage, nothing is ever mutated; this enables perfect replication and scaling of consumers of this event stream (i.e. other components may consume the event stream as a means to replicate the component’s state on a different continent or to react to changes). If the component’s state is lost—due to a machine failure or by being pushed out of a cache—it can easily be reconstructed by replaying the event stream (usually employing snapshots to speed up the process). *event-sourcing-scala* is supported by Akka Persistence.

Mailbox with Explicit Acknowledgement

By implementing a custom mailbox type it is possible to retry message processing at the receiving actor's end in order to handle temporary failures. This pattern is mostly useful in the local communication context where delivery guarantees are otherwise sufficient to fulfill the application's requirements.

Please note that the caveats for [The Rules for In-JVM \(Local\) Message Sends](#) do apply.

An example implementation of this pattern is shown at [Mailbox with Explicit Acknowledgement](#).

2.8.4 Dead Letters

Messages which cannot be delivered (and for which this can be ascertained) will be delivered to a synthetic actor called `/deadLetters`. This delivery happens on a best-effort basis; it may fail even within the local JVM (e.g. during actor termination). Messages sent via unreliable network transports will be lost without turning up as dead letters.

What Should I Use Dead Letters For?

The main use of this facility is for debugging, especially if an actor send does not arrive consistently (where usually inspecting the dead letters will tell you that the sender or recipient was set wrong somewhere along the way). In order to be useful for this purpose it is good practice to avoid sending to `deadLetters` where possible, i.e. run your application with a suitable dead letter logger (see more below) from time to time and clean up the log output. This exercise—like all else—requires judicious application of common sense: it may well be that avoiding to send to a terminated actor complicates the sender's code more than is gained in debug output clarity.

The dead letter service follows the same rules with respect to delivery guarantees as all other message sends, hence it cannot be used to implement guaranteed delivery.

How do I Receive Dead Letters?

An actor can subscribe to class `akka.actor.DeadLetter` on the event stream, see [Event Stream](#) (Java) or [event-stream-scala](#) (Scala) for how to do that. The subscribed actor will then receive all dead letters published in the (local) system from that point onwards. Dead letters are not propagated over the network, if you want to collect them in one place you will have to subscribe one actor per network node and forward them manually. Also consider that dead letters are generated at that node which can determine that a send operation is failed, which for a remote send can be the local system (if no network connection can be established) or the remote one (if the actor you are sending to does not exist at that point in time).

Dead Letters Which are (Usually) not Worrisome

Every time an actor does not terminate by its own decision, there is a chance that some messages which it sends to itself are lost. There is one which happens quite easily in complex shutdown scenarios that is usually benign: seeing a `akka.dispatch.Terminate` message dropped means that two termination requests were given, but of course only one can succeed. In the same vein, you might see `akka.actor.Terminated` messages from children while stopping a hierarchy of actors turning up in dead letters if the parent is still watching the child when the parent terminates.

2.9 Configuration

You can start using Akka without defining any configuration, since sensible default values are provided. Later on you might need to amend the settings to change the default behavior or adapt for specific runtime environments. Typical examples of settings that you might amend:

- log level and logger backend

- enable remoting
- message serializers
- definition of routers
- tuning of dispatchers

Akka uses the [Typesafe Config Library](#), which might also be a good choice for the configuration of your own application or library built with or without Akka. This library is implemented in Java with no external dependencies; you should have a look at its documentation (in particular about [ConfigFactory](#)), which is only summarized in the following.

Warning: If you use Akka from the Scala REPL from the 2.9.x series, and you do not provide your own `ClassLoader` to the `ActorSystem`, start the REPL with “-Yrepl-sync” to work around a deficiency in the REPLs provided `Context ClassLoader`.

2.9.1 Where configuration is read from

All configuration for Akka is held within instances of `ActorSystem`, or put differently, as viewed from the outside, `ActorSystem` is the only consumer of configuration information. While constructing an actor system, you can either pass in a `Config` object or not, where the second case is equivalent to passing `ConfigFactory.load()` (with the right class loader). This means roughly that the default is to parse all `application.conf`, `application.json` and `application.properties` found at the root of the class path—please refer to the aforementioned documentation for details. The actor system then merges in all `reference.conf` resources found at the root of the class path to form the fallback configuration, i.e. it internally uses

```
appConfig.withFallback(ConfigFactory.defaultReference(classLoader))
```

The philosophy is that code never contains default values, but instead relies upon their presence in the `reference.conf` supplied with the library in question.

Highest precedence is given to overrides given as system properties, see [the HOCON specification](#) (near the bottom). Also noteworthy is that the application configuration—which defaults to `application`—may be overridden using the `config.resource` property (there are more, please refer to the [Config docs](#)).

Note: If you are writing an Akka application, keep your configuration in `application.conf` at the root of the class path. If you are writing an Akka-based library, keep its configuration in `reference.conf` at the root of the JAR file.

2.9.2 When using JarJar, OneJar, Assembly or any jar-bundler

Warning: Akka’s configuration approach relies heavily on the notion of every module/jar having its own `reference.conf` file, all of these will be discovered by the configuration and loaded. Unfortunately this also means that if you put/merge multiple jars into the same jar, you need to merge all the `reference.conf`s as well. Otherwise all defaults will be lost and Akka will not function.

If you are using Maven to package your application, you can also make use of the [Apache Maven Shade Plugin](#) support for [Resource Transformers](#) to merge all the `reference.conf`s on the build classpath into one.

The plugin configuration might look like this:

```
<plugin>
<groupId>org.apache.maven.plugins</groupId>
<artifactId>maven-shade-plugin</artifactId>
<version>1.5</version>
<executions>
```

```

<execution>
  <phase>package</phase>
  <goals>
    <goal>shade</goal>
  </goals>
  <configuration>
    <shadedArtifactAttached>true</shadedArtifactAttached>
    <shadedClassifierName>allinone</shadedClassifierName>
    <artifactSet>
      <includes>
        <include>*:*</include>
      </includes>
    </artifactSet>
    <transformers>
      <transformer
        implementation="org.apache.maven.plugins.shade.resource.AppendingTransformer">
        <resource>reference.conf</resource>
      </transformer>
      <transformer
        implementation="org.apache.maven.plugins.shade.resource.ManifestResourceTransformer">
        <manifestEntries>
          <Main-Class>akka.Main</Main-Class>
        </manifestEntries>
      </transformer>
    </transformers>
  </configuration>
</execution>
</executions>
</plugin>

```

2.9.3 Custom application.conf

A custom `application.conf` might look like this:

```

# In this file you can override any option defined in the reference files.
# Copy in parts of the reference files and modify as you please.

akka {

  # Loggers to register at boot time (akka.event.Logging$DefaultLogger logs
  # to STDOUT)
  loggers = ["akka.event.slf4j.Slf4jLogger"]

  # Log level used by the configured loggers (see "loggers") as soon
  # as they have been started; before that, see "stdout-loglevel"
  # Options: OFF, ERROR, WARNING, INFO, DEBUG
  loglevel = "DEBUG"

  # Log level for the very basic logger activated during ActorSystem startup.
  # This logger prints the log messages to stdout (System.out).
  # Options: OFF, ERROR, WARNING, INFO, DEBUG
  stdout-loglevel = "DEBUG"

  # Filter of log events that is used by the LoggingAdapter before
  # publishing log events to the eventStream.
  logging-filter = "akka.event.slf4j.Slf4jLoggingFilter"

  actor {
    provider = "akka.cluster.ClusterActorRefProvider"

    default-dispatcher {

```

```

    # Throughput for default Dispatcher, set to 1 for as fair as possible
    throughput = 10
  }
}

remote {
  # The port clients should connect to. Default is 2552.
  netty.tcp.port = 4711
}
}

```

2.9.4 Including files

Sometimes it can be useful to include another configuration file, for example if you have one `application.conf` with all environment independent settings and then override some settings for specific environments.

Specifying system property with `-Dconfig.resource=/dev.conf` will load the `dev.conf` file, which includes the `application.conf`

`dev.conf`:

```

include "application"

akka {
  loglevel = "DEBUG"
}

```

More advanced include and substitution mechanisms are explained in the [HOCON](#) specification.

2.9.5 Logging of Configuration

If the system or config property `akka.log-config-on-start` is set to `on`, then the complete configuration at INFO level when the actor system is started. This is useful when you are uncertain of what configuration is used.

If in doubt, you can also easily and nicely inspect configuration objects before or after using them to construct an actor system:

```

Welcome to Scala version 2.11.7 (Java HotSpot(TM) 64-Bit Server VM, Java 1.8.0).
Type in expressions to have them evaluated.
Type :help for more information.

scala> import com.typesafe.config._
import com.typesafe.config._

scala> ConfigFactory.parseString("a.b=12")
res0: com.typesafe.config.Config = Config(SimpleConfigObject({"a" : {"b" : 12}}))

scala> res0.root.render
res1: java.lang.String =
{
  # String: 1
  "a" : {
    # String: 1
    "b" : 12
  }
}

```

The comments preceding every item give detailed information about the origin of the setting (file & line number) plus possible comments which were present, e.g. in the reference configuration. The settings as merged with the

reference and parsed by the actor system can be displayed like this:

```
final ActorSystem system = ActorSystem.create();
System.out.println(system.settings());
// this is a shortcut for system.settings().config().root().render()
```

2.9.6 A Word About ClassLoaders

In several places of the configuration file it is possible to specify the fully-qualified class name of something to be instantiated by Akka. This is done using Java reflection, which in turn uses a `ClassLoader`. Getting the right one in challenging environments like application containers or OSGi bundles is not always trivial, the current approach of Akka is that each `ActorSystem` implementation stores the current thread's context class loader (if available, otherwise just its own loader as in `this.getClass.getClassLoader`) and uses that for all reflective accesses. This implies that putting Akka on the boot class path will yield `NullPointerException` from strange places: this is simply not supported.

2.9.7 Application specific settings

The configuration can also be used for application specific settings. A good practice is to place those settings in an `Extension`, as described in:

- Scala API: *extending-akka-scala.settings*
- Java API: *Application specific settings*

2.9.8 Configuring multiple ActorSystem

If you have more than one `ActorSystem` (or you're writing a library and have an `ActorSystem` that may be separate from the application's) you may want to separate the configuration for each system.

Given that `ConfigFactory.load()` merges all resources with matching name from the whole class path, it is easiest to utilize that functionality and differentiate actor systems within the hierarchy of the configuration:

```
myapp1 {
  akka.loglevel = "WARNING"
  my.own.setting = 43
}
myapp2 {
  akka.loglevel = "ERROR"
  app2.setting = "appname"
}
my.own.setting = 42
my.other.setting = "hello"
```

```
val config = ConfigFactory.load()
val app1 = ActorSystem("MyApp1", config.getConfig("myapp1").withFallback(config))
val app2 = ActorSystem("MyApp2",
  config.getConfig("myapp2").withOnlyPath("akka").withFallback(config))
```

These two samples demonstrate different variations of the “lift-a-subtree” trick: in the first case, the configuration accessible from within the actor system is this

```
akka.loglevel = "WARNING"
my.own.setting = 43
my.other.setting = "hello"
// plus myapp1 and myapp2 subtrees
```

while in the second one, only the “akka” subtree is lifted, with the following result

```
akka.loglevel = "ERROR"
my.own.setting = 42
my.other.setting = "hello"
// plus myapp1 and myapp2 subtrees
```

Note: The configuration library is really powerful, explaining all features exceeds the scope affordable here. In particular not covered are how to include other configuration files within other files (see a small example at [Including files](#)) and copying parts of the configuration tree by way of path substitutions.

You may also specify and parse the configuration programmatically in other ways when instantiating the ActorSystem.

```
import akka.actor.ActorSystem
import com.typesafe.config.ConfigFactory
val customConf = ConfigFactory.parseString("""
  akka.actor.deployment {
    /my-service {
      router = round-robin-pool
      nr-of-instances = 3
    }
  }
  """)
// ConfigFactory.load sandwiches customConf between default reference
// config and default overrides, and then resolves it.
val system = ActorSystem("MySystem", ConfigFactory.load(customConf))
```

2.9.9 Reading configuration from a custom location

You can replace or supplement `application.conf` either in code or using system properties.

If you're using `ConfigFactory.load()` (which Akka does by default) you can replace `application.conf` by defining `-Dconfig.resource=whatever`, `-Dconfig.file=whatever`, or `-Dconfig.url=whatever`.

From inside your replacement file specified with `-Dconfig.resource` and friends, you can include "application" if you still want to use `application.{conf,json,properties}` as well. Settings specified before include "application" would be overridden by the included file, while those after would override the included file.

In code, there are many customization options.

There are several overloads of `ConfigFactory.load()`; these allow you to specify something to be sandwiched between system properties (which override) and the defaults (from `reference.conf`), replacing the usual `application.{conf,json,properties}` and replacing `-Dconfig.file` and friends.

The simplest variant of `ConfigFactory.load()` takes a resource basename (instead of application); `myname.conf`, `myname.json`, and `myname.properties` would then be used instead of `application.{conf,json,properties}`.

The most flexible variant takes a `Config` object, which you can load using any method in `ConfigFactory`. For example you could put a config string in code using `ConfigFactory.parseString()` or you could make a map and `ConfigFactory.parseMap()`, or you could load a file.

You can also combine your custom config with the usual config, that might look like:

```
// make a Config with just your special setting
Config myConfig =
  ConfigFactory.parseString("something=somethingElse");
// load the normal config stack (system props,
// then application.conf, then reference.conf)
Config regularConfig =
```

```

    ConfigFactory.load();
// override regular stack with myConfig
Config combined =
    myConfig.withFallback(regularConfig);
// put the result in between the overrides
// (system props) and defaults again
Config complete =
    ConfigFactory.load(combined);
// create ActorSystem
ActorSystem system =
    ActorSystem.create("myname", complete);

```

When working with `Config` objects, keep in mind that there are three “layers” in the cake:

- `ConfigFactory.defaultOverrides()` (system properties)
- the app’s settings
- `ConfigFactory.defaultReference()` (`reference.conf`)

The normal goal is to customize the middle layer while leaving the other two alone.

- `ConfigFactory.load()` loads the whole stack
- the overloads of `ConfigFactory.load()` let you specify a different middle layer
- the `ConfigFactory.parse()` variations load single files or resources

To stack two layers, use `override.withFallback(fallback)`; try to keep `system props` (`defaultOverrides()`) on top and `reference.conf` (`defaultReference()`) on the bottom.

Do keep in mind, you can often just add another `include` statement in `application.conf` rather than writing code. Includes at the top of `application.conf` will be overridden by the rest of `application.conf`, while those at the bottom will override the earlier stuff.

2.9.10 Actor Deployment Configuration

Deployment settings for specific actors can be defined in the `akka.actor.deployment` section of the configuration. In the deployment section it is possible to define things like dispatcher, mailbox, router settings, and remote deployment. Configuration of these features are described in the chapters detailing corresponding topics. An example may look like this:

```

akka.actor.deployment {

  # '/user/actorA/actorB' is a remote deployed actor
  /actorA/actorB {
    remote = "akka.tcp://sampleActorSystem@127.0.0.1:2553"
  }

  # all direct children of '/user/actorC' have a dedicated dispatcher
  "/actorC/*" {
    dispatcher = my-dispatcher
  }

  # '/user/actorD/actorE' has a special priority mailbox
  /actorD/actorE {
    mailbox = prio-mailbox
  }

  # '/user/actorF/actorG/actorH' is a random pool
  /actorF/actorG/actorH {
    router = random-pool
    nr-of-instances = 5
  }
}

```

```

}

my-dispatcher {
  fork-join-executor.parallelism-min = 10
  fork-join-executor.parallelism-max = 10
}
prio-mailbox {
  mailbox-type = "a.b.MyPrioMailbox"
}

```

Note: The deployment section for a specific actor is identified by the path of the actor relative to `/user`.

You can use asterisks as wildcard matches for the actor path sections, so you could specify: `/*/sampleActor` and that would match all `sampleActor` on that level in the hierarchy. You can also use wildcard in the last position to match all actors at a certain level: `/someParent/*`. Non-wildcard matches always have higher priority to match than wildcards, so: `/foo/bar` is considered **more specific** than `/foo/*` and only the highest priority match is used. Please note that it **cannot** be used to partially match section, like this: `/foo*/bar`, `/f*/o/bar` etc.

2.9.11 Listing of the Reference Configuration

Each Akka module has a reference configuration file with the default values.

akka-actor

```

#####
# Akka Actor Reference Config File #
#####

# This is the reference config file that contains all the default settings.
# Make your edits/overrides in your application.conf.

akka {
  # Akka version, checked against the runtime version of Akka.
  version = "2.4.0"

  # Home directory of Akka, modules in the deploy directory will be loaded
  home = ""

  # Loggers to register at boot time (akka.event.Logging$DefaultLogger logs
  # to STDOUT)
  loggers = ["akka.event.Logging$DefaultLogger"]

  # Filter of log events that is used by the LoggingAdapter before
  # publishing log events to the eventStream. It can perform
  # fine grained filtering based on the log source. The default
  # implementation filters on the `loglevel`.
  # FQCN of the LoggingFilter. The Class of the FQCN must implement
  # akka.event.LoggingFilter and have a public constructor with
  # (akka.actor.ActorSystem.Settings, akka.event.EventStream) parameters.
  logging-filter = "akka.event.DefaultLoggingFilter"

  # Loggers are created and registered synchronously during ActorSystem
  # start-up, and since they are actors, this timeout is used to bound the
  # waiting time
  logger-startup-timeout = 5s

  # Log level used by the configured loggers (see "loggers") as soon

```

```

# as they have been started; before that, see "stdout-loglevel"
# Options: OFF, ERROR, WARNING, INFO, DEBUG
loglevel = "INFO"

# Log level for the very basic logger activated during ActorSystem startup.
# This logger prints the log messages to stdout (System.out).
# Options: OFF, ERROR, WARNING, INFO, DEBUG
stdout-loglevel = "WARNING"

# Log the complete configuration at INFO level when the actor system is started.
# This is useful when you are uncertain of what configuration is used.
log-config-on-start = off

# Log at info level when messages are sent to dead letters.
# Possible values:
# on: all dead letters are logged
# off: no logging of dead letters
# n: positive integer, number of dead letters that will be logged
log-dead-letters = 10

# Possibility to turn off logging of dead letters while the actor system
# is shutting down. Logging is only done when enabled by 'log-dead-letters'
# setting.
log-dead-letters-during-shutdown = on

# List FQCN of extensions which shall be loaded at actor system startup.
# Should be on the format: 'extensions = ["foo", "bar"]' etc.
# See the Akka Documentation for more info about Extensions
extensions = []

# Toggles whether threads created by this ActorSystem should be daemons or not
daemonic = off

# JVM shutdown, System.exit(-1), in case of a fatal error,
# such as OutOfMemoryError
jvm-exit-on-fatal-error = on

actor {

  # FQCN of the ActorRefProvider to be used; the below is the built-in default,
  # another one is akka.remote.RemoteActorRefProvider in the akka-remote bundle.
  provider = "akka.actor.LocalActorRefProvider"

  # The guardian "/user" will use this class to obtain its supervisorStrategy.
  # It needs to be a subclass of akka.actor.SupervisorStrategyConfigurator.
  # In addition to the default there is akka.actor.StoppingSupervisorStrategy.
  guardian-supervisor-strategy = "akka.actor.DefaultSupervisorStrategy"

  # Timeout for ActorSystem.actorOf
  creation-timeout = 20s

  # Frequency with which stopping actors are prodded in case they had to be
  # removed from their parents
  reaper-interval = 5s

  # Serializes and deserializes (non-primitive) messages to ensure immutability,
  # this is only intended for testing.
  serialize-messages = off

  # Serializes and deserializes creators (in Props) to ensure that they can be
  # sent over the network, this is only intended for testing. Purely local deployments
  # as marked with deploy.scope == LocalScope are exempt from verification.
  serialize-creators = off

```

```

# Timeout for send operations to top-level actors which are in the process
# of being started. This is only relevant if using a bounded mailbox or the
# CallingThreadDispatcher for a top-level actor.
unstarted-push-timeout = 10s

typed {
  # Default timeout for typed actor methods with non-void return type
  timeout = 5s
}

# Mapping between 'deployment.router' short names to fully qualified class names
router.type-mapping {
  from-code = "akka.routing.NoRouter"
  round-robin-pool = "akka.routing.RoundRobinPool"
  round-robin-group = "akka.routing.RoundRobinGroup"
  random-pool = "akka.routing.RandomPool"
  random-group = "akka.routing.RandomGroup"
  balancing-pool = "akka.routing.BalancingPool"
  smallest-mailbox-pool = "akka.routing.SmallestMailboxPool"
  broadcast-pool = "akka.routing.BroadcastPool"
  broadcast-group = "akka.routing.BroadcastGroup"
  scatter-gather-pool = "akka.routing.ScatterGatherFirstCompletedPool"
  scatter-gather-group = "akka.routing.ScatterGatherFirstCompletedGroup"
  tail-chopping-pool = "akka.routing.TailChoppingPool"
  tail-chopping-group = "akka.routing.TailChoppingGroup"
  consistent-hashing-pool = "akka.routing.ConsistentHashingPool"
  consistent-hashing-group = "akka.routing.ConsistentHashingGroup"
}

deployment {

  # deployment id pattern - on the format: /parent/child etc.
  default {

    # The id of the dispatcher to use for this actor.
    # If undefined or empty the dispatcher specified in code
    # (Props.withDispatcher) is used, or default-dispatcher if not
    # specified at all.
    dispatcher = ""

    # The id of the mailbox to use for this actor.
    # If undefined or empty the default mailbox of the configured dispatcher
    # is used or if there is no mailbox configuration the mailbox specified
    # in code (Props.withMailbox) is used.
    # If there is a mailbox defined in the configured dispatcher then that
    # overrides this setting.
    mailbox = ""

    # routing (load-balance) scheme to use
    # - available: "from-code", "round-robin", "random", "smallest-mailbox",
    #               "scatter-gather", "broadcast"
    # - or:        Fully qualified class name of the router class.
    #               The class must extend akka.routing.CustomRouterConfig and
    #               have a public constructor with com.typesafe.config.Config
    #               and optional akka.actor.DynamicAccess parameter.
    # - default is "from-code";
    # Whether or not an actor is transformed to a Router is decided in code
    # only (Props.withRouter). The type of router can be overridden in the
    # configuration; specifying "from-code" means that the values specified
    # in the code shall be used.
    # In case of routing, the actors to be routed to can be specified
    # in several ways:

```

```

# - nr-of-instances: will create that many children
# - routees.paths: will route messages to these paths using ActorSelection,
#   i.e. will not create children
# - resizer: dynamically resizable number of routees as specified in
#   resizer below
router = "from-code"

# number of children to create in case of a router;
# this setting is ignored if routees.paths is given
nr-of-instances = 1

# within is the timeout used for routers containing future calls
within = 5 seconds

# number of virtual nodes per node for consistent-hashing router
virtual-nodes-factor = 10

tail-chopping-router {
  # interval is duration between sending message to next routee
  interval = 10 milliseconds
}

routees {
  # Alternatively to giving nr-of-instances you can specify the full
  # paths of those actors which should be routed to. This setting takes
  # precedence over nr-of-instances
  paths = []
}

# To use a dedicated dispatcher for the routees of the pool you can
# define the dispatcher configuration inline with the property name
# 'pool-dispatcher' in the deployment section of the router.
# For example:
# pool-dispatcher {
#   fork-join-executor.parallelism-min = 5
#   fork-join-executor.parallelism-max = 5
# }

# Routers with dynamically resizable number of routees; this feature is
# enabled by including (parts of) this section in the deployment
resizer {

  enabled = off

  # The fewest number of routees the router should ever have.
  lower-bound = 1

  # The most number of routees the router should ever have.
  # Must be greater than or equal to lower-bound.
  upper-bound = 10

  # Threshold used to evaluate if a routee is considered to be busy
  # (under pressure). Implementation depends on this value (default is 1).
  # 0: number of routees currently processing a message.
  # 1: number of routees currently processing a message has
  #   some messages in mailbox.
  # > 1: number of routees with at least the configured pressure-threshold
  #   messages in their mailbox. Note that estimating mailbox size of
  #   default UnboundedMailbox is O(N) operation.
  pressure-threshold = 1

  # Percentage to increase capacity whenever all routees are busy.
  # For example, 0.2 would increase 20% (rounded up), i.e. if current

```

```

# capacity is 6 it will request an increase of 2 more routees.
rampup-rate = 0.2

# Minimum fraction of busy routees before backing off.
# For example, if this is 0.3, then we'll remove some routees only when
# less than 30% of routees are busy, i.e. if current capacity is 10 and
# 3 are busy then the capacity is unchanged, but if 2 or less are busy
# the capacity is decreased.
# Use 0.0 or negative to avoid removal of routees.
backoff-threshold = 0.3

# Fraction of routees to be removed when the resizer reaches the
# backoffThreshold.
# For example, 0.1 would decrease 10% (rounded up), i.e. if current
# capacity is 9 it will request an decrease of 1 routee.
backoff-rate = 0.1

# Number of messages between resize operation.
# Use 1 to resize before each message.
messages-per-resize = 10
}
}

/IO-DNS/inet-address {
  mailbox = "unbounded"
  router = "consistent-hashing-pool"
  nr-of-instances = 4
}
}

default-dispatcher {
  # Must be one of the following
  # Dispatcher, PinnedDispatcher, or a FQCN to a class inheriting
  # MessageDispatcherConfigurator with a public constructor with
  # both com.typesafe.config.Config parameter and
  # akka.dispatch.DispatcherPrerequisites parameters.
  # PinnedDispatcher must be used together with executor=thread-pool-executor.
  type = "Dispatcher"

  # Which kind of ExecutorService to use for this dispatcher
  # Valid options:
  # - "default-executor" requires a "default-executor" section
  # - "fork-join-executor" requires a "fork-join-executor" section
  # - "thread-pool-executor" requires a "thread-pool-executor" section
  # - A FQCN of a class extending ExecutorServiceConfigurator
  executor = "default-executor"

  # This will be used if you have set "executor = "default-executor"".
  # If an ActorSystem is created with a given ExecutionContext, this
  # ExecutionContext will be used as the default executor for all
  # dispatchers in the ActorSystem configured with
  # executor = "default-executor". Note that "default-executor"
  # is the default value for executor, and therefore used if not
  # specified otherwise. If no ExecutionContext is given,
  # the executor configured in "fallback" will be used.
  default-executor {
    fallback = "fork-join-executor"
  }

  # This will be used if you have set "executor = "fork-join-executor""
  fork-join-executor {
    # Min number of threads to cap factor-based parallelism number to
    parallelism-min = 8
  }
}

```

```

# The parallelism factor is used to determine thread pool size using the
# following formula: ceil(available processors * factor). Resulting size
# is then bounded by the parallelism-min and parallelism-max values.
parallelism-factor = 3.0

# Max number of threads to cap factor-based parallelism number to
parallelism-max = 64

# Setting to "FIFO" to use queue like peeking mode which "poll" or "LIFO" to use stack
# like peeking mode which "pop".
task-peeking-mode = "FIFO"
}

# This will be used if you have set "executor = "thread-pool-executor""
thread-pool-executor {
  # Keep alive time for threads
  keep-alive-time = 60s

  # Min number of threads to cap factor-based core number to
  core-pool-size-min = 8

  # The core pool size factor is used to determine thread pool core size
  # using the following formula: ceil(available processors * factor).
  # Resulting size is then bounded by the core-pool-size-min and
  # core-pool-size-max values.
  core-pool-size-factor = 3.0

  # Max number of threads to cap factor-based number to
  core-pool-size-max = 64

  # Minimum number of threads to cap factor-based max number to
  # (if using a bounded task queue)
  max-pool-size-min = 8

  # Max no of threads (if using a bounded task queue) is determined by
  # calculating: ceil(available processors * factor)
  max-pool-size-factor = 3.0

  # Max number of threads to cap factor-based max number to
  # (if using a bounded task queue)
  max-pool-size-max = 64

  # Specifies the bounded capacity of the task queue (< 1 == unbounded)
  task-queue-size = -1

  # Specifies which type of task queue will be used, can be "array" or
  # "linked" (default)
  task-queue-type = "linked"

  # Allow core threads to time out
  allow-core-timeout = on
}

# How long time the dispatcher will wait for new actors until it shuts down
shutdown-timeout = 1s

# Throughput defines the number of messages that are processed in a batch
# before the thread is returned to the pool. Set to 1 for as fair as possible.
throughput = 5

# Throughput deadline for Dispatcher, set to 0 or negative for no deadline
throughput-deadline-time = 0ms

```

```

# For BalancingDispatcher: If the balancing dispatcher should attempt to
# schedule idle actors using the same dispatcher when a message comes in,
# and the dispatchers ExecutorService is not fully busy already.
attempt-teamwork = on

# If this dispatcher requires a specific type of mailbox, specify the
# fully-qualified class name here; the actually created mailbox will
# be a subtype of this type. The empty string signifies no requirement.
mailbox-requirement = ""
}

default-mailbox {
# FQCN of the MailboxType. The Class of the FQCN must have a public
# constructor with
# (akka.actor.ActorSystem.Settings, com.typesafe.config.Config) parameters.
mailbox-type = "akka.dispatch.UnboundedMailbox"

# If the mailbox is bounded then it uses this setting to determine its
# capacity. The provided value must be positive.
# NOTICE:
# Up to version 2.1 the mailbox type was determined based on this setting;
# this is no longer the case, the type must explicitly be a bounded mailbox.
mailbox-capacity = 1000

# If the mailbox is bounded then this is the timeout for enqueueing
# in case the mailbox is full. Negative values signify infinite
# timeout, which should be avoided as it bears the risk of dead-lock.
mailbox-push-timeout-time = 10s

# For Actor with Stash: The default capacity of the stash.
# If negative (or zero) then an unbounded stash is used (default)
# If positive then a bounded stash is used and the capacity is set using
# the property
stash-capacity = -1
}

mailbox {
# Mapping between message queue semantics and mailbox configurations.
# Used by akka.dispatch.RequiresMessageQueue[T] to enforce different
# mailbox types on actors.
# If your Actor implements RequiresMessageQueue[T], then when you create
# an instance of that actor its mailbox type will be decided by looking
# up a mailbox configuration via T in this mapping
requirements {
"akka.dispatch.UnboundedMessageQueueSemantics" =
akka.actor.mailbox.unbounded-queue-based
"akka.dispatch.BoundedMessageQueueSemantics" =
akka.actor.mailbox.bounded-queue-based
"akka.dispatch.DequeBasedMessageQueueSemantics" =
akka.actor.mailbox.unbounded-deque-based
"akka.dispatch.UnboundedDequeBasedMessageQueueSemantics" =
akka.actor.mailbox.unbounded-deque-based
"akka.dispatch.BoundedDequeBasedMessageQueueSemantics" =
akka.actor.mailbox.bounded-deque-based
"akka.dispatch.MultipleConsumerSemantics" =
akka.actor.mailbox.unbounded-queue-based
"akka.dispatch.ControlAwareMessageQueueSemantics" =
akka.actor.mailbox.unbounded-control-aware-queue-based
"akka.dispatch.UnboundedControlAwareMessageQueueSemantics" =
akka.actor.mailbox.unbounded-control-aware-queue-based
"akka.dispatch.BoundedControlAwareMessageQueueSemantics" =
akka.actor.mailbox.bounded-control-aware-queue-based
}
}

```

```

    "akka.event.LoggerMessageQueueSemantics" =
      akka.actor.mailbox.logger-queue
  }

  unbounded-queue-based {
    # FQCN of the MailboxType, The Class of the FQCN must have a public
    # constructor with (akka.actor.ActorSystem.Settings,
    # com.typesafe.config.Config) parameters.
    mailbox-type = "akka.dispatch.UnboundedMailbox"
  }

  bounded-queue-based {
    # FQCN of the MailboxType, The Class of the FQCN must have a public
    # constructor with (akka.actor.ActorSystem.Settings,
    # com.typesafe.config.Config) parameters.
    mailbox-type = "akka.dispatch.BoundedMailbox"
  }

  unbounded-deque-based {
    # FQCN of the MailboxType, The Class of the FQCN must have a public
    # constructor with (akka.actor.ActorSystem.Settings,
    # com.typesafe.config.Config) parameters.
    mailbox-type = "akka.dispatch.UnboundedDequeBasedMailbox"
  }

  bounded-deque-based {
    # FQCN of the MailboxType, The Class of the FQCN must have a public
    # constructor with (akka.actor.ActorSystem.Settings,
    # com.typesafe.config.Config) parameters.
    mailbox-type = "akka.dispatch.BoundedDequeBasedMailbox"
  }

  unbounded-control-aware-queue-based {
    # FQCN of the MailboxType, The Class of the FQCN must have a public
    # constructor with (akka.actor.ActorSystem.Settings,
    # com.typesafe.config.Config) parameters.
    mailbox-type = "akka.dispatch.UnboundedControlAwareMailbox"
  }

  bounded-control-aware-queue-based {
    # FQCN of the MailboxType, The Class of the FQCN must have a public
    # constructor with (akka.actor.ActorSystem.Settings,
    # com.typesafe.config.Config) parameters.
    mailbox-type = "akka.dispatch.BoundedControlAwareMailbox"
  }

  # The LoggerMailbox will drain all messages in the mailbox
  # when the system is shutdown and deliver them to the StandardOutLogger.
  # Do not change this unless you know what you are doing.
  logger-queue {
    mailbox-type = "akka.event.LoggerMailboxType"
  }
}

debug {
  # enable function of Actor.loggable(), which is to log any received message
  # at DEBUG level, see the "Testing Actor Systems" section of the Akka
  # Documentation at http://akka.io/docs
  receive = off

  # enable DEBUG logging of all AutoReceiveMessages (Kill, PoisonPill et.c.)
  autoreceive = off
}

```

```

# enable DEBUG logging of actor lifecycle changes
lifecycle = off

# enable DEBUG logging of all LoggingFSMs for events, transitions and timers
fsm = off

# enable DEBUG logging of subscription changes on the eventStream
event-stream = off

# enable DEBUG logging of unhandled messages
unhandled = off

# enable WARN logging of misconfigured routers
router-misconfiguration = off
}

# Entries for pluggable serializers and their bindings.
serializers {
  java = "akka.serialization.JavaSerializer"
  bytes = "akka.serialization.ByteArraySerializer"
}

# Class to Serializer binding. You only need to specify the name of an
# interface or abstract base class of the messages. In case of ambiguity it
# is using the most specific configured class, or giving a warning and
# choosing the "first" one.
#
# To disable one of the default serializers, assign its class to "none", like
# "java.io.Serializable" = none
serialization-bindings {
  "[B" = bytes
  "java.io.Serializable" = java
}

# Configuration namespace of serialization identifiers.
# Each serializer implementation must have an entry in the following format:
# `akka.actor.serialization-identifiers."FQCN" = ID`
# where `FQCN` is fully qualified class name of the serializer implementation
# and `ID` is globally unique serializer identifier number.
# Identifier values from 0 to 16 are reserved for Akka internal usage.
serialization-identifiers {
  "akka.serialization.JavaSerializer" = 1
  "akka.serialization.ByteArraySerializer" = 4
}

# Configuration items which are used by the akka.actor.ActorDSL._ methods
dsl {
  # Maximum queue size of the actor created by newInbox(); this protects
  # against faulty programs which use select() and consistently miss messages
  inbox-size = 1000

  # Default timeout to assume for operations like Inbox.receive et al
  default-timeout = 5s
}
}

# Used to set the behavior of the scheduler.
# Changing the default values may change the system behavior drastically so make
# sure you know what you're doing! See the Scheduler section of the Akka
# Documentation for more details.
scheduler {
  # The LightArrayRevolverScheduler is used as the default scheduler in the
  # system. It does not execute the scheduled tasks on exact time, but on every

```

```

# tick, it will run everything that is (over)due. You can increase or decrease
# the accuracy of the execution timing by specifying smaller or larger tick
# duration. If you are scheduling a lot of tasks you should consider increasing
# the ticks per wheel.
# Note that it might take up to 1 tick to stop the Timer, so setting the
# tick-duration to a high value will make shutting down the actor system
# take longer.
tick-duration = 10ms

# The timer uses a circular wheel of buckets to store the timer tasks.
# This should be set such that the majority of scheduled timeouts (for high
# scheduling frequency) will be shorter than one rotation of the wheel
# (ticks-per-wheel * ticks-duration)
# THIS MUST BE A POWER OF TWO!
ticks-per-wheel = 512

# This setting selects the timer implementation which shall be loaded at
# system start-up.
# The class given here must implement the akka.actor.Scheduler interface
# and offer a public constructor which takes three arguments:
# 1) com.typesafe.config.Config
# 2) akka.event.LoggingAdapter
# 3) java.util.concurrent.ThreadFactory
implementation = akka.actor.LightArrayRevolverScheduler

# When shutting down the scheduler, there will typically be a thread which
# needs to be stopped, and this timeout determines how long to wait for
# that to happen. In case of timeout the shutdown of the actor system will
# proceed without running possibly still enqueued tasks.
shutdown-timeout = 5s
}

io {

# By default the select loops run on dedicated threads, hence using a
# PinnedDispatcher
pinned-dispatcher {
  type = "PinnedDispatcher"
  executor = "thread-pool-executor"
  thread-pool-executor.allow-core-timeout = off
}

tcp {

# The number of selectors to stripe the served channels over; each of
# these will use one select loop on the selector-dispatcher.
nr-of-selectors = 1

# Maximum number of open channels supported by this TCP module; there is
# no intrinsic general limit, this setting is meant to enable DoS
# protection by limiting the number of concurrently connected clients.
# Also note that this is a "soft" limit; in certain cases the implementation
# will accept a few connections more or a few less than the number configured
# here. Must be an integer > 0 or "unlimited".
max-channels = 256000

# When trying to assign a new connection to a selector and the chosen
# selector is at full capacity, retry selector choosing and assignment
# this many times before giving up
selector-association-retries = 10

# The maximum number of connection that are accepted in one go,
# higher numbers decrease latency, lower numbers increase fairness on

```

```

# the worker-dispatcher
batch-accept-limit = 10

# The number of bytes per direct buffer in the pool used to read or write
# network data from the kernel.
direct-buffer-size = 128 KiB

# The maximal number of direct buffers kept in the direct buffer pool for
# reuse.
direct-buffer-pool-limit = 1000

# The duration a connection actor waits for a 'Register' message from
# its commander before aborting the connection.
register-timeout = 5s

# The maximum number of bytes delivered by a 'Received' message. Before
# more data is read from the network the connection actor will try to
# do other work.
# The purpose of this setting is to impose a smaller limit than the
# configured receive buffer size. When using value 'unlimited' it will
# try to read all from the receive buffer.
max-received-message-size = unlimited

# Enable fine grained logging of what goes on inside the implementation.
# Be aware that this may log more than once per message sent to the actors
# of the tcp implementation.
trace-logging = off

# Fully qualified config path which holds the dispatcher configuration
# to be used for running the select() calls in the selectors
selector-dispatcher = "akka.io.pinned-dispatcher"

# Fully qualified config path which holds the dispatcher configuration
# for the read/write worker actors
worker-dispatcher = "akka.actor.default-dispatcher"

# Fully qualified config path which holds the dispatcher configuration
# for the selector management actors
management-dispatcher = "akka.actor.default-dispatcher"

# Fully qualified config path which holds the dispatcher configuration
# on which file IO tasks are scheduled
file-io-dispatcher = "akka.actor.default-dispatcher"

# The maximum number of bytes (or "unlimited") to transfer in one batch
# when using 'WriteFile' command which uses 'FileChannel.transferTo' to
# pipe files to a TCP socket. On some OS like Linux 'FileChannel.transferTo'
# may block for a long time when network IO is faster than file IO.
# Decreasing the value may improve fairness while increasing may improve
# throughput.
file-io-transferTo-limit = 512 KiB

# The number of times to retry the 'finishConnect' call after being notified about
# OP_CONNECT. Retries are needed if the OP_CONNECT notification doesn't imply that
# 'finishConnect' will succeed, which is the case on Android.
finish-connect-retries = 5

# On Windows connection aborts are not reliably detected unless an OP_READ is
# registered on the selector _after_ the connection has been reset. This
# workaround enables an OP_CONNECT which forces the abort to be visible on Windows.
# Enabling this setting on other platforms than Windows will cause various failures
# and undefined behavior.
# Possible values of this key are on, off and auto where auto will enable the

```

```

# workaround if Windows is detected automatically.
windows-connection-abort-workaround-enabled = off
}

udp {

# The number of selectors to stripe the served channels over; each of
# these will use one select loop on the selector-dispatcher.
nr-of-selectors = 1

# Maximum number of open channels supported by this UDP module Generally
# UDP does not require a large number of channels, therefore it is
# recommended to keep this setting low.
max-channels = 4096

# The select loop can be used in two modes:
# - setting "infinite" will select without a timeout, hogging a thread
# - setting a positive timeout will do a bounded select call,
#   enabling sharing of a single thread between multiple selectors
#   (in this case you will have to use a different configuration for the
#   selector-dispatcher, e.g. using "type=Dispatcher" with size 1)
# - setting it to zero means polling, i.e. calling selectNow()
select-timeout = infinite

# When trying to assign a new connection to a selector and the chosen
# selector is at full capacity, retry selector choosing and assignment
# this many times before giving up
selector-association-retries = 10

# The maximum number of datagrams that are read in one go,
# higher numbers decrease latency, lower numbers increase fairness on
# the worker-dispatcher
receive-throughput = 3

# The number of bytes per direct buffer in the pool used to read or write
# network data from the kernel.
direct-buffer-size = 128 KiB

# The maximal number of direct buffers kept in the direct buffer pool for
# reuse.
direct-buffer-pool-limit = 1000

# The maximum number of bytes delivered by a 'Received' message. Before
# more data is read from the network the connection actor will try to
# do other work.
received-message-size-limit = unlimited

# Enable fine grained logging of what goes on inside the implementation.
# Be aware that this may log more than once per message sent to the actors
# of the tcp implementation.
trace-logging = off

# Fully qualified config path which holds the dispatcher configuration
# to be used for running the select() calls in the selectors
selector-dispatcher = "akka.io.pinned-dispatcher"

# Fully qualified config path which holds the dispatcher configuration
# for the read/write worker actors
worker-dispatcher = "akka.actor.default-dispatcher"

# Fully qualified config path which holds the dispatcher configuration
# for the selector management actors
management-dispatcher = "akka.actor.default-dispatcher"

```

```

}

udp-connected {

  # The number of selectors to stripe the served channels over; each of
  # these will use one select loop on the selector-dispatcher.
  nr-of-selectors = 1

  # Maximum number of open channels supported by this UDP module Generally
  # UDP does not require a large number of channels, therefore it is
  # recommended to keep this setting low.
  max-channels = 4096

  # The select loop can be used in two modes:
  # - setting "infinite" will select without a timeout, hogging a thread
  # - setting a positive timeout will do a bounded select call,
  #   enabling sharing of a single thread between multiple selectors
  #   (in this case you will have to use a different configuration for the
  #   selector-dispatcher, e.g. using "type=Dispatcher" with size 1)
  # - setting it to zero means polling, i.e. calling selectNow()
  select-timeout = infinite

  # When trying to assign a new connection to a selector and the chosen
  # selector is at full capacity, retry selector choosing and assignment
  # this many times before giving up
  selector-association-retries = 10

  # The maximum number of datagrams that are read in one go,
  # higher numbers decrease latency, lower numbers increase fairness on
  # the worker-dispatcher
  receive-throughput = 3

  # The number of bytes per direct buffer in the pool used to read or write
  # network data from the kernel.
  direct-buffer-size = 128 KiB

  # The maximal number of direct buffers kept in the direct buffer pool for
  # reuse.
  direct-buffer-pool-limit = 1000

  # The maximum number of bytes delivered by a 'Received' message. Before
  # more data is read from the network the connection actor will try to
  # do other work.
  received-message-size-limit = unlimited

  # Enable fine grained logging of what goes on inside the implementation.
  # Be aware that this may log more than once per message sent to the actors
  # of the tcp implementation.
  trace-logging = off

  # Fully qualified config path which holds the dispatcher configuration
  # to be used for running the select() calls in the selectors
  selector-dispatcher = "akka.io.pinned-dispatcher"

  # Fully qualified config path which holds the dispatcher configuration
  # for the read/write worker actors
  worker-dispatcher = "akka.actor.default-dispatcher"

  # Fully qualified config path which holds the dispatcher configuration
  # for the selector management actors
  management-dispatcher = "akka.actor.default-dispatcher"
}

```

```

dns {
  # Fully qualified config path which holds the dispatcher configuration
  # for the manager and resolver router actors.
  # For actual router configuration see akka.actor.deployment./IO-DNS/*
  dispatcher = "akka.actor.default-dispatcher"

  # Name of the subconfig at path akka.io.dns, see inet-address below
  resolver = "inet-address"

  inet-address {
    # Must implement akka.io.DnsProvider
    provider-object = "akka.io.InetAddressDnsProvider"

    # These TTLs are set to default java 6 values
    positive-ttl = 30s
    negative-ttl = 10s

    # How often to sweep out expired cache entries.
    # Note that this interval has nothing to do with TTLs
    cache-cleanup-interval = 120s
  }
}
}
}

```

akka-agent

```

#####
# Akka Agent Reference Config File #
#####

# This is the reference config file that contains all the default settings.
# Make your edits/overrides in your application.conf.

akka {
  agent {

    # The dispatcher used for agent-send-off actor
    send-off-dispatcher {
      executor = thread-pool-executor
      type = PinnedDispatcher
    }

    # The dispatcher used for agent-alter-off actor
    alter-off-dispatcher {
      executor = thread-pool-executor
      type = PinnedDispatcher
    }
  }
}

```

akka-camel

```

#####
# Akka Camel Reference Config File #
#####

# This is the reference config file that contains all the default settings.

```

```
# Make your edits/overrides in your application.conf.

akka {
  camel {
    # FQCN of the ContextProvider to be used to create or locate a CamelContext
    # it must implement akka.camel.ContextProvider and have a no-arg constructor
    # the built-in default create a fresh DefaultCamelContext
    context-provider = akka.camel.DefaultContextProvider

    # Whether JMX should be enabled or disabled for the Camel Context
    jmx = off
    # enable/disable streaming cache on the Camel Context
    streamingCache = on
    consumer {
      # Configured setting which determines whether one-way communications
      # between an endpoint and this consumer actor
      # should be auto-acknowledged or application-acknowledged.
      # This flag has only effect when exchange is in-only.
      auto-ack = on

      # When endpoint is out-capable (can produce responses) reply-timeout is the
      # maximum time the endpoint can take to send the response before the message
      # exchange fails. This setting is used for out-capable, in-only,
      # manually acknowledged communication.
      reply-timeout = 1m

      # The duration of time to await activation of an endpoint.
      activation-timeout = 10s
    }

    #Scheme to FQCN mappings for CamelMessage body conversions
    conversions {
      "file" = "java.io.InputStream"
    }
  }
}
}
```

akka-cluster

```
#####
# Akka Cluster Reference Config File #
#####

# This is the reference config file that contains all the default settings.
# Make your edits/overrides in your application.conf.

akka {

  cluster {
    # Initial contact points of the cluster.
    # The nodes to join automatically at startup.
    # Comma separated full URIs defined by a string on the form of
    # "akka.tcp://system@hostname:port"
    # Leave as empty if the node is supposed to be joined manually.
    seed-nodes = []

    # how long to wait for one of the seed nodes to reply to initial join request
    seed-node-timeout = 5s

    # If a join request fails it will be retried after this period.
    # Disable join retry by specifying "off".
  }
}
```

```

retry-unsuccessful-join-after = 10s

# Should the 'leader' in the cluster be allowed to automatically mark
# unreachable nodes as DOWN after a configured time of unreachability?
# Using auto-down implies that two separate clusters will automatically be
# formed in case of network partition.
# Disable with "off" or specify a duration to enable auto-down.
auto-down-unreachable-after = off

# Time margin after which shards or singletons that belonged to a downed/removed
# partition are created in surviving partition. The purpose of this margin is that
# in case of a network partition the persistent actors in the non-surviving partitions
# must be stopped before corresponding persistent actors are started somewhere else.
# This is useful if you implement downing strategies that handle network partitions,
# e.g. by keeping the larger side of the partition and shutting down the smaller side.
# It will not add any extra safety for auto-down-unreachable-after, since that is not
# handling network partitions.
# Disable with "off" or specify a duration to enable.
down-removal-margin = off

# By default, the leader will not move 'Joining' members to 'Up' during a network
# split. This feature allows the leader to accept 'Joining' members to be 'WeaklyUp'
# so they become part of the cluster even during a network split. The leader will
# move 'WeaklyUp' members to 'Up' status once convergence has been reached. This
# feature must be off if some members are running Akka 2.3.X.
# WeaklyUp is an EXPERIMENTAL feature.
allow-weakly-up-members = off

# The roles of this member. List of strings, e.g. roles = ["A", "B"].
# The roles are part of the membership information and can be used by
# routers or other services to distribute work to certain member types,
# e.g. front-end and back-end nodes.
roles = []

role {
  # Minimum required number of members of a certain role before the leader
  # changes member status of 'Joining' members to 'Up'. Typically used together
  # with 'Cluster.registerOnMemberUp' to defer some action, such as starting
  # actors, until the cluster has reached a certain size.
  # E.g. to require 2 nodes with role 'frontend' and 3 nodes with role 'backend':
  #   frontend.min-nr-of-members = 2
  #   backend.min-nr-of-members = 3
  #<role-name>.min-nr-of-members = 1
}

# Minimum required number of members before the leader changes member status
# of 'Joining' members to 'Up'. Typically used together with
# 'Cluster.registerOnMemberUp' to defer some action, such as starting actors,
# until the cluster has reached a certain size.
min-nr-of-members = 1

# Enable/disable info level logging of cluster events
log-info = on

# Enable or disable JMX MBeans for management of the cluster
jmx.enabled = on

# how long should the node wait before starting the periodic tasks
# maintenance tasks?
periodic-tasks-initial-delay = 1s

# how often should the node send out gossip information?
gossip-interval = 1s

```

```

# discard incoming gossip messages if not handled within this duration
gossip-time-to-live = 2s

# how often should the leader perform maintenance tasks?
leader-actions-interval = 1s

# how often should the node move nodes, marked as unreachable by the failure
# detector, out of the membership ring?
unreachable-nodes-reaper-interval = 1s

# How often the current internal stats should be published.
# A value of 0s can be used to always publish the stats, when it happens.
# Disable with "off".
publish-stats-interval = off

# The id of the dispatcher to use for cluster actors. If not specified
# default dispatcher is used.
# If specified you need to define the settings of the actual dispatcher.
use-dispatcher = ""

# Gossip to random node with newer or older state information, if any with
# this probability. Otherwise Gossip to any random live node.
# Probability value is between 0.0 and 1.0. 0.0 means never, 1.0 means always.
gossip-different-view-probability = 0.8

# Reduced the above probability when the number of nodes in the cluster
# greater than this value.
reduce-gossip-different-view-probability = 400

# Settings for the Phi accrual failure detector (http://ddg.jaist.ac.jp/pub/HDY+04.pdf
# [Hayashibara et al]) used by the cluster subsystem to detect unreachable
# members.
# The default PhiAccrualFailureDetector will trigger if there are no heartbeats within
# the duration heartbeat-interval + acceptable-heartbeat-pause + threshold_adjustment,
# i.e. around 5.5 seconds with default settings.
failure-detector {

  # FQCN of the failure detector implementation.
  # It must implement akka.remote.FailureDetector and have
  # a public constructor with a com.typesafe.config.Config and
  # akka.actor.EventStream parameter.
  implementation-class = "akka.remote.PhiAccrualFailureDetector"

  # How often keep-alive heartbeat messages should be sent to each connection.
  heartbeat-interval = 1 s

  # Defines the failure detector threshold.
  # A low threshold is prone to generate many wrong suspicions but ensures
  # a quick detection in the event of a real crash. Conversely, a high
  # threshold generates fewer mistakes but needs more time to detect
  # actual crashes.
  threshold = 8.0

  # Number of the samples of inter-heartbeat arrival times to adaptively
  # calculate the failure timeout for connections.
  max-sample-size = 1000

  # Minimum standard deviation to use for the normal distribution in
  # AccrualFailureDetector. Too low standard deviation might result in
  # too much sensitivity for sudden, but normal, deviations in heartbeat
  # inter arrival times.
  min-std-deviation = 100 ms
}

```

```

# Number of potentially lost/delayed heartbeats that will be
# accepted before considering it to be an anomaly.
# This margin is important to be able to survive sudden, occasional,
# pauses in heartbeat arrivals, due to for example garbage collect or
# network drop.
acceptable-heartbeat-pause = 3 s

# Number of member nodes that each member will send heartbeat messages to,
# i.e. each node will be monitored by this number of other nodes.
monitored-by-nr-of-members = 5

# After the heartbeat request has been sent the first failure detection
# will start after this period, even though no heartbeat message has
# been received.
expected-response-after = 1 s

}

metrics {
  # Enable or disable metrics collector for load-balancing nodes.
  enabled = on

  # FQCN of the metrics collector implementation.
  # It must implement akka.cluster.MetricsCollector and
  # have public constructor with akka.actor.ActorSystem parameter.
  # The default SigarMetricsCollector uses JMX and Hyperic SIGAR, if SIGAR
  # is on the classpath, otherwise only JMX.
  collector-class = "akka.cluster.SigarMetricsCollector"

  # How often metrics are sampled on a node.
  # Shorter interval will collect the metrics more often.
  collect-interval = 3s

  # How often a node publishes metrics information.
  gossip-interval = 3s

  # How quickly the exponential weighting of past data is decayed compared to
  # new data. Set lower to increase the bias toward newer values.
  # The relevance of each data sample is halved for every passing half-life
  # duration, i.e. after 4 times the half-life, a data sample's relevance is
  # reduced to 6% of its original relevance. The initial relevance of a data
  # sample is given by  $1 - 0.5^{(\text{collect-interval} / \text{half-life})}$ .
  # See http://en.wikipedia.org/wiki/Moving\_average#Exponential\_moving\_average
  moving-average-half-life = 12s
}

# If the tick-duration of the default scheduler is longer than the
# tick-duration configured here a dedicated scheduler will be used for
# periodic tasks of the cluster, otherwise the default scheduler is used.
# See akka.scheduler settings for more details.
scheduler {
  tick-duration = 33ms
  ticks-per-wheel = 512
}

}

# Default configuration for routers
actor.deployment.default {
  # MetricsSelector to use
  # - available: "mix", "heap", "cpu", "load"
  # - or: Fully qualified class name of the MetricsSelector class.

```

```

#       The class must extend akka.cluster.routing.MetricsSelector
#       and have a public constructor with com.typesafe.config.Config
#       parameter.
# - default is "mix"
metrics-selector = mix
}
actor.deployment.default.cluster {
  # enable cluster aware router that deploys to nodes in the cluster
  enabled = off

  # Maximum number of routees that will be deployed on each cluster
  # member node.
  # Note that max-total-nr-of-instances defines total number of routees, but
  # number of routees per node will not be exceeded, i.e. if you
  # define max-total-nr-of-instances = 50 and max-nr-of-instances-per-node = 2
  # it will deploy 2 routees per new member in the cluster, up to
  # 25 members.
  max-nr-of-instances-per-node = 1

  # Maximum number of routees that will be deployed, in total
  # on all nodes. See also description of max-nr-of-instances-per-node.
  # For backwards compatibility reasons, nr-of-instances
  # has the same purpose as max-total-nr-of-instances for cluster
  # aware routers and nr-of-instances (if defined by user) takes
  # precedence over max-total-nr-of-instances.
  max-total-nr-of-instances = 10000

  # Defines if routees are allowed to be located on the same node as
  # the head router actor, or only on remote nodes.
  # Useful for master-worker scenario where all routees are remote.
  allow-local-routees = on

  # Use members with specified role, or all members if undefined or empty.
  use-role = ""
}

# Protobuf serializer for cluster messages
actor {
  serializers {
    akka-cluster = "akka.cluster.protobuf.ClusterMessageSerializer"
  }

  serialization-bindings {
    "akka.cluster.ClusterMessage" = akka-cluster
  }

  serialization-identifiers {
    "akka.cluster.protobuf.ClusterMessageSerializer" = 5
  }

  router.type-mapping {
    adaptive-pool = "akka.cluster.routing.AdaptiveLoadBalancingPool"
    adaptive-group = "akka.cluster.routing.AdaptiveLoadBalancingGroup"
  }
}
}

```

akka-multi-node-testkit

```
#####
# Akka Remote Testing Reference Config File #
#####

# This is the reference config file that contains all the default settings.
# Make your edits/overrides in your application.conf.

akka {
  testconductor {

    # Timeout for joining a barrier: this is the maximum time any participants
    # waits for everybody else to join a named barrier.
    barrier-timeout = 30s

    # Timeout for interrogation of TestConductor's Controller actor
    query-timeout = 5s

    # Threshold for packet size in time unit above which the failure injector will
    # split the packet and deliver in smaller portions; do not give value smaller
    # than HashedWheelTimer resolution (would not make sense)
    packet-split-threshold = 100ms

    # amount of time for the ClientFSM to wait for the connection to the conductor
    # to be successful
    connect-timeout = 20s

    # Number of connect attempts to be made to the conductor controller
    client-reconnects = 30

    # minimum time interval which is to be inserted between reconnect attempts
    reconnect-backoff = 1s

  }

  netty {
    # (I&O) Used to configure the number of I/O worker threads on server sockets
    server-socket-worker-pool {
      # Min number of threads to cap factor-based number to
      pool-size-min = 1

      # The pool size factor is used to determine thread pool size
      # using the following formula: ceil(available processors * factor).
      # Resulting size is then bounded by the pool-size-min and
      # pool-size-max values.
      pool-size-factor = 1.0

      # Max number of threads to cap factor-based number to
      pool-size-max = 2
    }

    # (I&O) Used to configure the number of I/O worker threads on client sockets
    client-socket-worker-pool {
      # Min number of threads to cap factor-based number to
      pool-size-min = 1

      # The pool size factor is used to determine thread pool size
      # using the following formula: ceil(available processors * factor).
      # Resulting size is then bounded by the pool-size-min and
      # pool-size-max values.
      pool-size-factor = 1.0

      # Max number of threads to cap factor-based number to
      pool-size-max = 2
    }
  }
}
```

```

    }
  }
}
}

```

akka-persistence

```

#####
# Akka Persistence Extension Reference Configuration File #
#####

# This is the reference config file that contains all the default settings.
# Make your edits in your application.conf in order to override these settings.

# Directory of persistence journal and snapshot store plugins is available at the
# Akka Community Projects page http://akka.io/community/

# Default persistence extension settings.
akka.persistence {
  journal {
    # Absolute path to the journal plugin configuration entry used by
    # persistent actor or view by default.
    # Persistent actor or view can override `journalPluginId` method
    # in order to rely on a different journal plugin.
    plugin = ""
  }
  snapshot-store {
    # Absolute path to the snapshot plugin configuration entry used by
    # persistent actor or view by default.
    # Persistent actor or view can override `snapshotPluginId` method
    # in order to rely on a different snapshot plugin.
    # It is not mandatory to specify a snapshot store plugin.
    # If you don't use snapshots you don't have to configure it.
    # Note that Cluster Sharding is using snapshots, so if you
    # use Cluster Sharding you need to define a snapshot store plugin.
    plugin = ""
  }
  # used as default-snapshot store if no plugin configured
  # (see `akka.persistence.snapshot-store`)
  no-snapshot-store {
    class = "akka.persistence.snapshot.NoSnapshotStore"
  }
  # Default persistent view settings.
  view {
    # Automated incremental view update.
    auto-update = on
    # Interval between incremental updates.
    auto-update-interval = 5s
    # Maximum number of messages to replay per incremental view update.
    # Set to -1 for no upper limit.
    auto-update-replay-max = -1
  }
  # Default reliable delivery settings.
  at-least-once-delivery {
    # Interval between re-delivery attempts.
    redeliver-interval = 5s
    # Maximum number of unconfirmed messages that will be sent in one
    # re-delivery burst.
    redelivery-burst-limit = 10000
    # After this number of delivery attempts a
    # `ReliableRedelivery.UnconfirmedWarning`, message will be sent to the actor.

```

```

warn-after-number-of-unconfirmed-attempts = 5
# Maximum number of unconfirmed messages that an actor with
# AtLeastOnceDelivery is allowed to hold in memory.
max-unconfirmed-messages = 100000
}
# Default persistent extension thread pools.
dispatchers {
  # Dispatcher used by every plugin which does not declare explicit
  # `plugin-dispatcher` field.
  default-plugin-dispatcher {
    type = PinnedDispatcher
    executor = "thread-pool-executor"
  }
  # Default dispatcher for message replay.
  default-replay-dispatcher {
    type = Dispatcher
    executor = "fork-join-executor"
    fork-join-executor {
      parallelism-min = 2
      parallelism-max = 8
    }
  }
  # Default dispatcher for streaming snapshot IO
  default-stream-dispatcher {
    type = Dispatcher
    executor = "fork-join-executor"
    fork-join-executor {
      parallelism-min = 2
      parallelism-max = 8
    }
  }
}

# Fallback settings for journal plugin configurations.
# These settings are used if they are not defined in plugin config section.
journal-plugin-fallback {

  # Fully qualified class name providing journal plugin api implementation.
  # It is mandatory to specify this property.
  # The class must have a constructor without parameters or constructor with
  # one `com.typesafe.config.Config` parameter.
  class = ""

  # Dispatcher for the plugin actor.
  plugin-dispatcher = "akka.persistence.dispatchers.default-plugin-dispatcher"

  # Dispatcher for message replay.
  replay-dispatcher = "akka.persistence.dispatchers.default-replay-dispatcher"

  # Maximum size of a persistent message batch written to the journal.
  max-message-batch-size = 200

  circuit-breaker {
    max-failures = 10
    call-timeout = 10s
    reset-timeout = 30s
  }

  # The replay filter can detect a corrupt event stream by inspecting
  # sequence numbers and writerUuid when replaying events.
  replay-filter {
    # What the filter should do when detecting invalid events.
    # Supported values:

```

```

# `repair-by-discard-old` : discard events from old writers,
#                          warning is logged
# `fail` : fail the replay, error is logged
# `warn` : log warning but emit events untouched
# `off` : disable this feature completely
mode = repair-by-discard-old

# It uses a look ahead buffer for analyzing the events.
# This defines the size (in number of events) of the buffer.
window-size = 100

# How many old writerUuid to remember
max-old-writers = 10
}
}

# Fallback settings for snapshot store plugin configurations
# These settings are used if they are not defined in plugin config section.
snapshot-store-plugin-fallback {

# Fully qualified class name providing snapshot store plugin api
# implementation. It is mandatory to specify this property if
# snapshot store is enabled.
# The class must have a constructor without parameters or constructor with
# one `com.typesafe.config.Config` parameter.
class = ""

# Dispatcher for the plugin actor.
plugin-dispatcher = "akka.persistence.dispatchers.default-plugin-dispatcher"

circuit-breaker {
  max-failures = 5
  call-timeout = 20s
  reset-timeout = 60s
}
}
}

# Protobuf serialization for the persistent extension messages.
akka.actor {
  serializers {
    akka-persistence-message = "akka.persistence.serialization.MessageSerializer"
    akka-persistence-snapshot = "akka.persistence.serialization.SnapshotSerializer"
  }
  serialization-bindings {
    "akka.persistence.serialization.Message" = akka-persistence-message
    "akka.persistence.serialization.Snapshot" = akka-persistence-snapshot
  }
  serialization-identifiers {
    "akka.persistence.serialization.MessageSerializer" = 7
    "akka.persistence.serialization.SnapshotSerializer" = 8
  }
}

#####
# Persistence plugins included with the extension #
#####

# In-memory journal plugin.
akka.persistence.journal.inmem {
  # Class name of the plugin.
  class = "akka.persistence.journal.inmem.InmemJournal"
}

```

```

# Dispatcher for the plugin actor.
plugin-dispatcher = "akka.actor.default-dispatcher"
}

# Local file system snapshot store plugin.
akka.persistence.snapshot-store.local {
  # Class name of the plugin.
  class = "akka.persistence.snapshot.local.LocalSnapshotStore"
  # Dispatcher for the plugin actor.
  plugin-dispatcher = "akka.persistence.dispatchers.default-plugin-dispatcher"
  # Dispatcher for streaming snapshot IO.
  stream-dispatcher = "akka.persistence.dispatchers.default-stream-dispatcher"
  # Storage location of snapshot files.
  dir = "snapshots"
  # Number load attempts when recovering from the latest snapshot fails
  # yet older snapshot files are available. Each recovery attempt will try
  # to recover using an older than previously failed-on snapshot file
  # (if any are present).
  max-load-attempts = 3
}

# LevelDB journal plugin.
# Note: this plugin requires explicit LevelDB dependency, see below.
akka.persistence.journal.leveldb {
  # Class name of the plugin.
  class = "akka.persistence.journal.leveldb.LeveldbJournal"
  # Dispatcher for the plugin actor.
  plugin-dispatcher = "akka.persistence.dispatchers.default-plugin-dispatcher"
  # Dispatcher for message replay.
  replay-dispatcher = "akka.persistence.dispatchers.default-replay-dispatcher"
  # Storage location of LevelDB files.
  dir = "journal"
  # Use fsync on write.
  fsync = on
  # Verify checksum on read.
  checksum = off
  # Native LevelDB (via JNI) or LevelDB Java port.
  native = on
}

# Shared LevelDB journal plugin (for testing only).
# Note: this plugin requires explicit LevelDB dependency, see below.
akka.persistence.journal.leveldb-shared {
  # Class name of the plugin.
  class = "akka.persistence.journal.leveldb.SharedLeveldbJournal"
  # Dispatcher for the plugin actor.
  plugin-dispatcher = "akka.actor.default-dispatcher"
  # Timeout for async journal operations.
  timeout = 10s
  store {
    # Dispatcher for shared store actor.
    store-dispatcher = "akka.persistence.dispatchers.default-plugin-dispatcher"
    # Dispatcher for message replay.
    replay-dispatcher = "akka.persistence.dispatchers.default-replay-dispatcher"
    # Storage location of LevelDB files.
    dir = "journal"
    # Use fsync on write.
    fsync = on
    # Verify checksum on read.
    checksum = off
    # Native LevelDB (via JNI) or LevelDB Java port.
    native = on
  }
}

```

```

}

# LevelDB persistence requires the following dependency declarations:
#
# SBT:
#   "org.iq80.leveldb"           % "leveldb"           % "0.7"
#   "org.fusesource.leveldbjni" % "leveldbjni-all"    % "1.8"
#
# Maven:
#   <dependency>
#     <groupId>org.iq80.leveldb</groupId>
#     <artifactId>leveldb</artifactId>
#     <version>0.7</version>
#   </dependency>
#   <dependency>
#     <groupId>org.fusesource.leveldbjni</groupId>
#     <artifactId>leveldbjni-all</artifactId>
#     <version>1.8</version>
#   </dependency>

```

akka-remote

```

#####
# Akka Remote Reference Config File #
#####

# This is the reference config file that contains all the default settings.
# Make your edits/overrides in your application.conf.

# comments about akka.actor settings left out where they are already in akka-
# actor.jar, because otherwise they would be repeated in config rendering.

akka {

  actor {

    serializers {
      akka-containers = "akka.remote.serialization.MessageContainerSerializer"
      proto = "akka.remote.serialization.ProtobufSerializer"
      daemon-create = "akka.remote.serialization.DaemonMsgCreateSerializer"
    }

    serialization-bindings {
      "akka.actor.ActorSelectionMessage" = akka-containers
      "akka.remote.DaemonMsgCreate" = daemon-create

      # Since akka.protobuf.Message does not extend Serializable but
      # GeneratedMessage does, need to use the more specific one here in order
      # to avoid ambiguity.
      "akka.protobuf.GeneratedMessage" = proto

      # Since com.google.protobuf.Message does not extend Serializable but
      # GeneratedMessage does, need to use the more specific one here in order
      # to avoid ambiguity.
      # This com.google.protobuf serialization binding is only used if the class can be loaded,
      # i.e. com.google.protobuf dependency has been added in the application project.
      "com.google.protobuf.GeneratedMessage" = proto
    }

    serialization-identifiers {

```

```

"akka.remote.serialization.ProtobufSerializer" = 2
"akka.remote.serialization.DaemonMsgCreateSerializer" = 3
"akka.remote.serialization.MessageContainerSerializer" = 6
}

deployment {

  default {

    # if this is set to a valid remote address, the named actor will be
    # deployed at that node e.g. "akka.tcp://sys@host:port"
    remote = ""

    target {

      # A list of hostnames and ports for instantiating the children of a
      # router
      # The format should be on "akka.tcp://sys@host:port", where:
      # - sys is the remote actor system name
      # - hostname can be either hostname or IP address the remote actor
      # should connect to
      # - port should be the port for the remote server on the other node
      # The number of actor instances to be spawned is still taken from the
      # nr-of-instances setting as for local routers; the instances will be
      # distributed round-robin among the given nodes.
      nodes = []

    }
  }
}

remote {

  ### General settings

  # Timeout after which the startup of the remoting subsystem is considered
  # to be failed. Increase this value if your transport drivers (see the
  # enabled-transport section) need longer time to be loaded.
  startup-timeout = 10 s

  # Timeout after which the graceful shutdown of the remoting subsystem is
  # considered to be failed. After the timeout the remoting system is
  # forcefully shut down. Increase this value if your transport drivers
  # (see the enabled-transport section) need longer time to stop properly.
  shutdown-timeout = 10 s

  # Before shutting down the drivers, the remoting subsystem attempts to flush
  # all pending writes. This setting controls the maximum time the remoting is
  # willing to wait before moving on to shut down the drivers.
  flush-wait-on-shutdown = 2 s

  # Reuse inbound connections for outbound messages
  use-passive-connections = on

  # Controls the backoff interval after a refused write is reattempted.
  # (Transports may refuse writes if their internal buffer is full)
  backoff-interval = 5 ms

  # Acknowledgment timeout of management commands sent to the transport stack.
  command-ack-timeout = 30 s

  # If set to a nonempty string remoting will use the given dispatcher for

```

```

# its internal actors otherwise the default dispatcher is used. Please note
# that since remoting can load arbitrary 3rd party drivers (see
# "enabled-transport" and "adapters" entries) it is not guaranteed that
# every module will respect this setting.
use-dispatcher = "akka.remote.default-remote-dispatcher"

### Security settings

# Enable untrusted mode for full security of server managed actors, prevents
# system messages to be send by clients, e.g. messages like 'Create',
# 'Suspend', 'Resume', 'Terminate', 'Supervise', 'Link' etc.
untrusted-mode = off

# When 'untrusted-mode=on' inbound actor selections are by default discarded.
# Actors with paths defined in this white list are granted permission to receive actor
# selections messages.
# E.g. trusted-selection-paths = ["/user/receptionist", "/user/namingService"]
trusted-selection-paths = []

# Should the remote server require that its peers share the same
# secure-cookie (defined in the 'remote' section)? Secure cookies are passed
# between during the initial handshake. Connections are refused if the initial
# message contains a mismatching cookie or the cookie is missing.
require-cookie = off

# Deprecated since 2.4-M1
secure-cookie = ""

### Logging

# If this is "on", Akka will log all inbound messages at DEBUG level,
# if off then they are not logged
log-received-messages = off

# If this is "on", Akka will log all outbound messages at DEBUG level,
# if off then they are not logged
log-sent-messages = off

# Sets the log granularity level at which Akka logs remoting events. This setting
# can take the values OFF, ERROR, WARNING, INFO, DEBUG, or ON. For compatibility
# reasons the setting "on" will default to "debug" level. Please note that the effective
# logging level is still determined by the global logging level of the actor system:
# for example debug level remoting events will be only logged if the system
# is running with debug level logging.
# Failures to deserialize received messages also fall under this flag.
log-remote-lifecycle-events= on

# Logging of message types with payload size in bytes larger than
# this value. Maximum detected size per message type is logged once,
# with an increase threshold of 10%.
# By default this feature is turned off. Activate it by setting the property to
# a value in bytes, such as 1000b. Note that for all messages larger than this
# limit there will be extra performance and scalability cost.
log-frame-size-exceeding = off

# Log warning if the number of messages in the backoff buffer in the endpoint
# writer exceeds this limit. It can be disabled by setting the value to off.
log-buffer-size-exceeding = 50000

### Failure detection and recovery

# Settings for the failure detector to monitor connections.
# For TCP it is not important to have fast failure detection, since

```

```

# most connection failures are captured by TCP itself.
# The default DeadlineFailureDetector will trigger if there are no heartbeats within
# the duration heartbeat-interval + acceptable-heartbeat-pause, i.e. 20 seconds
# with the default settings.
transport-failure-detector {

    # FQCN of the failure detector implementation.
    # It must implement akka.remote.FailureDetector and have
    # a public constructor with a com.typesafe.config.Config and
    # akka.actor.EventStream parameter.
    implementation-class = "akka.remote.DeadlineFailureDetector"

    # How often keep-alive heartbeat messages should be sent to each connection.
    heartbeat-interval = 4 s

    # Number of potentially lost/delayed heartbeats that will be
    # accepted before considering it to be an anomaly.
    # A margin to the 'heartbeat-interval' is important to be able to survive sudden,
    # occasional, pauses in heartbeat arrivals, due to for example garbage collect or
    # network drop.
    acceptable-heartbeat-pause = 16 s
}

# Settings for the Phi accrual failure detector (http://ddg.jaist.ac.jp/pub/HDY+04.pdf
# [Hayashibara et al]) used for remote death watch.
# The default PhiAccrualFailureDetector will trigger if there are no heartbeats within
# the duration heartbeat-interval + acceptable-heartbeat-pause + threshold_adjustment,
# i.e. around 12.5 seconds with default settings.
watch-failure-detector {

    # FQCN of the failure detector implementation.
    # It must implement akka.remote.FailureDetector and have
    # a public constructor with a com.typesafe.config.Config and
    # akka.actor.EventStream parameter.
    implementation-class = "akka.remote.PhiAccrualFailureDetector"

    # How often keep-alive heartbeat messages should be sent to each connection.
    heartbeat-interval = 1 s

    # Defines the failure detector threshold.
    # A low threshold is prone to generate many wrong suspicions but ensures
    # a quick detection in the event of a real crash. Conversely, a high
    # threshold generates fewer mistakes but needs more time to detect
    # actual crashes.
    threshold = 10.0

    # Number of the samples of inter-heartbeat arrival times to adaptively
    # calculate the failure timeout for connections.
    max-sample-size = 200

    # Minimum standard deviation to use for the normal distribution in
    # AccrualFailureDetector. Too low standard deviation might result in
    # too much sensitivity for sudden, but normal, deviations in heartbeat
    # inter arrival times.
    min-std-deviation = 100 ms

    # Number of potentially lost/delayed heartbeats that will be
    # accepted before considering it to be an anomaly.
    # This margin is important to be able to survive sudden, occasional,
    # pauses in heartbeat arrivals, due to for example garbage collect or
    # network drop.
    acceptable-heartbeat-pause = 10 s
}

```

```

# How often to check for nodes marked as unreachable by the failure
# detector
unreachable-nodes-reaper-interval = 1s

# After the heartbeat request has been sent the first failure detection
# will start after this period, even though no heartbeat message has
# been received.
expected-response-after = 1 s

}

# After failed to establish an outbound connection, the remoting will mark the
# address as failed. This configuration option controls how much time should
# be elapsed before reattempting a new connection. While the address is
# gated, all messages sent to the address are delivered to dead-letters.
# Since this setting limits the rate of reconnects setting it to a
# very short interval (i.e. less than a second) may result in a storm of
# reconnect attempts.
retry-gate-closed-for = 5 s

# After catastrophic communication failures that result in the loss of system
# messages or after the remote DeathWatch triggers the remote system gets
# quarantined to prevent inconsistent behavior.
# This setting controls how long the Quarantine marker will be kept around
# before being removed to avoid long-term memory leaks.
# WARNING: DO NOT change this to a small value to re-enable communication with
# quarantined nodes. Such feature is not supported and any behavior between
# the affected systems after lifting the quarantine is undefined.
prune-quarantine-marker-after = 5 d

# This setting defines the maximum number of unacknowledged system messages
# allowed for a remote system. If this limit is reached the remote system is
# declared to be dead and its UID marked as tainted.
system-message-buffer-size = 20000

# This setting defines the maximum idle time after an individual
# acknowledgement for system messages is sent. System message delivery
# is guaranteed by explicit acknowledgement messages. These acks are
# piggybacked on ordinary traffic messages. If no traffic is detected
# during the time period configured here, the remoting will send out
# an individual ack.
system-message-ack-piggyback-timeout = 0.3 s

# This setting defines the time after internal management signals
# between actors (used for DeathWatch and supervision) that have not been
# explicitly acknowledged or negatively acknowledged are resent.
# Messages that were negatively acknowledged are always immediately
# resent.
resent-interval = 2 s

# Maximum number of unacknowledged system messages that will be resent
# each 'resent-interval'. If you watch many (> 1000) remote actors you can
# increase this value to for example 600, but a too large limit (e.g. 10000)
# may flood the connection and might cause false failure detection to trigger.
# Test such a configuration by watching all actors at the same time and stop
# all watched actors at the same time.
resent-limit = 200

# WARNING: this setting should not be not changed unless all of its consequences
# are properly understood which assumes experience with remoting internals
# or expert advice.
# This setting defines the time after redelivery attempts of internal management

```

```

# signals are stopped to a remote system that has been not confirmed to be alive by
# this system before.
initial-system-message-delivery-timeout = 3 m

### Transports and adapters

# List of the transport drivers that will be loaded by the remoting.
# A list of fully qualified config paths must be provided where
# the given configuration path contains a transport-class key
# pointing to an implementation class of the Transport interface.
# If multiple transports are provided, the address of the first
# one will be used as a default address.
enabled-transports = ["akka.remote.netty.tcp"]

# Transport drivers can be augmented with adapters by adding their
# name to the applied-adapters setting in the configuration of a
# transport. The available adapters should be configured in this
# section by providing a name, and the fully qualified name of
# their corresponding implementation. The class given here
# must implement akka.akka.remote.transport.TransportAdapterProvider
# and have public constructor without parameters.
adapters {
  gremlin = "akka.remote.transport.FailureInjectorProvider"
  trttl = "akka.remote.transport.ThrottlerProvider"
}

### Default configuration for the Netty based transport drivers

netty.tcp {
  # The class given here must implement the akka.remote.transport.Transport
  # interface and offer a public constructor which takes two arguments:
  # 1) akka.actor.ExtendedActorSystem
  # 2) com.typesafe.config.Config
  transport-class = "akka.remote.transport.netty.NettyTransport"

  # Transport drivers can be augmented with adapters by adding their
  # name to the applied-adapters list. The last adapter in the
  # list is the adapter immediately above the driver, while
  # the first one is the top of the stack below the standard
  # Akka protocol
  applied-adapters = []

  transport-protocol = tcp

  # The default remote server port clients should connect to.
  # Default is 2552 (AKKA), use 0 if you want a random available port
  # This port needs to be unique for each actor system on the same machine.
  port = 2552

  # The hostname or ip clients should connect to.
  # InetAddress.getLocalHost.getHostAddress is used if empty
  hostname = ""

  # Use this setting to bind a network interface to a different port
  # than remoting protocol expects messages at. This may be used
  # when running akka nodes in a separated networks (under NATs or docker containers).
  # Use 0 if you want a random available port. Examples:
  #
  # akka.remote.netty.tcp.port = 2552
  # akka.remote.netty.tcp.bind-port = 2553
  # Network interface will be bound to the 2553 port, but remoting protocol will
  # expect messages sent to port 2552.
  #

```

```

# akka.remote.netty.tcp.port = 0
# akka.remote.netty.tcp.bind-port = 0
# Network interface will be bound to a random port, and remoting protocol will
# expect messages sent to the bound port.
#
# akka.remote.netty.tcp.port = 2552
# akka.remote.netty.tcp.bind-port = 0
# Network interface will be bound to a random port, but remoting protocol will
# expect messages sent to port 2552.
#
# akka.remote.netty.tcp.port = 0
# akka.remote.netty.tcp.bind-port = 2553
# Network interface will be bound to the 2553 port, and remoting protocol will
# expect messages sent to the bound port.
#
# akka.remote.netty.tcp.port = 2552
# akka.remote.netty.tcp.bind-port = ""
# Network interface will be bound to the 2552 port, and remoting protocol will
# expect messages sent to the bound port.
#
# akka.remote.netty.tcp.port if empty
bind-port = ""

# Use this setting to bind a network interface to a different hostname or ip
# than remoting protocol expects messages at.
# Use "0.0.0.0" to bind to all interfaces.
# akka.remote.netty.tcp.hostname if empty
bind-hostname = ""

# Enables SSL support on this transport
enable-ssl = false

# Sets the connectTimeoutMillis of all outbound connections,
# i.e. how long a connect may take until it is timed out
connection-timeout = 15 s

# If set to "<id.of.dispatcher>" then the specified dispatcher
# will be used to accept inbound connections, and perform IO. If "" then
# dedicated threads will be used.
# Please note that the Netty driver only uses this configuration and does
# not read the "akka.remote.use-dispatcher" entry. Instead it has to be
# configured manually to point to the same dispatcher if needed.
use-dispatcher-for-io = ""

# Sets the high water mark for the in and outbound sockets,
# set to 0b for platform default
write-buffer-high-water-mark = 0b

# Sets the low water mark for the in and outbound sockets,
# set to 0b for platform default
write-buffer-low-water-mark = 0b

# Sets the send buffer size of the Sockets,
# set to 0b for platform default
send-buffer-size = 256000b

# Sets the receive buffer size of the Sockets,
# set to 0b for platform default
receive-buffer-size = 256000b

# Maximum message size the transport will accept, but at least
# 32000 bytes.
# Please note that UDP does not support arbitrary large datagrams,

```

```

# so this setting has to be chosen carefully when using UDP.
# Both send-buffer-size and receive-buffer-size settings has to
# be adjusted to be able to buffer messages of maximum size.
maximum-frame-size = 128000b

# Sets the size of the connection backlog
backlog = 4096

# Enables the TCP_NODELAY flag, i.e. disables Nagle's algorithm
tcp-nodelay = on

# Enables TCP Keepalive, subject to the O/S kernel's configuration
tcp-keepalive = on

# Enables SO_REUSEADDR, which determines when an ActorSystem can open
# the specified listen port (the meaning differs between *nix and Windows)
# Valid values are "on", "off" and "off-for-windows"
# due to the following Windows bug: http://bugs.sun.com/bugdatabase/view\_bug.do?bug\_id=4476
# "off-for-windows" of course means that it's "on" for all other platforms
tcp-reuse-addr = off-for-windows

# Used to configure the number of I/O worker threads on server sockets
server-socket-worker-pool {
  # Min number of threads to cap factor-based number to
  pool-size-min = 2

  # The pool size factor is used to determine thread pool size
  # using the following formula: ceil(available processors * factor).
  # Resulting size is then bounded by the pool-size-min and
  # pool-size-max values.
  pool-size-factor = 1.0

  # Max number of threads to cap factor-based number to
  pool-size-max = 2
}

# Used to configure the number of I/O worker threads on client sockets
client-socket-worker-pool {
  # Min number of threads to cap factor-based number to
  pool-size-min = 2

  # The pool size factor is used to determine thread pool size
  # using the following formula: ceil(available processors * factor).
  # Resulting size is then bounded by the pool-size-min and
  # pool-size-max values.
  pool-size-factor = 1.0

  # Max number of threads to cap factor-based number to
  pool-size-max = 2
}

}

netty.udp = ${akka.remote.netty.tcp}
netty.udp {
  transport-protocol = udp
}

netty.ssl = ${akka.remote.netty.tcp}
netty.ssl = {
  # Enable SSL/TLS encryption.
  # This must be enabled on both the client and server to work.

```

```

enable-ssl = true

security {
  # This is the Java Key Store used by the server connection
  key-store = "keystore"

  # This password is used for decrypting the key store
  key-store-password = "changeme"

  # This password is used for decrypting the key
  key-password = "changeme"

  # This is the Java Key Store used by the client connection
  trust-store = "truststore"

  # This password is used for decrypting the trust store
  trust-store-password = "changeme"

  # Protocol to use for SSL encryption, choose from:
  # Java 6 & 7:
  #   'SSLv3', 'TLSv1'
  # Java 7:
  #   'TLSv1.1', 'TLSv1.2'
  protocol = "TLSv1"

  # Example: ["TLS_RSA_WITH_AES_128_CBC_SHA", "TLS_RSA_WITH_AES_256_CBC_SHA"]
  # You need to install the JCE Unlimited Strength Jurisdiction Policy
  # Files to use AES 256.
  # More info here:
  # http://docs.oracle.com/javase/7/docs/technotes/guides/security/SunProviders.html#SunJCE
  enabled-algorithms = ["TLS_RSA_WITH_AES_128_CBC_SHA"]

  # There are three options, in increasing order of security:
  # "" or SecureRandom => (default)
  # "SHA1PRNG" => Can be slow because of blocking issues on Linux
  # "AES128CounterSecureRNG" => fastest startup and based on AES encryption
  # algorithm
  # "AES256CounterSecureRNG"
  #
  # The following are deprecated in Akka 2.4. They use one of 3 possible
  # seed sources, depending on availability: /dev/random, random.org and
  # SecureRandom (provided by Java)
  # "AES128CounterInetRNG"
  # "AES256CounterInetRNG" (Install JCE Unlimited Strength Jurisdiction
  # Policy Files first)
  # Setting a value here may require you to supply the appropriate cipher
  # suite (see enabled-algorithms section above)
  random-number-generator = ""
}
}

### Default configuration for the failure injector transport adapter

gremlin {
  # Enable debug logging of the failure injector transport adapter
  debug = off
}

### Default dispatcher for the remoting subsystem

default-remote-dispatcher {
  type = Dispatcher
  executor = "fork-join-executor"
}

```

```

fork-join-executor {
  # Min number of threads to cap factor-based parallelism number to
  parallelism-min = 2
  parallelism-max = 2
}

backoff-remote-dispatcher {
  type = Dispatcher
  executor = "fork-join-executor"
  fork-join-executor {
    # Min number of threads to cap factor-based parallelism number to
    parallelism-min = 2
    parallelism-max = 2
  }
}

}
}

```

akka-testkit

```

#####
# Akka Testkit Reference Config File #
#####

# This is the reference config file that contains all the default settings.
# Make your edits/overrides in your application.conf.

akka {
  test {
    # factor by which to scale timeouts during tests, e.g. to account for shared
    # build system load
    timefactor = 1.0

    # duration of EventFilter.intercept waits after the block is finished until
    # all required messages are received
    filter-leeway = 3s

    # duration to wait in expectMsg and friends outside of within() block
    # by default
    single-expect-default = 3s

    # The timeout that is added as an implicit by DefaultTimeout trait
    default-timeout = 5s

    calling-thread-dispatcher {
      type = akka.testkit.CallingThreadDispatcherConfigurator
    }
  }
}

```

akka-cluster-metrics ~~~~~

```

#####
# Akka Cluster Metrics Reference Config File #
#####

# This is the reference config file that contains all the default settings.
# Make your edits in your application.conf in order to override these settings.

```

```

# Sigar provisioning:
#
# User can provision sigar classes and native library in one of the following ways:
#
# 1) Use https://github.com/kamon-io/sigar-loader Kamon sigar-loader as a project dependency for
# Metrics extension will extract and load sigar library on demand with help of Kamon sigar provi
#
# 2) Use https://github.com/kamon-io/sigar-loader Kamon sigar-loader as java agent: `java -javaa
# Kamon sigar loader agent will extract and load sigar library during JVM start.
#
# 3) Place `sigar.jar` on the `classpath` and sigar native library for the o/s on the `java.libra
# User is required to manage both project dependency and library deployment manually.

# Cluster metrics extension.
# Provides periodic statistics collection and publication throughout the cluster.
akka.cluster.metrics {
  # Full path of dispatcher configuration key.
  # Use "" for default key `akka.actor.default-dispatcher`.
  dispatcher = ""
  # How long should any actor wait before starting the periodic tasks.
  periodic-tasks-initial-delay = 1s
  # Sigar native library extract location.
  # Use per-application-instance scoped location, such as program working directory.
  native-library-extract-folder = ${user.dir}/"native"
  # Metrics supervisor actor.
  supervisor {
    # Actor name. Example name space: /system/cluster-metrics
    name = "cluster-metrics"
    # Supervision strategy.
    strategy {
      #
      # FQCN of class providing `akka.actor.SupervisorStrategy`.
      # Must have a constructor with signature `(com.typesafe.config.Config)`.
      # Default metrics strategy provider is a configurable extension of `OneForOneStrategy`
      provider = "akka.cluster.metrics.ClusterMetricsStrategy"
      #
      # Configuration of the default strategy provider.
      # Replace with custom settings when overriding the provider.
      configuration = {
        # Log restart attempts.
        loggingEnabled = true
        # Child actor restart-on-failure window.
        withinTimeRange = 3s
        # Maximum number of restart attempts before child actor is stopped.
        maxNrOfRetries = 3
      }
    }
  }
}
# Metrics collector actor.
collector {
  # Enable or disable metrics collector for load-balancing nodes.
  # Metrics collection can also be controlled at runtime by sending control messages
  # to /system/cluster-metrics actor: `akka.cluster.metrics.{CollectionStartMessage,Collectio
  enabled = on
  # FQCN of the metrics collector implementation.
  # It must implement `akka.cluster.metrics.MetricsCollector` and
  # have public constructor with akka.actor.ActorSystem parameter.
  # Will try to load in the following order of priority:
  # 1) configured custom collector 2) internal `SigarMetricsCollector` 3) internal `JmxMetric
  provider = ""
  # Try all 3 available collector providers, or else fail on the configured custom collector.
  fallback = true
}

```

```

# How often metrics are sampled on a node.
# Shorter interval will collect the metrics more often.
# Also controls frequency of the metrics publication to the node system event bus.
sample-interval = 3s
# How often a node publishes metrics information to the other nodes in the cluster.
# Shorter interval will publish the metrics gossip more often.
gossip-interval = 3s
# How quickly the exponential weighting of past data is decayed compared to
# new data. Set lower to increase the bias toward newer values.
# The relevance of each data sample is halved for every passing half-life
# duration, i.e. after 4 times the half-life, a data sample's relevance is
# reduced to 6% of its original relevance. The initial relevance of a data
# sample is given by 1 - 0.5 ^ (collect-interval / half-life).
# See http://en.wikipedia.org/wiki/Moving_average#Exponential_moving_average
moving-average-half-life = 12s
}
}

# Cluster metrics extension serializers and routers.
akka.actor {
  # Protobuf serializer for remote cluster metrics messages.
  serializers {
    akka-cluster-metrics = "akka.cluster.metrics.protobuf.MessageSerializer"
  }
  # Interface binding for remote cluster metrics messages.
  serialization-bindings {
    "akka.cluster.metrics.ClusterMetricsMessage" = akka-cluster-metrics
  }
  # Globally unique metrics extension serializer identifier.
  serialization-identifiers {
    "akka.cluster.metrics.protobuf.MessageSerializer" = 10
  }
  # Provide routing of messages based on cluster metrics.
  router.type-mapping {
    cluster-metrics-adaptive-pool = "akka.cluster.metrics.AdaptiveLoadBalancingPool"
    cluster-metrics-adaptive-group = "akka.cluster.metrics.AdaptiveLoadBalancingGroup"
  }
}
}

```

akka-cluster-tools ~~~~~

```

#####
# Akka Cluster Tools Reference Config File #
#####

# This is the reference config file that contains all the default settings.
# Make your edits/overrides in your application.conf.

# //#pub-sub-ext-config
# Settings for the DistributedPubSub extension
akka.cluster.pub-sub {
  # Actor name of the mediator actor, /system/distributedPubSubMediator
  name = distributedPubSubMediator

  # Start the mediator on members tagged with this role.
  # All members are used if undefined or empty.
  role = ""

  # The routing logic to use for 'Send'
  # Possible values: random, round-robin, broadcast
  routing-logic = random

  # How often the DistributedPubSubMediator should send out gossip information

```

```

gossip-interval = 1s

# Removed entries are pruned after this duration
removed-time-to-live = 120s

# Maximum number of elements to transfer in one message when synchronizing the registries.
# Next chunk will be transferred in next round of gossip.
max-delta-elements = 3000

# The id of the dispatcher to use for DistributedPubSubMediator actors.
# If not specified default dispatcher is used.
# If specified you need to define the settings of the actual dispatcher.
use-dispatcher = ""

}
# //#pub-sub-ext-config

# Protobuf serializer for cluster DistributedPubSubMediator messages
akka.actor {
  serializers {
    akka-pubsub = "akka.cluster.pubsub.protobuf.DistributedPubSubMessageSerializer"
  }
  serialization-bindings {
    "akka.cluster.pubsub.DistributedPubSubMessage" = akka-pubsub
  }
  serialization-identifiers {
    "akka.cluster.pubsub.protobuf.DistributedPubSubMessageSerializer" = 9
  }
}

# //#receptionist-ext-config
# Settings for the ClusterClientReceptionist extension
akka.cluster.client.receptionist {
  # Actor name of the ClusterReceptionist actor, /system/receptionist
  name = receptionist

  # Start the receptionist on members tagged with this role.
  # All members are used if undefined or empty.
  role = ""

  # The receptionist will send this number of contact points to the client
  number-of-contacts = 3

  # The actor that tunnel response messages to the client will be stopped
  # after this time of inactivity.
  response-tunnel-receive-timeout = 30s

  # The id of the dispatcher to use for ClusterReceptionist actors.
  # If not specified default dispatcher is used.
  # If specified you need to define the settings of the actual dispatcher.
  use-dispatcher = ""
}
# //#receptionist-ext-config

# //#cluster-client-config
# Settings for the ClusterClient
akka.cluster.client {
  # Actor paths of the ClusterReceptionist actors on the servers (cluster nodes)
  # that the client will try to contact initially. It is mandatory to specify
  # at least one initial contact.
  # Comma separated full actor paths defined by a string on the form of
  # "akka.tcp://system@hostname:port/system/receptionist"

```

```

initial-contacts = []

# Interval at which the client retries to establish contact with one of
# ClusterReceptionist on the servers (cluster nodes)
establishing-get-contacts-interval = 3s

# Interval at which the client will ask the ClusterReceptionist for
# new contact points to be used for next reconnect.
refresh-contacts-interval = 60s

# How often failure detection heartbeat messages should be sent
heartbeat-interval = 2s

# Number of potentially lost/delayed heartbeats that will be
# accepted before considering it to be an anomaly.
# The ClusterClient is using the akka.remote.DeadlineFailureDetector, which
# will trigger if there are no heartbeats within the duration
# heartbeat-interval + acceptable-heartbeat-pause, i.e. 15 seconds with
# the default settings.
acceptable-heartbeat-pause = 13s

# If connection to the receptionist is not established the client will buffer
# this number of messages and deliver them the connection is established.
# When the buffer is full old messages will be dropped when new messages are sent
# via the client. Use 0 to disable buffering, i.e. messages will be dropped
# immediately if the location of the singleton is unknown.
# Maximum allowed buffer size is 10000.
buffer-size = 1000
}
# //#cluster-client-config

# Protobuf serializer for ClusterClient messages
akka.actor {
  serializers {
    akka-cluster-client = "akka.cluster.client.protobuf.ClusterClientMessageSerializer"
  }
  serialization-bindings {
    "akka.cluster.client.ClusterClientMessage" = akka-cluster-client
  }
  serialization-identifiers {
    "akka.cluster.client.protobuf.ClusterClientMessageSerializer" = 15
  }
}

# //#singleton-config
akka.cluster.singleton {
  # The actor name of the child singleton actor.
  singleton-name = "singleton"

  # Singleton among the nodes tagged with specified role.
  # If the role is not specified it's a singleton among all nodes in the cluster.
  role = ""

  # When a node is becoming oldest it sends hand-over request to previous oldest,
  # that might be leaving the cluster. This is retried with this interval until
  # the previous oldest confirms that the hand over has started or the previous
  # oldest member is removed from the cluster (+ akka.cluster.down-removal-margin).
  hand-over-retry-interval = 1s

  # The number of retries are derived from hand-over-retry-interval and
  # akka.cluster.down-removal-margin (or ClusterSingletonManagerSettings.removalMargin),
  # but it will never be less than this property.
  min-number-of-hand-over-retries = 10
}

```

```

}
# //#singleton-config

# //#singleton-proxy-config
akka.cluster.singleton-proxy {
  # The actor name of the singleton actor that is started by the ClusterSingletonManager
  singleton-name = ${akka.cluster.singleton.singleton-name}

  # The role of the cluster nodes where the singleton can be deployed.
  # If the role is not specified then any node will do.
  role = ""

  # Interval at which the proxy will try to resolve the singleton instance.
  singleton-identification-interval = 1s

  # If the location of the singleton is unknown the proxy will buffer this
  # number of messages and deliver them when the singleton is identified.
  # When the buffer is full old messages will be dropped when new messages are
  # sent via the proxy.
  # Use 0 to disable buffering, i.e. messages will be dropped immediately if
  # the location of the singleton is unknown.
  # Maximum allowed buffer size is 10000.
  buffer-size = 1000
}
# //#singleton-proxy-config

# Serializer for cluster ClusterSingleton messages
akka.actor {
  serializers {
    akka-singleton = "akka.cluster.singleton.protobuf.ClusterSingletonMessageSerializer"
  }
  serialization-bindings {
    "akka.cluster.singleton.ClusterSingletonMessage" = akka-singleton
  }
  serialization-identifiers {
    "akka.cluster.singleton.protobuf.ClusterSingletonMessageSerializer" = 14
  }
}

```

akka-cluster-sharding ~~~~~

```

#####
# Akka Cluster Sharding Reference Config File #
#####

# This is the reference config file that contains all the default settings.
# Make your edits/overrides in your application.conf.

# //#sharding-ext-config
# Settings for the ClusterShardingExtension
akka.cluster.sharding {

  # The extension creates a top level actor with this name in top level system scope,
  # e.g. '/system/sharding'
  guardian-name = sharding

  # Specifies that entities runs on cluster nodes with a specific role.
  # If the role is not specified (or empty) all nodes in the cluster are used.
  role = ""

  # When this is set to 'on' the active entity actors will automatically be restarted
  # upon Shard restart. i.e. if the Shard is started on a different ShardRegion

```

```

# due to rebalance or crash.
remember-entities = off

# If the coordinator can't store state changes it will be stopped
# and started again after this duration, with an exponential back-off
# of up to 5 times this duration.
coordinator-failure-backoff = 5 s

# The ShardRegion retries registration and shard location requests to the
# ShardCoordinator with this interval if it does not reply.
retry-interval = 2 s

# Maximum number of messages that are buffered by a ShardRegion actor.
buffer-size = 100000

# Timeout of the shard rebalancing process.
handoff-timeout = 60 s

# Time given to a region to acknowledge it's hosting a shard.
shard-start-timeout = 10 s

# If the shard is remembering entities and can't store state changes
# will be stopped and then started again after this duration. Any messages
# sent to an affected entity may be lost in this process.
shard-failure-backoff = 10 s

# If the shard is remembering entities and an entity stops itself without
# using passivate. The entity will be restarted after this duration or when
# the next message for it is received, which ever occurs first.
entity-restart-backoff = 10 s

# Rebalance check is performed periodically with this interval.
rebalance-interval = 10 s

# Absolute path to the journal plugin configuration entity that is to be
# used for the internal persistence of ClusterSharding. If not defined
# the default journal plugin is used. Note that this is not related to
# persistence used by the entity actors.
journal-plugin-id = ""

# Absolute path to the snapshot plugin configuration entity that is to be
# used for the internal persistence of ClusterSharding. If not defined
# the default snapshot plugin is used. Note that this is not related to
# persistence used by the entity actors.
snapshot-plugin-id = ""

# Parameter which determines how the coordinator will be store a state
# valid values either "persistence" or "ddata"
# The "ddata" mode is experimental, since it depends on the experimental
# module akka-distributed-data-experimental.
state-store-mode = "persistence"

# The shard saves persistent snapshots after this number of persistent
# events. Snapshots are used to reduce recovery times.
snapshot-after = 1000

# Setting for the default shard allocation strategy
least-shard-allocation-strategy {
  # Threshold of how large the difference between most and least number of
  # allocated shards must be to begin the rebalancing.
  rebalance-threshold = 10

  # The number of ongoing rebalancing processes is limited to this number.

```

```

    max-simultaneous-rebalance = 3
  }

  # Timeout of waiting the initial distributed state (an initial state will be queried again if t
  # works only for state-store-mode = "ddata"
  waiting-for-state-timeout = 5 s

  # Timeout of waiting for update the distributed state (update will be retried if the timeout hap
  # works only for state-store-mode = "ddata"
  updating-state-timeout = 5 s

  # Settings for the coordinator singleton. Same layout as akka.cluster.singleton.
  coordinator-singleton = ${akka.cluster.singleton}

  # The id of the dispatcher to use for ClusterSharding actors.
  # If not specified default dispatcher is used.
  # If specified you need to define the settings of the actual dispatcher.
  # This dispatcher for the entity actors is defined by the user provided
  # Props, i.e. this dispatcher is not used for the entity actors.
  use-dispatcher = ""
}
# //#sharding-ext-config

# Protobuf serializer for Cluster Sharding messages
akka.actor {
  serializers {
    akka-sharding = "akka.cluster.sharding.protobuf.ClusterShardingMessageSerializer"
  }
  serialization-bindings {
    "akka.cluster.sharding.ClusterShardingSerializable" = akka-sharding
  }
  serialization-identifiers {
    "akka.cluster.sharding.protobuf.ClusterShardingMessageSerializer" = 13
  }
}

```

akka-distributed-data ~~~~~

```

#####
# Akka Distributed DataReference Config File #
#####

# This is the reference config file that contains all the default settings.
# Make your edits/overrides in your application.conf.

#//#distributed-data
# Settings for the DistributedData extension
akka.cluster.distributed-data {
  # Actor name of the Replicator actor, /system/ddataReplicator
  name = ddataReplicator

  # Replicas are running on members tagged with this role.
  # All members are used if undefined or empty.
  role = ""

  # How often the Replicator should send out gossip information
  gossip-interval = 2 s

  # How often the subscribers will be notified of changes, if any
  notify-subscribers-interval = 500 ms
}

```

```
# Maximum number of entries to transfer in one gossip message when synchronizing
# the replicas. Next chunk will be transferred in next round of gossip.
max-delta-elements = 1000

# The id of the dispatcher to use for Replicator actors. If not specified
# default dispatcher is used.
# If specified you need to define the settings of the actual dispatcher.
use-dispatcher = ""

# How often the Replicator checks for pruning of data associated with
# removed cluster nodes.
pruning-interval = 30 s

# How long time it takes (worst case) to spread the data to all other replica nodes.
# This is used when initiating and completing the pruning process of data associated
# with removed cluster nodes. The time measurement is stopped when any replica is
# unreachable, so it should be configured to worst case in a healthy cluster.
max-pruning-dissemination = 60 s

}
###distributed-data

# Protobuf serializer for cluster DistributedData messages
akka.actor {
  serializers {
    akka-data-replication = "akka.cluster.ddata.protobuf.ReplicatorMessageSerializer"
    akka-replicated-data = "akka.cluster.ddata.protobuf.ReplicatedDataSerializer"
  }
  serialization-bindings {
    "akka.cluster.ddata.Replicator$ReplicatorMessage" = akka-data-replication
    "akka.cluster.ddata.ReplicatedDataSerialization" = akka-replicated-data
  }
  serialization-identifiers {
    "akka.cluster.ddata.protobuf.ReplicatedDataSerializer" = 11
    "akka.cluster.ddata.protobuf.ReplicatorMessageSerializer" = 12
  }
}
```

ACTORS

3.1 Actors

The *Actor Model* provides a higher level of abstraction for writing concurrent and distributed systems. It alleviates the developer from having to deal with explicit locking and thread management, making it easier to write correct concurrent and parallel systems. Actors were defined in the 1973 paper by Carl Hewitt but have been popularized by the Erlang language, and used for example at Ericsson with great success to build highly concurrent and reliable telecom systems.

The API of Akka's Actors is similar to Scala Actors which has borrowed some of its syntax from Erlang.

3.1.1 Creating Actors

Note: Since Akka enforces parental supervision every actor is supervised and (potentially) the supervisor of its children, it is advisable that you familiarize yourself with *Actor Systems* and *Supervision and Monitoring* and it may also help to read *Actor References, Paths and Addresses*.

Defining an Actor class

Actors in Java are implemented by extending the `UntypedActor` class and implementing the `onReceive` method. This method takes the message as a parameter.

Here is an example:

```
import akka.actor.UntypedActor;
import akka.event.Logging;
import akka.event.LoggingAdapter;

public class MyUntypedActor extends UntypedActor {
    LoggingAdapter log = Logging.getLogger(getContext().system(), this);

    public void onReceive(Object message) throws Exception {
        if (message instanceof String) {
            log.info("Received String message: {}", message);
            getSender().tell(message, getSelf());
        } else {
            unhandled(message);
        }
    }
}
```

Props

Props is a configuration class to specify options for the creation of actors, think of it as an immutable and thus freely shareable recipe for creating an actor including associated deployment information (e.g. which dispatcher to use, see more below). Here are some examples of how to create a Props instance.

```
import akka.actor.Props;
import akka.japi.Creator;

static class MyActorC implements Creator<MyActor> {
    @Override public MyActor create() {
        return new MyActor("...");
    }
}

Props props1 = Props.create(MyUntypedActor.class);
Props props2 = Props.create(MyActor.class, "...");
Props props3 = Props.create(new MyActorC());
```

The second line shows how to pass constructor arguments to the Actor being created. The presence of a matching constructor is verified during construction of the Props object, resulting in an `IllegalArgumentException` if no or multiple matching constructors are found.

The third line demonstrates the use of a Creator. The creator class must be static, which is verified during Props construction. The type parameter's upper bound is used to determine the produced actor class, falling back to Actor if fully erased. An example of a parametric factory could be:

```
static class ParametricCreator<T extends MyActor> implements Creator<T> {
    @Override public T create() {
        // ... fabricate actor here
    }
}
```

Note: In order for mailbox requirements—like using a deque-based mailbox for actors using the stash—to be picked up, the actor type needs to be known before creating it, which is what the Creator type argument allows. Therefore make sure to use the specific type for your actors wherever possible.

Recommended Practices

It is a good idea to provide static factory methods on the `UntypedActor` which help keeping the creation of suitable Props as close to the actor definition as possible. This also allows usage of the Creator-based methods which statically verify that the used constructor actually exists instead relying only on a runtime check.

```
public class DemoActor extends UntypedActor {

    /**
     * Create Props for an actor of this type.
     * @param magicNumber The magic number to be passed to this actor's constructor.
     * @return a Props for creating this actor, which can then be further configured
     *         (e.g. calling .withDispatcher() on it)
     */
    public static Props props(final int magicNumber) {
        return Props.create(new Creator<DemoActor>() {
            private static final long serialVersionUID = 1L;

            @Override
            public DemoActor create() throws Exception {
                return new DemoActor(magicNumber);
            }
        });
    }
}
```

```

}

final int magicNumber;

public DemoActor(int magicNumber) {
    this.magicNumber = magicNumber;
}

@Override
public void onReceive(Object msg) {
    // some behavior here
}

}

system.actorOf(DemoActor.props(42), "demo");

```

Another good practice is to declare what messages an Actor can receive as close to the actor definition as possible (e.g. as static classes inside the Actor or using other suitable class), which makes it easier to know what it can receive.

```

public class DemoMessagesActor extends UntypedActor {

    static public class Greeting {
        private final String from;

        public Greeting(String from) {
            this.from = from;
        }

        public String getGreeter() {
            return from;
        }
    }

    public void onReceive(Object message) throws Exception {
        if (message instanceof Greeting) {
            getSender().tell("Hello " + ((Greeting) message).getGreeter(), getSelf());
        } else {
            unhandled(message);
        }
    }
}

```

Creating Actors with Props

Actors are created by passing a Props instance into the actorOf factory method which is available on ActorSystem and ActorContext.

```

import akka.actor.ActorRef;
import akka.actor.ActorSystem;

// ActorSystem is a heavy object: create only one per application
final ActorSystem system = ActorSystem.create("MySystem");
final ActorRef myActor = system.actorOf(Props.create(MyUntypedActor.class),
    "myactor");

```

Using the ActorSystem will create top-level actors, supervised by the actor system's provided guardian actor, while using an actor's context will create a child actor.

```

class A extends UntypedActor {
    final ActorRef child =
        getContext().actorOf(Props.create(MyUntypedActor.class), "myChild");
}

```

```
// plus some behavior ...
}
```

It is recommended to create a hierarchy of children, grand-children and so on such that it fits the logical failure-handling structure of the application, see [Actor Systems](#).

The call to `actorOf` returns an instance of `ActorRef`. This is a handle to the actor instance and the only way to interact with it. The `ActorRef` is immutable and has a one to one relationship with the Actor it represents. The `ActorRef` is also serializable and network-aware. This means that you can serialize it, send it over the wire and use it on a remote host and it will still be representing the same Actor on the original node, across the network.

The name parameter is optional, but you should preferably name your actors, since that is used in log messages and for identifying actors. The name must not be empty or start with \$, but it may contain URL encoded characters (eg. %20 for a blank space). If the given name is already in use by another child to the same parent an `InvalidActorNameException` is thrown.

Actors are automatically started asynchronously when created.

Dependency Injection

If your `UntypedActor` has a constructor that takes parameters then those need to be part of the `Props` as well, as described above. But there are cases when a factory method must be used, for example when the actual constructor arguments are determined by a dependency injection framework.

```
import akka.actor.Actor;
import akka.actor.IndirectActorProducer;

class DependencyInjector implements IndirectActorProducer {
    final Object applicationContext;
    final String beanName;

    public DependencyInjector(Object applicationContext, String beanName) {
        this.applicationContext = applicationContext;
        this.beanName = beanName;
    }

    @Override
    public Class<? extends Actor> actorClass() {
        return MyActor.class;
    }

    @Override
    public MyActor produce() {
        MyActor result;
        // obtain fresh Actor instance from DI framework ...
        return result;
    }
}

final ActorRef myActor = getContext().actorOf(
    Props.create(DependencyInjector.class, applicationContext, "MyActor"),
    "myactor3");
```

Warning: You might be tempted at times to offer an `IndirectActorProducer` which always returns the same instance, e.g. by using a static field. This is not supported, as it goes against the meaning of an actor restart, which is described here: [What Restarting Means](#).
When using a dependency injection framework, actor beans *MUST NOT* have singleton scope.

Techniques for dependency injection and integration with dependency injection frameworks are described in more depth in the [Using Akka with Dependency Injection](#) guideline and the [Akka Java Spring](#) tutorial in `Typesafe Activator`.

The Inbox

When writing code outside of actors which shall communicate with actors, the `ask` pattern can be a solution (see below), but there are two things it cannot do: receiving multiple replies (e.g. by subscribing an `ActorRef` to a notification service) and watching other actors' lifecycle. For these purposes there is the `Inbox` class:

```
final Inbox inbox = Inbox.create(system);
inbox.send(target, "hello");
try {
    assert inbox.receive(Duration.create(1, TimeUnit.SECONDS)).equals("world");
} catch (java.util.concurrent.TimeoutException e) {
    // timeout
}
```

The `send` method wraps a normal `tell` and supplies the internal actor's reference as the sender. This allows the reply to be received on the last line. Watching an actor is quite simple as well:

```
final Inbox inbox = Inbox.create(system);
inbox.watch(target);
target.tell(PoisonPill.getInstance(), ActorRef.noSender());
try {
    assert inbox.receive(Duration.create(1, TimeUnit.SECONDS)) instanceof Terminated;
} catch (java.util.concurrent.TimeoutException e) {
    // timeout
}
```

3.1.2 UntypedActor API

The `UntypedActor` class defines only one abstract method, the above mentioned `onReceive(Object message)`, which implements the behavior of the actor.

If the current actor behavior does not match a received message, it's recommended that you call the `unhandled` method, which by default publishes a new `akka.actor.UnhandledMessage(message, sender, recipient)` on the actor system's event stream (set configuration item `akka.actor.debug.unhandled` to `on` to have them converted into actual `Debug` messages).

In addition, it offers:

- `getSelf` reference to the `ActorRef` of the actor
- `getSender` reference sender `Actor` of the last received message, typically used as described in [Reply to messages](#)
- `supervisorStrategy` user overridable definition the strategy to use for supervising child actors

This strategy is typically declared inside the actor in order to have access to the actor's internal state within the decider function: since failure is communicated as a message sent to the supervisor and processed like other messages (albeit outside of the normal behavior), all values and variables within the actor are available, as is the `getSender()` reference (which will be the immediate child reporting the failure; if the original failure occurred within a distant descendant it is still reported one level up at a time).

- `getContext` exposes contextual information for the actor and the current message, such as:
 - factory methods to create child actors (`actorOf`)
 - system that the actor belongs to
 - parent supervisor
 - supervised children
 - lifecycle monitoring
 - hotswap behavior stack as described in [HotSwap](#)

The remaining visible methods are user-overridable life-cycle hooks which are described in the following:

```

public void preStart() {
}

public void preRestart(Throwable reason, scala.Option<Object> message) {
  for (ActorRef each : getContext().getChildren()) {
    getContext().unwatch(each);
    getContext().stop(each);
  }
  postStop();
}


public void postRestart(Throwable reason) {
  preStart();
}

public void postStop() {
}

```

The implementations shown above are the defaults provided by the `UntypedActor` class.

Actor Lifecycle



A path in an actor system represents a “place” which might be occupied by a living actor. Initially (apart from system initialized actors) a path is empty. When `actorOf()` is called it assigns an *incarnation* of the actor described by the passed `Props` to the given path. An actor incarnation is identified by the path *and* a *UID*. A restart only swaps the `Actor` instance defined by the `Props` but the incarnation and hence the *UID* remains the same.

The lifecycle of an incarnation ends when the actor is stopped. At that point the appropriate lifecycle events are called and watching actors are notified of the termination. After the incarnation is stopped, the path can be reused

again by creating an actor with `actorOf()`. In this case the name of the new incarnation will be the same as the previous one but the UIDs will differ. An actor can be stopped by the actor itself, another actor or the `ActorSystem` (see *Stopping actors*).

An `ActorRef` always represents an incarnation (path and UID) not just a given path. Therefore if an actor is stopped and a new one with the same name is created an `ActorRef` of the old incarnation will not point to the new one.

`ActorSelection` on the other hand points to the path (or multiple paths if wildcards are used) and is completely oblivious to which incarnation is currently occupying it. `ActorSelection` cannot be watched for this reason. It is possible to resolve the current incarnation's `ActorRef` living under the path by sending an `Identify` message to the `ActorSelection` which will be replied to with an `ActorIdentity` containing the correct reference (see *Identifying Actors via Actor Selection*). This can also be done with the `resolveOne` method of the `ActorSelection`, which returns a `Future` of the matching `ActorRef`.

Lifecycle Monitoring aka DeathWatch

In order to be notified when another actor terminates (i.e. stops permanently, not temporary failure and restart), an actor may register itself for reception of the `Terminated` message dispatched by the other actor upon termination (see *Stopping Actors*). This service is provided by the `DeathWatch` component of the actor system.

Registering a monitor is easy (see fourth line, the rest is for demonstrating the whole functionality):

```
import akka.actor.Terminated;

public class WatchActor extends UntypedActor {
    final ActorRef child = this.getContext().actorOf(Props.empty(), "child");
    {
        this.getContext().watch(child); // <-- the only call needed for registration
    }
    ActorRef lastSender = getContext().system().deadLetters();

    @Override
    public void onReceive(Object message) {
        if (message.equals("kill")) {
            getContext().stop(child);
            lastSender = getSender();
        } else if (message instanceof Terminated) {
            final Terminated t = (Terminated) message;
            if (t.getActor() == child) {
                lastSender.tell("finished", getSelf());
            }
        } else {
            unhandled(message);
        }
    }
}
```

It should be noted that the `Terminated` message is generated independent of the order in which registration and termination occur. In particular, the watching actor will receive a `Terminated` message even if the watched actor has already been terminated at the time of registration.

Registering multiple times does not necessarily lead to multiple messages being generated, but there is no guarantee that only exactly one such message is received: if termination of the watched actor has generated and queued the message, and another registration is done before this message has been processed, then a second message will be queued, because registering for monitoring of an already terminated actor leads to the immediate generation of the `Terminated` message.

It is also possible to deregister from watching another actor's liveliness using `getContext().unwatch(target)`. This works even if the `Terminated` message has already been enqueued in the mailbox; after calling `unwatch` no `Terminated` message for that actor will be processed anymore.

Start Hook

Right after starting the actor, its `preStart` method is invoked.

```
@Override
public void preStart() {
    child = getContext().actorOf(Props.empty());
}
```

This method is called when the actor is first created. During restarts it is called by the default implementation of `postRestart`, which means that by overriding that method you can choose whether the initialization code in this method is called only exactly once for this actor or for every restart. Initialization code which is part of the actor's constructor will always be called when an instance of the actor class is created, which happens at every restart.

Restart Hooks

All actors are supervised, i.e. linked to another actor with a fault handling strategy. Actors may be restarted in case an exception is thrown while processing a message (see *Supervision and Monitoring*). This restart involves the hooks mentioned above:

1. The old actor is informed by calling `preRestart` with the exception which caused the restart and the message which triggered that exception; the latter may be `None` if the restart was not caused by processing a message, e.g. when a supervisor does not trap the exception and is restarted in turn by its supervisor, or if an actor is restarted due to a sibling's failure. If the message is available, then that message's sender is also accessible in the usual way (i.e. by calling `getSender()`).

This method is the best place for cleaning up, preparing hand-over to the fresh actor instance, etc. By default it stops all children and calls `postStop`.

2. The initial factory from the `actorOf` call is used to produce the fresh instance.
3. The new actor's `postRestart` method is invoked with the exception which caused the restart. By default the `preStart` is called, just as in the normal start-up case.

An actor restart replaces only the actual actor object; the contents of the mailbox is unaffected by the restart, so processing of messages will resume after the `postRestart` hook returns. The message that triggered the exception will not be received again. Any message sent to an actor while it is being restarted will be queued to its mailbox as usual.

Warning: Be aware that the ordering of failure notifications relative to user messages is not deterministic. In particular, a parent might restart its child before it has processed the last messages sent by the child before the failure. See *Discussion: Message Ordering* for details.

Stop Hook

After stopping an actor, its `postStop` hook is called, which may be used e.g. for deregistering this actor from other services. This hook is guaranteed to run after message queuing has been disabled for this actor, i.e. messages sent to a stopped actor will be redirected to the `deadLetters` of the `ActorSystem`.

3.1.3 Identifying Actors via Actor Selection

As described in *Actor References, Paths and Addresses*, each actor has a unique logical path, which is obtained by following the chain of actors from child to parent until reaching the root of the actor system, and it has a physical path, which may differ if the supervision chain includes any remote supervisors. These paths are used by the system to look up actors, e.g. when a remote message is received and the recipient is searched, but they are also useful more directly: actors may look up other actors by specifying absolute or relative paths—logical or physical—and receive back an `ActorSelection` with the result:

```
// will look up this absolute path
getContext().actorSelection("/user/serviceA/actor");
// will look up sibling beneath same supervisor
getContext().actorSelection("../joe");
```

Note: It is always preferable to communicate with other Actors using their ActorRef instead of relying upon ActorSelection. Exceptions are

- sending messages using the *At-Least-Once Delivery* facility
- initiating first contact with a remote system

In all other cases ActorRefs can be provided during Actor creation or initialization, passing them from parent to child or introducing Actors by sending their ActorRefs to other Actors within messages.

The supplied path is parsed as a `java.net.URI`, which basically means that it is split on `/` into path elements. If the path starts with `/`, it is absolute and the look-up starts at the root guardian (which is the parent of `/user`); otherwise it starts at the current actor. If a path element equals `..`, the look-up will take a step “up” towards the supervisor of the currently traversed actor, otherwise it will step “down” to the named child. It should be noted that the `..` in actor paths here always means the logical structure, i.e. the supervisor.

The path elements of an actor selection may contain wildcard patterns allowing for broadcasting of messages to that section:

```
// will look all children to serviceB with names starting with worker
getContext().actorSelection("/user/serviceB/worker*");
// will look up all siblings beneath same supervisor
getContext().actorSelection("../*");
```

Messages can be sent via the ActorSelection and the path of the ActorSelection is looked up when delivering each message. If the selection does not match any actors the message will be dropped.

To acquire an ActorRef for an ActorSelection you need to send a message to the selection and use the `getSender` reference of the reply from the actor. There is a built-in `Identify` message that all Actors will understand and automatically reply to with a `ActorIdentity` message containing the ActorRef. This message is handled specially by the actors which are traversed in the sense that if a concrete name lookup fails (i.e. a non-wildcard path element does not correspond to a live actor) then a negative result is generated. Please note that this does not mean that delivery of that reply is guaranteed, it still is a normal message.

```
import akka.actor.ActorIdentity;
import akka.actor.ActorSelection;
import akka.actor.Identify;

public class Follower extends UntypedActor {
    final String identifyId = "1";
    {
        ActorSelection selection =
            getContext().actorSelection("/user/another");
        selection.tell(new Identify(identifyId), getSelf());
    }
    ActorRef another;

    final ActorRef probe;
    public Follower(ActorRef probe) {
        this.probe = probe;
    }

    @Override
    public void onReceive(Object message) {
        if (message instanceof ActorIdentity) {
            ActorIdentity identity = (ActorIdentity) message;
            if (identity.correlationId().equals(identifyId)) {
```

```

ActorRef ref = identity.getRef();
if (ref == null)
    getContext().stop(getSelf());
else {
    another = ref;
    getContext().watch(another);
    probe.tell(ref, getSelf());
}
}
} else if (message instanceof Terminated) {
    final Terminated t = (Terminated) message;
    if (t.getActor().equals(another)) {
        getContext().stop(getSelf());
    }
} else {
    unhandled(message);
}
}
}
}

```

You can also acquire an ActorRef for an ActorSelection with the `resolveOne` method of the ActorSelection. It returns a Future of the matching ActorRef if such an actor exists. It is completed with failure `[[akka.actor.ActorNotFound]]` if no such actor exists or the identification didn't complete within the supplied *timeout*.

Remote actor addresses may also be looked up, if *remoting* is enabled:

```
getContext().actorSelection("akka.tcp://app@otherhost:1234/user/serviceB");
```

An example demonstrating remote actor look-up is given in [Remoting Sample](#).

3.1.4 Messages and immutability

IMPORTANT: Messages can be any kind of object but have to be immutable. Akka can't enforce immutability (yet) so this has to be by convention.

Here is an example of an immutable message:

```

public class ImmutableMessage {
    private final int sequenceNumber;
    private final List<String> values;

    public ImmutableMessage(int sequenceNumber, List<String> values) {
        this.sequenceNumber = sequenceNumber;
        this.values = Collections.unmodifiableList(new ArrayList<String>(values));
    }

    public int getSequenceNumber() {
        return sequenceNumber;
    }

    public List<String> getValues() {
        return values;
    }
}

```

3.1.5 Send messages

Messages are sent to an Actor through one of the following methods.

- `tell` means “fire-and-forget”, e.g. send a message asynchronously and return immediately.

- `ask` sends a message asynchronously and returns a `Future` representing a possible reply.

Message ordering is guaranteed on a per-sender basis.

Note: There are performance implications of using `ask` since something needs to keep track of when it times out, there needs to be something that bridges a `Promise` into an `ActorRef` and it also needs to be reachable through remoting. So always prefer `tell` for performance, and only `ask` if you must.

In all these methods you have the option of passing along your own `ActorRef`. Make it a practice of doing so because it will allow the receiver actors to be able to respond to your message, since the sender reference is sent along with the message.

Tell: Fire-forget

This is the preferred way of sending messages. No blocking waiting for a message. This gives the best concurrency and scalability characteristics.

```
// don't forget to think about who is the sender (2nd argument)
target.tell(message, getSelf());
```

The sender reference is passed along with the message and available within the receiving actor via its `getSender` method while processing this message. Inside of an actor it is usually `getSelf` who shall be the sender, but there can be cases where replies shall be routed to some other actor—e.g. the parent—in which the second argument to `tell` would be a different one. Outside of an actor and if no reply is needed the second argument can be `null`; if a reply is needed outside of an actor you can use the `ask`-pattern described next..

Ask: Send-And-Receive-Future

The `ask` pattern involves actors as well as futures, hence it is offered as a use pattern rather than a method on `ActorRef`:

```
import static akka.pattern.Patterns.ask;
import static akka.pattern.Patterns.pipe;
import scala.concurrent.Future;
import scala.concurrent.duration.Duration;
import akka.dispatch.Futures;
import akka.dispatch.Mapper;
import akka.util.Timeout;

final Timeout t = new Timeout(Duration.create(5, TimeUnit.SECONDS));

final ArrayList<Future<Object>> futures = new ArrayList<Future<Object>>();
futures.add(ask(actorA, "request", 1000)); // using 1000ms timeout
futures.add(ask(actorB, "another request", t)); // using timeout from
// above

final Future<Iterable<Object>> aggregate = Futures.sequence(futures,
    system.dispatcher());

final Future<Result> transformed = aggregate.map(
    new Mapper<Iterable<Object>, Result>() {
        public Result apply(Iterable<Object> coll) {
            final Iterator<Object> it = coll.iterator();
            final String x = (String) it.next();
            final String s = (String) it.next();
            return new Result(x, s);
        }
    }, system.dispatcher());

pipe(transformed, system.dispatcher()).to(actorC);
```

This example demonstrates `ask` together with the `pipe` pattern on futures, because this is likely to be a common combination. Please note that all of the above is completely non-blocking and asynchronous: `ask` produces a `Future`, two of which are composed into a new future using the `Futures.sequence` and `map` methods and then `pipe` installs an `onComplete`-handler on the future to effect the submission of the aggregated `Result` to another actor.

Using `ask` will send a message to the receiving Actor as with `tell`, and the receiving actor must reply with `getSender().tell(reply, getSelf())` in order to complete the returned `Future` with a value. The `ask` operation involves creating an internal actor for handling this reply, which needs to have a timeout after which it is destroyed in order not to leak resources; see more below.

Warning: To complete the future with an exception you need send a `Failure` message to the sender. This is *not done automatically* when an actor throws an exception while processing a message.

```
try {
    String result = operation();
    getSender().tell(result, getSelf());
} catch (Exception e) {
    getSender().tell(new akka.actor.Status.Failure(e), getSelf());
    throw e;
}
```

If the actor does not complete the future, it will expire after the timeout period, specified as parameter to the `ask` method; this will complete the `Future` with an `AskTimeoutException`.

See [Futures](#) for more information on how to await or query a future.

The `onComplete`, `onSuccess`, or `onFailure` methods of the `Future` can be used to register a callback to get a notification when the `Future` completes. Gives you a way to avoid blocking.

Warning: When using future callbacks, inside actors you need to carefully avoid closing over the containing actor's reference, i.e. do not call methods or access mutable state on the enclosing actor from within the callback. This would break the actor encapsulation and may introduce synchronization bugs and race conditions because the callback will be scheduled concurrently to the enclosing actor. Unfortunately there is not yet a way to detect these illegal accesses at compile time. See also: [Actors and shared mutable state](#)

Forward message

You can forward a message from one actor to another. This means that the original sender address/reference is maintained even though the message is going through a 'mediator'. This can be useful when writing actors that work as routers, load-balancers, replicators etc. You need to pass along your context variable as well.

```
target.forward(result, getContext());
```

3.1.6 Receive messages

When an actor receives a message it is passed into the `onReceive` method, this is an abstract method on the `UntypedActor` base class that needs to be defined.

Here is an example:

```
import akka.actor.UntypedActor;
import akka.event.Logging;
import akka.event.LoggingAdapter;

public class MyUntypedActor extends UntypedActor {
    LoggingAdapter log = Logging.getLogger(getContext().system(), this);

    public void onReceive(Object message) throws Exception {
```

```

if (message instanceof String) {
    log.info("Received String message: {}", message);
    getSender().tell(message, getSelf());
} else
    unhandled(message);
}
}

```

An alternative to using `if-instanceof` checks is to use [Apache Commons MethodUtils](#) to invoke a named method whose parameter type matches the message type.

3.1.7 Reply to messages

If you want to have a handle for replying to a message, you can use `getSender()`, which gives you an `ActorRef`. You can reply by sending to that `ActorRef` with `getSender().tell(replyMsg, getSelf())`. You can also store the `ActorRef` for replying later, or passing on to other actors. If there is no sender (a message was sent without an actor or future context) then the sender defaults to a ‘dead-letter’ actor ref.

```

@Override
public void onReceive(Object msg) {
    Object result =
        // calculate result ...

    // do not forget the second argument!
    getSender().tell(result, getSelf());
}

```

3.1.8 Receive timeout

The `UntypedActorContext` `setReceiveTimeout` defines the inactivity timeout after which the sending of a `ReceiveTimeout` message is triggered. When specified, the receive function should be able to handle an `akka.actor.ReceiveTimeout` message. 1 millisecond is the minimum supported timeout.

Please note that the receive timeout might fire and enqueue the `ReceiveTimeout` message right after another message was enqueued; hence it is **not guaranteed** that upon reception of the receive timeout there must have been an idle period beforehand as configured via this method.

Once set, the receive timeout stays in effect (i.e. continues firing repeatedly after inactivity periods). Pass in `Duration.Undefined` to switch off this feature.

```

import akka.actor.ActorRef;
import akka.actor.ReceiveTimeout;
import akka.actor.UntypedActor;
import scala.concurrent.duration.Duration;

public class MyReceiveTimeoutUntypedActor extends UntypedActor {

    public MyReceiveTimeoutUntypedActor() {
        // To set an initial delay
        getContext().setReceiveTimeout(Duration.create("30 seconds"));
    }

    public void onReceive(Object message) {
        if (message.equals("Hello")) {
            // To set in a response to a message
            getContext().setReceiveTimeout(Duration.create("1 second"));
        } else if (message instanceof ReceiveTimeout) {
            // To turn it off
            getContext().setReceiveTimeout(Duration.Undefined());
        } else {

```

```

    unhandled(message);
  }
}
}

```

Messages marked with `NotInfluenceReceiveTimeout` will not reset the timer. This can be useful when `ReceiveTimeout` should be fired by external inactivity but not influenced by internal activity, e.g. scheduled tick messages.

3.1.9 Stopping actors

Actors are stopped by invoking the `stop` method of a `ActorRefFactory`, i.e. `ActorContext` or `ActorSystem`. Typically the context is used for stopping the actor itself or child actors and the system for stopping top level actors. The actual termination of the actor is performed asynchronously, i.e. `stop` may return before the actor is stopped.

```

import akka.actor.ActorRef;
import akka.actor.UntypedActor;
import akka.event.Logging;
import akka.event.LoggingAdapter;

public class MyStoppingActor extends UntypedActor {

    ActorRef child = null;

    // ... creation of child ...

    public void onReceive(Object message) throws Exception {
        if (message.equals("interrupt-child")) {
            context().stop(child);
        } else if (message.equals("done")) {
            context().stop(getSelf());
        } else {
            unhandled(message);
        }
    }
}

```

Processing of the current message, if any, will continue before the actor is stopped, but additional messages in the mailbox will not be processed. By default these messages are sent to the `deadLetters` of the `ActorSystem`, but that depends on the mailbox implementation.

Termination of an actor proceeds in two steps: first the actor suspends its mailbox processing and sends a stop command to all its children, then it keeps processing the internal termination notifications from its children until the last one is gone, finally terminating itself (invoking `postStop`, dumping mailbox, publishing `Terminated` on the *DeathWatch*, telling its supervisor). This procedure ensures that actor system sub-trees terminate in an orderly fashion, propagating the stop command to the leaves and collecting their confirmation back to the stopped supervisor. If one of the actors does not respond (i.e. processing a message for extended periods of time and therefore not receiving the stop command), this whole process will be stuck.

Upon `ActorSystem.terminate`, the system guardian actors will be stopped, and the aforementioned process will ensure proper termination of the whole system.

The `postStop` hook is invoked after an actor is fully stopped. This enables cleaning up of resources:

```

@Override
public void postStop() {
    // clean up resources here ...
}

```

Note: Since stopping an actor is asynchronous, you cannot immediately reuse the name of the child you just stopped; this will result in an `InvalidActorNameException`. Instead, watch the terminating actor and create its replacement in response to the `Terminated` message which will eventually arrive.

PoisonPill

You can also send an actor the `akka.actor.PoisonPill` message, which will stop the actor when the message is processed. `PoisonPill` is enqueued as ordinary messages and will be handled after messages that were already queued in the mailbox.

Use it like this:

```
myActor.tell(akka.actor.PoisonPill.getInstance(), sender);
```

Graceful Stop

`gracefulStop` is useful if you need to wait for termination or compose ordered termination of several actors:

```
import static akka.pattern.Patterns.gracefulStop;
import scala.concurrent.Await;
import scala.concurrent.Future;
import scala.concurrent.duration.Duration;
import akka.pattern.AskTimeoutException;
```

```
try {
    Future<Boolean> stopped =
        gracefulStop(actorRef, Duration.create(5, TimeUnit.SECONDS), Manager.SHUTDOWN);
    Await.result(stopped, Duration.create(6, TimeUnit.SECONDS));
    // the actor has been stopped
} catch (AskTimeoutException e) {
    // the actor wasn't stopped within 5 seconds
}
```

```
public class Manager extends UntypedActor {

    public static final String SHUTDOWN = "shutdown";

    ActorRef worker = getContext().watch(getContext().actorOf(
        Props.create(Cruncher.class), "worker"));

    public void onReceive(Object message) {
        if (message.equals("job")) {
            worker.tell("crunch", getSelf());
        } else if (message.equals(SHUTDOWN)) {
            worker.tell(PoisonPill.getInstance(), getSelf());
            getContext().become(shuttingDown);
        }
    }

    Procedure<Object> shuttingDown = new Procedure<Object>() {
        @Override
        public void apply(Object message) {
            if (message.equals("job")) {
                getSender().tell("service unavailable, shutting down", getSelf());
            } else if (message instanceof Terminated) {
                getContext().stop(getSelf());
            }
        }
    };
}
```

When `gracefulStop()` returns successfully, the actor's `postStop()` hook will have been executed: there exists a happens-before edge between the end of `postStop()` and the return of `gracefulStop()`.

In the above example a custom `Manager.SHUTDOWN` message is sent to the target actor to initiate the process of stopping the actor. You can use `PoisonPill` for this, but then you have limited possibilities to perform interactions with other actors before stopping the target actor. Simple cleanup tasks can be handled in `postStop`.

Warning: Keep in mind that an actor stopping and its name being deregistered are separate events which happen asynchronously from each other. Therefore it may be that you will find the name still in use after `gracefulStop()` returned. In order to guarantee proper deregistration, only reuse names from within a supervisor you control and only in response to a `Terminated` message, i.e. not for top-level actors.

3.1.10 HotSwap

Upgrade

Akka supports hotswapping the Actor's message loop (e.g. its implementation) at runtime. Use the `getContext().become` method from within the Actor. The hotswapped code is kept in a Stack which can be pushed (replacing or adding at the top) and popped.

Warning: Please note that the actor will revert to its original behavior when restarted by its Supervisor.

To hotswap the Actor using `getContext().become`:

```
import akka.japi.Procedure;

public class HotSwapActor extends UntypedActor {

    Procedure<Object> angry = new Procedure<Object>() {
        @Override
        public void apply(Object message) {
            if (message.equals("bar")) {
                getSender().tell("I am already angry?", getSelf());
            } else if (message.equals("foo")) {
                getContext().become(happy);
            }
        }
    };

    Procedure<Object> happy = new Procedure<Object>() {
        @Override
        public void apply(Object message) {
            if (message.equals("bar")) {
                getSender().tell("I am already happy :-)", getSelf());
            } else if (message.equals("foo")) {
                getContext().become(angry);
            }
        }
    };

    public void onReceive(Object message) {
        if (message.equals("bar")) {
            getContext().become(angry);
        } else if (message.equals("foo")) {
            getContext().become(happy);
        } else {
            unhandled(message);
        }
    }
}
```

This variant of the `become` method is useful for many different things, such as to implement a Finite State Machine (FSM). It will replace the current behavior (i.e. the top of the behavior stack), which means that you do not use `unbecome`, instead always the next behavior is explicitly installed.

The other way of using `become` does not replace but add to the top of the behavior stack. In this case care must be taken to ensure that the number of “pop” operations (i.e. `unbecome`) matches the number of “push” ones in the long run, otherwise this amounts to a memory leak (which is why this behavior is not the default).

```
public class UntypedActorSwapper {

    public static class Swap {
        public static Swap SWAP = new Swap();

        private Swap() {
        }
    }

    public static class Swapper extends UntypedActor {
        LoggingAdapter log = Logging.getLogger(getContext().system(), this);

        public void onReceive(Object message) {
            if (message == SWAP) {
                log.info("Hi");
                getContext().become(new Procedure<Object>() {
                    @Override
                    public void apply(Object message) {
                        log.info("Ho");
                        getContext().unbecome(); // resets the latest 'become'
                    }
                }, false); // this signals stacking of the new behavior
            } else {
                unhandled(message);
            }
        }
    }

    public static void main(String... args) {
        ActorSystem system = ActorSystem.create("MySystem");
        ActorRef swap = system.actorOf(Props.create(Swapper.class));
        swap.tell(SWAP, ActorRef.noSender()); // logs Hi
        swap.tell(SWAP, ActorRef.noSender()); // logs Ho
        swap.tell(SWAP, ActorRef.noSender()); // logs Hi
        swap.tell(SWAP, ActorRef.noSender()); // logs Ho
        swap.tell(SWAP, ActorRef.noSender()); // logs Hi
        swap.tell(SWAP, ActorRef.noSender()); // logs Ho
    }
}
```

3.1.11 Stash

The `UntypedActorWithStash` class enables an actor to temporarily stash away messages that can not or should not be handled using the actor’s current behavior. Upon changing the actor’s message handler, i.e., right before invoking `getContext().become()` or `getContext().unbecome()`, all stashed messages can be “unstashed”, thereby prepending them to the actor’s mailbox. This way, the stashed messages can be processed in the same order as they have been received originally. An actor that extends `UntypedActorWithStash` will automatically get a deque-based mailbox.

Note: The abstract class `UntypedActorWithStash` implements the marker interface `RequiresMessageQueue<DequeBasedMessageQueueSemantics>` which requests the system to automatically choose a deque based mailbox implementation for the actor. If you want more control over the mailbox, see the documentation on mailboxes: [Mailboxes](#).

Here is an example of the `UntypedActorWithStash` class in action:

```
import akka.actor.UntypedActorWithStash;

public class ActorWithProtocol extends UntypedActorWithStash {
    public void onReceive(Object msg) {
        if (msg.equals("open")) {
            unstashAll();
            getContext().become(new Procedure<Object>() {
                public void apply(Object msg) throws Exception {
                    if (msg.equals("write")) {
                        // do writing...
                    } else if (msg.equals("close")) {
                        unstashAll();
                        getContext().unbecome();
                    } else {
                        stash();
                    }
                }
            }, false); // add behavior on top instead of replacing
        } else {
            stash();
        }
    }
}
```

Invoking `stash()` adds the current message (the message that the actor received last) to the actor's stash. It is typically invoked when handling the default case in the actor's message handler to stash messages that aren't handled by the other cases. It is illegal to stash the same message twice; to do so results in an `IllegalStateException` being thrown. The stash may also be bounded in which case invoking `stash()` may lead to a capacity violation, which results in a `StashOverflowException`. The capacity of the stash can be configured using the `stash-capacity` setting (an `Int`) of the mailbox's configuration.

Invoking `unstashAll()` enqueues messages from the stash to the actor's mailbox until the capacity of the mailbox (if any) has been reached (note that messages from the stash are prepended to the mailbox). In case a bounded mailbox overflows, a `MessageQueueAppendFailedException` is thrown. The stash is guaranteed to be empty after calling `unstashAll()`.

The stash is backed by a `scala.collection.immutable.Vector`. As a result, even a very large number of messages may be stashed without a major impact on performance.

Note that the stash is part of the ephemeral actor state, unlike the mailbox. Therefore, it should be managed like other parts of the actor's state which have the same property. The `UntypedActorWithStash` implementation of `preRestart` will call `unstashAll()`, which is usually the desired behavior.

Note: If you want to enforce that your actor can only work with an unbounded stash, then you should use the `UntypedActorWithUnboundedStash` class instead.

3.1.12 Killing an Actor

You can kill an actor by sending a `Kill` message. This will cause the actor to throw a `ActorKilledException`, triggering a failure. The actor will suspend operation and its supervisor will be asked how to handle the failure, which may mean resuming the actor, restarting it or terminating it completely. See [What Supervision Means](#) for more information.

Use `Kill` like this:

```
victim.tell(akka.actor.Kill.getInstance(), ActorRef.noSender());
```

3.1.13 Actors and exceptions

It can happen that while a message is being processed by an actor, that some kind of exception is thrown, e.g. a database exception.

What happens to the Message

If an exception is thrown while a message is being processed (i.e. taken out of its mailbox and handed over to the current behavior), then this message will be lost. It is important to understand that it is not put back on the mailbox. So if you want to retry processing of a message, you need to deal with it yourself by catching the exception and retry your flow. Make sure that you put a bound on the number of retries since you don't want a system to livelock (so consuming a lot of cpu cycles without making progress). Another possibility would be to have a look at the *PeekMailbox pattern*.

What happens to the mailbox

If an exception is thrown while a message is being processed, nothing happens to the mailbox. If the actor is restarted, the same mailbox will be there. So all messages on that mailbox will be there as well.

What happens to the actor

If code within an actor throws an exception, that actor is suspended and the supervision process is started (see *Supervision and Monitoring*). Depending on the supervisor's decision the actor is resumed (as if nothing happened), restarted (wiping out its internal state and starting from scratch) or terminated.

3.1.14 Initialization patterns

The rich lifecycle hooks of Actors provide a useful toolkit to implement various initialization patterns. During the lifetime of an `ActorRef`, an actor can potentially go through several restarts, where the old instance is replaced by a fresh one, invisibly to the outside observer who only sees the `ActorRef`.

One may think about the new instances as "incarnations". Initialization might be necessary for every incarnation of an actor, but sometimes one needs initialization to happen only at the birth of the first instance when the `ActorRef` is created. The following sections provide patterns for different initialization needs.

Initialization via constructor

Using the constructor for initialization has various benefits. First of all, it makes it possible to use `val` fields to store any state that does not change during the life of the actor instance, making the implementation of the actor more robust. The constructor is invoked for every incarnation of the actor, therefore the internals of the actor can always assume that proper initialization happened. This is also the drawback of this approach, as there are cases when one would like to avoid reinitializing internals on restart. For example, it is often useful to preserve child actors across restarts. The following section provides a pattern for this case.

Initialization via `preStart`

The method `preStart()` of an actor is only called once directly during the initialization of the first instance, that is, at creation of its `ActorRef`. In the case of restarts, `preStart()` is called from `postRestart()`, therefore if not overridden, `preStart()` is called on every incarnation. However, overriding `postRestart()` one can disable this behavior, and ensure that there is only one call to `preStart()`.

One useful usage of this pattern is to disable creation of new `ActorRefs` for children during restarts. This can be achieved by overriding `preRestart()`:

```
@Override
public void preStart() {
    // Initialize children here
}

// Overriding postRestart to disable the call to preStart()
// after restarts
@Override
public void postRestart(Throwable reason) {
}

// The default implementation of preRestart() stops all the children
// of the actor. To opt-out from stopping the children, we
// have to override preRestart()
@Override
public void preRestart(Throwable reason, Option<Object> message)
    throws Exception {
    // Keep the call to postStop(), but no stopping of children
    postStop();
}
}
```

Please note, that the child actors are *still restarted*, but no new `ActorRef` is created. One can recursively apply the same principles for the children, ensuring that their `preStart()` method is called only at the creation of their refs.

For more information see [What Restarting Means](#).

Initialization via message passing

There are cases when it is impossible to pass all the information needed for actor initialization in the constructor, for example in the presence of circular dependencies. In this case the actor should listen for an initialization message, and use `become()` or a finite state-machine state transition to encode the initialized and uninitialized states of the actor.

```
private String initializeMe = null;

@Override
public void onReceive(Object message) throws Exception {
    if (message.equals("init")) {
        initializeMe = "Up and running";
        getContext().become(new Procedure<Object>() {
            @Override
            public void apply(Object message) throws Exception {
                if (message.equals("U OK?"))
                    getSender().tell(initializeMe, getSelf());
            }
        });
    }
}
}
```

If the actor may receive messages before it has been initialized, a useful tool can be the `Stash` to save messages until the initialization finishes, and replaying them after the actor became initialized.

Warning: This pattern should be used with care, and applied only when none of the patterns above are applicable. One of the potential issues is that messages might be lost when sent to remote actors. Also, publishing an `ActorRef` in an uninitialized state might lead to the condition that it receives a user message before the initialization has been done.

3.2 Typed Actors

Akka Typed Actors is an implementation of the [Active Objects](#) pattern. Essentially turning method invocations into asynchronous dispatch instead of synchronous that has been the default way since Smalltalk came out.

Typed Actors consist of 2 “parts”, a public interface and an implementation, and if you’ve done any work in “enterprise” Java, this will be very familiar to you. As with normal Actors you have an external API (the public interface instance) that will delegate method calls asynchronously to a private instance of the implementation.

The advantage of Typed Actors vs. Actors is that with TypedActors you have a static contract, and don’t need to define your own messages, the downside is that it places some limitations on what you can do and what you can’t, i.e. you can’t use become/unbecome.

Typed Actors are implemented using [JDK Proxies](#) which provide a pretty easy-worked API to intercept method calls.

Note: Just as with regular Akka Untyped Actors, Typed Actors process one call at a time.

3.2.1 When to use Typed Actors

Typed actors are nice for bridging between actor systems (the “inside”) and non-actor code (the “outside”), because they allow you to write normal OO-looking code on the outside. Think of them like doors: their practicality lies in interfacing between private sphere and the public, but you don’t want that many doors inside your house, do you? For a longer discussion see [this blog post](#).

A bit more background: TypedActors can very easily be abused as RPC, and that is an abstraction which is [well-known](#) to be leaky. Hence TypedActors are not what we think of first when we talk about making highly scalable concurrent software easier to write correctly. They have their niche, use them sparingly.

3.2.2 The tools of the trade

Before we create our first Typed Actor we should first go through the tools that we have at our disposal, it’s located in `akka.actor.TypedActor`.

```
//Returns the Typed Actor Extension
TypedActorExtension extension =
    TypedActor.get(system); //system is an instance of ActorSystem

//Returns whether the reference is a Typed Actor Proxy or not
TypedActor.get(system).isTypedActor(someReference);

//Returns the backing Akka Actor behind an external Typed Actor Proxy
TypedActor.get(system).getActorRefFor(someReference);

//Returns the current ActorContext,
// method only valid within methods of a TypedActor implementation
ActorContext context = TypedActor.context();

//Returns the external proxy of the current Typed Actor,
// method only valid within methods of a TypedActor implementation
Squarer sq = TypedActor.<Squarer>self();

//Returns a contextual instance of the Typed Actor Extension
//this means that if you create other Typed Actors with this,
//they will become children to the current Typed Actor.
TypedActor.get(TypedActor.context());
```

Warning: Same as not exposing `this` of an Akka Actor, it's important not to expose `this` of a Typed Actor, instead you should pass the external proxy reference, which is obtained from within your Typed Actor as `TypedActor.self()`, this is your external identity, as the `ActorRef` is the external identity of an Akka Actor.

3.2.3 Creating Typed Actors

To create a Typed Actor you need to have one or more interfaces, and one implementation.

The following imports are assumed:

```
import akka.actor.TypedActor;
import akka.actor.*;
import akka.japi.*;
import akka.dispatch.Futures;

import scala.concurrent.Await;
import scala.concurrent.Future;
import scala.concurrent.duration.Duration;
import java.util.concurrent.TimeUnit;

import java.util.List;
import java.util.ArrayList;
import java.util.Random;
import akka.routing.RoundRobinGroup;
public class TypedActorDocTest {
    Object someReference = null;
    ActorSystem system = null;

    static
    public interface Squarer {
        void squareDontCare(int i); //fire-forget

        Future<Integer> square(int i); //non-blocking send-request-reply

        Option<Integer> squareNowPlease(int i); //blocking send-request-reply

        int squareNow(int i); //blocking send-request-reply
    }

    static
    class SquarerImpl implements Squarer {
        private String name;

        public SquarerImpl() {
            this.name = "default";
        }

        public SquarerImpl(String name) {
            this.name = name;
        }

        public void squareDontCare(int i) {
            int sq = i * i; //Nobody cares :(
        }

        public Future<Integer> square(int i) {
            return Futures.successful(i * i);
        }
    }
}
```

```

public Option<Integer> squareNowPlease(int i) {
    return Option.some(i * i);
}

public int squareNow(int i) {
    return i * i;
}
}

@Test public void mustGetTheTypedActorExtension() {

    try {

        //Returns the Typed Actor Extension
        TypedActorExtension extension =
            TypedActor.get(system); //system is an instance of ActorSystem

        //Returns whether the reference is a Typed Actor Proxy or not
        TypedActor.get(system).isTypedActor(someReference);

        //Returns the backing Akka Actor behind an external Typed Actor Proxy
        TypedActor.get(system).getActorRefFor(someReference);

        //Returns the current ActorContext,
        // method only valid within methods of a TypedActor implementation
        ActorContext context = TypedActor.context();

        //Returns the external proxy of the current Typed Actor,
        // method only valid within methods of a TypedActor implementation
        Squarer sq = TypedActor.<Squarer>self();

        //Returns a contextual instance of the Typed Actor Extension
        //this means that if you create other Typed Actors with this,
        //they will become children to the current Typed Actor.
        TypedActor.get(TypedActor.context());

    } catch (Exception e) {
        //dun care
    }
}

@Test public void createATypedActor() {
    try {
        Squarer mySquarer =
            TypedActor.get(system).typedActorOf(
                new TypedProps<SquarerImpl>(Squarer.class, SquarerImpl.class));
        Squarer otherSquarer =
            TypedActor.get(system).typedActorOf(
                new TypedProps<SquarerImpl>(Squarer.class,
                    new Creator<SquarerImpl>() {
                        public SquarerImpl create() { return new SquarerImpl("foo"); }
                    },
                    "name");

        mySquarer.squareDontCare(10);

        Future<Integer> fSquare = mySquarer.square(10); //A Future[Int]

        Option<Integer> oSquare = mySquarer.squareNowPlease(10); //Option[Int]

        int iSquare = mySquarer.squareNow(10); //Int

        assertEquals(100, Await.result(fSquare,
            Duration.create(3, TimeUnit.SECONDS)).intValue());
    }
}

```

```

assertEquals(100, oSquare.get().intValue());

assertEquals(100, iSquare);

TypedActor.get(system).stop(mySquarer);

TypedActor.get(system).poisonPill(otherSquarer);
} catch (Exception e) {
    //Ignore
}
}

@Test public void createHierarchies() {
    try {
        Squarer childSquarer =
            TypedActor.get(TypedActor.context()).
                typedActorOf(
                    new TypedProps<SquarerImpl>(Squarer.class, SquarerImpl.class)
                );
        //Use "childSquarer" as a Squarer
    } catch (Exception e) {
        //dun care
    }
}

@Test public void proxyAnyActorRef() {
    try {
        final ActorRef actorRefToRemoteActor = system.deadLetters();
        Squarer typedActor =
            TypedActor.get(system).
                typedActorOf(
                    new TypedProps<Squarer>(Squarer.class),
                    actorRefToRemoteActor
                );
        //Use "typedActor" as a FooBar
    } catch (Exception e) {
        //dun care
    }
}

interface HasName {
    String name();
}

class Named implements HasName {
    private int id = new Random().nextInt(1024);

    @Override public String name() { return "name-" + id; }
}

@Test public void typedRouterPattern() {
    try {
        // prepare routees
        TypedActorExtension typed = TypedActor.get(system);

        Named named1 =
            typed.typedActorOf(new TypedProps<Named>(Named.class));

        Named named2 =
            typed.typedActorOf(new TypedProps<Named>(Named.class));
    }
}

```

```

List<Named> routees = new ArrayList<Named>();
routees.add(named1);
routees.add(named2);

List<String> routeePaths = new ArrayList<String>();
routeePaths.add(typed.getActorRefFor(named1).path().toStringWithoutAddress());
routeePaths.add(typed.getActorRefFor(named2).path().toStringWithoutAddress());

// prepare untyped router
ActorRef router = system.actorOf(new RoundRobinGroup(routeePaths).props(), "router");

// prepare typed proxy, forwarding MethodCall messages to `router`
Named typedRouter = typed.typedActorOf(new TypedProps<Named>(Named.class), router);

System.out.println("actor was: " + typedRouter.name()); // name-243
System.out.println("actor was: " + typedRouter.name()); // name-614
System.out.println("actor was: " + typedRouter.name()); // name-243
System.out.println("actor was: " + typedRouter.name()); // name-614

typed.poisonPill(named1);
typed.poisonPill(named2);
typed.poisonPill(typedRouter);

} catch (Exception e) {
    //dun care
}
}
}

```

Our example interface:

```

public interface Squarer {
    // typed actor iface methods ...
}

```

Our example implementation of that interface:

```

class SquarerImpl implements Squarer {
    private String name;

    public SquarerImpl() {
        this.name = "default";
    }

    public SquarerImpl(String name) {
        this.name = name;
    }

    // typed actor impl methods ...
}

```

The most trivial way of creating a Typed Actor instance of our Squarer:

```

Squarer mySquarer =
    TypedActor.get(system).typedActorOf(
        new TypedProps<SquarerImpl>(Squarer.class, SquarerImpl.class));

```

First type is the type of the proxy, the second type is the type of the implementation. If you need to call a specific constructor you do it like this:

```

Squarer otherSquarer =
    TypedActor.get(system).typedActorOf(
        new TypedProps<SquarerImpl>(Squarer.class,
            new Creator<SquarerImpl>() {

```

```

        public SquarerImpl create() { return new SquarerImpl("foo"); }
    },
    "name");

```

Since you supply a `Props`, you can specify which dispatcher to use, what the default timeout should be used and more. Now, our `Squarer` doesn't have any methods, so we'd better add those.

```

public interface Squarer {
    void squareDontCare(int i); //fire-forget

    Future<Integer> square(int i); //non-blocking send-request-reply

    Option<Integer> squareNowPlease(int i); //blocking send-request-reply

    int squareNow(int i); //blocking send-request-reply
}

```

Alright, now we've got some methods we can call, but we need to implement those in `SquarerImpl`.

```

class SquarerImpl implements Squarer {
    private String name;

    public SquarerImpl() {
        this.name = "default";
    }

    public SquarerImpl(String name) {
        this.name = name;
    }

    public void squareDontCare(int i) {
        int sq = i * i; //Nobody cares :(
    }

    public Future<Integer> square(int i) {
        return Futures.successful(i * i);
    }

    public Option<Integer> squareNowPlease(int i) {
        return Option.some(i * i);
    }

    public int squareNow(int i) {
        return i * i;
    }
}

```

Excellent, now we have an interface and an implementation of that interface, and we know how to create a `Typed Actor` from that, so let's look at calling these methods.

3.2.4 Method dispatch semantics

Methods returning:

- `void` will be dispatched with `fire-and-forget` semantics, exactly like `ActorRef.tell`
- `scala.concurrent.Future<?>` will use `send-request-reply` semantics, exactly like `ActorRef.ask`
- `akka.japi.Option<?>` will use `send-request-reply` semantics, but *will* block to wait for an answer, and return `akka.japi.Option.None` if no answer was produced within the timeout, or

`akka.japi.Option.Some<?>` containing the result otherwise. Any exception that was thrown during this call will be rethrown.

- Any other type of value will use `send-request-reply` semantics, but *will* block to wait for an answer, throwing `java.util.concurrent.TimeoutException` if there was a timeout or rethrow any exception that was thrown during this call. Note that due to the Java exception and reflection mechanisms, such a `TimeoutException` will be wrapped in a `java.lang.reflect.UndeclaredThrowableException` unless the interface method explicitly declares the `TimeoutException` as a thrown checked exception.

3.2.5 Messages and immutability

While Akka cannot enforce that the parameters to the methods of your Typed Actors are immutable, we *strongly* recommend that parameters passed are immutable.

One-way message send

```
mySquarer.squareDontCare(10);
```

As simple as that! The method will be executed on another thread; asynchronously.

Request-reply message send

```
Option<Integer> oSquare = mySquarer.squareNowPlease(10); //Option[Int]
```

This will block for as long as the timeout that was set in the `Props` of the Typed Actor, if needed. It will return `None` if a timeout occurs.

```
int iSquare = mySquarer.squareNow(10); //Int
```

This will block for as long as the timeout that was set in the `Props` of the Typed Actor, if needed. It will throw a `java.util.concurrent.TimeoutException` if a timeout occurs. Note that here, such a `TimeoutException` will be wrapped in a `java.lang.reflect.UndeclaredThrowableException` by the Java reflection mechanism, because the interface method does not explicitly declare the `TimeoutException` as a thrown checked exception. To get the `TimeoutException` directly, declare `throws java.util.concurrent.TimeoutException` at the interface method.

Request-reply-with-future message send

```
Future<Integer> fSquare = mySquarer.square(10); //A Future[Int]
```

This call is asynchronous, and the `Future` returned can be used for asynchronous composition.

3.2.6 Stopping Typed Actors

Since Akka's Typed Actors are backed by Akka Actors they must be stopped when they aren't needed anymore.

```
TypedActor.get(system).stop(mySquarer);
```

This asynchronously stops the Typed Actor associated with the specified proxy ASAP.

```
TypedActor.get(system).poisonPill(otherSquarer);
```

This asynchronously stops the Typed Actor associated with the specified proxy after it's done with all calls that were made prior to this call.

3.2.7 Typed Actor Hierarchies

Since you can obtain a contextual Typed Actor Extension by passing in an ActorContext you can create child Typed Actors by invoking `typedActorOf(..)` on that.

```
Squarer childSquarer =
  TypedActor.get(TypedActor.context()).
    typedActorOf(
      new TypedProps<SquarerImpl>(Squarer.class, SquarerImpl.class)
    );
//Use "childSquarer" as a Squarer
```

You can also create a child Typed Actor in regular Akka Actors by giving the `UntypedActorContext` as an input parameter to `TypedActor.get(...)`.

3.2.8 Supervisor Strategy

By having your Typed Actor implementation class implement `TypedActor.Supervisor` you can define the strategy to use for supervising child actors, as described in *Supervision and Monitoring* and *Fault Tolerance*.

3.2.9 Receive arbitrary messages

If your implementation class of your TypedActor extends `akka.actor.TypedActor.Receiver`, all messages that are not `MethodCall` instances will be passed into the `onReceive`-method.

This allows you to react to `DeathWatch Terminated`-messages and other types of messages, e.g. when interfacing with untyped actors.

3.2.10 Lifecycle callbacks

By having your Typed Actor implementation class implement any and all of the following:

- `TypedActor.PreStart`
- `TypedActor.PostStop`
- `TypedActor.PreRestart`
- `TypedActor.PostRestart`

You can hook into the lifecycle of your Typed Actor.

3.2.11 Proxying

You can use the `typedActorOf` that takes a `TypedProps` and an `ActorRef` to proxy the given `ActorRef` as a `TypedActor`. This is usable if you want to communicate remotely with `TypedActors` on other machines, just pass the `ActorRef` to `typedActorOf`.

3.2.12 Lookup & Remoting

Since `TypedActors` are backed by Akka Actors, you can use `typedActorOf` to proxy `ActorRefs` potentially residing on remote nodes.

```
Squarer typedActor =
  TypedActor.get(system).
    typedActorOf(
      new TypedProps<Squarer>(Squarer.class),
      actorRefToRemoteActor
```

```
);
//Use "typedActor" as a FooBar
```

3.2.13 Typed Router pattern

Sometimes you want to spread messages between multiple actors. The easiest way to achieve this in Akka is to use a *Router*, which can implement a specific routing logic, such as `smallest-mailbox` or `consistent-hashing` etc.

Routers are not provided directly for typed actors, but it is really easy to leverage an untyped router and use a typed proxy in front of it. To showcase this let's create typed actors that assign themselves some random id, so we know that in fact, the router has sent the message to different actors:

```
@Test public void typedRouterPattern() {
    try {
        // prepare routees
        TypedActorExtension typed = TypedActor.get(system);

        Named named1 =
            typed.typedActorOf(new TypedProps<Named>(Named.class));

        Named named2 =
            typed.typedActorOf(new TypedProps<Named>(Named.class));

        List<Named> routees = new ArrayList<Named>();
        routees.add(named1);
        routees.add(named2);

        List<String> routeePaths = new ArrayList<String>();
        routeePaths.add(typed.getActorRefFor(named1).path().toStringWithoutAddress());
        routeePaths.add(typed.getActorRefFor(named2).path().toStringWithoutAddress());

        // prepare untyped router
        ActorRef router = system.actorOf(new RoundRobinGroup(routeePaths).props(), "router");

        // prepare typed proxy, forwarding MethodCall messages to `router`
        Named typedRouter = typed.typedActorOf(new TypedProps<Named>(Named.class), router);

        System.out.println("actor was: " + typedRouter.name()); // name-243
        System.out.println("actor was: " + typedRouter.name()); // name-614
        System.out.println("actor was: " + typedRouter.name()); // name-243
        System.out.println("actor was: " + typedRouter.name()); // name-614

        typed.poisonPill(named1);
        typed.poisonPill(named2);
        typed.poisonPill(typedRouter);

    } catch (Exception e) {
        //dun care
    }
}
```

In order to round robin among a few instances of such actors, you can simply create a plain untyped router, and then facade it with a `TypedActor` like shown in the example below. This works because typed actors of course communicate using the same mechanisms as normal actors, and methods calls on them get transformed into message sends of `MethodCall` messages.

```
// prepare routees
TypedActorExtension typed = TypedActor.get(system);
```

```
Named named1 =
    typed.typedActorOf(new TypedProps<Named>(Named.class));

Named named2 =
    typed.typedActorOf(new TypedProps<Named>(Named.class));

List<Named> routees = new ArrayList<Named>();
routees.add(named1);
routees.add(named2);

List<String> routeePaths = new ArrayList<String>();
routeePaths.add(typed.getActorRefFor(named1).path().toStringWithoutAddress());
routeePaths.add(typed.getActorRefFor(named2).path().toStringWithoutAddress());

// prepare untyped router
ActorRef router = system.actorOf(new RoundRobinGroup(routeePaths).props(), "router");

// prepare typed proxy, forwarding MethodCall messages to `router`
Named typedRouter = typed.typedActorOf(new TypedProps<Named>(Named.class), router);

System.out.println("actor was: " + typedRouter.name()); // name-243
System.out.println("actor was: " + typedRouter.name()); // name-614
System.out.println("actor was: " + typedRouter.name()); // name-243
System.out.println("actor was: " + typedRouter.name()); // name-614
```

3.3 Fault Tolerance


As explained in *Actor Systems* each actor is the supervisor of its children, and as such each actor defines fault handling supervisor strategy. This strategy cannot be changed afterwards as it is an integral part of the actor system's structure.

3.3.1 Fault Handling in Practice

First, let us look at a sample that illustrates one way to handle data store errors, which is a typical source of failure in real world applications. Of course it depends on the actual application what is possible to do when the data store is unavailable, but in this sample we use a best effort re-connect approach.

Read the following source code. The inlined comments explain the different pieces of the fault handling and why they are added. It is also highly recommended to run this sample as it is easy to follow the log output to understand what is happening in runtime.

Diagrams of the Fault Tolerance Sample



The above diagram illustrates the normal message flow.

Normal flow:

Step	Description
1	The progress Listener starts the work.
2	The Worker schedules work by sending Do messages periodically to itself
3, 4, 5	When receiving Do the Worker tells the CounterService to increment the counter, three times. The Increment message is forwarded to the Counter, which updates its counter variable and sends current value to the Storage.
6, 7	The Worker asks the CounterService of current value of the counter and pipes the result back to the Listener.

Step	Description
1	The Storage throws StorageException.
2	The CounterService is supervisor of the Storage and restarts the Storage when StorageException is thrown.
3, 4, 5, 6	The Storage continues to fail and is restarted.
7	After 3 failures and restarts within 5 seconds the Storage is stopped by its supervisor, i.e. the CounterService.
8	The CounterService is also watching the Storage for termination and receives the Terminated message when the Storage has been stopped ...
9, 10, 11	and tells the Counter that there is no Storage.
12	The CounterService schedules a Reconnect message to itself.
13, 14	When it receives the Reconnect message it creates a new Storage ...
15, 16	and tells the Counter to use the new Storage

Full Source Code of the Fault Tolerance Sample

```
import java.util.ArrayList;
import java.util.HashMap;
import java.util.List;
import java.util.Map;

import akka.actor.*;
import akka.dispatch.Mapper;
import akka.japi.Function;
import scala.concurrent.duration.Duration;
import akka.util.Timeout;
import akka.event.Logging;
import akka.event.LoggingAdapter;
import com.typesafe.config.Config;
import com.typesafe.config.ConfigFactory;

import static akka.japi.Util.classTag;

import static akka.actor.SupervisorStrategy.restart;
import static akka.actor.SupervisorStrategy.stop;
import static akka.actor.SupervisorStrategy.escalate;
import akka.actor.SupervisorStrategy.Directive;
import static akka.pattern.Patterns.ask;
import static akka.pattern.Patterns.pipe;

import static docs.actor.japi.FaultHandlingDocSample.WorkerApi.*;
import static docs.actor.japi.FaultHandlingDocSample.CounterServiceApi.*;
import static docs.actor.japi.FaultHandlingDocSample.CounterApi.*;
import static docs.actor.japi.FaultHandlingDocSample.StorageApi.*;

public class FaultHandlingDocSample {

    /**
     * Runs the sample
     */
    public static void main(String[] args) {
        Config config = ConfigFactory.parseString("akka.loglevel = DEBUG \n" +
            "akka.actor.debug.lifecycle = on");

        ActorSystem system = ActorSystem.create("FaultToleranceSample", config);
```

```

ActorRef worker = system.actorOf(Props.create(Worker.class), "worker");
ActorRef listener = system.actorOf(Props.create(Listener.class), "listener");
// start the work and listen on progress
// note that the listener is used as sender of the tell,
// i.e. it will receive replies from the worker
worker.tell(Start, listener);
}

/**
 * Listens on progress from the worker and shuts down the system when enough
 * work has been done.
 */
public static class Listener extends UntypedActor {
    final LoggingAdapter log = Logging.getLogger(getContext().system(), this);

    @Override
    public void preStart() {
        // If we don't get any progress within 15 seconds then the service
        // is unavailable
        getContext().setReceiveTimeout(Duration.create("15 seconds"));
    }

    public void onReceive(Object msg) {
        log.debug("received message {}", msg);
        if (msg instanceof Progress) {
            Progress progress = (Progress) msg;
            log.info("Current progress: {} %", progress.percent);
            if (progress.percent >= 100.0) {
                log.info("That's all, shutting down");
                getContext().system().terminate();
            }
        } else if (msg == ReceiveTimeout.getInstance()) {
            // No progress within 15 seconds, ServiceUnavailable
            log.error("Shutting down due to unavailable service");
            getContext().system().terminate();
        } else {
            unhandled(msg);
        }
    }
}

public interface WorkerApi {
    public static final Object Start = "Start";
    public static final Object Do = "Do";

    public static class Progress {
        public final double percent;

        public Progress(double percent) {
            this.percent = percent;
        }

        public String toString() {
            return String.format("%s(%s)", getClass().getSimpleName(), percent);
        }
    }
}

/**
 * Worker performs some work when it receives the Start message. It will
 * continuously notify the sender of the Start message of current Progress.
 * The Worker supervise the CounterService.

```

```

*/
public static class Worker extends UntypedActor {
    final LoggingAdapter log = Logging.getLogger(getContext().system(), this);
    final Timeout askTimeout = new Timeout(Duration.create(5, "seconds"));

    // The sender of the initial Start message will continuously be notified
    // about progress
    ActorRef progressListener;
    final ActorRef counterService = getContext().actorOf(
        Props.create(CounterService.class), "counter");
    final int totalCount = 51;

    // Stop the CounterService child if it throws ServiceUnavailable
    private static SupervisorStrategy strategy = new OneForOneStrategy(-1,
        Duration.Inf(), new Function<Throwable, Directive>() {
            @Override
            public Directive apply(Throwable t) {
                if (t instanceof ServiceUnavailable) {
                    return stop();
                } else {
                    return escalate();
                }
            }
        });

    @Override
    public SupervisorStrategy supervisorStrategy() {
        return strategy;
    }

    public void onReceive(Object msg) {
        log.debug("received message {}", msg);
        if (msg.equals(Start) && progressListener == null) {
            progressListener = getSender();
            getContext().system().scheduler().schedule(
                Duration.Zero(), Duration.create(1, "second"), getSelf(), Do,
                getContext().dispatcher(), null
            );
        } else if (msg.equals(Do)) {
            counterService.tell(new Increment(1), getSelf());
            counterService.tell(new Increment(1), getSelf());
            counterService.tell(new Increment(1), getSelf());

            // Send current progress to the initial sender
            pipe(ask(counterService, GetCurrentCount, askTimeout)
                .mapTo(classTag(CurrentCount.class))
                .map(new Mapper<CurrentCount, Progress>() {
                    public Progress apply(CurrentCount c) {
                        return new Progress(100.0 * c.count / totalCount);
                    }
                }, getContext().dispatcher(), getContext().dispatcher())
                .to(progressListener);
        } else {
            unhandled(msg);
        }
    }
}

public interface CounterServiceApi {

    public static final Object GetCurrentCount = "GetCurrentCount";

    public static class CurrentCount {

```

```

public final String key;
public final long count;

public CurrentCount(String key, long count) {
    this.key = key;
    this.count = count;
}

public String toString() {
    return String.format("%s(%s, %s)", getClass().getSimpleName(), key, count);
}
}

public static class Increment {
    public final long n;

    public Increment(long n) {
        this.n = n;
    }

    public String toString() {
        return String.format("%s(%s)", getClass().getSimpleName(), n);
    }
}

public static class ServiceUnavailable extends RuntimeException {
    private static final long serialVersionUID = 1L;
    public ServiceUnavailable(String msg) {
        super(msg);
    }
}
}

/**
 * Adds the value received in Increment message to a persistent counter.
 * Replies with CurrentCount when it is asked for CurrentCount. CounterService
 * supervise Storage and Counter.
 */
public static class CounterService extends UntypedActor {

    // Reconnect message
    static final Object Reconnect = "Reconnect";

    private static class SenderMsgPair {
        final ActorRef sender;
        final Object msg;

        SenderMsgPair(ActorRef sender, Object msg) {
            this.msg = msg;
            this.sender = sender;
        }
    }

    final LoggingAdapter log = Logging.getLogger(getContext().system(), this);
    final String key = getSelf().path().name();
    ActorRef storage;
    ActorRef counter;
    final List<SenderMsgPair> backlog = new ArrayList<SenderMsgPair>();
    final int MAX_BACKLOG = 10000;

    // Restart the storage child when StorageException is thrown.

```

```

// After 3 restarts within 5 seconds it will be stopped.
private static SupervisorStrategy strategy = new OneForOneStrategy(3,
    Duration.create("5 seconds"), new Function<Throwable, Directive>() {
    @Override
    public Directive apply(Throwable t) {
        if (t instanceof StorageException) {
            return restart();
        } else {
            return escalate();
        }
    }
});

@Override
public SupervisorStrategy supervisorStrategy() {
    return strategy;
}

@Override
public void preStart() {
    initStorage();
}

/**
 * The child storage is restarted in case of failure, but after 3 restarts,
 * and still failing it will be stopped. Better to back-off than
 * continuously failing. When it has been stopped we will schedule a
 * Reconnect after a delay. Watch the child so we receive Terminated message
 * when it has been terminated.
 */
void initStorage() {
    storage = getContext().watch(getContext().actorOf(
        Props.create(Storage.class), "storage"));
    // Tell the counter, if any, to use the new storage
    if (counter != null)
        counter.tell(new UseStorage(storage), getSelf());
    // We need the initial value to be able to operate
    storage.tell(new Get(key), getSelf());
}

@Override
public void onReceive(Object msg) {
    log.debug("received message {}", msg);
    if (msg instanceof Entry && ((Entry) msg).key.equals(key) &&
        counter == null) {
        // Reply from Storage of the initial value, now we can create the Counter
        final long value = ((Entry) msg).value;
        counter = getContext().actorOf(Props.create(Counter.class, key, value));
        // Tell the counter to use current storage
        counter.tell(new UseStorage(storage), getSelf());
        // and send the buffered backlog to the counter
        for (SenderMsgPair each : backlog) {
            counter.tell(each.msg, each.sender);
        }
        backlog.clear();
    } else if (msg instanceof Increment) {
        forwardOrPlaceInBacklog(msg);
    } else if (msg.equals(GetCurrentCount)) {
        forwardOrPlaceInBacklog(msg);
    } else if (msg instanceof Terminated) {
        // After 3 restarts the storage child is stopped.
        // We receive Terminated because we watch the child, see initStorage.
        storage = null;
    }
}

```

```

    // Tell the counter that there is no storage for the moment
    counter.tell(new UseStorage(null), getSelf());
    // Try to re-establish storage after while
    getContext().system().scheduler().scheduleOnce(
        Duration.create(10, "seconds"), getSelf(), Reconnect,
        getContext().dispatcher(), null);
} else if (msg.equals(Reconnect)) {
    // Re-establish storage after the scheduled delay
    initStorage();
} else {
    unhandled(msg);
}
}

void forwardOrPlaceInBacklog(Object msg) {
    // We need the initial value from storage before we can start delegate to
    // the counter. Before that we place the messages in a backlog, to be sent
    // to the counter when it is initialized.
    if (counter == null) {
        if (backlog.size() >= MAX_BACKLOG)
            throw new ServiceUnavailable("CounterService not available," +
                " lack of initial value");
        backlog.add(new SenderMsgPair(getSender(), msg));
    } else {
        counter.forward(msg, getContext());
    }
}

public interface CounterApi {
    public static class UseStorage {
        public final ActorRef storage;

        public UseStorage(ActorRef storage) {
            this.storage = storage;
        }

        public String toString() {
            return String.format("%s(%s)", getClass().getSimpleName(), storage);
        }
    }
}

/**
 * The in memory count variable that will send current value to the Storage,
 * if there is any storage available at the moment.
 */
public static class Counter extends UntypedActor {
    final LoggingAdapter log = Logging.getLogger(getContext().system(), this);
    final String key;
    long count;
    ActorRef storage;

    public Counter(String key, long initialValue) {
        this.key = key;
        this.count = initialValue;
    }

    @Override
    public void onReceive(Object msg) {
        log.debug("received message {}", msg);
        if (msg instanceof UseStorage) {

```

```

        storage = ((UseStorage) msg).storage;
        storeCount();
    } else if (msg instanceof Increment) {
        count += ((Increment) msg).n;
        storeCount();
    } else if (msg.equals(GetCurrentCount)) {
        getSender().tell(new CurrentCount(key, count), getSelf());
    } else {
        unhandled(msg);
    }
}

void storeCount() {
    // Delegate dangerous work, to protect our valuable state.
    // We can continue without storage.
    if (storage != null) {
        storage.tell(new Store(new Entry(key, count)), getSelf());
    }
}

public interface StorageApi {

    public static class Store {
        public final Entry entry;

        public Store(Entry entry) {
            this.entry = entry;
        }

        public String toString() {
            return String.format("%s(%s)", getClass().getSimpleName(), entry);
        }
    }

    public static class Entry {
        public final String key;
        public final long value;

        public Entry(String key, long value) {
            this.key = key;
            this.value = value;
        }

        public String toString() {
            return String.format("%s(%s, %s)", getClass().getSimpleName(), key, value);
        }
    }

    public static class Get {
        public final String key;

        public Get(String key) {
            this.key = key;
        }

        public String toString() {
            return String.format("%s(%s)", getClass().getSimpleName(), key);
        }
    }

    public static class StorageException extends RuntimeException {
        private static final long serialVersionUID = 1L;
    }
}

```

```

    public StorageException(String msg) {
        super(msg);
    }
}

/**
 * Saves key/value pairs to persistent storage when receiving Store message.
 * Replies with current value when receiving Get message. Will throw
 * StorageException if the underlying data store is out of order.
 */
public static class Storage extends UntypedActor {

    final LoggingAdapter log = Logging.getLogger(getContext().system(), this);
    final DummyDB db = DummyDB.instance;

    @Override
    public void onReceive(Object msg) {
        log.debug("received message {}", msg);
        if (msg instanceof Store) {
            Store store = (Store) msg;
            db.save(store.entry.key, store.entry.value);
        } else if (msg instanceof Get) {
            Get get = (Get) msg;
            Long value = db.load(get.key);
            getSender().tell(new Entry(get.key, value == null ?
                Long.valueOf(0L) : value), getSelf());
        } else {
            unhandled(msg);
        }
    }
}

public static class DummyDB {
    public static final DummyDB instance = new DummyDB();
    private final Map<String, Long> db = new HashMap<String, Long>();

    private DummyDB() {
    }

    public synchronized void save(String key, Long value) throws StorageException {
        if (11 <= value && value <= 14)
            throw new StorageException("Simulated store failure " + value);
        db.put(key, value);
    }

    public synchronized Long load(String key) throws StorageException {
        return db.get(key);
    }
}
}

```

3.3.2 Creating a Supervisor Strategy

The following sections explain the fault handling mechanism and alternatives in more depth.

For the sake of demonstration let us consider the following strategy:

```

private static SupervisorStrategy strategy =
    new OneForOneStrategy(10, Duration.create("1 minute"),
        new Function<Throwable, Directive>() {

```

```

@Override
public Directive apply(Throwable t) {
    if (t instanceof ArithmeticException) {
        return resume();
    } else if (t instanceof NullPointerException) {
        return restart();
    } else if (t instanceof IllegalArgumentException) {
        return stop();
    } else {
        return escalate();
    }
}
});

@Override
public SupervisorStrategy supervisorStrategy() {
    return strategy;
}

```

I have chosen a few well-known exception types in order to demonstrate the application of the fault handling directives described in *Supervision and Monitoring*. First off, it is a one-for-one strategy, meaning that each child is treated separately (an all-for-one strategy works very similarly, the only difference is that any decision is applied to all children of the supervisor, not only the failing one). There are limits set on the restart frequency, namely maximum 10 restarts per minute. `-1` and `Duration.Inf()` means that the respective limit does not apply, leaving the possibility to specify an absolute upper limit on the restarts or to make the restarts work infinitely. The child actor is stopped if the limit is exceeded.

Note: If the strategy is declared inside the supervising actor (as opposed to a separate class) its decider has access to all internal state of the actor in a thread-safe fashion, including obtaining a reference to the currently failed child (available as the `getSender` of the failure message).

Default Supervisor Strategy

`Escalate` is used if the defined strategy doesn't cover the exception that was thrown.

When the supervisor strategy is not defined for an actor the following exceptions are handled by default:

- `ActorInitializationException` will stop the failing child actor
- `ActorKilledException` will stop the failing child actor
- `Exception` will restart the failing child actor
- Other types of `Throwable` will be escalated to parent actor

If the exception escalate all the way up to the root guardian it will handle it in the same way as the default strategy defined above.

Stopping Supervisor Strategy

Closer to the Erlang way is the strategy to just stop children when they fail and then take corrective action in the supervisor when `DeathWatch` signals the loss of the child. This strategy is also provided pre-packaged as `SupervisorStrategy.stoppingStrategy` with an accompanying `StoppingSupervisorStrategy` configurator to be used when you want the `"/user"` guardian to apply it.

Logging of Actor Failures

By default the `SupervisorStrategy` logs failures unless they are escalated. Escalated failures are supposed to be handled, and potentially logged, at a level higher in the hierarchy.

You can mute the default logging of a `SupervisorStrategy` by setting `loggingEnabled` to `false` when instantiating it. Customized logging can be done inside the `Decider`. Note that the reference to the currently failed child is available as the `getSender` when the `SupervisorStrategy` is declared inside the supervising actor.

You may also customize the logging in your own `SupervisorStrategy` implementation by overriding the `logFailure` method.

3.3.3 Supervision of Top-Level Actors

Toplevel actors means those which are created using `system.actorOf()`, and they are children of the *User Guardian*. There are no special rules applied in this case, the guardian simply applies the configured strategy.

3.3.4 Test Application

The following section shows the effects of the different directives in practice, where a test setup is needed. First off, we need a suitable supervisor:

```
public class Supervisor extends UntypedActor {

    private static SupervisorStrategy strategy =
        new OneForOneStrategy(10, Duration.create("1 minute"),
            new Function<Throwable, Directive>() {
                @Override
                public Directive apply(Throwable t) {
                    if (t instanceof ArithmeticException) {
                        return resume();
                    } else if (t instanceof NullPointerException) {
                        return restart();
                    } else if (t instanceof IllegalArgumentException) {
                        return stop();
                    } else {
                        return escalate();
                    }
                }
            });

    @Override
    public SupervisorStrategy supervisorStrategy() {
        return strategy;
    }

    public void onReceive(Object o) {
        if (o instanceof Props) {
            getSender().tell(getContext().actorOf((Props) o), getSelf());
        } else {
            unhandled(o);
        }
    }
}
```

This supervisor will be used to create a child, with which we can experiment:

```
public class Child extends UntypedActor {
    int state = 0;

    public void onReceive(Object o) throws Exception {
        if (o instanceof Exception) {
            throw (Exception) o;
        } else if (o instanceof Integer) {
```

```

    state = (Integer) o;
  } else if (o.equals("get")) {
    getSender().tell(state, getSelf());
  } else {
    unhandled(o);
  }
}
}
}

```

The test is easier by using the utilities described in *akka-testkit*, where `TestProbe` provides an actor ref useful for receiving and inspecting replies.

```

import akka.actor.ActorRef;
import akka.actor.ActorSystem;
import akka.actor.SupervisorStrategy;
import static akka.actor.SupervisorStrategy.resume;
import static akka.actor.SupervisorStrategy.restart;
import static akka.actor.SupervisorStrategy.stop;
import static akka.actor.SupervisorStrategy.escalate;
import akka.actor.SupervisorStrategy.Directive;
import akka.actor.OneForOneStrategy;
import akka.actor.Props;
import akka.actor.Terminated;
import akka.actor.UntypedActor;
import scala.collection.immutable.Seq;
import scala.concurrent.Await;
import static akka.pattern.Patterns.ask;
import scala.concurrent.duration.Duration;
import akka.testkit.TestProbe;

public class FaultHandlingTest {
    static ActorSystem system;
    Duration timeout = Duration.create(5, SECONDS);

    @BeforeClass
    public static void start() {
        system = ActorSystem.create("test");
    }

    @AfterClass
    public static void cleanup() {
        JavaTestKit.shutdownActorSystem(system);
        system = null;
    }

    @Test
    public void mustEmploySupervisorStrategy() throws Exception {
        // code here
    }
}

```

Let us create actors:

```

Props superprops = Props.create(Supervisor.class);
ActorRef supervisor = system.actorOf(superprops, "supervisor");
ActorRef child = (ActorRef) Await.result(ask(supervisor,
    Props.create(Child.class), 5000), timeout);

```

The first test shall demonstrate the `Resume` directive, so we try it out by setting some non-initial state in the actor and have it fail:

```
child.tell(42, ActorRef.noSender());
assert Await.result(ask(child, "get", 5000), timeout).equals(42);
child.tell(new ArithmeticException(), ActorRef.noSender());
assert Await.result(ask(child, "get", 5000), timeout).equals(42);
```

As you can see the value 42 survives the fault handling directive. Now, if we change the failure to a more serious `NullPointerException`, that will no longer be the case:

```
child.tell(new NullPointerException(), ActorRef.noSender());
assert Await.result(ask(child, "get", 5000), timeout).equals(0);
```

And finally in case of the fatal `IllegalArgumentException` the child will be terminated by the supervisor:

```
final TestProbe probe = new TestProbe(system);
probe.watch(child);
child.tell(new IllegalArgumentException(), ActorRef.noSender());
probe.expectMsgClass(Terminated.class);
```

Up to now the supervisor was completely unaffected by the child's failure, because the directives set did handle it. In case of an `Exception`, this is not true anymore and the supervisor escalates the failure.

```
child = (ActorRef) Await.result(ask(supervisor,
    Props.create(Child.class), 5000), timeout);
probe.watch(child);
assert Await.result(ask(child, "get", 5000), timeout).equals(0);
child.tell(new Exception(), ActorRef.noSender());
probe.expectMsgClass(Terminated.class);
```

The supervisor itself is supervised by the top-level actor provided by the `ActorSystem`, which has the default policy to restart in case of all `Exception` cases (with the notable exceptions of `ActorInitializationException` and `ActorKilledException`). Since the default directive in case of a restart is to kill all children, we expected our poor child not to survive this failure.

In case this is not desired (which depends on the use case), we need to use a different supervisor which overrides this behavior.

```
public class Supervisor2 extends UntypedActor {

    private static SupervisorStrategy strategy = new OneForOneStrategy(10,
        Duration.create("1 minute"),
        new Function<Throwable, Directive>() {
            @Override
            public Directive apply(Throwable t) {
                if (t instanceof ArithmeticException) {
                    return resume();
                } else if (t instanceof NullPointerException) {
                    return restart();
                } else if (t instanceof IllegalArgumentException) {
                    return stop();
                } else {
                    return escalate();
                }
            }
        }
    );

    @Override
    public SupervisorStrategy supervisorStrategy() {
        return strategy;
    }

    public void onReceive(Object o) {
        if (o instanceof Props) {
            getSender().tell(getContext().actorOf((Props) o), getSelf());
        }
    }
}
```

```

    } else {
        unhandled(o);
    }
}

@Override
public void preRestart(Throwable cause, Option<Object> msg) {
    // do not kill all children, which is the default here
}
}

```

With this parent, the child survives the escalated restart, as demonstrated in the last test:

```

superprops = Props.create(Supervisor2.class);
supervisor = system.actorOf(superprops);
child = (ActorRef) Await.result(ask(supervisor,
    Props.create(Child.class), 5000), timeout);
child.tell(23, ActorRef.noSender());
assert Await.result(ask(child, "get", 5000), timeout).equals(23);
child.tell(new Exception(), ActorRef.noSender());
assert Await.result(ask(child, "get", 5000), timeout).equals(0);

```

3.4 Dispatchers

An Akka MessageDispatcher is what makes Akka Actors “tick”, it is the engine of the machine so to speak. All MessageDispatcher implementations are also an ExecutionContext, which means that they can be used to execute arbitrary code, for instance *Futures*.

3.4.1 Default dispatcher

Every ActorSystem will have a default dispatcher that will be used in case nothing else is configured for an Actor. The default dispatcher can be configured, and is by default a Dispatcher with the specified default-executor. If an ActorSystem is created with an ExecutionContext passed in, this ExecutionContext will be used as the default executor for all dispatchers in this ActorSystem. If no ExecutionContext is given, it will fallback to the executor specified in akka.actor.default-dispatcher.default-executor.fallback. By default this is a “fork-join-executor”, which gives excellent performance in most cases.

3.4.2 Looking up a Dispatcher

Dispatchers implement the ExecutionContext interface and can thus be used to run Future invocations etc.

```

// this is scala.concurrent.ExecutionContext
// for use with Futures, Scheduler, etc.
final ExecutionContext ex = system.dispatchers().lookup("my-dispatcher");

```

3.4.3 Setting the dispatcher for an Actor

So in case you want to give your Actor a different dispatcher than the default, you need to do two things, of which the first is to configure the dispatcher:

```

my-dispatcher {
    # Dispatcher is the name of the event-based dispatcher
    type = Dispatcher
    # What kind of ExecutionService to use
    executor = "fork-join-executor"
}

```

```
# Configuration for the fork join pool
fork-join-executor {
  # Min number of threads to cap factor-based parallelism number to
  parallelism-min = 2
  # Parallelism (threads) ... ceil(available processors * factor)
  parallelism-factor = 2.0
  # Max number of threads to cap factor-based parallelism number to
  parallelism-max = 10
}
# Throughput defines the maximum number of messages to be
# processed per actor before the thread jumps to the next actor.
# Set to 1 for as fair as possible.
throughput = 100
}
```

And here's another example that uses the "thread-pool-executor":

```
my-thread-pool-dispatcher {
  # Dispatcher is the name of the event-based dispatcher
  type = Dispatcher
  # What kind of ExecutionService to use
  executor = "thread-pool-executor"
  # Configuration for the thread pool
  thread-pool-executor {
    # minimum number of threads to cap factor-based core number to
    core-pool-size-min = 2
    # No of core threads ... ceil(available processors * factor)
    core-pool-size-factor = 2.0
    # maximum number of threads to cap factor-based number to
    core-pool-size-max = 10
  }
  # Throughput defines the maximum number of messages to be
  # processed per actor before the thread jumps to the next actor.
  # Set to 1 for as fair as possible.
  throughput = 100
}
```

For more options, see the default-dispatcher section of the [Configuration](#).

Then you create the actor as usual and define the dispatcher in the deployment configuration.

```
ActorRef myActor =
  system.actorOf(Props.create(MyUntypedActor.class),
    "myactor");
```

```
akka.actor.deployment {
  /myactor {
    dispatcher = my-dispatcher
  }
}
```

An alternative to the deployment configuration is to define the dispatcher in code. If you define the dispatcher in the deployment configuration then this value will be used instead of programmatically provided parameter.

```
ActorRef myActor =
  system.actorOf(Props.create(MyUntypedActor.class).withDispatcher("my-dispatcher"),
    "myactor3");
```

Note: The dispatcher you specify in `withDispatcher` and the dispatcher property in the deployment configuration is in fact a path into your configuration. So in this example it's a top-level section, but you could for instance put it as a sub-section, where you'd use periods to denote sub-sections, like this: `"foo.bar.my-dispatcher"`

3.4.4 Types of dispatchers

There are 4 different types of message dispatchers:

- Dispatcher
 - This is an event-based dispatcher that binds a set of Actors to a thread pool. It is the default dispatcher used if one is not specified.
 - Sharability: Unlimited
 - Mailboxes: Any, creates one per Actor
 - Use cases: Default dispatcher, Bulkheading
 - **Driven by:** `java.util.concurrent.ExecutorService` specify using “executor” using “fork-join-executor”, “thread-pool-executor” or the FQCN of an `akka.dispatcher.ExecutorServiceConfigurator`
- PinnedDispatcher
 - This dispatcher dedicates a unique thread for each actor using it; i.e. each actor will have its own thread pool with only one thread in the pool.
 - Sharability: None
 - Mailboxes: Any, creates one per Actor
 - Use cases: Bulkheading
 - **Driven by:** Any `akka.dispatch.ThreadPoolExecutorConfigurator` by default a “thread-pool-executor”
- CallingThreadDispatcher
 - This dispatcher runs invocations on the current thread only. This dispatcher does not create any new threads, but it can be used from different threads concurrently for the same actor. See [CallingThreadDispatcher](#) for details and restrictions.
 - Sharability: Unlimited
 - Mailboxes: Any, creates one per Actor per Thread (on demand)
 - Use cases: Testing
 - Driven by: The calling thread (duh)

More dispatcher configuration examples

Configuring a PinnedDispatcher:

```
my-pinned-dispatcher {
  executor = "thread-pool-executor"
  type = PinnedDispatcher
}
```

And then using it:

```
ActorRef myActor = system.actorOf(Props.create(MyUntypedActor.class)
  .withDispatcher("my-pinned-dispatcher"));
```

Note that `thread-pool-executor` configuration as per the above `my-thread-pool-dispatcher` example is NOT applicable. This is because every actor will have its own thread pool when using `PinnedDispatcher`, and that pool will have only one thread.

Note that it's not guaranteed that the *same* thread is used over time, since the core pool timeout is used for `PinnedDispatcher` to keep resource usage down in case of idle actors. To use the same thread all the time you need to add `thread-pool-executor.allow-core-timeout=off` to the configuration of the `PinnedDispatcher`.

3.5 Mailboxes

An Akka `Mailbox` holds the messages that are destined for an `Actor`. Normally each `Actor` has its own mailbox, but with for example a `BalancingPool` all routees will share a single mailbox instance.

3.5.1 Mailbox Selection

Requiring a Message Queue Type for an Actor

It is possible to require a certain type of message queue for a certain type of actor by having that actor implement the parameterized interface `RequiresMessageQueue`. Here is an example:

```
import akka.dispatch.BoundedMessageQueueSemantics;
import akka.dispatch.RequiresMessageQueue;

public class MyBoundedUntypedActor extends MyUntypedActor
    implements RequiresMessageQueue<BoundedMessageQueueSemantics> {
}
```

The type parameter to the `RequiresMessageQueue` interface needs to be mapped to a mailbox in configuration like this:

```
bounded-mailbox {
  mailbox-type = "akka.dispatch.BoundedMailbox"
  mailbox-capacity = 1000
  mailbox-push-timeout-time = 10s
}

akka.actor.mailbox.requirements {
  "akka.dispatch.BoundedMessageQueueSemantics" = bounded-mailbox
}
```

Now every time you create an actor of type `MyBoundedUntypedActor` it will try to get a bounded mailbox. If the actor has a different mailbox configured in deployment, either directly or via a dispatcher with a specified mailbox type, then that will override this mapping.

Note: The type of the queue in the mailbox created for an actor will be checked against the required type in the interface and if the queue doesn't implement the required type then actor creation will fail.

Requiring a Message Queue Type for a Dispatcher

A dispatcher may also have a requirement for the mailbox type used by the actors running on it. An example is the `BalancingDispatcher` which requires a message queue that is thread-safe for multiple concurrent consumers. Such a requirement is formulated within the dispatcher configuration section like this:

```
my-dispatcher {
  mailbox-requirement = org.example.MyInterface
}
```

The given requirement names a class or interface which will then be ensured to be a supertype of the message queue's implementation. In case of a conflict—e.g. if the actor requires a mailbox type which does not satisfy this requirement—then actor creation will fail.

How the Mailbox Type is Selected

When an actor is created, the `ActorRefProvider` first determines the dispatcher which will execute it. Then the mailbox is determined as follows:

1. If the actor's deployment configuration section contains a `mailbox` key then that names a configuration section describing the mailbox type to be used.
2. If the actor's `Props` contains a mailbox selection—i.e. `withMailbox` was called on it—then that names a configuration section describing the mailbox type to be used.
3. If the dispatcher's configuration section contains a `mailbox-type` key the same section will be used to configure the mailbox type.
4. If the actor requires a mailbox type as described above then the mapping for that requirement will be used to determine the mailbox type to be used; if that fails then the dispatcher's requirement—if any—will be tried instead.
5. If the dispatcher requires a mailbox type as described above then the mapping for that requirement will be used to determine the mailbox type to be used.
6. The default mailbox `akka.actor.default-mailbox` will be used.

Default Mailbox

When the mailbox is not specified as described above the default mailbox is used. By default it is an unbounded mailbox, which is backed by a `java.util.concurrent.ConcurrentLinkedQueue`.

`SingleConsumerOnlyUnboundedMailbox` is an even more efficient mailbox, and it can be used as the default mailbox, but it cannot be used with a `BalancingDispatcher`.

Configuration of `SingleConsumerOnlyUnboundedMailbox` as default mailbox:

```
akka.actor.default-mailbox {
  mailbox-type = "akka.dispatch.SingleConsumerOnlyUnboundedMailbox"
}
```

Which Configuration is passed to the Mailbox Type

Each mailbox type is implemented by a class which extends `MailboxType` and takes two constructor arguments: a `ActorSystem.Settings` object and a `Config` section. The latter is computed by obtaining the named configuration section from the actor system's configuration, overriding its `id` key with the configuration path of the mailbox type and adding a fall-back to the default mailbox configuration section.

3.5.2 Builtin Mailbox Implementations

Akka comes shipped with a number of mailbox implementations:

- **UnboundedMailbox** (default)
 - The default mailbox
 - Backed by a `java.util.concurrent.ConcurrentLinkedQueue`
 - Blocking: No
 - Bounded: No
 - Configuration name: `"unbounded"` or `"akka.dispatch.UnboundedMailbox"`
- **SingleConsumerOnlyUnboundedMailbox**
 - Backed by a very efficient Multiple-Producer Single-Consumer queue, cannot be used with `BalancingDispatcher`
 - Blocking: No
 - Bounded: No
 - Configuration name: `"akka.dispatch.SingleConsumerOnlyUnboundedMailbox"`

- **NonBlockingBoundedMailbox**
 - Backed by a very efficient Multiple-Producer Single-Consumer queue
 - Blocking: No (discards overflowing messages into deadLetters)
 - Bounded: Yes
 - Configuration name: "akka.dispatch.NonBlockingBoundedMailbox"
- **UnboundedControlAwareMailbox**
 - Delivers messages that extend akka.dispatch.ControlMessage with higher priority
 - Backed by two java.util.concurrent.ConcurrentLinkedQueue
 - Blocking: No
 - Bounded: No
 - Configuration name: "akka.dispatch.UnboundedControlAwareMailbox"

Additional implementations exist which can be potentially blocking on pushing into the queue, however they can be configured using

- **BoundedMailbox**
 - Backed by a java.util.concurrent.LinkedBlockingQueue
 - Blocking: Yes
 - Bounded: Yes
 - Configuration name: "bounded" or "akka.dispatch.BoundedMailbox"
- **UnboundedPriorityMailbox**
 - Backed by a java.util.concurrent.PriorityBlockingQueue
 - Delivery order for messages of equal priority is undefined - contrast with the UnboundedStablePriorityMailbox
 - Blocking: Yes
 - Bounded: No
 - Configuration name: "akka.dispatch.UnboundedPriorityMailbox"
- **BoundedPriorityMailbox**
 - Backed by a java.util.PriorityQueue wrapped in an akka.util.BoundedBlockingQueue
 - Delivery order for messages of equal priority is undefined - contrast with the BoundedStablePriorityMailbox
 - Blocking: Yes
 - Bounded: Yes
 - Configuration name: "akka.dispatch.BoundedPriorityMailbox"
- **UnboundedStablePriorityMailbox**
 - Backed by a java.util.concurrent.PriorityBlockingQueue wrapped in an akka.util.PriorityQueueStabilizer
 - FIFO order is preserved for messages of equal priority - contrast with the UnboundedPriorityMailbox
 - Blocking: Yes
 - Bounded: No
 - Configuration name: "akka.dispatch.UnboundedStablePriorityMailbox"
- **BoundedStablePriorityMailbox**

- Backed by a `java.util.PriorityQueue` wrapped in an `akka.util.PriorityQueueStabilizer` and an `akka.util.BoundedBlockingQueue`
- FIFO order is preserved for messages of equal priority - contrast with the `BoundedPriorityMailbox`
- Blocking: Yes
- Bounded: Yes
- Configuration name: `"akka.dispatch.BoundedStablePriorityMailbox"`

- **BoundedControlAwareMailbox**

- Delivers messages that extend `akka.dispatch.ControlMessage` with higher priority
- Backed by two `java.util.concurrent.ConcurrentLinkedQueue` and blocking on enqueue if capacity has been reached
- Blocking: Yes
- Bounded: Yes
- Configuration name: `"akka.dispatch.BoundedControlAwareMailbox"`

3.5.3 Mailbox configuration examples

PriorityMailbox

How to create a `PriorityMailbox`:

```
public class MyPrioMailbox extends UnboundedStablePriorityMailbox {
    // needed for reflective instantiation
    public MyPrioMailbox(ActorSystem.Settings settings, Config config) {
        // Create a new PriorityGenerator, lower prio means more important
        super(new PriorityGenerator() {
            @Override
            public int gen(Object message) {
                if (message.equals("highpriority"))
                    return 0; // 'highpriority messages should be treated first if possible
                else if (message.equals("lowpriority"))
                    return 2; // 'lowpriority messages should be treated last if possible
                else if (message.equals(PoisonPill.getInstance()))
                    return 3; // PoisonPill when no other left
                else
                    return 1; // By default they go between high and low prio
            }
        });
    }
}
```

And then add it to the configuration:

```
prio-dispatcher {
    mailbox-type = "docs.dispatcher.DispatcherDocSpec$MyPrioMailbox"
    //Other dispatcher configuration goes here
}
```

And then an example on how you would use it:

```
class Demo extends UntypedActor {
    LoggingAdapter log = Logging.getLogger(getContext().system(), this);
    {
        for (Object msg : new Object[] { "lowpriority", "lowpriority",
            "highpriority", "pigdog", "pigdog2", "pigdog3", "highpriority",
            PoisonPill.getInstance() }) {
            getSelf().tell(msg, getSelf());
        }
    }
}
```

```

    }
  }

  public void onReceive(Object message) {
    log.info(message.toString());
  }
}

// We create a new Actor that just prints out what it processes
ActorRef myActor = system.actorOf(Props.create(Demo.class, this)
    .withDispatcher("prio-dispatcher"));

/*
Logs:
'highpriority
'highpriority
'pigdog
'pigdog2
'pigdog3
'lowpriority
'lowpriority
*/

```

It is also possible to configure a mailbox type directly like this:

```

prio-mailbox {
  mailbox-type = "docs.dispatcher.DispatcherDocSpec$MyPrioMailbox"
  //Other mailbox configuration goes here
}

akka.actor.deployment {
  /priomailboxactor {
    mailbox = prio-mailbox
  }
}

```

And then use it either from deployment like this:

```

ActorRef myActor =
  system.actorOf(Props.create(MyUntypedActor.class),
    "priomailboxactor");

```

Or code like this:

```

ActorRef myActor =
  system.actorOf(Props.create(MyUntypedActor.class)
    .withMailbox("prio-mailbox"));

```

ControlAwareMailbox

A `ControlAwareMailbox` can be very useful if an actor needs to be able to receive control messages immediately no matter how many other messages are already in its mailbox.

It can be configured like this:

```

control-aware-dispatcher {
  mailbox-type = "akka.dispatch.UnboundedControlAwareMailbox"
  //Other dispatcher configuration goes here
}

```

Control messages need to extend the `ControlMessage` trait:

```
public class MyControlMessage implements ControlMessage {}
```

And then an example on how you would use it:

```
class Demo extends UntypedActor {
  LoggingAdapter log = Logging.getLogger(getContext().system(), this);
  {
    for (Object msg : new Object[] { "foo", "bar", new MyControlMessage(),
      PoisonPill.getInstance() }) {
      getSelf().tell(msg, getSelf());
    }
  }

  public void onReceive(Object message) {
    log.info(message.toString());
  }
}

// We create a new Actor that just prints out what it processes
ActorRef myActor = system.actorOf(Props.create(Demo.class, this)
  .withDispatcher("control-aware-dispatcher"));

/*
Logs:
  'MyControlMessage
  'foo
  'bar
*/
```

3.5.4 Creating your own Mailbox type

An example is worth a thousand quacks:

```
import akka.actor.ActorRef;
import akka.actor.ActorSystem;
import akka.dispatch.Envelope;
import akka.dispatch.MailboxType;
import akka.dispatch.MessageQueue;
import akka.dispatch.ProducesMessageQueue;
import com.typesafe.config.Config;
import java.util.concurrent.ConcurrentLinkedQueue;
import java.util.Queue;
import scala.Option;

public class MyUnboundedJMailbox implements MailboxType,
  ProducesMessageQueue<MyUnboundedJMailbox.MyMessageQueue> {

  // This is the MessageQueue implementation
  public static class MyMessageQueue implements MessageQueue,
    MyUnboundedJMessageQueueSemantics {
    private final Queue<Envelope> queue =
      new ConcurrentLinkedQueue<Envelope>();

    // these must be implemented; queue used as example
    public void enqueue(ActorRef receiver, Envelope handle) {
      queue.offer(handle);
    }

    public Envelope dequeue() { return queue.poll(); }
    public int numberOfMessages() { return queue.size(); }
    public boolean hasMessages() { return !queue.isEmpty(); }
    public void cleanUp(ActorRef owner, MessageQueue deadLetters) {
      for (Envelope handle: queue) {
```

```

        deadLetters.enqueue(owner, handle);
    }
}

// This constructor signature must exist, it will be called by Akka
public MyUnboundedJMailbox(ActorSystem.Settings settings, Config config) {
    // put your initialization code here
}

// The create method is called to create the MessageQueue
public MessageQueue create(Option<ActorRef> owner, Option<ActorSystem> system) {
    return new MyMessageQueue();
}
}

// Marker interface used for mailbox requirements mapping
public interface MyUnboundedJMessageQueueSemantics {
}

```

And then you just specify the FQCN of your MailboxType as the value of the “mailbox-type” in the dispatcher configuration, or the mailbox configuration.

Note: Make sure to include a constructor which takes `akka.actor.ActorSystem.Settings` and `com.typesafe.config.Config` arguments, as this constructor is invoked reflectively to construct your mailbox type. The config passed in as second argument is that section from the configuration which describes the dispatcher or mailbox setting using this mailbox type; the mailbox type will be instantiated once for each dispatcher or mailbox setting using it.

You can also use the mailbox as a requirement on the dispatcher like this:

```

custom-dispatcher {
  mailbox-requirement =
    "docs.dispatcher.MyUnboundedJMessageQueueSemantics"
}

akka.actor.mailbox.requirements {
  "docs.dispatcher.MyUnboundedJMessageQueueSemantics" =
    custom-dispatcher-mailbox
}

custom-dispatcher-mailbox {
  mailbox-type = "docs.dispatcher.MyUnboundedJMailbox"
}

```

Or by defining the requirement on your actor class like this:

```

public class MySpecialActor extends UntypedActor implements
  RequiresMessageQueue<MyUnboundedJMessageQueueSemantics> {
  // ...
}

```

3.5.5 Special Semantics of `system.actorOf`

In order to make `system.actorOf` both synchronous and non-blocking while keeping the return type `ActorRef` (and the semantics that the returned ref is fully functional), special handling takes place for this case. Behind the scenes, a hollow kind of actor reference is constructed, which is sent to the system’s guardian actor who actually creates the actor and its context and puts those inside the reference. Until that has happened, messages sent to the `ActorRef` will be queued locally, and only upon swapping the real filling in will they be transferred into the real mailbox. Thus,

```
final Props props = ...
// this actor uses MyCustomMailbox, which is assumed to be a singleton
system.actorOf(props.withDispatcher("myCustomMailbox").tell("bang", sender);
assert(MyCustomMailbox.getInstance().getLastEnqueued().equals("bang"));
```

will probably fail; you will have to allow for some time to pass and retry the check à la `TestKit.awaitCond`.

3.6 Routing

Messages can be sent via a router to efficiently route them to destination actors, known as its *routees*. A Router can be used inside or outside of an actor, and you can manage the routees yourselves or use a self contained router actor with configuration capabilities.

Different routing strategies can be used, according to your application's needs. Akka comes with several useful routing strategies right out of the box. But, as you will see in this chapter, it is also possible to *create your own*.

3.6.1 A Simple Router

The following example illustrates how to use a Router and manage the routees from within an actor.

```
public final class Work implements Serializable {
    private static final long serialVersionUID = 1L;
    public final String payload;
    public Work(String payload) {
        this.payload = payload;
    }
}

public class Master extends UntypedActor {

    Router router;
    {
        List<Routee> routees = new ArrayList<Routee>();
        for (int i = 0; i < 5; i++) {
            ActorRef r = getContext().actorOf(Props.create(Worker.class));
            getContext().watch(r);
            routees.add(new ActorRefRoutee(r));
        }
        router = new Router(new RoundRobinRoutingLogic(), routees);
    }

    public void onReceive(Object msg) {
        if (msg instanceof Work) {
            router.route(msg, getSender());
        } else if (msg instanceof Terminated) {
            router = router.removeRoutee(((Terminated) msg).actor());
            ActorRef r = getContext().actorOf(Props.create(Worker.class));
            getContext().watch(r);
            router = router.addRoutee(new ActorRefRoutee(r));
        }
    }
}
```

We create a Router and specify that it should use `RoundRobinRoutingLogic` when routing the messages to the routees.

The routing logic shipped with Akka are:

- `akka.routing.RoundRobinRoutingLogic`
- `akka.routing.RandomRoutingLogic`

- `akka.routing.SmallestMailboxRoutingLogic`
- `akka.routing.BroadcastRoutingLogic`
- `akka.routing.ScatterGatherFirstCompletedRoutingLogic`
- `akka.routing.TailChoppingRoutingLogic`
- `akka.routing.ConsistentHashingRoutingLogic`

We create the routees as ordinary child actors wrapped in `ActorRefRoutee`. We watch the routees to be able to replace them if they are terminated.

Sending messages via the router is done with the `route` method, as is done for the `Work` messages in the example above.

The `Router` is immutable and the `RoutingLogic` is thread safe; meaning that they can also be used outside of actors.

Note: In general, any message sent to a router will be sent onwards to its routees, but there is one exception. The special *Broadcast Messages* will send to *all* of a router's routees

3.6.2 A Router Actor

A router can also be created as a self contained actor that manages the routees itself and loads routing logic and other settings from configuration.

This type of router actor comes in two distinct flavors:

- **Pool** - The router creates routees as child actors and removes them from the router if they terminate.
- **Group** - The routee actors are created externally to the router and the router sends messages to the specified path using actor selection, without watching for termination.

The settings for a router actor can be defined in configuration or programmatically. Although router actors can be defined in the configuration file, they must still be created programmatically, i.e. you cannot make a router through external configuration alone. If you define the router actor in the configuration file then these settings will be used instead of any programmatically provided parameters.

You send messages to the routees via the router actor in the same way as for ordinary actors, i.e. via its `ActorRef`. The router actor forwards messages onto its routees without changing the original sender. When a routee replies to a routed message, the reply will be sent to the original sender, not to the router actor.

Note: In general, any message sent to a router will be sent onwards to its routees, but there are a few exceptions. These are documented in the *Specially Handled Messages* section below.

Pool

The following code and configuration snippets show how to create a *round-robin* router that forwards messages to five `Worker` routees. The routees will be created as the router's children.

```
akka.actor.deployment {
  /parent/router1 {
    router = round-robin-pool
    nr-of-instances = 5
  }
}
```

```
ActorRef router1 =
  getContext().actorOf(FromConfig.getInstance().props(Props.create(Worker.class)),
    "router1");
```

Here is the same example, but with the router configuration provided programmatically instead of from configuration.

```
ActorRef router2 =
  getContext().actorOf(new RoundRobinPool(5).props(Props.create(Worker.class)),
    "router2");
```

Remote Deployed Routees

In addition to being able to create local actors as routees, you can instruct the router to deploy its created children on a set of remote hosts. Routees will be deployed in round-robin fashion. In order to deploy routees remotely, wrap the router configuration in a `RemoteRouterConfig`, attaching the remote addresses of the nodes to deploy to. Remote deployment requires the `akka-remote` module to be included in the classpath.

```
Address[] addresses = {
  new Address("akka.tcp", "remotesys", "otherhost", 1234),
  AddressFromURIStrng.parse("akka.tcp://othersys@anotherhost:1234")};
ActorRef routerRemote = system.actorOf(
  new RemoteRouterConfig(new RoundRobinPool(5), addresses).props(
    Props.create(Echo.class)));
```

Senders

When a routee sends a message, it can *set itself as the sender*.

```
getSender().tell("reply", getSelf());
```

However, it is often useful for routees to set the *router* as a sender. For example, you might want to set the router as the sender if you want to hide the details of the routees behind the router. The following code snippet shows how to set the parent router as sender.

```
getSender().tell("reply", getContext().parent());
```

Supervision

Routees that are created by a pool router will be created as the router's children. The router is therefore also the children's supervisor.

The supervision strategy of the router actor can be configured with the `supervisorStrategy` property of the `Pool`. If no configuration is provided, routers default to a strategy of "always escalate". This means that errors are passed up to the router's supervisor for handling. The router's supervisor will decide what to do about any errors.

Note the router's supervisor will treat the error as an error with the router itself. Therefore a directive to stop or restart will cause the router *itself* to stop or restart. The router, in turn, will cause its children to stop and restart.

It should be mentioned that the router's restart behavior has been overridden so that a restart, while still re-creating the children, will still preserve the same number of actors in the pool.

This means that if you have not specified `supervisorStrategy` of the router or its parent a failure in a routee will escalate to the parent of the router, which will by default restart the router, which will restart all routees (it uses `Escalate` and does not stop routees during restart). The reason is to make the default behave such that adding `withRouter` to a child's definition does not change the supervision strategy applied to the child. This might be an inefficiency that you can avoid by specifying the strategy when defining the router.

Setting the strategy is easily done:

```
final SupervisorStrategy strategy =
  new OneForOneStrategy(5, Duration.create(1, TimeUnit.MINUTES),
    Collections.<Class<? extends Throwable>>singletonList(Exception.class));
```

```
final ActorRef router = system.actorOf(new RoundRobinPool(5).
    withSupervisorStrategy(strategy).props(Props.create(Echo.class)));
```

Note: If the child of a pool router terminates, the pool router will not automatically spawn a new child. In the event that all children of a pool router have terminated the router will terminate itself unless it is a dynamic router, e.g. using a resizer.

Group

Sometimes, rather than having the router actor create its routees, it is desirable to create routees separately and provide them to the router for its use. You can do this by passing an paths of the routees to the router's configuration. Messages will be sent with `ActorSelection` to these paths.

The example below shows how to create a router by providing it with the path strings of three routee actors.

```
akka.actor.deployment {
  /parent/router3 {
    router = round-robin-group
    routees.paths = ["/user/workers/w1", "/user/workers/w2", "/user/workers/w3"]
  }
}
```

```
ActorRef router3 =
  getContext().actorOf(FromConfig.getInstance().props(), "router3");
```

Here is the same example, but with the router configuration provided programmatically instead of from configuration.

```
ActorRef router4 =
  getContext().actorOf(new RoundRobinGroup(paths).props(), "router4");
```

The routee actors are created externally from the router:

```
system.actorOf(Props.create(Workers.class), "workers");
```

```
public class Workers extends UntypedActor {
  @Override public void preStart() {
    getContext().actorOf(Props.create(Worker.class), "w1");
    getContext().actorOf(Props.create(Worker.class), "w2");
    getContext().actorOf(Props.create(Worker.class), "w3");
  }
  // ...
}
```

The paths may contain protocol and address information for actors running on remote hosts. Remoting requires the `akka-remote` module to be included in the classpath.

```
akka.actor.deployment {
  /parent/remoteGroup {
    router = round-robin-group
    routees.paths = [
      "akka.tcp://app@10.0.0.1:2552/user/workers/w1",
      "akka.tcp://app@10.0.0.2:2552/user/workers/w1",
      "akka.tcp://app@10.0.0.3:2552/user/workers/w1"
    ]
  }
}
```

3.6.3 Router usage

In this section we will describe how to create the different types of router actors.

The router actors in this section are created from within a top level actor named `parent`. Note that deployment paths in the configuration starts with `/parent/` followed by the name of the router actor.

```
system.actorOf(Props.create(Parent.class), "parent");
```

RoundRobinPool and RoundRobinGroup

Routes in a [round-robin](#) fashion to its routees.

RoundRobinPool defined in configuration:

```
akka.actor.deployment {
  /parent/router1 {
    router = round-robin-pool
    nr-of-instances = 5
  }
}
```

```
ActorRef router1 =
  getContext().actorOf(FromConfig.getInstance().props(Props.create(Worker.class)),
    "router1");
```

RoundRobinPool defined in code:

```
ActorRef router2 =
  getContext().actorOf(new RoundRobinPool(5).props(Props.create(Worker.class)),
    "router2");
```

RoundRobinGroup defined in configuration:

```
akka.actor.deployment {
  /parent/router3 {
    router = round-robin-group
    routees.paths = ["/user/workers/w1", "/user/workers/w2", "/user/workers/w3"]
  }
}
```

```
ActorRef router3 =
  getContext().actorOf(FromConfig.getInstance().props(), "router3");
```

RoundRobinGroup defined in code:

```
List<String> paths = Arrays.asList("/user/workers/w1", "/user/workers/w2",
  "/user/workers/w3");
ActorRef router4 =
  getContext().actorOf(new RoundRobinGroup(paths).props(), "router4");
```

RandomPool and RandomGroup

This router type selects one of its routees randomly for each message.

RandomPool defined in configuration:

```
akka.actor.deployment {
  /parent/router5 {
    router = random-pool
    nr-of-instances = 5
  }
}
```

```
ActorRef router5 =
  getContext().actorOf(FromConfig.getInstance().props(
    Props.create(Worker.class)), "router5");
```

RandomPool defined in code:

```
ActorRef router6 =
  getContext().actorOf(new RandomPool(5).props(Props.create(Worker.class)),
    "router6");
```

RandomGroup defined in configuration:

```
akka.actor.deployment {
  /parent/router7 {
    router = random-group
    routees.paths = ["/user/workers/w1", "/user/workers/w2", "/user/workers/w3"]
  }
}
```

```
ActorRef router7 =
  getContext().actorOf(FromConfig.getInstance().props(), "router7");
```

RandomGroup defined in code:

```
List<String> paths = Arrays.asList("/user/workers/w1", "/user/workers/w2",
  "/user/workers/w3");
ActorRef router8 =
  getContext().actorOf(new RandomGroup(paths).props(), "router8");
```

BalancingPool

A Router that will try to redistribute work from busy routees to idle routees. All routees share the same mailbox.

BalancingPool defined in configuration:

```
akka.actor.deployment {
  /parent/router9 {
    router = balancing-pool
    nr-of-instances = 5
  }
}
```

```
ActorRef router9 =
  getContext().actorOf(FromConfig.getInstance().props(
    Props.create(Worker.class)), "router9");
```

BalancingPool defined in code:

```
ActorRef router10 =
  getContext().actorOf(new BalancingPool(5).props(
    Props.create(Worker.class)), "router10");
```

Additional configuration for the balancing dispatcher, which is used by the pool, can be configured in the pool-dispatcher section of the router deployment configuration.

```
akka.actor.deployment {
  /parent/router9b {
    router = balancing-pool
    nr-of-instances = 5
    pool-dispatcher {
      attempt-teamwork = off
    }
  }
}
```

The BalancingPool automatically uses a special BalancingDispatcher for its routees - disregarding any dispatcher that is set on the the routee Props object. This is needed in order to implement the balancing semantics via sharing the same mailbox by all the routees.

While it is not possible to change the dispatcher used by the routees, it is possible to fine tune the used *executor*. By default the `fork-join-dispatcher` is used and can be configured as explained in *Dispatchers*. In situations where the routees are expected to perform blocking operations it may be useful to replace it with a `thread-pool-executor` hinting the number of allocated threads explicitly:

```
akka.actor.deployment {
  /parent/router10b {
    router = balancing-pool
    nr-of-instances = 5
    pool-dispatcher {
      executor = "thread-pool-executor"

      # allocate exactly 5 threads for this pool
      thread-pool-executor {
        core-pool-size-min = 5
        core-pool-size-max = 5
      }
    }
  }
}
```

There is no Group variant of the `BalancingPool`.

SmallestMailboxPool

A Router that tries to send to the non-suspended child routee with fewest messages in mailbox. The selection is done in this order:

- pick any idle routee (not processing message) with empty mailbox
- pick any routee with empty mailbox
- pick routee with fewest pending messages in mailbox
- pick any remote routee, remote actors are consider lowest priority, since their mailbox size is unknown

`SmallestMailboxPool` defined in configuration:

```
akka.actor.deployment {
  /parent/router11 {
    router = smallest-mailbox-pool
    nr-of-instances = 5
  }
}
```

```
ActorRef router11 =
  getContext().actorOf(FromConfig.getInstance().props(
    Props.create(Worker.class)), "router11");
```

`SmallestMailboxPool` defined in code:

```
ActorRef router12 =
  getContext().actorOf(new SmallestMailboxPool(5).props(
    Props.create(Worker.class)), "router12");
```

There is no Group variant of the `SmallestMailboxPool` because the size of the mailbox and the internal dispatching state of the actor is not practically available from the paths of the routees.

BroadcastPool and BroadcastGroup

A broadcast router forwards the message it receives to *all* its routees.

`BroadcastPool` defined in configuration:

```
akka.actor.deployment {
  /parent/router13 {
    router = broadcast-pool
    nr-of-instances = 5
  }
}
```

```
ActorRef router13 =
  getContext().actorOf(FromConfig.getInstance().props(
    Props.create(Worker.class)), "router13");
```

BroadcastPool defined in code:

```
ActorRef router14 =
  getContext().actorOf(new BroadcastPool(5).props(Props.create(Worker.class)),
    "router14");
```

BroadcastGroup defined in configuration:

```
akka.actor.deployment {
  /parent/router15 {
    router = broadcast-group
    routees.paths = ["/user/workers/w1", "/user/workers/w2", "/user/workers/w3"]
  }
}
```

```
ActorRef router15 =
  getContext().actorOf(FromConfig.getInstance().props(), "router15");
```

BroadcastGroup defined in code:

```
List<String> paths = Arrays.asList("/user/workers/w1", "/user/workers/w2",
  "/user/workers/w3");
ActorRef router16 =
  getContext().actorOf(new BroadcastGroup(paths).props(), "router16");
```

Note: Broadcast routers always broadcast *every* message to their routees. If you do not want to broadcast every message, then you can use a non-broadcasting router and use *Broadcast Messages* as needed.

ScatterGatherFirstCompletedPool and ScatterGatherFirstCompletedGroup

The `ScatterGatherFirstCompletedRouter` will send the message on to all its routees. It then waits for first reply it gets back. This result will be sent back to original sender. Other replies are discarded.

It is expecting at least one reply within a configured duration, otherwise it will reply with `akka.pattern.AskTimeoutException` in a `akka.actor.Status.Failure`.

ScatterGatherFirstCompletedPool defined in configuration:

```
akka.actor.deployment {
  /parent/router17 {
    router = scatter-gather-pool
    nr-of-instances = 5
    within = 10 seconds
  }
}
```

```
ActorRef router17 =
  getContext().actorOf(FromConfig.getInstance().props(
    Props.create(Worker.class)), "router17");
```

ScatterGatherFirstCompletedPool defined in code:

```
FiniteDuration within = FiniteDuration.create(10, TimeUnit.SECONDS);
ActorRef router18 =
  getContext().actorOf(new ScatterGatherFirstCompletedPool(5, within).props(
    Props.create(Worker.class)), "router18");
```

ScatterGatherFirstCompletedGroup defined in configuration:

```
akka.actor.deployment {
  /parent/router19 {
    router = scatter-gather-group
    routees.paths = ["/user/workers/w1", "/user/workers/w2", "/user/workers/w3"]
    within = 10 seconds
  }
}
```

```
ActorRef router19 =
  getContext().actorOf(FromConfig.getInstance().props(), "router19");
```

ScatterGatherFirstCompletedGroup defined in code:

```
List<String> paths = Arrays.asList("/user/workers/w1", "/user/workers/w2",
  "/user/workers/w3");
FiniteDuration within2 = FiniteDuration.create(10, TimeUnit.SECONDS);
ActorRef router20 =
  getContext().actorOf(new ScatterGatherFirstCompletedGroup(paths, within2).props(),
  "router20");
```

TailChoppingPool and TailChoppingGroup

The TailChoppingRouter will first send the message to one, randomly picked, routee and then after a small delay to a second routee (picked randomly from the remaining routees) and so on. It waits for first reply it gets back and forwards it back to original sender. Other replies are discarded.

The goal of this router is to decrease latency by performing redundant queries to multiple routees, assuming that one of the other actors may still be faster to respond than the initial one.

This optimisation was described nicely in a blog post by Peter Bailis: [Doing redundant work to speed up distributed queries](#).

TailChoppingPool defined in configuration:

```
akka.actor.deployment {
  /parent/router21 {
    router = tail-chopping-pool
    nr-of-instances = 5
    within = 10 seconds
    tail-chopping-router.interval = 20 milliseconds
  }
}
```

```
ActorRef router21 =
  getContext().actorOf(FromConfig.getInstance().props(
    Props.create(Worker.class)), "router21");
```

TailChoppingPool defined in code:

```
FiniteDuration within3 = FiniteDuration.create(10, TimeUnit.SECONDS);
FiniteDuration interval = FiniteDuration.create(20, TimeUnit.MILLISECONDS);
ActorRef router22 =
  getContext().actorOf(new TailChoppingPool(5, within3, interval).props(
    Props.create(Worker.class)), "router22");
```

TailChoppingGroup defined in configuration:

```
akka.actor.deployment {
  /parent/router23 {
    router = tail-chopping-group
    routees.paths = ["/user/workers/w1", "/user/workers/w2", "/user/workers/w3"]
    within = 10 seconds
    tail-chopping-router.interval = 20 milliseconds
  }
}
```

```
ActorRef router23 =
  getContext().actorOf(FromConfig.getInstance().props(), "router23");
```

TailChoppingGroup defined in code:

```
List<String> paths = Arrays.asList("/user/workers/w1", "/user/workers/w2",
  "/user/workers/w3");
FiniteDuration within4 = FiniteDuration.create(10, TimeUnit.SECONDS);
FiniteDuration interval2 = FiniteDuration.create(20, TimeUnit.MILLISECONDS);
ActorRef router24 =
  getContext().actorOf(new TailChoppingGroup(paths, within4, interval2).props(),
  "router24");
```

ConsistentHashingPool and ConsistentHashingGroup

The ConsistentHashingPool uses [consistent hashing](#) to select a routee based on the sent message. This article gives good insight into how consistent hashing is implemented.

There is 3 ways to define what data to use for the consistent hash key.

- You can define `withHashMapper` of the router to map incoming messages to their consistent hash key. This makes the decision transparent for the sender.
- The messages may implement `akka.routing.ConsistentHashingRouter.ConsistentHashable`. The key is part of the message and it's convenient to define it together with the message definition.
- The messages can be wrapped in a `akka.routing.ConsistentHashingRouter.ConsistentHashableEnvelope` to define what data to use for the consistent hash key. The sender knows the key to use.

These ways to define the consistent hash key can be use together and at the same time for one router. The `withHashMapper` is tried first.

Code example:

```
public class Cache extends UntypedActor {
  Map<String, String> cache = new HashMap<String, String>();

  public void onReceive(Object msg) {
    if (msg instanceof Entry) {
      Entry entry = (Entry) msg;
      cache.put(entry.key, entry.value);
    } else if (msg instanceof Get) {
      Get get = (Get) msg;
      Object value = cache.get(get.key);
      getSender().tell(value == null ? NOT_FOUND : value,
        getContext().self());
    } else if (msg instanceof Evict) {
      Evict evict = (Evict) msg;
      cache.remove(evict.key);
    } else {
      unhandled(msg);
    }
  }
}
```

```

}

public final class Evict implements Serializable {
    private static final long serialVersionUID = 1L;
    public final String key;
    public Evict(String key) {
        this.key = key;
    }
}

public final class Get implements Serializable, ConsistentHashable {
    private static final long serialVersionUID = 1L;
    public final String key;
    public Get(String key) {
        this.key = key;
    }
    public Object consistentHashKey() {
        return key;
    }
}

public final class Entry implements Serializable {
    private static final long serialVersionUID = 1L;
    public final String key;
    public final String value;
    public Entry(String key, String value) {
        this.key = key;
        this.value = value;
    }
}

public final String NOT_FOUND = "NOT_FOUND";

final ConsistentHashMapMapper hashMapMapper = new ConsistentHashMapMapper() {
    @Override
    public Object hashKey(Object message) {
        if (message instanceof Evict) {
            return ((Evict) message).key;
        } else {
            return null;
        }
    }
};

ActorRef cache = system.actorOf(
    new ConsistentHashingPool(10).withHashMapMapper(hashMapper).props(
        Props.create(Cache.class)),
    "cache");

cache.tell(new ConsistentHashableEnvelope(
    new Entry("hello", "HELLO"), "hello"), getRef());
cache.tell(new ConsistentHashableEnvelope(
    new Entry("hi", "HI"), "hi"), getRef());

cache.tell(new Get("hello"), getRef());
expectMsgEquals("HELLO");

cache.tell(new Get("hi"), getRef());
expectMsgEquals("HI");

cache.tell(new Evict("hi"), getRef());
cache.tell(new Get("hi"), getRef());
expectMsgEquals(NOT_FOUND);

```

In the above example you see that the `Get` message implements `ConsistentHashable` itself, while the `Entry` message is wrapped in a `ConsistentHashableEnvelope`. The `Evict` message is handled by the `hashMapping` partial function.

`ConsistentHashingPool` defined in configuration:

```
akka.actor.deployment {
  /parent/router25 {
    router = consistent-hashing-pool
    nr-of-instances = 5
    virtual-nodes-factor = 10
  }
}
```

```
ActorRef router25 =
  getContext().actorOf(FromConfig.getInstance().props(Props.create(Worker.class)),
    "router25");
```

`ConsistentHashingPool` defined in code:

```
ActorRef router26 =
  getContext().actorOf(new ConsistentHashingPool(5).props(
    Props.create(Worker.class)), "router26");
```

`ConsistentHashingGroup` defined in configuration:

```
akka.actor.deployment {
  /parent/router27 {
    router = consistent-hashing-group
    routees.paths = ["/user/workers/w1", "/user/workers/w2", "/user/workers/w3"]
    virtual-nodes-factor = 10
  }
}
```

```
ActorRef router27 =
  getContext().actorOf(FromConfig.getInstance().props(), "router27");
```

`ConsistentHashingGroup` defined in code:

```
List<String> paths = Arrays.asList("/user/workers/w1", "/user/workers/w2",
  "/user/workers/w3");
ActorRef router28 =
  getContext().actorOf(new ConsistentHashingGroup(paths).props(), "router28");
```

`virtual-nodes-factor` is the number of virtual nodes per routee that is used in the consistent hash node ring to make the distribution more uniform.

3.6.4 Specially Handled Messages

Most messages sent to router actors will be forwarded according to the routers' routing logic. However there are a few types of messages that have special behavior.

Note that these special messages, except for the `Broadcast` message, are only handled by self contained router actors and not by the `akka.routing.Router` component described in [A Simple Router](#).

Broadcast Messages

A `Broadcast` message can be used to send a message to *all* of a router's routees. When a router receives a `Broadcast` message, it will broadcast that message's *payload* to all routees, no matter how that router would normally route its messages.

The example below shows how you would use a `Broadcast` message to send a very important message to every routee of a router.

```
router.tell(new Broadcast("Watch out for Davy Jones' locker"), getTestActor());
```

In this example the router receives the `Broadcast` message, extracts its payload ("Watch out for Davy Jones' locker"), and then sends the payload on to all of the router's routees. It is up to each each routee actor to handle the received payload message.

PoisonPill Messages

A `PoisonPill` message has special handling for all actors, including for routers. When any actor receives a `PoisonPill` message, that actor will be stopped. See the [PoisonPill](#) documentation for details.

```
router.tell(PoisonPill.getInstance(), getTestActor());
```

For a router, which normally passes on messages to routees, it is important to realise that `PoisonPill` messages are processed by the router only. `PoisonPill` messages sent to a router will *not* be sent on to routees.

However, a `PoisonPill` message sent to a router may still affect its routees, because it will stop the router and when the router stops it also stops its children. Stopping children is normal actor behavior. The router will stop routees that it has created as children. Each child will process its current message and then stop. This may lead to some messages being unprocessed. See the documentation on [Stopping actors](#) for more information.

If you wish to stop a router and its routees, but you would like the routees to first process all the messages currently in their mailboxes, then you should not send a `PoisonPill` message to the router. Instead you should wrap a `PoisonPill` message inside a `Broadcast` message so that each routee will receive the `PoisonPill` message. Note that this will stop all routees, even if the routees aren't children of the router, i.e. even routees programmatically provided to the router.

```
router.tell(new Broadcast(PoisonPill.getInstance()), getTestActor());
```

With the code shown above, each routee will receive a `PoisonPill` message. Each routee will continue to process its messages as normal, eventually processing the `PoisonPill`. This will cause the routee to stop. After all routees have stopped the router will itself be *stopped automatically* unless it is a dynamic router, e.g. using a resizer.

Note: Brendan W McAdams' excellent blog post [Distributing Akka Workloads - And Shutting Down Afterwards](#) discusses in more detail how `PoisonPill` messages can be used to shut down routers and routees.

Kill Messages

Kill messages are another type of message that has special handling. See [Killing an Actor](#) for general information about how actors handle `Kill` messages.

When a `Kill` message is sent to a router the router processes the message internally, and does *not* send it on to its routees. The router will throw an `ActorKilledException` and fail. It will then be either resumed, restarted or terminated, depending how it is supervised.

Routees that are children of the router will also be suspended, and will be affected by the supervision directive that is applied to the router. Routees that are not the routers children, i.e. those that were created externally to the router, will not be affected.

```
router.tell(Kill.getInstance(), getTestActor());
```

As with the `PoisonPill` message, there is a distinction between killing a router, which indirectly kills its children (who happen to be routees), and killing routees directly (some of whom may not be children.) To kill routees directly the router should be sent a `Kill` message wrapped in a `Broadcast` message.

```
router.tell(new Broadcast(Kill.getInstance()), getTestActor());
```

Management Messages

- Sending `akka.routing.GetRoutees` to a router actor will make it send back its currently used routees in a `akka.routing.Routees` message.
- Sending `akka.routing.AddRoutee` to a router actor will add that routee to its collection of routees.
- Sending `akka.routing.RemoveRoutee` to a router actor will remove that routee to its collection of routees.
- Sending `akka.routing.AdjustPoolSize` to a pool router actor will add or remove that number of routees to its collection of routees.

These management messages may be handled after other messages, so if you send `AddRoutee` immediately followed an ordinary message you are not guaranteed that the routees have been changed when the ordinary message is routed. If you need to know when the change has been applied you can send `AddRoutee` followed by `GetRoutees` and when you receive the `Routees` reply you know that the preceding change has been applied.

3.6.5 Dynamically Resizable Pool

All pools can be used with a fixed number of routees or with a resize strategy to adjust the number of routees dynamically.

Pool with resizer defined in configuration:

```
akka.actor.deployment {
  /parent/router29 {
    router = round-robin-pool
    resizer {
      lower-bound = 2
      upper-bound = 15
      messages-per-resize = 100
    }
  }
}
```

```
ActorRef router29 =
  getContext().actorOf(FromConfig.getInstance().props(
    Props.create(Worker.class)), "router29");
```

Several more configuration options are available and described in `akka.actor.deployment.default.resizer` section of the reference [Configuration](#).

Pool with resizer defined in code:

```
DefaultResizer resizer = new DefaultResizer(2, 15);
ActorRef router30 =
  getContext().actorOf(new RoundRobinPool(5).withResizer(resizer).props(
    Props.create(Worker.class)), "router30");
```

It is also worth pointing out that if you define the “router” in the configuration file then this value will be used instead of any programmatically sent parameters.

Note: Resizing is triggered by sending messages to the actor pool, but it is not completed synchronously; instead a message is sent to the “head” `RouterActor` to perform the size change. Thus you cannot rely on resizing to instantaneously create new workers when all others are busy, because the message just sent will be queued to the mailbox of a busy actor. To remedy this, configure the pool to use a balancing dispatcher, see [Configuring Dispatchers](#) for more information.

3.6.6 How Routing is Designed within Akka

On the surface routers look like normal actors, but they are actually implemented differently. Routers are designed to be extremely efficient at receiving messages and passing them quickly on to routees.

A normal actor can be used for routing messages, but an actor's single-threaded processing can become a bottleneck. Routers can achieve much higher throughput with an optimization to the usual message-processing pipeline that allows concurrent routing. This is achieved by embedding routers' routing logic directly in their `ActorRef` rather than in the router actor. Messages sent to a router's `ActorRef` can be immediately routed to the routee, bypassing the single-threaded router actor entirely.

The cost to this is, of course, that the internals of routing code are more complicated than if routers were implemented with normal actors. Fortunately all of this complexity is invisible to consumers of the routing API. However, it is something to be aware of when implementing your own routers.

3.6.7 Custom Router

You can create your own router should you not find any of the ones provided by Akka sufficient for your needs. In order to roll your own router you have to fulfill certain criteria which are explained in this section.

Before creating your own router you should consider whether a normal actor with router-like behavior might do the job just as well as a full-blown router. As explained *above*, the primary benefit of routers over normal actors is their higher performance. But they are somewhat more complicated to write than normal actors. Therefore if lower maximum throughput is acceptable in your application you may wish to stick with traditional actors. This section, however, assumes that you wish to get maximum performance and so demonstrates how you can create your own router.

The router created in this example is replicating each message to a few destinations.

Start with the routing logic:

```
public class RedundancyRoutingLogic implements RoutingLogic {
    private final int nbrCopies;

    public RedundancyRoutingLogic(int nbrCopies) {
        this.nbrCopies = nbrCopies;
    }
    RoundRobinRoutingLogic roundRobin = new RoundRobinRoutingLogic();

    @Override
    public Routee select(Object message, IndexedSeq<Routee> routees) {
        List<Routee> targets = new ArrayList<Routee>();
        for (int i = 0; i < nbrCopies; i++) {
            targets.add(roundRobin.select(message, routees));
        }
        return new SeveralRoutees(targets);
    }
}
```

`select` will be called for each message and in this example pick a few destinations by round-robin, by reusing the existing `RoundRobinRoutingLogic` and wrap the result in a `SeveralRoutees` instance. `SeveralRoutees` will send the message to all of the supplied routes.

The implementation of the routing logic must be thread safe, since it might be used outside of actors.

A unit test of the routing logic:

```
public final class TestRoutee implements Routee {
    public final int n;

    public TestRoutee(int n) {
        this.n = n;
    }
}
```

```

@Override
public void send(Object message, ActorRef sender) {
}

@Override
public int hashCode() {
    return n;
}

@Override
public boolean equals(Object obj) {
    return (obj instanceof TestRoutee) &&
        n == ((TestRoutee) obj).n;
}
}

RedundancyRoutingLogic logic = new RedundancyRoutingLogic(3);

List<Routee> routeeList = new ArrayList<Routee>();
for (int n = 1; n <= 7; n++) {
    routeeList.add(new TestRoutee(n));
}
IndexedSeq<Routee> routees = immutableIndexedSeq(routeeList);

SeveralRoutees r1 = (SeveralRoutees) logic.select("msg", routees);
assertEquals(r1.getRoutees().get(0), routeeList.get(0));
assertEquals(r1.getRoutees().get(1), routeeList.get(1));
assertEquals(r1.getRoutees().get(2), routeeList.get(2));

SeveralRoutees r2 = (SeveralRoutees) logic.select("msg", routees);
assertEquals(r2.getRoutees().get(0), routeeList.get(3));
assertEquals(r2.getRoutees().get(1), routeeList.get(4));
assertEquals(r2.getRoutees().get(2), routeeList.get(5));

SeveralRoutees r3 = (SeveralRoutees) logic.select("msg", routees);
assertEquals(r3.getRoutees().get(0), routeeList.get(6));
assertEquals(r3.getRoutees().get(1), routeeList.get(0));
assertEquals(r3.getRoutees().get(2), routeeList.get(1));

```

You could stop here and use the `RedundancyRoutingLogic` with a `akka.routing.Router` as described in [A Simple Router](#).

Let us continue and make this into a self contained, configurable, router actor.

Create a class that extends `PoolBase`, `GroupBase` or `CustomRouterConfig`. That class is a factory for the routing logic and holds the configuration for the router. Here we make it a `Group`.

```

import java.util.List;

import scala.Option;
import scala.collection.immutable.Iterable;
import akka.actor.ActorContext;
import akka.actor.ActorPath;
import akka.actor.ActorSystem;
import akka.actor.Props;
import akka.dispatch.Dispatchers;
import akka.routing.Group;
import akka.routing.Routee;
import akka.routing.Router;
import akka.routing.RouterActor;
import akka.routing.RouterConfig;
import akka.routing.RoutingLogic;

```

```

import com.typesafe.config.Config;

import akka.routing.GroupBase;
import static docs.jrouting.CustomRouterDocTest.RedundancyRoutingLogic;

public class RedundancyGroup extends GroupBase {
    private final List<String> paths;
    private final int nbrCopies;

    public RedundancyGroup(List<String> paths, int nbrCopies) {
        this.paths = paths;
        this.nbrCopies = nbrCopies;
    }

    public RedundancyGroup(Config config) {
        this(config.getStringList("routees.paths"),
            config.getInt("nbr-copies"));
    }

    @Override
    public java.lang.Iterable<String> getPaths(ActorSystem system) {
        return paths;
    }

    @Override
    public Router createRouter(ActorSystem system) {
        return new Router(new RedundancyRoutingLogic(nbrCopies));
    }

    @Override
    public String routerDispatcher() {
        return Dispatchers.DefaultDispatcherId();
    }
}

```

This can be used exactly as the router actors provided by Akka.

```

for (int n = 1; n <= 10; n++) {
    system.actorOf(Props.create(Storage.class), "s" + n);
}

List<String> paths = new ArrayList<String>();
for (int n = 1; n <= 10; n++) {
    paths.add("/user/s" + n);
}

ActorRef redundancy1 =
    system.actorOf(new RedundancyGroup(paths, 3).props(),
        "redundancy1");
redundancy1.tell("important", getTestActor());

```

Note that we added a constructor in `RedundancyGroup` that takes a `Config` parameter. That makes it possible to define it in configuration.

```

akka.actor.deployment {
  /redundancy2 {
    router = "docs.jrouting.RedundancyGroup"
    routees.paths = ["/user/s1", "/user/s2", "/user/s3"]
    nbr-copies = 5
  }
}

```

Note the fully qualified class name in the `router` property. The router class must extend

`akka.routing.RouterConfig` (`Pool`, `Group` or `CustomRouterConfig`) and have constructor with one `com.typesafe.config.Config` parameter. The deployment section of the configuration is passed to the constructor.

```
ActorRef redundancy2 = system.actorOf(FromConfig.getInstance().props(),
    "redundancy2");
redundancy2.tell("very important", getTestActor());
```

3.6.8 Configuring Dispatchers

The dispatcher for created children of the pool will be taken from `Props` as described in *dispatchers-scala*.

To make it easy to define the dispatcher of the routees of the pool you can define the dispatcher inline in the deployment section of the config.

```
akka.actor.deployment {
  /poolWithDispatcher {
    router = random-pool
    nr-of-instances = 5
    pool-dispatcher {
      fork-join-executor.parallelism-min = 5
      fork-join-executor.parallelism-max = 5
    }
  }
}
```

That is the only thing you need to do enable a dedicated dispatcher for a pool.

Note: If you use a group of actors and route to their paths, then they will still use the same dispatcher that was configured for them in their `Props`, it is not possible to change an actors dispatcher after it has been created.

The “head” router cannot always run on the same dispatcher, because it does not process the same type of messages, hence this special actor does not use the dispatcher configured in `Props`, but takes the `routerDispatcher` from the `RouterConfig` instead, which defaults to the actor system’s default dispatcher. All standard routers allow setting this property in their constructor or factory method, custom routers have to implement the method in a suitable way.

```
Props props =
  // "head" router actor will run on "router-dispatcher" dispatcher
  // Worker routees will run on "pool-dispatcher" dispatcher
  new RandomPool(5).withDispatcher("router-dispatcher").props(
    Props.create(Worker.class));
ActorRef router = system.actorOf(props, "poolWithDispatcher");
```

Note: It is not allowed to configure the `routerDispatcher` to be a `akka.dispatch.BalancingDispatcherConfigurator` since the messages meant for the special router actor cannot be processed by any other actor.

3.7 Building Finite State Machine Actors

3.7.1 Overview

The FSM (Finite State Machine) pattern is best described in the [Erlang design principles](#). In short, it can be seen as a set of relations of the form:

$$\text{State(S)} \times \text{Event(E)} \rightarrow \text{Actions (A), State(S')}$$

These relations are interpreted as meaning:

If we are in state S and the event E occurs, we should perform the actions A and make a transition to the state S' .

While the Scala programming language enables the formulation of a nice internal DSL (domain specific language) for formulating finite state machines (see *fsm-scala*), Java's verbosity does not lend itself well to the same approach. This chapter describes ways to effectively achieve the same separation of concerns through self-discipline.

3.7.2 How State should be Handled

All mutable fields (or transitively mutable data structures) referenced by the FSM actor's implementation should be collected in one place and only mutated using a small well-defined set of methods. One way to achieve this is to assemble all mutable state in a superclass which keeps it private and offers protected methods for mutating it.

```
import java.util.ArrayList;
import java.util.List;
import akka.actor.ActorRef;

public abstract class MyFSMBase extends UntypedActor {

    /*
     * This is the mutable state of this state machine.
     */
    protected enum State {
        IDLE, ACTIVE;
    }

    private State state = State.IDLE;
    private ActorRef target;
    private List<Object> queue;

    /*
     * Then come all the mutator methods:
     */
    protected void init(ActorRef target) {
        this.target = target;
        queue = new ArrayList<Object>();
    }

    protected void setState(State s) {
        if (state != s) {
            transition(state, s);
            state = s;
        }
    }

    protected void enqueue(Object o) {
        if (queue != null)
            queue.add(o);
    }

    protected List<Object> drainQueue() {
        final List<Object> q = queue;
        if (q == null)
            throw new IllegalStateException("drainQueue(): not yet initialized");
        queue = new ArrayList<Object>();
        return q;
    }

    /*
```

```

    * Here are the interrogation methods:
    */
    protected boolean isInitialized() {
        return target != null;
    }

    protected State getState() {
        return state;
    }

    protected ActorRef getTarget() {
        if (target == null)
            throw new IllegalStateException("getTarget(): not yet initialized");
        return target;
    }

    /*
    * And finally the callbacks (only one in this example: react to state change)
    */
    abstract protected void transition(State old, State next);
}

```

The benefit of this approach is that state changes can be acted upon in one central place, which makes it impossible to forget inserting code for reacting to state transitions when adding to the FSM's machinery.

3.7.3 Message Buncher Example

The base class shown above is designed to support a similar example as for the Scala FSM documentation: an actor which receives and queues messages, to be delivered in batches to a configurable target actor. The messages involved are:

```

public final class SetTarget {
    final ActorRef ref;

    public SetTarget(ActorRef ref) {
        this.ref = ref;
    }
}

public final class Queue {
    final Object o;

    public Queue(Object o) {
        this.o = o;
    }
}

public final Object flush = new Object();

public final class Batch {
    final List<Object> objects;

    public Batch(List<Object> objects) {
        this.objects = objects;
    }
}

```

This actor has only the two states `IDLE` and `ACTIVE`, making their handling quite straight-forward in the concrete actor derived from the base class:

```

import akka.event.LoggingAdapter;
import akka.event.Logging;
import akka.actor.UntypedActor;

public class MyFSM extends MyFSMBase {

    private final LoggingAdapter log =
        Logging.getLogger(getContext().system(), this);

    @Override
    public void onReceive(Object o) {

        if (getState() == State.IDLE) {

            if (o instanceof SetTarget)
                init(((SetTarget) o).ref);

            else
                whenUnhandled(o);

        } else if (getState() == State.ACTIVE) {

            if (o == flush)
                setState(State.IDLE);

            else
                whenUnhandled(o);

        }

    }

    @Override
    public void transition(State old, State next) {
        if (old == State.ACTIVE) {
            getTarget().tell(new Batch(drainQueue()), getSelf());
        }
    }

    private void whenUnhandled(Object o) {
        if (o instanceof Queue && isInitialized()) {
            enqueue(((Queue) o).o);
            setState(State.ACTIVE);
        } else {
            log.warning("received unknown message {} in state {}", o, getState());
        }
    }
}

```

The trick here is to factor out common functionality like `whenUnhandled` and `transition` in order to obtain a few well-defined points for reacting to change or insert logging.

3.7.4 State-Centric vs. Event-Centric

In the example above, the subjective complexity of state and events was roughly equal, making it a matter of taste whether to choose primary dispatch on either; in the example a state-based dispatch was chosen. Depending on how evenly the matrix of possible states and events is populated, it may be more practical to handle different events first and distinguish the states in the second tier. An example would be a state machine which has a multitude of internal states but handles only very few distinct events.

3.8 Persistence

Akka persistence enables stateful actors to persist their internal state so that it can be recovered when an actor is started, restarted after a JVM crash or by a supervisor, or migrated in a cluster. The key concept behind Akka persistence is that only changes to an actor's internal state are persisted but never its current state directly (except for optional snapshots). These changes are only ever appended to storage, nothing is ever mutated, which allows for very high transaction rates and efficient replication. Stateful actors are recovered by replaying stored changes to these actors from which they can rebuild internal state. This can be either the full history of changes or starting from a snapshot which can dramatically reduce recovery times. Akka persistence also provides point-to-point communication with at-least-once message delivery semantics.

Note: Java 8 lambda expressions are also supported. (See section *Persistence (Java with Lambda Support)*)

Akka persistence is inspired by and the official replacement of the `eventsourced` library. It follows the same concepts and architecture of `eventsourced` but significantly differs on API and implementation level. See also *Migration Guide Eventsourced to Akka Persistence 2.3.x*

3.8.1 Dependencies

Akka persistence is a separate jar file. Make sure that you have the following dependency in your project:

```
<dependency>
  <groupId>com.typesafe.akka</groupId>
  <artifactId>akka-persistence_2.11</artifactId>
  <version>2.4.0</version>
</dependency>
```

Akka persistence extension comes with few built-in persistence plugins, including in-memory heap based journal, local file-system based snapshot-store and LevelDB based journal.

LevelDB based plugins will require the following additional dependency declaration:

```
<dependency>
  <groupId>org.iq80.leveldb</groupId>
  <artifactId>leveldb</artifactId>
  <version>0.7</version>
</dependency>
<dependency>
  <groupId>org.fusesource.leveldbjni</groupId>
  <artifactId>leveldbjni-all</artifactId>
  <version>1.8</version>
</dependency>
```

3.8.2 Architecture

- *UntypedPersistentActor*: Is a persistent, stateful actor. It is able to persist events to a journal and can react to them in a thread-safe manner. It can be used to implement both *command* as well as *event sourced* actors. When a persistent actor is started or restarted, journaled messages are replayed to that actor, so that it can recover internal state from these messages.
- *UntypedPersistentView*: A view is a persistent, stateful actor that receives journaled messages that have been written by another persistent actor. A view itself does not journal new messages, instead, it updates internal state only from a persistent actor's replicated message stream.
- *UntypedPersistentActorAtLeastOnceDelivery*: To send messages with at-least-once delivery semantics to destinations, also in case of sender and receiver JVM crashes.
- *AsyncWriteJournal*: A journal stores the sequence of messages sent to a persistent actor. An application can control which messages are journaled and which are received by the persistent actor without being

journalled. The storage backend of a journal is pluggable. Persistence extension comes with a “leveldb” journal plugin, which writes to the local filesystem, and replicated journals are available as [Community plugins](#).

- *Snapshot store*: A snapshot store persists snapshots of a persistent actor’s or a view’s internal state. Snapshots are used for optimizing recovery times. The storage backend of a snapshot store is pluggable. Persistence extension comes with a “local” snapshot storage plugin, which writes to the local filesystem, and replicated snapshot stores are available as [Community plugins](#).

3.8.3 Event sourcing

The basic idea behind [Event Sourcing](#) is quite simple. A persistent actor receives a (non-persistent) command which is first validated if it can be applied to the current state. Here, validation can mean anything, from simple inspection of a command message’s fields up to a conversation with several external services, for example. If validation succeeds, events are generated from the command, representing the effect of the command. These events are then persisted and, after successful persistence, used to change the actor’s state. When the persistent actor needs to be recovered, only the persisted events are replayed of which we know that they can be successfully applied. In other words, events cannot fail when being replayed to a persistent actor, in contrast to commands. Event sourced actors may of course also process commands that do not change application state, such as query commands, for example.

Akka persistence supports event sourcing with the `UntypedPersistentActor` abstract class. An actor that extends this class uses the `persist` method to persist and handle events. The behavior of an `UntypedPersistentActor` is defined by implementing `receiveRecover` and `receiveCommand`. This is demonstrated in the following example.

```
import akka.actor.ActorRef;
import akka.actor.ActorSystem;
import akka.actor.Props;
import akka.japi.Procedure;
import akka.persistence.SnapshotOffer;
import akka.persistence.UntypedPersistentActor;

import java.io.Serializable;
import java.util.ArrayList;
import static java.util.Arrays.asList;

class Cmd implements Serializable {
    private static final long serialVersionUID = 1L;
    private final String data;

    public Cmd(String data) {
        this.data = data;
    }

    public String getData() {
        return data;
    }
}

class Evt implements Serializable {
    private static final long serialVersionUID = 1L;
    private final String data;

    public Evt(String data) {
        this.data = data;
    }

    public String getData() {
        return data;
    }
}
```

```

}

class ExampleState implements Serializable {
    private static final long serialVersionUID = 1L;
    private final ArrayList<String> events;

    public ExampleState() {
        this(new ArrayList<String>());
    }

    public ExampleState(ArrayList<String> events) {
        this.events = events;
    }

    public ExampleState copy() {
        return new ExampleState(new ArrayList<String>(events));
    }

    public void update(Evt evt) {
        events.add(evt.getData());
    }

    public int size() {
        return events.size();
    }

    @Override
    public String toString() {
        return events.toString();
    }
}

class ExamplePersistentActor extends UntypedPersistentActor {
    @Override
    public String persistenceId() { return "sample-id-1"; }

    private ExampleState state = new ExampleState();

    public int getNumEvents() {
        return state.size();
    }

    @Override
    public void onReceiveRecover(Object msg) {
        if (msg instanceof Evt) {
            state.update((Evt) msg);
        } else if (msg instanceof SnapshotOffer) {
            state = (ExampleState)((SnapshotOffer)msg).snapshot();
        } else {
            unhandled(msg);
        }
    }

    @Override
    public void onReceiveCommand(Object msg) {
        if (msg instanceof Cmd) {
            final String data = ((Cmd)msg).getData();
            final Evt evt1 = new Evt(data + "-" + getNumEvents());
            final Evt evt2 = new Evt(data + "-" + (getNumEvents() + 1));
            persistAll(asList(evt1, evt2), new Procedure<Evt>() {
                public void apply(Evt evt) throws Exception {
                    state.update(evt);
                    if (evt.equals(evt2)) {

```

```

        getContext().system().eventStream().publish(evt);
    }
}
});
} else if (msg.equals("snap")) {
    // IMPORTANT: create a copy of snapshot
    // because ExampleState is mutable !!!
    saveSnapshot(state.copy());
} else if (msg.equals("print")) {
    System.out.println(state);
} else {
    unhandled(msg);
}
}
}
}

```

The example defines two data types, `Cmd` and `Evt` to represent commands and events, respectively. The state of the `ExamplePersistentActor` is a list of persisted event data contained in `ExampleState`.

The persistent actor's `onReceiveRecover` method defines how state is updated during recovery by handling `Evt` and `SnapshotOffer` messages. The persistent actor's `onReceiveCommand` method is a command handler. In this example, a command is handled by generating two events which are then persisted and handled. Events are persisted by calling `persist` with an event (or a sequence of events) as first argument and an event handler as second argument.

The `persist` method persists events asynchronously and the event handler is executed for successfully persisted events. Successfully persisted events are internally sent back to the persistent actor as individual messages that trigger event handler executions. An event handler may close over persistent actor state and mutate it. The sender of a persisted event is the sender of the corresponding command. This allows event handlers to reply to the sender of a command (not shown).

The main responsibility of an event handler is changing persistent actor state using event data and notifying others about successful state changes by publishing events.

When persisting events with `persist` it is guaranteed that the persistent actor will not receive further commands between the `persist` call and the execution(s) of the associated event handler. This also holds for multiple `persist` calls in context of a single command. Incoming messages are *stashed* until the `persist` is completed. You should be careful to not send more messages to a persistent actor than it can keep up with, otherwise the number of stashed messages will grow. It can be wise to protect against `OutOfMemoryError` by defining a maximum stash capacity in the mailbox configuration:

```
akka.actor.default-mailbox.stash-capacity=10000
```

If the stash capacity is exceeded for an actor the stashed messages are discarded and a `MessageQueueAppendFailedException` is thrown, causing actor restart if default supervision strategy is used.

Note that the stash capacity is per actor. If you have many persistent actors, e.g. when using cluster sharding, you may need to define a small stash capacity to ensure that the total number of stashed messages in the system don't consume too much memory.

If persistence of an event fails, `onPersistFailure` will be invoked (logging the error by default) and the actor will unconditionally be stopped. If persistence of an event is rejected before it is stored, e.g. due to serialization error, `onPersistRejected` will be invoked (logging a warning by default) and the actor continues with next message.

The easiest way to run this example yourself is to download [Typesafe Activator](#) and open the tutorial named [Akka Persistence Samples with Java](#). It contains instructions on how to run the `PersistentActorExample`.

Note: It's also possible to switch between different command handlers during normal processing and recovery with `getContext().become()` and `getContext().unbecome()`. To get the actor into the same state after recovery you need to take special care to perform the same state transitions with `become` and `unbecome` in the `receiveRecover` method as you would have done in the command handler. Note that when using

become from `receiveRecover` it will still only use the `receiveRecover` behavior when replaying the events. When replay is completed it will use the new behavior.

Identifiers

A persistent actor must have an identifier that doesn't change across different actor incarnations. The identifier must be defined with the `persistenceId` method.

```
@Override
public String persistenceId() {
    return "my-stable-persistence-id";
}
```

Recovery

By default, a persistent actor is automatically recovered on start and on restart by replaying journaled messages. New messages sent to a persistent actor during recovery do not interfere with replayed messages. They are cached and received by a persistent actor after recovery phase completes.

Note: Accessing the `sender()` for replayed messages will always result in a `deadLetters` reference, as the original sender is presumed to be long gone. If you indeed have to notify an actor during recovery in the future, store its `ActorPath` explicitly in your persisted events.

Recovery customization

Applications may also customise how recovery is performed by returning a customised `Recovery` object in the `recovery` method of a `UntypedPersistentActor`, for example setting an upper bound to the replay, which allows the actor to be replayed to a certain point “in the past” instead to its most up to date state:

```
@Override
public Recovery recovery() {
    return Recovery.create(457L);
}
```

Recovery can be disabled by returning `Recovery.none()` in the `recovery` method of a `PersistentActor`:

```
@Override
public Recovery recovery() {
    return Recovery.none();
}
```

Recovery status

A persistent actor can query its own recovery status via the methods

```
public boolean recoveryRunning();
public boolean recoveryFinished();
```

Sometimes there is a need for performing additional initialization when the recovery has completed, before processing any other message sent to the persistent actor. The persistent actor will receive a special `RecoveryCompleted` message right after recovery and before any other received messages.

```

@Override
public void onReceiveRecover(Object message) {
    if (message instanceof RecoveryCompleted) {
        // perform init after recovery, before any other messages
    }
}

@Override
public void onReceiveCommand(Object message) throws Exception {
    if (message instanceof String) {
        // ...
    } else {
        unhandled(message);
    }
}
}

```

If there is a problem with recovering the state of the actor from the journal, `onRecoveryFailure` is called (logging the error by default) and the actor will be stopped.

Relaxed local consistency requirements and high throughput use-cases

If faced with relaxed local consistency requirements and high throughput demands sometimes `PersistentActor` and its `persist` may not be enough in terms of consuming incoming `Commands` at a high rate, because it has to wait until all `Events` related to a given `Command` are processed in order to start processing the next `Command`. While this abstraction is very useful for most cases, sometimes you may be faced with relaxed requirements about consistency – for example you may want to process commands as fast as you can, assuming that `Event` will eventually be persisted and handled properly in the background and retroactively reacting to persistence failures if needed.

The `persistAsync` method provides a tool for implementing high-throughput persistent actors. It will *not* stash incoming `Commands` while the `Journal` is still working on persisting and/or user code is executing event callbacks.

In the below example, the event callbacks may be called “at any time”, even after the next `Command` has been processed. The ordering between events is still guaranteed (“`evt-b-1`” will be sent after “`evt-a-2`”, which will be sent after “`evt-a-1`” etc.).

```

class MyPersistentActor extends UntypedPersistentActor {
    @Override
    public String persistenceId() { return "some-persistence-id"; }

    @Override
    public void onReceiveRecover(Object msg) {
        // handle recovery here
    }

    @Override
    public void onReceiveCommand(Object msg) {
        sender().tell(msg, self());

        persistAsync(String.format("evt-%s-1", msg), new Procedure<String>() {
            @Override
            public void apply(String event) throws Exception {
                sender().tell(event, self());
            }
        });
        persistAsync(String.format("evt-%s-2", msg), new Procedure<String>() {
            @Override
            public void apply(String event) throws Exception {
                sender().tell(event, self());
            }
        });
    }
}

```

```
}
}
```

Note: In order to implement the pattern known as “*command sourcing*” simply `persistAsync` all incoming messages right away, and handle them in the callback.

Warning: The callback will not be invoked if the actor is restarted (or stopped) in between the call to `persistAsync` and the journal has confirmed the write.

Deferring actions until preceding persist handlers have executed

Sometimes when working with `persistAsync` you may find that it would be nice to define some actions in terms of “happens-after the previous `persistAsync` handlers have been invoked”. `PersistentActor` provides an utility method called `deferAsync`, which works similarly to `persistAsync` yet does not persist the passed in event. It is recommended to use it for *read* operations, and actions which do not have corresponding events in your domain model.

Using this method is very similar to the `persist` family of methods, yet it does **not** persist the passed in event. It will be kept in memory and used when invoking the handler.

```
class MyPersistentActor extends UntypedPersistentActor {
  @Override
  public String persistenceId() { return "some-persistence-id"; }

  @Override
  public void onReceiveRecover(Object msg) {
    // handle recovery here
  }

  @Override
  public void onReceiveCommand(Object msg) {
    final Procedure<String> replyToSender = new Procedure<String>() {
      @Override
      public void apply(String event) throws Exception {
        sender().tell(event, self());
      }
    };

    persistAsync(String.format("evt-%s-1", msg), replyToSender);
    persistAsync(String.format("evt-%s-2", msg), replyToSender);
    deferAsync(String.format("evt-%s-3", msg), replyToSender);
  }
}
```

Notice that the `sender()` is **safe** to access in the handler callback, and will be pointing to the original sender of the command for which this `deferAsync` handler was called.

```
final ActorRef persistentActor = system.actorOf(Props.create(MyPersistentActor.class));
persistentActor.tell("a", null);
persistentActor.tell("b", null);

// order of received messages:
// a
// b
// evt-a-1
// evt-a-2
// evt-a-3
// evt-b-1
```

```
// evt-b-2
// evt-b-3
```

Warning: The callback will not be invoked if the actor is restarted (or stopped) in between the call to `deferAsync` and the journal has processed and confirmed all preceding writes.

Nested persist calls

It is possible to call `persist` and `persistAsync` inside their respective callback blocks and they will properly retain both the thread safety (including the right value of `sender()`) as well as stashing guarantees.

In general it is encouraged to create command handlers which do not need to resort to nested event persisting, however there are situations where it may be useful. It is important to understand the ordering of callback execution in those situations, as well as their implication on the stashing behaviour (that `persist()` enforces). In the following example two `persist` calls are issued, and each of them issues another `persist` inside its callback:

```
@Override
public void onReceiveCommand(Object msg) {
    final Procedure<String> replyToSender = new Procedure<String>() {
        @Override
        public void apply(String event) throws Exception {
            sender().tell(event, self());
        }
    };

    final Procedure<String> outer1Callback = new Procedure<String>() {
        @Override
        public void apply(String event) throws Exception {
            sender().tell(event, self());
            persist(String.format("%s-inner-1", msg), replyToSender);
        }
    };

    final Procedure<String> outer2Callback = new Procedure<String>() {
        @Override
        public void apply(String event) throws Exception {
            sender().tell(event, self());
            persist(String.format("%s-inner-2", msg), replyToSender);
        }
    };

    persist(String.format("%s-outer-1", msg), outer1Callback);
    persist(String.format("%s-outer-2", msg), outer2Callback);
}
```

When sending two commands to this `PersistentActor`, the `persist` handlers will be executed in the following order:

```
persistentActor.tell("a", self());
persistentActor.tell("b", self());

// order of received messages:
// a
// a-outer-1
// a-outer-2
// a-inner-1
// a-inner-2
// and only then process "b"
// b
// b-outer-1
// b-outer-2
```

```
// b-inner-1
// b-inner-2
```

First the “outer layer” of persist calls is issued and their callbacks applied, after these have successfully completed the inner callbacks will be invoked (once the events they are persisting have been confirmed to be persisted by the journal). And only after all these handlers have been successfully invoked, the next command will be delivered to the persistent Actor. In other words, the stashing of incoming commands that is guaranteed by initially calling `persist()` on the outer layer is extended until all nested `persist` callbacks have been handled.

It is also possible to nest `persistAsync` calls, using the same pattern:

```
@Override
public void onReceiveCommand(Object msg) {
    final Procedure<String> replyToSender = new Procedure<String>() {
        @Override
        public void apply(String event) throws Exception {
            sender().tell(event, self());
        }
    };

    final Procedure<String> outer1Callback = new Procedure<String>() {
        @Override
        public void apply(String event) throws Exception {
            sender().tell(event, self());
            persistAsync(String.format("%s-inner-1", msg), replyToSender);
        }
    };
    final Procedure<String> outer2Callback = new Procedure<String>() {
        @Override
        public void apply(String event) throws Exception {
            sender().tell(event, self());
            persistAsync(String.format("%s-inner-1", msg), replyToSender);
        }
    };

    persistAsync(String.format("%s-outer-1", msg), outer1Callback);
    persistAsync(String.format("%s-outer-2", msg), outer2Callback);
}
```

In this case no stashing is happening, yet the events are still persisted and callbacks executed in the expected order:

```
persistentActor.tell("a", ActorRef.noSender());
persistentActor.tell("b", ActorRef.noSender());

// order of received messages:
// a
// b
// a-outer-1
// a-outer-2
// b-outer-1
// b-outer-2
// a-inner-1
// a-inner-2
// b-inner-1
// b-inner-2

// which can be seen as the following causal relationship:
// a -> a-outer-1 -> a-outer-2 -> a-inner-1 -> a-inner-2
// b -> b-outer-1 -> b-outer-2 -> b-inner-1 -> b-inner-2
```

While it is possible to nest mixed `persist` and `persistAsync` with keeping their respective semantics it is not a recommended practice as it may lead to overly complex nesting.

Failures

If persistence of an event fails, `onPersistFailure` will be invoked (logging the error by default) and the actor will unconditionally be stopped.

The reason that it cannot resume when persist fails is that it is unknown if the even was actually persisted or not, and therefore it is in an inconsistent state. Restarting on persistent failures will most likely fail anyway, since the journal is probably unavailable. It is better to stop the actor and after a back-off timeout start it again. The `akka.pattern.BackoffSupervisor` actor is provided to support such restarts.

```
@Override
public void preStart() throws Exception {
    final Props childProps = Props.create(MyPersistentActor1.class);
    final Props props = BackoffSupervisor.props(
        childProps,
        "myActor",
        Duration.create(3, TimeUnit.SECONDS),
        Duration.create(30, TimeUnit.SECONDS),
        0.2);
    getContext().actorOf(props, "mySupervisor");
    super.preStart();
}
```

If persistence of an event is rejected before it is stored, e.g. due to serialization error, `onPersistRejected` will be invoked (logging a warning by default) and the actor continues with next message.

If there is a problem with recovering the state of the actor from the journal when the actor is started, `onRecoveryFailure` is called (logging the error by default) and the actor will be stopped.

Atomic writes

Each event is of course stored atomically, but it is also possible to store several events atomically by using the `persistAll` or `persistAllAsync` method. That means that all events passed to that method are stored or none of them are stored if there is an error.

The recovery of a persistent actor will therefore never be done partially with only a subset of events persisted by `persistAll`.

Some journals may not support atomic writes of several events and they will then reject the `persistAll` command, i.e. `onPersistRejected` is called with an exception (typically `UnsupportedOperationException`).

Batch writes

To optimize throughput, a persistent actor internally batches events to be stored under high load before writing them to the journal (as a single batch). The batch size dynamically grows from 1 under low and moderate loads to a configurable maximum size (default is 200) under high load. When using `persistAsync` this increases the maximum throughput dramatically.

```
akka.persistence.journal.leveldb.max-message-batch-size = 200
```

A new batch write is triggered by a persistent actor as soon as a batch reaches the maximum size or if the journal completed writing the previous batch. Batch writes are never timer-based which keeps latencies at a minimum.

Message deletion

It is possible to delete all messages (journalled by a single persistent actor) up to a specified sequence number, persistent actors may call the `deleteMessages` method.

Deleting messages in event sourcing based applications is typically either not used at all, or used in conjunction with *snapshotting*, i.e. after a snapshot has been successfully stored, a `deleteMessages(toSequenceNr)`

up until the sequence number of the data held by that snapshot can be issued, to safely delete the previous events, while still having access to the accumulated state during replays - by loading the snapshot.

The result of the `deleteMessages` request is signaled to the persistent actor with a `DeleteMessagesSuccess` message if the delete was successful or a `DeleteMessagesFailure` message if it failed.

Persistence status handling

Persisting, deleting and replaying messages can either succeed or fail.

Method	Success	Failure / Rejection	After failure handler invoked
<code>persist / persistAsync</code> <code>onPersistRejected</code>	<code>persist</code> handler invoked	<code>onPersistFailure</code> No automatic actions.	Actor is stopped.
<code>recovery</code>	<code>RecoverySuccess</code>	<code>onRecoveryFailure</code>	Actor is stopped.
<code>deleteMessages</code>	<code>DeleteMessagesSuccess</code>	<code>DeleteMessagesFailure</code>	No automatic actions.

The most important operations (`persist` and `recovery`) have failure handlers modelled as explicit callbacks which the user can override in the `PersistentActor`. The default implementations of these handlers emit a log message (`error` for `persist/recovery` failures, and `warning` for others), logging the failure cause and information about which message caused the failure.

For critical failures, such as `recovery` or persisting events failing, the persistent actor will be stopped after the failure handler is invoked. This is because if the underlying journal implementation is signalling persistence failures it is most likely either failing completely or overloaded and restarting right-away and trying to persist the event again will most likely not help the journal recover – as it would likely cause a [Thundering herd problem](#), as many persistent actors would restart and try to persist their events again. Instead, using a `BackoffSupervisor` (as described in [Failures](#)) which implements an exponential-backoff strategy which allows for more breathing room for the journal to recover between restarts of the persistent actor.

Note: Journal implementations may choose to implement a retry mechanisms, e.g. such that only after a write fails N number of times a persistence failure is signalled back to the user. In other words, once a journal returns a failure, it is considered *fatal* by Akka Persistence, and the persistent actor which caused the failure will be stopped.

Check the documentation of the journal implementation you are using for details if/how it is using this technique.

Safely shutting down persistent actors

Special care should be given when when shutting down persistent actors from the outside. With normal Actors it is often acceptable to use the special `PoisonPill` message to signal to an Actor that it should stop itself once it receives this message – in fact this message is handled automatically by Akka, leaving the target actor no way to refuse stopping itself when given a poison pill.

This can be dangerous when used with `PersistentActor` due to the fact that incoming commands are *stashed* while the persistent actor is awaiting confirmation from the Journal that events have been written when `persist()` was used. Since the incoming commands will be drained from the Actor's mailbox and put into its internal stash while awaiting the confirmation (thus, before calling the `persist` handlers) the Actor **may receive and (auto)handle the `PoisonPill` before it processes the other messages which have been put into its stash**, causing a pre-mature shutdown of the Actor.

Warning: Consider using explicit shut-down messages instead of `PoisonPill` when working with persistent actors.

The example below highlights how messages arrive in the Actor's mailbox and how they interact with its internal stashing mechanism when `persist()` is used, notice the early stop behaviour that occurs when `PoisonPill` is used:

```

final class Shutdown {}

class MyPersistentActor extends UntypedPersistentActor {
    @Override
    public String persistenceId() {
        return "some-persistence-id";
    }

    @Override
    public void onReceiveCommand(Object msg) throws Exception {
        if (msg instanceof Shutdown) {
            context().stop(self());
        } else if (msg instanceof String) {
            System.out.println(msg);
            persist("handle-" + msg, new Procedure<String>() {
                @Override
                public void apply(String param) throws Exception {
                    System.out.println(param);
                }
            });
        } else unhandled(msg);
    }

    @Override
    public void onReceiveRecover(Object msg) throws Exception {
        // handle recovery...
    }
}

// UN-SAFE, due to PersistentActor's command stashing:
persistentActor.tell("a", ActorRef.noSender());
persistentActor.tell("b", ActorRef.noSender());
persistentActor.tell(PoisonPill.getInstance(), ActorRef.noSender());
// order of received messages:
// a
// # b arrives at mailbox, stashing;           internal-stash = [b]
// # PoisonPill arrives at mailbox, stashing; internal-stash = [b, Shutdown]
// PoisonPill is an AutoReceivedMessage, is handled automatically
// !! stop !!
// Actor is stopped without handling 'b' nor the 'a' handler!

// SAFE:
persistentActor.tell("a", ActorRef.noSender());
persistentActor.tell("b", ActorRef.noSender());
persistentActor.tell(new Shutdown(), ActorRef.noSender());
// order of received messages:
// a
// # b arrives at mailbox, stashing;           internal-stash = [b]
// # Shutdown arrives at mailbox, stashing; internal-stash = [b, Shutdown]
// handle-a
// # unstashing;                             internal-stash = [Shutdown]
// b
// handle-b
// # unstashing;                             internal-stash = []
// Shutdown
// -- stop --

```

3.8.4 Persistent Views

Warning: `UntypedPersistentView` is deprecated. Use *Persistence Query* instead. The corresponding query type is `EventsByPersistenceId`. There are several alternatives for connecting the `Source` to an actor corresponding to a previous `UntypedPersistentView` actor:

- `Sink.actorRef` is simple, but has the disadvantage that there is no back-pressure signal from the destination actor, i.e. if the actor is not consuming the messages fast enough the mailbox of the actor will grow
- `mapAsync` combined with *Ask: Send-And-Receive-Future* is almost as simple with the advantage of back-pressure being propagated all the way
- `ActorSubscriber` in case you need more fine grained control

The consuming actor may be a plain `UntypedActor` or an `UntypedPersistentActor` if it needs to store its own state (e.g. `fromSequenceNr` offset).

Persistent views can be implemented by extending the `UntypedPersistentView` trait and implementing the `onReceive` and the `persistenceId` methods.

```
class MyView extends UntypedPersistentView {
  @Override
  public String persistenceId() { return "some-persistence-id"; }

  @Override
  public String viewId() { return "my-stable-persistence-view-id"; }

  @Override
  public void onReceive(Object message) throws Exception {
    if (isPersistent()) {
      // handle message from Journal...
    } else if (message instanceof String) {
      // handle message from user...
    } else {
      unhandled(message);
    }
  }
}
```

The `persistenceId` identifies the persistent actor from which the view receives journaled messages. It is not necessary the referenced persistent actor is actually running. Views read messages from a persistent actor's journal directly. When a persistent actor is started later and begins to write new messages, the corresponding view is updated automatically, by default.

It is possible to determine if a message was sent from the `Journal` or from another actor in user-land by calling the `isPersistent` method. Having that said, very often you don't need this information at all and can simply apply the same logic to both cases (skip the `if isPersistent` check).

Updates

The default update interval of all persistent views of an actor system is configurable:

```
akka.persistence.view.auto-update-interval = 5s
```

`UntypedPersistentView` implementation classes may also override the `autoUpdateInterval` method to return a custom update interval for a specific view class or view instance. Applications may also trigger additional updates at any time by sending a view an `Update` message.

```
final ActorRef view = system.actorOf(Props.create(MyView.class));
view.tell(Update.create(true), null);
```

If the `await` parameter is set to `true`, messages that follow the `Update` request are processed when the incremental message replay, triggered by that update request, completed. If set to `false` (default), messages following

the update request may interleave with the replayed message stream. Automated updates always run with `await = false`.

Automated updates of all persistent views of an actor system can be turned off by configuration:

```
akka.persistence.view.auto-update = off
```

Implementation classes may override the configured default value by overriding the `autoUpdate` method. To limit the number of replayed messages per update request, applications can configure a custom `akka.persistence.view.auto-update-replay-max` value or override the `autoUpdateReplayMax` method. The number of replayed messages for manual updates can be limited with the `replayMax` parameter of the `Update` message.

Recovery

Initial recovery of persistent views works in the very same way as for a persistent actor (i.e. by sending a `Recover` message to self). The maximum number of replayed messages during initial recovery is determined by `autoUpdateReplayMax`. Further possibilities to customize initial recovery are explained in section [Recovery](#).

Identifiers

A persistent view must have an identifier that doesn't change across different actor incarnations. The identifier must be defined with the `viewId` method.

The `viewId` must differ from the referenced `persistenceId`, unless *Snapshots* of a view and its persistent actor shall be shared (which is what applications usually do not want).

3.8.5 Snapshots

Snapshots can dramatically reduce recovery times of persistent actor and views. The following discusses snapshots in context of persistent actor but this is also applicable to persistent views.

Persistent actor can save snapshots of internal state by calling the `saveSnapshot` method. If saving of a snapshot succeeds, the persistent actor receives a `SaveSnapshotSuccess` message, otherwise a `SaveSnapshotFailure` message

```
private Object state;

@Override
public void onReceiveCommand(Object message) {
    if (message.equals("snap")) {
        saveSnapshot(state);
    } else if (message instanceof SaveSnapshotSuccess) {
        SnapshotMetadata metadata = ((SaveSnapshotSuccess)message).metadata();
        // ...
    } else if (message instanceof SaveSnapshotFailure) {
        SnapshotMetadata metadata = ((SaveSnapshotFailure)message).metadata();
        // ...
    }
}
```

During recovery, the persistent actor is offered a previously saved snapshot via a `SnapshotOffer` message from which it can initialize internal state.

```
private Object state;

@Override
public void onReceiveRecover(Object message) {
    if (message instanceof SnapshotOffer) {
        state = ((SnapshotOffer)message).snapshot();
    }
}
```

```

    // ...
  } else if (message instanceof RecoveryCompleted) {
    // ...
  } else {
    // ...
  }
}

```

The replayed messages that follow the `SnapshotOffer` message, if any, are younger than the offered snapshot. They finally recover the persistent actor to its current (i.e. latest) state.

In general, a persistent actor is only offered a snapshot if that persistent actor has previously saved one or more snapshots and at least one of these snapshots matches the `SnapshotSelectionCriteria` that can be specified for recovery.

```

@Override
public Recovery recovery() {
  return Recovery.create(
    SnapshotSelectionCriteria
      .create(457L, System.currentTimeMillis()));
}

```

If not specified, they default to `SnapshotSelectionCriteria.latest()` which selects the latest (= youngest) snapshot. To disable snapshot-based recovery, applications should use `SnapshotSelectionCriteria.none()`. A recovery where no saved snapshot matches the specified `SnapshotSelectionCriteria` will replay all journaled messages.

Note: In order to use snapshots a default snapshot-store (`akka.persistence.snapshot-store.plugin`) must be configured, or the persistent actor can pick a snapshot store explicitly by overriding `String snapshotPluginId()`.

Since it is acceptable for some applications to not use any snapshotting, it is legal to not configure a snapshot store, however Akka will log a warning message when this situation is detected and then continue to operate until an actor tries to store a snapshot, at which point the the operation will fail (by replying with an `SaveSnapshotFailure` for example).

Note that *Cluster Sharding* is using snapshots, so if you use Cluster Sharding you need to define a snapshot store plugin.

Snapshot deletion

A persistent actor can delete individual snapshots by calling the `deleteSnapshot` method with the sequence number of when the snapshot was taken.

To bulk-delete a range of snapshots matching `SnapshotSelectionCriteria`, persistent actors should use the `deleteSnapshots` method.

Snapshot status handling

Saving or deleting snapshots can either succeed or fail – this information is reported back to the persistent actor via status messages as illustrated in the following table.

Method	Success	Failure message
<code>saveSnapshot (Any)</code>	<code>SaveSnapshotSuccess</code>	<code>SaveSnapshotFailure</code>
<code>deleteSnapshot (Long)</code>	<code>DeleteSnapshotSuccess</code>	<code>DeleteSnapshotFailure</code>
<code>deleteSnapshots (SnapshotSelectionCriteria)</code>	<code>DeleteSnapshotsSuccess</code>	<code>DeleteSnapshotsFailure</code>

3.8.6 At-Least-Once Delivery

To send messages with at-least-once delivery semantics to destinations you can extend the `UntypedPersistentActorWithAtLeastOnceDelivery` class instead of `UntypedPersistentActor` on the sending side. It takes care of re-sending messages when they have not been confirmed within a configurable timeout.

The state of the sending actor, including which messages that have been sent and still not been confirmed by the recipient, must be persistent so that it can survive a crash of the sending actor or JVM. The `UntypedPersistentActorWithAtLeastOnceDelivery` class does not persist anything by itself. It is your responsibility to persist the intent that a message is sent and that a confirmation has been received.

Note: At-least-once delivery implies that original message send order is not always preserved and the destination may receive duplicate messages. That means that the semantics do not match those of a normal `ActorRef` send operation:

- it is not at-most-once delivery
- message order for the same sender–receiver pair is not preserved due to possible resends
- after a crash and restart of the destination messages are still delivered—to the new actor incarnation

These semantics is similar to what an `ActorPath` represents (see *actor-lifecycle-scala*), therefore you need to supply a path and not a reference when delivering messages. The messages are sent to the path with an actor selection.

Use the `deliver` method to send a message to a destination. Call the `confirmDelivery` method when the destination has replied with a confirmation message.

Relationship between `deliver` and `confirmDelivery`

To send messages to the destination path, use the `deliver` method after you have persisted the intent to send the message.

The destination actor must send back a confirmation message. When the sending actor receives this confirmation message you should persist the fact that the message was delivered successfully and then call the `confirmDelivery` method.

If the persistent actor is not currently recovering, the `deliver` method will send the message to the destination actor. When recovering, messages will be buffered until they have been confirmed using `confirmDelivery`. Once recovery has completed, if there are outstanding messages that have not been confirmed (during the message replay), the persistent actor will resend these before sending any other messages.

`Deliver` requires a `deliveryIdToMessage` function to pass the provided `deliveryId` into the message so that correlation between `deliver` and `confirmDelivery` is possible. The `deliveryId` must do the round trip. Upon receipt of the message, destination actor will send the same “`deliveryId`” wrapped in a confirmation message back to the sender. The sender will then use it to call `confirmDelivery` method to complete delivery routine.

```
class Msg implements Serializable {
    public final long deliveryId;
    public final String s;

    public Msg(long deliveryId, String s) {
        this.deliveryId = deliveryId;
        this.s = s;
    }
}

class Confirm implements Serializable {
    public final long deliveryId;
```

```

public Confirm(long deliveryId) {
    this.deliveryId = deliveryId;
}
}

class MsgSent implements Serializable {
    public final String s;

    public MsgSent(String s) {
        this.s = s;
    }
}

class MsgConfirmed implements Serializable {
    public final long deliveryId;

    public MsgConfirmed(long deliveryId) {
        this.deliveryId = deliveryId;
    }
}

class MyPersistentActor extends UntypedPersistentActorWithAtLeastOnceDelivery {
    private final ActorSelection destination;

    @Override
    public String persistenceId() { return "persistence-id"; }

    public MyPersistentActor(ActorSelection destination) {
        this.destination = destination;
    }

    @Override
    public void onReceiveCommand(Object message) {
        if (message instanceof String) {
            String s = (String) message;
            persist(new MsgSent(s), new Procedure<MsgSent>() {
                public void apply(MsgSent evt) {
                    updateState(evt);
                }
            });
        } else if (message instanceof Confirm) {
            Confirm confirm = (Confirm) message;
            persist(new MsgConfirmed(confirm.deliveryId), new Procedure<MsgConfirmed>() {
                public void apply(MsgConfirmed evt) {
                    updateState(evt);
                }
            });
        } else {
            unhandled(message);
        }
    }

    @Override
    public void onReceiveRecover(Object event) {
        updateState(event);
    }

    void updateState(Object event) {
        if (event instanceof MsgSent) {
            final MsgSent evt = (MsgSent) event;
            deliver(destination, new Function<Long, Object>() {
                public Object apply(Long deliveryId) {
                    return new Msg(deliveryId, evt.s);
                }
            });
        }
    }
}

```

```

    }
  });
} else if (event instanceof MsgConfirmed) {
  final MsgConfirmed evt = (MsgConfirmed) event;
  confirmDelivery(evt.deliveryId);
}
}
}

class MyDestination extends UntypedActor {
  public void onReceive(Object message) throws Exception {
    if (message instanceof Msg) {
      Msg msg = (Msg) message;
      // ...
      getSender().tell(new Confirm(msg.deliveryId), getSelf());
    } else {
      unhandled(message);
    }
  }
}
}

```

The `deliveryId` generated by the persistence module is a strictly monotonically increasing sequence number without gaps. The same sequence is used for all destinations of the actor, i.e. when sending to multiple destinations the destinations will see gaps in the sequence. It is not possible to use custom `deliveryId`. However, you can send a custom correlation identifier in the message to the destination. You must then retain a mapping between the internal `deliveryId` (passed into the `deliveryIdToMessage` function) and your custom correlation id (passed into the message). You can do this by storing such mapping in a `Map(correlationId -> deliveryId)` from which you can retrieve the `deliveryId` to be passed into the `confirmDelivery` method once the receiver of your message has replied with your custom correlation id.

The `UntypedPersistentActorWithAtLeastOnceDelivery` class has a state consisting of unconfirmed messages and a sequence number. It does not store this state itself. You must persist events corresponding to the `deliver` and `confirmDelivery` invocations from your `PersistentActor` so that the state can be restored by calling the same methods during the recovery phase of the `PersistentActor`. Sometimes these events can be derived from other business level events, and sometimes you must create separate events. During recovery calls to `deliver` will not send out the message, but it will be sent later if no matching `confirmDelivery` was performed.

Support for snapshots is provided by `getDeliverySnapshot` and `setDeliverySnapshot`. The `AtLeastOnceDeliverySnapshot` contains the full delivery state, including unconfirmed messages. If you need a custom snapshot for other parts of the actor state you must also include the `AtLeastOnceDeliverySnapshot`. It is serialized using `protobuf` with the ordinary Akka serialization mechanism. It is easiest to include the bytes of the `AtLeastOnceDeliverySnapshot` as a blob in your custom snapshot.

The interval between redelivery attempts is defined by the `redeliverInterval` method. The default value can be configured with the `akka.persistence.at-least-once-delivery.redeliver-interval` configuration key. The method can be overridden by implementation classes to return non-default values.

The maximum number of messages that will be sent at each redelivery burst is defined by the `redeliveryBurstLimit` method (burst frequency is half of the redelivery interval). If there's a lot of unconfirmed messages (e.g. if the destination is not available for a long time), this helps to prevent an overwhelming amount of messages to be sent at once. The default value can be configured with the `akka.persistence.at-least-once-delivery.redelivery-burst-limit` configuration key. The method can be overridden by implementation classes to return non-default values.

After a number of delivery attempts a `AtLeastOnceDelivery.UnconfirmedWarning` message will be sent to `self`. The re-sending will still continue, but you can choose to call `confirmDelivery` to cancel the re-sending. The number of delivery attempts before emitting the warning is defined by the `warnAfterNumberOfUnconfirmedAttempts` method. The default value can be configured with the `akka.persistence.at-least-once-delivery.warn-after-number-of-unconfirmed-attempts` configuration key. The method can be overridden by implementation classes to return non-default values.

The `UntypedPersistentActorWithAtLeastOnceDelivery` class holds messages in memory until their successful delivery has been confirmed. The limit of maximum number of unconfirmed messages that the actor is allowed to hold in memory is defined by the `maxUnconfirmedMessages` method. If this limit is exceeded the `deliver` method will not accept more messages and it will throw `AtLeastOnceDelivery.MaxUnconfirmedMessagesExceededException`. The default value can be configured with the `akka.persistence.at-least-once-delivery.max-unconfirmed-messages` configuration key. The method can be overridden by implementation classes to return non-default values.

3.8.7 Event Adapters

In long running projects using event sourcing sometimes the need arises to detach the data model from the domain model completely.

Event Adapters help in situations where:

- **Version Migrations** – existing events stored in *Version 1* should be “upcasted” to a new *Version 2* representation, and the process of doing so involves actual code, not just changes on the serialization layer. For these scenarios the `toJournal` function is usually an identity function, however the `fromJournal` is implemented as `v1.Event=>v2.Event`, performing the necessary mapping inside the `fromJournal` method. This technique is sometimes referred to as “upcasting” in other CQRS libraries.
- **Separating Domain and Data models** – thanks to `EventAdapters` it is possible to completely separate the domain model from the model used to persist data in the Journals. For example one may want to use case classes in the domain model, however persist their protocol-buffer (or any other binary serialization format) counter-parts to the Journal. A simple `toJournal:MyModel=>MyDataModel` and `fromJournal:MyDataModel=>MyModel` adapter can be used to implement this feature.
- **Journal Specialized Data Types** – exposing data types understood by the underlying Journal, for example for data stores which understand JSON it is possible to write an `EventAdapter toJournal:Any=>JSON` such that the Journal can *directly* store the json instead of serializing the object to its binary representation.

Implementing an `EventAdapter` is rather stright forward:

```
class MyEventAdapter implements EventAdapter {
  @Override
  public String manifest(Object event) {
    return ""; // if no manifest needed, return ""
  }

  @Override
  public Object toJournal(Object event) {
    return event; // identity
  }

  @Override
  public EventSeq fromJournal(Object event, String manifest) {
    return EventSeq.single(event); // identity
  }
}
```

Then in order for it to be used on events coming to and from the journal you must bind it using the below configuration syntax:

```
akka.persistence.journal {
  inmem {
    event-adapters {
      tagging = "docs.persistence.MyTaggingEventAdapter"
      user-upcasting = "docs.persistence.UserUpcastingEventAdapter"
      item-upcasting = "docs.persistence.ItemUpcastingEventAdapter"
    }

    event-adapter-bindings {
      "docs.persistence.Item" = tagging
    }
  }
}
```

```

    "docs.persistence.TaggedEvent" = tagging
    "docs.persistence.v1.Event"   = [user-upcasting, item-upcasting]
  }
}
}

```

It is possible to bind multiple adapters to one class *for recovery*, in which case the `fromJournal` methods of all bound adapters will be applied to a given matching event (in order of definition in the configuration). Since each adapter may return from 0 to n adapted events (called as `EventSeq`), each adapter can investigate the event and if it should indeed adapt it return the adapted event(s) for it, other adapters which do not have anything to contribute during this adaptation simply return `EventSeq.empty`. The adapted events are then delivered in-order to the `PersistentActor` during replay.

Note: For more advanced schema evolution techniques refer to the *persistence-schema-evolution-scala* documentation.

3.8.8 Storage plugins

Storage backends for journals and snapshot stores are pluggable in the Akka persistence extension.

Directory of persistence journal and snapshot store plugins is available at the Akka Community Projects page, see [Community plugins](#)

Plugins can be selected either by “default”, for all persistent actors and views, or “individually”, when persistent actor or view defines it’s own set of plugins.

When persistent actor or view does NOT override `journalPluginId` and `snapshotPluginId` methods, persistence extension will use “default” journal and snapshot-store plugins configured in the `reference.conf`:

```

akka.persistence.journal.plugin = ""
akka.persistence.snapshot-store.plugin = ""

```

However, these entries are provided as empty “”, and require explicit user configuration via override in the user `application.conf`. For an example of journal plugin which writes messages to LevelDB see [Local LevelDB journal](#). For an example of snapshot store plugin which writes snapshots as individual files to the local filesystem see [Local snapshot store](#).

Applications can provide their own plugins by implementing a plugin API and activate them by configuration. Plugin development requires the following imports:

```

import akka.persistence.*;
import akka.persistence.journal.japi.AsyncWriteJournal;
import akka.persistence.snapshot.japi.SnapshotStore;

```

Journal plugin API

A journal plugin extends `AsyncWriteJournal`.

`AsyncWriteJournal` is an actor and the methods to be implemented are:

```

/**
 * Java API, Plugin API: asynchronously writes a batch (`Iterable`) of
 * persistent messages to the journal.
 *
 * The batch is only for performance reasons, i.e. all messages don't have to
 * be written atomically. Higher throughput can typically be achieved by using
 * batch inserts of many records compared to inserting records one-by-one, but
 * this aspect depends on the underlying data store and a journal
 * implementation can implement it as efficient as possible. Journals should
 * aim to persist events in-order for a given `persistenceId` as otherwise in

```

```

* case of a failure, the persistent state may be end up being inconsistent.
*
* Each 'AtomicWrite' message contains the single 'PersistentRepr' that
* corresponds to the event that was passed to the 'persist' method of the
* 'PersistentActor', or it contains several 'PersistentRepr' that corresponds
* to the events that were passed to the 'persistAll' method of the
* 'PersistentActor'. All 'PersistentRepr' of the 'AtomicWrite' must be
* written to the data store atomically, i.e. all or none must be stored. If
* the journal (data store) cannot support atomic writes of multiple events it
* should reject such writes with an 'Optional' with an
* 'UnsupportedOperationException' describing the issue. This limitation
* should also be documented by the journal plugin.
*
* If there are failures when storing any of the messages in the batch the
* returned 'Future' must be completed with failure. The 'Future' must only be
* completed with success when all messages in the batch have been confirmed
* to be stored successfully, i.e. they will be readable, and visible, in a
* subsequent replay. If there is uncertainty about if the messages were
* stored or not the 'Future' must be completed with failure.
*
* Data store connection problems must be signaled by completing the 'Future'
* with failure.
*
* The journal can also signal that it rejects individual messages
* ('AtomicWrite') by the returned
* 'Iterable<Optional<Exception>>'. The returned 'Iterable' must
* have as many elements as the input 'messages' 'Iterable'. Each 'Optional'
* element signals if the corresponding 'AtomicWrite' is rejected or not, with
* an exception describing the problem. Rejecting a message means it was not
* stored, i.e. it must not be included in a later replay. Rejecting a message
* is typically done before attempting to store it, e.g. because of
* serialization error.
*
* Data store connection problems must not be signaled as rejections.
*
* Note that it is possible to reduce number of allocations by caching some
* result 'Iterable' for the happy path, i.e. when no messages are rejected.
*/
Future<Iterable<Optional<Exception>>> doAsyncWriteMessages(Iterable<AtomicWrite> messages);

/**
 * Java API, Plugin API: synchronously deletes all persistent messages up to
 * 'toSequenceNr'.
 *
 * @see AsyncRecoveryPlugin
 */
Future<Void> doAsyncDeleteMessagesTo(String persistenceId, long toSequenceNr);

```

If the storage backend API only supports synchronous, blocking writes, the methods should be implemented as:

```

@Override
public Future<Iterable<Optional<Exception>>> doAsyncWriteMessages(
    Iterable<AtomicWrite> messages) {
    try {
        Iterable<Optional<Exception>> result = new ArrayList<Optional<Exception>>();
        // blocking call here...
        // result.add(..)
        return Futures.successful(result);
    } catch (Exception e) {
        return Futures.failed(e);
    }
}

```

A journal plugin must also implement the methods defined in `AsyncRecovery` for replays and sequence number recovery:

```
/**
 * Java API, Plugin API: asynchronously replays persistent messages.
 * Implementations replay a message by calling 'replayCallback'. The returned
 * future must be completed when all messages (matching the sequence number
 * bounds) have been replayed. The future must be completed with a failure if
 * any of the persistent messages could not be replayed.
 *
 * The 'replayCallback' must also be called with messages that have been
 * marked as deleted. In this case a replayed message's 'deleted' method must
 * return 'true'.
 *
 * The 'toSequenceNr' is the lowest of what was returned by
 * {@link #doAsyncReadHighestSequenceNr} and what the user specified as
 * recovery {@link akka.persistence.Recovery} parameter.
 *
 * @param persistenceId
 *         id of the persistent actor.
 * @param fromSequenceNr
 *         sequence number where replay should start (inclusive).
 * @param toSequenceNr
 *         sequence number where replay should end (inclusive).
 * @param max
 *         maximum number of messages to be replayed.
 * @param replayCallback
 *         called to replay a single message. Can be called from any thread.
 */
Future<Void> doAsyncReplayMessages(String persistenceId, long fromSequenceNr,
    long toSequenceNr, long max, Consumer<PersistentRepr> replayCallback);

/**
 * Java API, Plugin API: asynchronously reads the highest stored sequence
 * number for the given 'persistenceId'. The persistent actor will use the
 * highest sequence number after recovery as the starting point when
 * persisting new events. This sequence number is also used as 'toSequenceNr'
 * in subsequent call to [[#asyncReplayMessages]] unless the user has
 * specified a lower 'toSequenceNr'.
 *
 * @param persistenceId
 *         id of the persistent actor.
 * @param fromSequenceNr
 *         hint where to start searching for the highest sequence number.
 */
Future<Long> doAsyncReadHighestSequenceNr(String persistenceId, long fromSequenceNr);
```

A journal plugin can be activated with the following minimal configuration:

```
# Path to the journal plugin to be used
akka.persistence.journal.plugin = "my-journal"

# My custom journal plugin
my-journal {
  # Class name of the plugin.
  class = "docs.persistence.MyJournal"
  # Dispatcher for the plugin actor.
  plugin-dispatcher = "akka.actor.default-dispatcher"
}
```

The specified plugin class must have a no-arg constructor. The plugin-dispatcher is the dispatcher used for the plugin actor. If not specified, it defaults to `akka.persistence.dispatchers.default-plugin-dispatcher`.

The journal plugin instance is an actor so the methods corresponding to requests from persistent actors are executed sequentially. It may delegate to asynchronous libraries, spawn futures, or delegate to other actors to achieve parallelism.

The journal plugin class must have a constructor without parameters or constructor with one `com.typesafe.config.Config` parameter. The plugin section of the actor system's config will be passed in the config constructor parameter.

Don't run journal tasks/futures on the system default dispatcher, since that might starve other tasks.

Snapshot store plugin API

A snapshot store plugin must extend the `SnapshotStore` actor and implement the following methods:

```
/**
 * Java API, Plugin API: asynchronously loads a snapshot.
 *
 * @param persistenceId
 *         id of the persistent actor.
 * @param criteria
 *         selection criteria for loading.
 */
Future<Optional<SelectedSnapshot>> doLoadAsync(String persistenceId,
      SnapshotSelectionCriteria criteria);

/**
 * Java API, Plugin API: asynchronously saves a snapshot.
 *
 * @param metadata
 *         snapshot metadata.
 * @param snapshot
 *         snapshot.
 */
Future<Void> doSaveAsync(SnapshotMetadata metadata, Object snapshot);

/**
 * Java API, Plugin API: deletes the snapshot identified by 'metadata'.
 *
 * @param metadata
 *         snapshot metadata.
 */
Future<Void> doDeleteAsync(SnapshotMetadata metadata);

/**
 * Java API, Plugin API: deletes all snapshots matching 'criteria'.
 *
 * @param persistenceId
 *         id of the persistent actor.
 * @param criteria
 *         selection criteria for deleting.
 */
Future<Void> doDeleteAsync(String persistenceId, SnapshotSelectionCriteria criteria);
```

A snapshot store plugin can be activated with the following minimal configuration:

```
# Path to the snapshot store plugin to be used
akka.persistence.snapshot-store.plugin = "my-snapshot-store"

# My custom snapshot store plugin
my-snapshot-store {
  # Class name of the plugin.
  class = "docs.persistence.MySnapshotStore"
  # Dispatcher for the plugin actor.
```

```
plugin-dispatcher = "akka.persistence.dispatchers.default-plugin-dispatcher"
}
```

The specified plugin class must have a no-arg constructor. The `plugin-dispatcher` is the dispatcher used for the plugin actor. If not specified, it defaults to `akka.persistence.dispatchers.default-plugin-dispatcher`.

The snapshot store instance is an actor so the methods corresponding to requests from persistent actors are executed sequentially. It may delegate to asynchronous libraries, spawn futures, or delegate to other actors to achieve parallelism.

The snapshot store plugin class must have a constructor without parameters or constructor with one `com.typesafe.config.Config` parameter. The plugin section of the actor system's config will be passed in the config constructor parameter.

Don't run snapshot store tasks/futures on the system default dispatcher, since that might starve other tasks.

Plugin TCK

In order to help developers build correct and high quality storage plugins, we provide an Technology Compatibility Kit (TCK for short).

The TCK is usable from Java as well as Scala projects, for Java you need to include the `akka-persistence-tck-experimental` dependency:

```
<dependency>
  <groupId>com.typesafe.akka</groupId>
  <artifactId>akka-persistence-tck-experimental_${scala.version}</artifactId>
  <version>2.3.5</version>
  <scope>test</scope>
</dependency>
```

To include the Journal TCK tests in your test suite simply extend the provided `JavaJournalSpec`:

```
class MyJournalSpecTest extends JavaJournalSpec {

  public MyJournalSpecTest() {
    super(ConfigFactory.parseString(
      "persistence.journal.plugin = " +
      "\"akka.persistence.journal.leveldb-shared\""));
  }
}
```

We also provide a simple benchmarking class `JavaJournalPerfSpec` which includes all the tests that `JavaJournalSpec` has, and also performs some longer operations on the Journal while printing it's performance stats. While it is NOT aimed to provide a proper benchmarking environment it can be used to get a rough feel about your journals performance in the most typical scenarios.

In order to include the `SnapshotStore` TCK tests in your test suite simply extend the `SnapshotStoreSpec`:

```
class MySnapshotStoreTest extends JavaSnapshotStoreSpec {

  public MySnapshotStoreTest() {
    super(ConfigFactory.parseString(
      "akka.persistence.snapshot-store.plugin = " +
      "\"akka.persistence.snapshot-store.local\""));
  }
}
```

In case your plugin requires some setting up (starting a mock database, removing temporary files etc.) you can override the `beforeAll` and `afterAll` methods to hook into the tests lifecycle:

```

class MyJournalSpecTest extends JavaJournalSpec {

    List<File> storageLocations = new ArrayList<File>();

    public MyJournalSpecTest() {
        super(ConfigFactory.parseString(
            "persistence.journal.plugin = " +
            "\"akka.persistence.journal.leveldb-shared\""));

        Config config = system().settings().config();
        storageLocations.add(new File(
            config.getString("akka.persistence.journal.leveldb.dir")));
        storageLocations.add(new File(
            config.getString("akka.persistence.snapshot-store.local.dir")));
    }

    @Override
    public void beforeAll() {
        for (File storageLocation : storageLocations) {
            FileUtils.deleteRecursively(storageLocation);
        }
        super.beforeAll();
    }

    @Override
    public void afterAll() {
        super.afterAll();
        for (File storageLocation : storageLocations) {
            FileUtils.deleteRecursively(storageLocation);
        }
    }
}

```

We *highly recommend* including these specifications in your test suite, as they cover a broad range of cases you might have otherwise forgotten to test for when writing a plugin from scratch.

3.8.9 Pre-packaged plugins

Local LevelDB journal

LevelDB journal plugin config entry is `akka.persistence.journal.leveldb` and it writes messages to a local LevelDB instance. Enable this plugin by defining config property:

```
# Path to the journal plugin to be used
akka.persistence.journal.plugin = "akka.persistence.journal.leveldb"
```

LevelDB based plugins will also require the following additional dependency declaration:

```
<dependency>
  <groupId>org.iq80.leveldb</groupId>
  <artifactId>leveldb</artifactId>
  <version>0.7</version>
</dependency>
<dependency>
  <groupId>org.fusesource.leveldbjni</groupId>
  <artifactId>leveldbjni-all</artifactId>
  <version>1.8</version>
</dependency>
```

The default location of the LevelDB files is a directory named `journal` in the current working directory. This location can be changed by configuration where the specified path can be relative or absolute:

```
akka.persistence.journal.leveldb.dir = "target/journal"
```

With this plugin, each actor system runs its own private LevelDB instance.

Shared LevelDB journal

A LevelDB instance can also be shared by multiple actor systems (on the same or on different nodes). This, for example, allows persistent actors to failover to a backup node and continue using the shared journal instance from the backup node.

Warning: A shared LevelDB instance is a single point of failure and should therefore only be used for testing purposes. Highly-available, replicated journal are available as [Community plugins](#).

A shared LevelDB instance is started by instantiating the `SharedLeveldbStore` actor.

```
final ActorRef store = system.actorOf(Props.create(SharedLeveldbStore.class), "store");
```

By default, the shared instance writes journaled messages to a local directory named `journal` in the current working directory. The storage location can be changed by configuration:

```
akka.persistence.journal.leveldb-shared.store.dir = "target/shared"
```

Actor systems that use a shared LevelDB store must activate the `akka.persistence.journal.leveldb-shared` plugin.

```
akka.persistence.journal.plugin = "akka.persistence.journal.leveldb-shared"
```

This plugin must be initialized by injecting the (remote) `SharedLeveldbStore` actor reference. Injection is done by calling the `SharedLeveldbJournal.setStore` method with the actor reference as argument.

```
class SharedStorageUsage extends UntypedActor {
  @Override
  public void preStart() throws Exception {
    String path = "akka.tcp://example@127.0.0.1:2552/user/store";
    ActorSelection selection = getContext().actorSelection(path);
    selection.tell(new Identify(1), getSelf());
  }

  @Override
  public void onReceive(Object message) throws Exception {
    if (message instanceof ActorIdentity) {
      ActorIdentity identity = (ActorIdentity) message;
      if (identity.correlationId().equals(1)) {
        ActorRef store = identity.getRef();
        if (store != null) {
          SharedLeveldbJournal.setStore(store, getContext().system());
        }
      }
    }
  }
}
```

Internal journal commands (sent by persistent actors) are buffered until injection completes. Injection is idempotent i.e. only the first injection is used.

Local snapshot store

Local snapshot store plugin config entry is `akka.persistence.snapshot-store.local` and it writes snapshot files to the local filesystem. Enable this plugin by defining config property:

```
# Path to the snapshot store plugin to be used
akka.persistence.snapshot-store.plugin = "akka.persistence.snapshot-store.local"
```

The default storage location is a directory named `snapshots` in the current working directory. This can be changed by configuration where the specified path can be relative or absolute:

```
akka.persistence.snapshot-store.local.dir = "target/snapshots"
```

Note that it is not mandatory to specify a snapshot store plugin. If you don't use snapshots you don't have to configure it.

3.8.10 Custom serialization

Serialization of snapshots and payloads of `Persistent` messages is configurable with Akka's *Serialization* infrastructure. For example, if an application wants to serialize

- payloads of type `MyPayload` with a custom `MyPayloadSerializer` and
- snapshots of type `MySnapshot` with a custom `MySnapshotSerializer`

it must add

```
akka.actor {
  serializers {
    my-payload = "docs.persistence.MyPayloadSerializer"
    my-snapshot = "docs.persistence.MySnapshotSerializer"
  }
  serialization-bindings {
    "docs.persistence.MyPayload" = my-payload
    "docs.persistence.MySnapshot" = my-snapshot
  }
}
```

to the application configuration. If not specified, a default serializer is used.

For more advanced schema evolution techniques refer to the *persistence-schema-evolution-scala* documentation.

3.8.11 Testing

When running tests with LevelDB default settings in `sbt`, make sure to set `fork := true` in your `sbt` project otherwise, you'll see an `UnsatisfiedLinkError`. Alternatively, you can switch to a LevelDB Java port by setting

```
akka.persistence.journal.leveldb.native = off
```

or

```
akka.persistence.journal.leveldb-shared.store.native = off
```

in your Akka configuration. The LevelDB Java port is for testing purposes only.

Warning: It is not possible to test persistence provided classes (i.e. *PersistentActor* and *AtLeastOnceDelivery*) using `TestActorRef` due to its *synchronous* nature. These traits need to be able to perform asynchronous tasks in the background in order to handle internal persistence related events. When testing Persistence based projects always rely on *asynchronous messaging using the TestKit*.

3.8.12 Configuration

There are several configuration properties for the persistence module, please refer to the *reference configuration*.

3.8.13 Multiple persistence plugin configurations

By default, persistent actor or view will use “default” journal and snapshot store plugins configured in the following sections of the `reference.conf` configuration resource:

```
# Absolute path to the default journal plugin configuration entry.
akka.persistence.journal.plugin = "akka.persistence.journal.inmem"
# Absolute path to the default snapshot store plugin configuration entry.
akka.persistence.snapshot-store.plugin = "akka.persistence.snapshot-store.local"
```

Note that in this case actor or view overrides only `persistenceId` method:

```
abstract class ActorWithDefaultPlugins extends UntypedPersistentActor {
  @Override
  public String persistenceId() { return "123"; }
}
```

When persistent actor or view overrides `journalPluginId` and `snapshotPluginId` methods, the actor or view will be serviced by these specific persistence plugins instead of the defaults:

```
abstract class ActorWithOverridePlugins extends UntypedPersistentActor {
  @Override
  public String persistenceId() { return "123"; }
  // Absolute path to the journal plugin configuration entry in the `reference.conf`
  @Override
  public String journalPluginId() { return "akka.persistence.chronicle.journal"; }
  // Absolute path to the snapshot store plugin configuration entry in the `reference.conf`
  @Override
  public String snapshotPluginId() { return "akka.persistence.chronicle.snapshot-store"; }
}
```

Note that `journalPluginId` and `snapshotPluginId` must refer to properly configured `reference.conf` plugin entries with standard class property as well as settings which are specific for those plugins, i.e.:

```
# Configuration entry for the custom journal plugin, see `journalPluginId`.
akka.persistence.chronicle.journal {
  # Standard persistence extension property: provider FQCN.
  class = "akka.persistence.chronicle.ChronicleSyncJournal"
  # Custom setting specific for the journal `ChronicleSyncJournal`.
  folder = ${user.dir}/store/journal
}
# Configuration entry for the custom snapshot store plugin, see `snapshotPluginId`.
akka.persistence.chronicle.snapshot-store {
  # Standard persistence extension property: provider FQCN.
  class = "akka.persistence.chronicle.ChronicleSnapshotStore"
  # Custom setting specific for the snapshot store `ChronicleSnapshotStore`.
  folder = ${user.dir}/store/snapshot
}
```

3.9 Persistence - Schema Evolution

When working on long running projects using *Persistence*, or any kind of *Event Sourcing* architectures, schema evolution becomes one of the more important technical aspects of developing your application. The requirements as well as our own understanding of the business domain may (and will) change in time.

In fact, if a project matures to the point where you need to evolve its schema to adapt to changing business requirements you can view this as first signs of its success – if you wouldn’t need to adapt anything over an apps lifecycle that could mean that no-one is really using it actively.

In this chapter we will investigate various schema evolution strategies and techniques from which you can pick and choose the ones that match your domain and challenge at hand.

Note: This page proposes a number of possible solutions to the schema evolution problem and explains how some of the utilities Akka provides can be used to achieve this, it is by no means a complete (closed) set of solutions.

Sometimes, based on the capabilities of your serialization formats, you may be able to evolve your schema in different ways than outlined in the sections below. If you discover useful patterns or techniques for schema evolution feel free to submit Pull Requests to this page to extend it.

3.9.1 Schema evolution in event-sourced systems

In recent years we have observed a tremendous move towards immutable append-only datastores, with event-sourcing being the prime technique successfully being used in these settings. For an excellent overview why and how immutable data makes scalability and systems design much simpler you may want to read Pat Helland's excellent [Immutability Changes Everything](#) whitepaper.

Since with [Event Sourcing](#) the **events are immutable** and usually never deleted – the way schema evolution is handled differs from how one would go about it in a mutable database setting (e.g. in typical CRUD database applications). The system needs to be able to continue to work in the presence of “old” events which were stored under the “old” schema. We also want to limit complexity in the business logic layer, exposing a consistent view over all of the events of a given type to `PersistentActor`s and *persistence queries*. This allows the business logic layer to focus on solving business problems instead of having to explicitly deal with different schemas.

The system needs to be able to continue to work in the presence of “old” events which were stored under the “old” schema, and we want to limit the complexity to the data layer, exposing a consistent view over all of the events of a given type to `PersistentActor`s and persistence queries, which allows these layers to focus on the business problems instead handling the different schemas explicitly in the business logic layers.

In summary, schema evolution in event sourced systems exposes the following characteristics:

- Allow the system to continue operating without large scale migrations to be applied,
- Allow the system to read “old” events from the underlying storage, however present them in a “new” view to the application logic,
- Transparently promote events to the latest versions during recovery (or queries) such that the business logic need not consider multiple versions of events

Types of schema evolution

Before we explain the various techniques that can be used to safely evolve the schema of your persistent events over time, we first need to define what the actual problem is, and what the typical styles of changes are.

Since events are never deleted, we need to have a way to be able to replay (read) old events, in such way that does not force the `PersistentActor` to be aware of all possible versions of an event that it may have persisted in the past. Instead, we want the Actors to work on some form of “latest” version of the event and provide some means of either converting old “versions” of stored events into this “latest” event type, or constantly evolve the event definition - in a backwards compatible way - such that the new deserialization code can still read old events.

The most common schema changes you will likely are:

- *adding a field to an event type,*
- *remove or rename field in event type,*
- *remove event type,*
- *split event into multiple smaller events.*

The following sections will explain some patterns which can be used to safely evolve your schema when facing those changes.

3.9.2 Picking the right serialization format

Picking the serialization format is a very important decision you will have to make while building your application. It affects which kind of evolutions are simple (or hard) to do, how much work is required to add a new datatype, and, last but not least, serialization performance.

If you find yourself realising you have picked “the wrong” serialization format, it is always possible to change the format used for storing new events, however you would have to keep the old deserialization code in order to be able to replay events that were persisted using the old serialization scheme. It is possible to “rebuild” an event-log from one serialization format to another one, however it may be a more involved process if you need to perform this on a live system.

Binary serialization formats that we have seen work well for long-lived applications include the very flexible IDL based: [Google Protobuf](#), [Apache Thrift](#) or [Apache Avro](#). Avro schema evolution is more “entire schema” based, instead of single fields focused like in protobuf or thrift, and usually requires using some kind of schema registry.

Users who want their data to be human-readable directly in the write-side datastore may opt to use plain-old JSON as the storage format, though that comes at a cost of lacking support for schema evolution and relatively large marshalling latency.

There are plenty excellent blog posts explaining the various trade-offs between popular serialization formats, one post we would like to highlight is the very well illustrated [Schema evolution in Avro, Protocol Buffers and Thrift](#) by Martin Kleppmann.

Provided default serializers

Akka Persistence provides [Google Protocol Buffers](#) based serializers (using [Akka Serialization](#)) for its own message types such as `PersistentRepr`, `AtomicWrite` and snapshots. Journal plugin implementations *may* choose to use those provided serializers, or pick a serializer which suits the underlying database better.

Note: Serialization is **NOT** handled automatically by Akka Persistence itself. Instead, it only provides the above described serializers, and in case a `AsyncWriteJournal` plugin implementation chooses to use them directly, the above serialization scheme will be used.

Please refer to your write journal’s documentation to learn more about how it handles serialization!

For example, some journals may choose to not use Akka Serialization *at all* and instead store the data in a format that is more “native” for the underlying datastore, e.g. using JSON or some other kind of format that the target datastore understands directly.

The below figure explains how the default serialization scheme works, and how it fits together with serializing the user provided message itself, which we will from here on refer to as the `payload` (highlighted in yellow):

The blue colored regions of the `PersistentMessage` indicate what is serialized using the generated protocol buffers serializers, and the yellow payload indicates the user provided event (by calling `persist(payload) (...)`). As you can see, the `PersistentMessage` acts as an envelope around the payload, adding various fields related to the origin of the event (`persistenceId`, `sequenceNr` and more).

More advanced techniques (e.g. [Remove event class and ignore events](#)) will dive into using the manifests for increasing the flexibility of the persisted vs. exposed types even more. However for now we will focus on the simpler evolution techniques, concerning simply configuring the payload serializers.

By default the `payload` will be serialized using Java Serialization. This is fine for testing and initial phases of your development (while you’re still figuring out things and the data will not need to stay persisted forever). However, once you move to production you should really *pick a different serializer for your payloads*.

Warning: Do not rely on Java serialization (which will be picked by Akka by default if you don’t specify any serializers) for *serious* application development! It does not lean itself well to evolving schemas over long periods of time, and its performance is also not very high (it never was designed for high-throughput scenarios).




Figure 3.1: Akka Persistence provided serializers wrap the user payload in an envelope containing all persistence-relevant information. **If the Journal uses provided Protobuf serializers for the wrapper types (e.g. PersistentRepr), then the payload will be serialized using the user configured serializer, and if none is provided explicitly, Java serialization will be used for it.**

Configuring payload serializers

This section aims to highlight the complete basics on how to define custom serializers using *Akka Serialization*. Many journal plugin implementations use Akka Serialization, thus it is tremendously important to understand how to configure it to work with your event classes.

Note: Read the *Akka Serialization* docs to learn more about defining custom serializers, to improve performance and maintainability of your system. Do not depend on Java serialization for production deployments.

The below snippet explains in the minimal amount of lines how a custom serializer can be registered. For more in-depth explanations on how serialization picks the serializer to use etc, please refer to its documentation.

First we start by defining our domain model class, here representing a person:

```

public class Person {
    public final String name;
    public final String surname;
    public Person(String name, String surname) {
        this.name = name;
        this.surname = surname;
    }
}
  
```

Next we implement a serializer (or extend an existing one to be able to handle the new Person class):

```

/**
 * Simplest possible serializer, uses a string representation of the Person class.
 *
 * Usually a serializer like this would use a library like:
 * protobuf, kryo, avro, cap'n proto, flatbuffers, SBE or some other dedicated serializer backend
 * to perform the actual to/from bytes marshalling.
 */
public class SimplestPossiblePersonSerializer extends SerializerWithStringManifest {
    private final Charset utf8 = Charset.forName("UTF-8");

    private final String personManifest = Person.class.getName();

    // unique identifier of the serializer
    @Override public int identifier() {
  
```

```

    return 1234567;
  }

  // extract manifest to be stored together with serialized object
  @Override public String manifest(Object o) {
    return o.getClass().getName();
  }

  // serialize the object
  @Override public byte[] toBinary(Object obj) {
    if (obj instanceof Person) {
      Person p = (Person) obj;
      return (p.name + "|" + p.surname).getBytes(utf8);
    } else {
      throw new IllegalArgumentException(
        "Unable to serialize to bytes, clazz was: " + obj.getClass().getName());
    }
  }

  // deserialize the object, using the manifest to indicate which logic to apply
  @Override public Object fromBinary(byte[] bytes, String manifest) {
    if (personManifest.equals(manifest)) {
      String nameAndSurname = new String(bytes, utf8);
      String[] parts = nameAndSurname.split("[|]");
      return new Person(parts[0], parts[1]);
    } else {
      throw new IllegalArgumentException(
        "Unable to deserialize from bytes, manifest was: " + manifest +
        "! Bytes length: " + bytes.length);
    }
  }
}

```

And finally we register the serializer and bind it to handle the `docs.persistence.Person` class:

```

# application.conf
akka {
  actor {
    serializers {
      person = "docs.persistence.SimplestPossiblePersonSerializer"
    }

    serialization-bindings {
      "docs.persistence.Person" = person
    }
  }
}

```

Deserialization will be performed by the same serializer which serialized the message initially because of the `identifier` being stored together with the message.

Please refer to the [Akka Serialization](#) documentation for more advanced use of serializers, especially the [Serializer with String Manifest](#) section since it is very useful for Persistence based applications dealing with schema evolutions, as we will see in some of the examples below.

3.9.3 Schema evolution in action

In this section we will discuss various schema evolution techniques using concrete examples and explaining some of the various options one might go about handling the described situation. The list below is by no means a complete guide, so feel free to adapt these techniques depending on your serializer's capabilities and/or other domain specific limitations.

Add fields

Situation: You need to add a field to an existing message type. For example, a `SeatReservation(String letter, int row)` now needs to have an associated code which indicates if it is a window or aisle seat.

Solution: Adding fields is the most common change you'll need to apply to your messages so make sure the serialization format you picked for your payloads can handle it appropriately, i.e. such changes should be *binary compatible*. This is easily achieved using the right serializer toolkit – we recommend something like [Google Protocol Buffers](#) or [Apache Thrift](#) however other tools may fit your needs just as well – picking a serializer backend is something you should research before picking one to run with. In the following examples we will be using protobuf, mostly because we are familiar with it, it does its job well and Akka is using it internally as well.

While being able to read messages with missing fields is half of the solution, you also need to deal with the missing values somehow. This is usually modeled as some kind of default value, or by representing the field as an `Optional<T>` See below for an example how reading an optional field from from a serialized protocol buffers message might look like.

```
public enum SeatType {
    Window("W"), Aisle("A"), Other("O"), Unknown("");

    private final String code;

    private SeatType(String code) {
        this.code = code;
    }

    public static SeatType fromCode(String c) {
        if (Window.code.equals(c))
            return Window;
        else if (Aisle.code.equals(c))
            return Aisle;
        else if (Other.code.equals(c))
            return Other;
        else
            return Unknown;
    }
}

public class SeatReserved {
    public final String letter;
    public final int row;
    public final SeatType seatType;

    public SeatReserved(String letter, int row, SeatType seatType) {
        this.letter = letter;
        this.row = row;
        this.seatType = seatType;
    }
}
```

Next we prepare an protocol definition using the protobuf Interface Description Language, which we'll use to generate the serializer code to be used on the Akka Serialization layer (notice that the schema approach allows us to easily rename fields, as long as the numeric identifiers of the fields do not change):

```
// FlightAppModels.proto
option java_package = "docs.persistence.proto";
option optimize_for = SPEED;

message SeatReserved {
    required string letter = 1;
    required uint32 row = 2;
    optional string seatType = 3; // the new field
}
```

The serializer implementation uses the protobuf generated classes to marshal the payloads. Optional fields can

be handled explicitly or missing values by calling the `has...` methods on the protobuf object, which we do for `seatType` in order to use a `Unknown` type in case the event was stored before we had introduced the field to this event type:

```
/**
 * Example serializer impl which uses protocol buffers generated classes (proto.*)
 * to perform the to/from binary marshalling.
 */
public class AddedFieldsSerializerWithProtobuf extends SerializerWithStringManifest {
    @Override public int identifier() {
        return 67876;
    }

    private final String seatReservedManifest = SeatReserved.class.getName();

    @Override public String manifest(Object o) {
        return o.getClass().getName();
    }

    @Override public Object fromBinary(byte[] bytes, String manifest) {
        if (seatReservedManifest.equals(manifest)) {
            // use generated protobuf serializer
            try {
                return seatReserved(FlightAppModels.SeatReserved.parseFrom(bytes));
            } catch (InvalidProtocolBufferException e) {
                throw new IllegalArgumentException(e.getMessage());
            }
        } else {
            throw new IllegalArgumentException("Unable to handle manifest: " + manifest);
        }
    }

    @Override public byte[] toBinary(Object o) {
        if (o instanceof SeatReserved) {
            SeatReserved s = (SeatReserved) o;
            return FlightAppModels.SeatReserved.newBuilder()
                .setRow(s.row)
                .setLetter(s.letter)
                .setSeatType(s.seatType.code)
                .build().toByteArray();
        } else {
            throw new IllegalArgumentException("Unable to handle: " + o);
        }
    }

    // -- fromBinary helpers --

    private SeatReserved seatReserved(FlightAppModels.SeatReserved p) {
        return new SeatReserved(p.getLetter(), p.getRow(), seatType(p));
    }

    // handle missing field by assigning "Unknown" value
    private SeatType seatType(FlightAppModels.SeatReserved p) {
        if (p.hasSeatType())
            return SeatType.fromCode(p.getSeatType());
        else
            return SeatType.Unknown;
    }
}
```


Rename fields

Situation: When first designing the system the `SeatReserved` event featured an `code` field. After some time you discover that what was originally called `code` actually means `seatNr`, thus the model should be changed to reflect this concept more accurately.

Solution 1 - using IDL based serializers: First, we will discuss the most efficient way of dealing with such kinds of schema changes – IDL based serializers.

IDL stands for Interface Description Language, and means that the schema of the messages that will be stored is based on this description. Most IDL based serializers also generate the serializer / deserializer code so that using them is not too hard. Examples of such serializers are protobuf or thrift.

Using these libraries rename operations are “free”, because the field name is never actually stored in the binary representation of the message. This is one of the advantages of schema based serializers, even though that they add the overhead of having to maintain the schema. When using serializers like this, no additional code change (except renaming the field and method used during serialization) is needed to perform such evolution:



This is how such a rename would look in protobuf:

```

// protobuf message definition, BEFORE:
message SeatReserved {
  required string code = 1;
}

// protobuf message definition, AFTER:
message SeatReserved {
  required string seatNr = 1; // field renamed, id remains the same
}
  
```


It is important to learn about the strengths and limitations of your serializers, in order to be able to move swiftly and refactor your models fearlessly as you go on with the project.

Note: Learn in-depth about the serialization engine you’re using as it will impact how you can approach schema evolution.

Some operations are “free” in certain serialization formats (more often than not: removing/adding optional fields, sometimes renaming fields etc.), while some other operations are strictly not possible.

Solution 2 - by manually handling the event versions: Another solution, in case your serialization format does not support renames as easily as the above mentioned formats, is versioning your schema. For example, you could have made your events carry an additional field called `_version` which was set to 1 (because it was the initial schema), and once you change the schema you bump this number to 2, and write an adapter which can perform the rename.

This approach is popular when your serialization format is something like JSON, where renames can not be performed automatically by the serializer. You can do these kinds of “promotions” either manually (as shown in the example below) or using a library like [Stamina](#) which helps to create those $V1 \rightarrow V2 \rightarrow V3 \rightarrow \dots \rightarrow Vn$ promotion chains without much boilerplate.



The following snippet showcases how one could apply renames if working with plain JSON (using a `JsonObject` as an example JSON representation):

```

public class JsonRenamedFieldAdapter implements EventAdapter {
    // use your favorite json library
    private final ExampleJsonMarshaller marshaller = new ExampleJsonMarshaller();

    private final String V1 = "v1";
    private final String V2 = "v2";

    // this could be done independently for each event type
    @Override public String manifest(Object event) {
        return V2;
    }

    @Override public JsonObject toJournal(Object event) {
        return marshaller.toJson(event);
    }

    @Override public EventSeq fromJournal(Object event, String manifest) {
        if (event instanceof JsonObject) {
            JsonObject json = (JsonObject) event;
            if (V1.equals(manifest))
                json = rename(json, "code", "seatNr");
            return EventSeq.single(json);
        } else {
            throw new IllegalArgumentException("Can only work with JSON, was: " +
                event.getClass().getName());
        }
    }

    private JsonObject rename(JsonObject json, String from, String to) {
        // use your favorite json library to rename the field
        JsonObject renamed = json;
        return renamed;
    }
}

```

As you can see, manually handling renames induces some boilerplate onto the `EventAdapter`, however much of it you will find is common infrastructure code that can be either provided by an external library (for promotion management) or put together in a simple helper class.

Note: The technique of versioning events and then promoting them to the latest version using JSON transformations can of course be applied to more than just field renames – it also applies to adding fields and all kinds of

changes in the message format.

Remove event class and ignore events

Situation: While investigating app performance you notice that insane amounts of `CustomerBlinked` events are being stored for every customer each time he/she blinks. Upon investigation you decide that the event does not add any value and should be deleted. You still have to be able to replay from a journal which contains those old `CustomerBlinked` events though.

Naive solution - drop events in `EventAdapter`:

The problem of removing an event type from the domain model is not as much its removal, as the implications for the recovery mechanisms that this entails. For example, a naive way of filtering out certain kinds of events from being delivered to a recovering `PersistentActor` is pretty simple, as one can simply filter them out in an `EventAdapter`:




Figure 3.2: The `EventAdapter` can drop old events (O) by emitting an empty `EventSeq`. Other events can simply be passed through (E).

This however does not address the underlying cost of having to deserialize all the events during recovery, even those which will be filtered out by the adapter. In the next section we will improve the above explained mechanism to avoid deserializing events which would be filtered out by the adapter anyway, thus allowing to save precious time during a recovery containing lots of such events (without actually having to delete them).

Improved solution - deserialize into tombstone:

In the just described technique we have saved the `PersistentActor` from receiving un-wanted events by filtering them out in the `EventAdapter`, however the event itself still was deserialized and loaded into memory. This has two notable *downsides*:

- first, that the deserialization was actually performed, so we spent some of our time budget on the deserialization, even though the event does not contribute anything to the persistent actors state.
- second, that we are *unable to remove the event class* from the system – since the serializer still needs to create the actual instance of it, as it does not know it will not be used.

The solution to these problems is to use a serializer that is aware of that event being no longer needed, and can notice this before starting to deserialize the object.

This approach allows us to *remove the original class from our classpath*, which makes for less “old” classes lying around in the project. This can for example be implemented by using an `SerializerWithStringManifest` (documented in depth in *Serializer with String Manifest*). By looking at the string manifest, the serializer can notice that the type is no longer needed, and skip the deserialization all-together:

The serializer detects that the string manifest points to a removed event type and skips attempting to deserialize it:

```
public class EventDeserializationSkipped {
    public static EventDeserializationSkipped instance =
        new EventDeserializationSkipped();
}
```




Figure 3.3: The serializer is aware of the old event types that need to be skipped (O), and can skip deserializing them altogether by simply returning a “tombstone” (T), which the EventAdapter converts into an empty EventSeq. Other events (E) can simply be passed through.

```

private EventDeserializationSkipped() {
}

public class RemovedEventsAwareSerializer extends SerializerWithStringManifest {
    private final Charset utf8 = Charset.forName("UTF-8");
    private final String customerBlinkedManifest = "blinked";

    // unique identifier of the serializer
    @Override public int identifier() {
        return 8337;
    }

    // extract manifest to be stored together with serialized object
    @Override public String manifest(Object o) {
        if (o instanceof CustomerBlinked)
            return customerBlinkedManifest;
        else
            return o.getClass().getName();
    }

    @Override public byte[] toBinary(Object o) {
        return o.toString().getBytes(utf8); // example serialization
    }

    @Override public Object fromBinary(byte[] bytes, String manifest) {
        if (customerBlinkedManifest.equals(manifest))
            return EventDeserializationSkipped.instance;
        else
            return new String(bytes, utf8);
    }
}

```

The EventAdapter we implemented is aware of EventDeserializationSkipped events (our “Tombstones”), and emits an empty EventSeq whenever such object is encountered:

```

public class SkippedEventsAwareAdapter implements EventAdapter {
    @Override public String manifest(Object event) {
        return "";
    }

    @Override public Object toJournal(Object event) {
        return event;
    }

    @Override public EventSeq fromJournal(Object event, String manifest) {
        if (event == EventDeserializationSkipped.instance)

```

```

return EventSeq.empty();
else
return EventSeq.single(event);
}
}

```

Detach domain model from data model

Situation: You want to separate the application model (often called the “*domain model*”) completely from the models used to persist the corresponding events (the “*data model*”). For example because the data representation may change independently of the domain model.

Another situation where this technique may be useful is when your serialization tool of choice requires generated classes to be used for serialization and deserialization of objects, like for example [Google Protocol Buffers](#) do, yet you do not want to leak this implementation detail into the domain model itself, which you’d like to model as plain Java classes.

Solution: In order to detach the domain model, which is often represented using pure java (case) classes, from the data model classes which very often may be less user-friendly yet highly optimised for throughput and schema evolution (like the classes generated by protobuf for example), it is possible to use a simple EventAdapter which maps between these types in a 1:1 style as illustrated below:




Figure 3.4: Domain events (**A**) are adapted to the data model events (**D**) by the `EventAdapter`. The data model can be a format natively understood by the journal, such that it can store it more efficiently or include additional data for the event (e.g. tags), for ease of later querying.

We will use the following domain and data models to showcase how the separation can be implemented by the adapter:

```

// Domain model - highly optimised for domain language and maybe "fluent" usage
public class Customer {
    public final String name;

    public Customer(String name) {
        this.name = name;
    }
}

public class Seat {
    public final String code;

    public Seat(String code) {
        this.code = code;
    }

    public SeatBooked bookFor(Customer customer) {
        return new SeatBooked(code, customer);
    }
}

```

```

}

public class SeatBooked {
    public final String code;
    public final Customer customer;

    public SeatBooked(String code, Customer customer) {
        this.code = code;
        this.customer = customer;
    }
}

// Data model - highly optimised for schema evolution and persistence
public class SeatBookedData {
    public final String code;
    public final String customerName;

    public SeatBookedData(String code, String customerName) {
        this.code = code;
        this.customerName = customerName;
    }
}

```

The `EventAdapter` takes care of converting from one model to the other one (in both directions), allowing the models to be completely detached from each other, such that they can be optimised independently as long as the mapping logic is able to convert between them:

```

class DetachedModelsAdapter implements EventAdapter {
    @Override public String manifest(Object event) {
        return "";
    }

    @Override public Object toJournal(Object event) {
        if (event instanceof SeatBooked) {
            SeatBooked s = (SeatBooked) event;
            return new SeatBookedData(s.code, s.customer.name);
        } else {
            throw new IllegalArgumentException("Unsupported: " + event.getClass());
        }
    }

    @Override public EventSeq fromJournal(Object event, String manifest) {
        if (event instanceof SeatBookedData) {
            SeatBookedData d = (SeatBookedData) event;
            return EventSeq.single(new SeatBooked(d.code, new Customer(d.customerName)));
        } else {
            throw new IllegalArgumentException("Unsupported: " + event.getClass());
        }
    }
}

```

The same technique could also be used directly in the `Serializer` if the end result of marshalling is bytes. Then the serializer can simply convert the bytes to the domain object by using the generated protobuf builders.

Store events as human-readable data model

Situation: You want to keep your persisted events in a human-readable format, for example JSON.

Solution: This is a special case of the *Detach domain model from data model* pattern, and thus requires some co-operation from the `Journal` implementation to achieve this.

An example of a `Journal` which may implement this pattern is MongoDB, however other databases such as PostgreSQL and Cassandra could also do it because of their built-in JSON capabilities.

In this approach, the `EventAdapter` is used as the marshalling layer: it serializes the events to/from JSON. The journal plugin notices that the incoming event type is JSON (for example by performing a `match` on the incoming event) and stores the incoming object directly.

```
public class JsonDataModelAdapter implements EventAdapter {

    // use your favorite json library
    private final ExampleJsonMarshaller marshaller =
        new ExampleJsonMarshaller();

    @Override public String manifest(Object event) {
        return "";
    }

    @Override public JsObject toJournal(Object event) {
        return marshaller.toJson(event);
    }

    @Override public EventSeq fromJournal(Object event, String manifest) {
        if (event instanceof JsObject) {
            JsObject json = (JsObject) event;
            return EventSeq.single(marshaller.fromJson(json));
        } else {
            throw new IllegalArgumentException(
                "Unable to fromJournal a non-JSON object! Was: " + event.getClass());
        }
    }
}
```

Note: This technique only applies if the Akka Persistence plugin you are using provides this capability. Check the documentation of your favourite plugin to see if it supports this style of persistence.

If it doesn't, you may want to skim the [list of existing journal plugins](#), just in case some other plugin for your favourite datastore *does* provide this capability.

Alternative solution:

In fact, an `AsyncWriteJournal` implementation could natively decide to not use binary serialization at all, and *always* serialize the incoming messages as JSON - in which case the `toJournal` implementation of the `EventAdapter` would be an identity function, and the `fromJournal` would need to de-serialize messages from JSON.

Note: If in need of human-readable events on the *write-side* of your application reconsider whether preparing materialized views using *Persistence Query* would not be an efficient way to go about this, without compromising the write-side's throughput characteristics.

If indeed you want to use a human-readable representation on the write-side, pick a Persistence plugin that provides that functionality, or – implement one yourself.

Split large event into fine-grained events

Situation: While refactoring your domain events, you find that one of the events has become too large (coarse-grained) and needs to be split up into multiple fine-grained events.

Solution: Let us consider a situation where an event represents “user details changed”. After some time we discover that this event is too coarse, and needs to be split into “user name changed” and “user address changed”, because somehow users keep changing their usernames a lot and we'd like to keep this as a separate event.

The write side change is very simple, we simply persist `UserNameChanged` or `UserAddressChanged` depending on what the user actually intended to change (instead of the composite `UserDetailsChanged` that

we had in version 1 of our model).




Figure 3.5: The EventAdapter splits the incoming event into smaller more fine grained events during recovery.

During recovery however, we now need to convert the old V1 model into the V2 representation of the change. Depending if the old event contains a name change, we either emit the `UserNameChanged` or we don't, and the address change is handled similarly:

```

interface V1 {};
interface V2 {}

// V1 event:
public class UserDetailsChanged implements V1 {
    public final String name;
    public final String address;
    public UserDetailsChanged(String name, String address) {
        this.name = name;
        this.address = address;
    }
}

// corresponding V2 events:
public class UserNameChanged implements V2 {
    public final String name;

    public UserNameChanged(String name) {
        this.name = name;
    }
}

public class UserAddressChanged implements V2 {
    public final String address;

    public UserAddressChanged(String address) {
        this.address = address;
    }
}

// event splitting adapter:
public class UserEventsAdapter implements EventAdapter {
    @Override public String manifest(Object event) {
        return "";
    }

    @Override public EventSeq fromJournal(Object event, String manifest) {
        if (event instanceof UserDetailsChanged) {
            UserDetailsChanged c = (UserDetailsChanged) event;
            if (c.name == null)
                return EventSeq.single(new UserAddressChanged(c.address));
            else if (c.address == null)
                return EventSeq.single(new UserNameChanged(c.name));
        }
    }
}
  
```

```

else
    return EventSeq.create(
        new UserNameChanged(c.name),
        new UserAddressChanged(c.address));
} else {
    return EventSeq.single(event);
}
}

@Override public Object toJournal(Object event) {
    return event;
}
}

```

By returning an `EventSeq` from the event adapter, the recovered event can be converted to multiple events before being delivered to the persistent actor.

3.10 Persistence Query

Akka persistence query complements *Persistence* by providing a universal asynchronous stream based query interface that various journal plugins can implement in order to expose their query capabilities.

The most typical use case of persistence query is implementing the so-called query side (also known as “read side”) in the popular CQRS architecture pattern - in which the writing side of the application (e.g. implemented using akka persistence) is completely separated from the “query side”. Akka Persistence Query itself is *not* directly the query side of an application, however it can help to migrate data from the write side to the query side database. In very simple scenarios Persistence Query may be powerful enough to fulfill the query needs of your app, however we highly recommend (in the spirit of CQRS) of splitting up the write/read sides into separate datastores as the need arises.

Warning: This module is marked as “**experimental**” as of its introduction in Akka 2.4.0. We will continue to improve this API based on our users’ feedback, which implies that while we try to keep incompatible changes to a minimum the binary compatibility guarantee for maintenance releases does not apply to the contents of the `akka.persistence.query` package.

3.10.1 Dependencies

Akka persistence query is a separate jar file. Make sure that you have the following dependency in your project:

```

<dependency>
  <groupId>com.typesafe.akka</groupId>
  <artifactId>akka-persistence-query-experimental_2.11</artifactId>
  <version>2.4.0</version>
</dependency>

```

3.10.2 Design overview

Akka persistence query is purposely designed to be a very loosely specified API. This is in order to keep the provided APIs general enough for each journal implementation to be able to expose its best features, e.g. a SQL journal can use complex SQL queries or if a journal is able to subscribe to a live event stream this should also be possible to expose the same API - a typed stream of events.

Each read journal must explicitly document which types of queries it supports. Refer to the your journal’s plugins documentation for details on which queries and semantics it supports.

While Akka Persistence Query does not provide actual implementations of `ReadJournal`, it defines a number of pre-defined query types for the most common query scenarios, that most journals are likely to implement (however they are not required to).

3.10.3 Read Journals

In order to issue queries one has to first obtain an instance of a `ReadJournal`. Read journals are implemented as [Community plugins](#), each targeting a specific datastore (for example Cassandra or JDBC databases). For example, given a library that provides a `akka.persistence.query.my-read-journal` obtaining the related journal is as simple as:

```
// obtain read journal by plugin id
final MyJavadslReadJournal readJournal =
    PersistenceQuery.get(system).getReadJournalFor(MyJavadslReadJournal.class,
        "akka.persistence.query.my-read-journal");

// issue query to journal
Source<EventEnvelope, BoxedUnit> source =
    readJournal.eventsByPersistenceId("user-1337", 0, Long.MAX_VALUE);

// materialize stream, consuming events
ActorMaterializer mat = ActorMaterializer.create(system);
source.runForeach(event -> System.out.println("Event: " + event), mat);
```

Journal implementers are encouraged to put this identifier in a variable known to the user, such that one can access it via `getJournalFor(NoopJournal.class, NoopJournal.identifier)`, however this is not enforced.

Read journal implementations are available as [Community plugins](#).

Predefined queries

Akka persistence query comes with a number of query interfaces built in and suggests Journal implementors to implement them according to the semantics described below. It is important to notice that while these query types are very common a journal is not obliged to implement all of them - for example because in a given journal such query would be significantly inefficient.

Note: Refer to the documentation of the `ReadJournal` plugin you are using for a specific list of supported query types. For example, Journal plugins should document their stream completion strategies.

The predefined queries are:

`AllPersistenceIdsQuery` and `CurrentPersistenceIdsQuery`

`allPersistenceIds` which is designed to allow users to subscribe to a stream of all persistent ids in the system. By default this stream should be assumed to be a “live” stream, which means that the journal should keep emitting new persistence ids as they come into the system:

```
readJournal.allPersistenceIds();
```

If your usage does not require a live stream, you can use the `currentPersistenceIds` query:

```
readJournal.currentPersistenceIds();
```

EventsByPersistenceIdQuery and CurrentEventsByPersistenceIdQuery

`eventsByPersistenceId` is a query equivalent to replaying a *PersistentActor*, however, since it is a stream it is possible to keep it alive and watch for additional incoming events persisted by the persistent actor identified by the given `persistenceId`.

```
readJournal.eventsByPersistenceId("user-us-1337", 0L, Long.MAX_VALUE);
```

Most journals will have to revert to polling in order to achieve this, which can typically be configured with a `refresh-interval` configuration property.

If your usage does not require a live stream, you can use the `currentEventsByPersistenceId` query.

EventsByTag and CurrentEventsByTag

`eventsByTag` allows querying events regardless of which `persistenceId` they are associated with. This query is hard to implement in some journals or may need some additional preparation of the used data store to be executed efficiently. The goal of this query is to allow querying for all events which are “tagged” with a specific tag. That includes the use case to query all domain events of an Aggregate Root type. Please refer to your read journal plugin’s documentation to find out if and how it is supported.

Some journals may support tagging of events via an *Event Adapters* that wraps the events in a `akka.persistence.journal.Tagged` with the given tags. The journal may support other ways of doing tagging - again, how exactly this is implemented depends on the used journal. Here is an example of such a tagging event adapter:

```
public class MyTaggingEventAdapter implements WriteEventAdapter {

    @Override
    public Object toJournal(Object event) {
        if (event instanceof String) {
            String s = (String) event;
            Set<String> tags = new HashSet<String>();
            if (s.contains("green")) tags.add("green");
            if (s.contains("black")) tags.add("black");
            if (s.contains("blue")) tags.add("blue");
            if (tags.isEmpty())
                return event;
            else
                return new Tagged(event, tags);
        } else {
            return event;
        }
    }

    @Override
    public String manifest(Object event) {
        return "";
    }
}
```

Note: A very important thing to keep in mind when using queries spanning multiple `persistenceIds`, such as `EventsByTag` is that the order of events at which the events appear in the stream rarely is guaranteed (or stable between materializations).

Journals *may* choose to opt for strict ordering of the events, and should then document explicitly what kind of ordering guarantee they provide - for example “*ordered by timestamp ascending, independently of persistenceId*” is easy to achieve on relational databases, yet may be hard to implement efficiently on plain key-value datastores.

In the example below we query all events which have been tagged (we assume this was performed by the write-side using an *EventAdapter*, or that the journal is smart enough that it can figure out what we mean by this tag - for example if the journal stored the events as json it may try to find those with the field `tag` set to this value etc.).

```
// assuming journal is able to work with numeric offsets we can:
final Source<EventEnvelope, BoxedUnit> blueThings =
    readJournal.eventsByTag("blue", 0L);

// find top 10 blue things:
final Future<List<Object>> top10BlueThings =
    (Future<List<Object>>) blueThings
        .map(t -> t.event())
        .take(10) // cancels the query stream after pulling 10 elements
        .<List<Object>>runFold(new ArrayList<>(10), (acc, e) -> {
            acc.add(e);
            return acc;
        }, mat);

// start another query, from the known offset
Source<EventEnvelope, BoxedUnit> blue = readJournal.eventsByTag("blue", 10);
```

As you can see, we can use all the usual stream combinators available from [Akka Streams](#) on the resulting query stream, including for example taking the first 10 and cancelling the stream. It is worth pointing out that the built-in `EventsByTag` query has an optionally supported offset parameter (of type `Long`) which the journals can use to implement resumable-streams. For example a journal may be able to use a `WHERE` clause to begin the read starting from a specific row, or in a datastore that is able to order events by insertion time it could treat the `Long` as a timestamp and select only older events.

If your usage does not require a live stream, you can use the `currentEventsByTag` query.

Materialized values of queries

Journals are able to provide additional information related to a query by exposing [materialized values](#), which are a feature of [Akka Streams](#) that allows to expose additional values at stream materialization time.

More advanced query journals may use this technique to expose information about the character of the materialized stream, for example if it's finite or infinite, strictly ordered or not ordered at all. The materialized value type is defined as the second type parameter of the returned `Source`, which allows journals to provide users with their specialised query object, as demonstrated in the sample below:

```
public class RichEvent {
    public final Set<String> tags;
    public final Object payload;

    public RichEvent(Set<String> tags, Object payload) {
        this.tags = tags;
        this.payload = payload;
    }
}

// a plugin can provide:
public final class QueryMetadata{
    public final boolean deterministicOrder;
    public final boolean infinite;

    public QueryMetadata(boolean deterministicOrder, boolean infinite) {
        this.deterministicOrder = deterministicOrder;
        this.infinite = infinite;
    }
}

public Source<RichEvent, QueryMetadata> byTagsWithMeta(Set<String> tags) {
```

```

Set<String> tags = new HashSet<String>();
tags.add("red");
tags.add("blue");
final Source<RichEvent, QueryMetadata> events = readJournal.byTagsWithMeta(tags)
  .mapMaterializedValue(meta -> {
    System.out.println("The query is: " +
      "ordered deterministically: " + meta.deterministicOrder + " " +
      "infinite: " + meta.infinite);
    return meta;
  });

events.map(event -> {
  System.out.println("Event payload: " + event.payload);
  return event.payload;
}).runWith(Sink.ignore(), mat);

```

3.10.4 Performance and denormalization

When building systems using *Event sourcing* and CQRS (Command & Query Responsibility Segregation) techniques it is tremendously important to realise that the write-side has completely different needs from the read-side, and separating those concerns into datastores that are optimised for either side makes it possible to offer the best experience for the write and read sides independently.

For example, in a bidding system it is important to “take the write” and respond to the bidder that we have accepted the bid as soon as possible, which means that write-throughput is of highest importance for the write-side – often this means that data stores which are able to scale to accommodate these requirements have a less expressive query side.

On the other hand the same application may have some complex statistics view or we may have analysts working with the data to figure out best bidding strategies and trends – this often requires some kind of expressive query capabilities like for example SQL or writing Spark jobs to analyse the data. Therefore the data stored in the write-side needs to be projected into the other read-optimised datastore.

Note: When referring to **Materialized Views** in Akka Persistence think of it as “some persistent storage of the result of a Query”. In other words, it means that the view is created once, in order to be afterwards queried multiple times, as in this format it may be more efficient or interesting to query it (instead of the source events directly).

Materialize view to Reactive Streams compatible datastore

If the read datastore exposes a *Reactive Streams* interface then implementing a simple projection is as simple as, using the read-journal and feeding it into the databases driver interface, for example like so:

```

final ReactiveStreamsCompatibleDBDriver driver = new ReactiveStreamsCompatibleDBDriver();
final Subscriber<List<Object>> dbBatchWriter = driver.batchWriter();

// Using an example (Reactive Streams) Database driver
readJournal
  .eventsByPersistenceId("user-1337", 0L, Long.MAX_VALUE)
  .map(envelope -> envelope.event())
  .grouped(20) // batch inserts into groups of 20
  .runWith(Sink.create(dbBatchWriter), mat); // write batches to read-side database

```

Materialize view using mapAsync

If the target database does not provide a reactive streams *Subscriber* that can perform writes, you may have to implement the write logic using plain functions or Actors instead.

In case your write logic is state-less and you just need to convert the events from one data type to another before writing into the alternative datastore, then the projection is as simple as:

```
class ExampleStore {
  Future<Void> save(Object any) {
    // ...
  }
}

final ExampleStore store = new ExampleStore();

readJournal
  .eventsByTag("bid", 0L)
  .mapAsync(1, store::save)
  .runWith(Sink.ignore(), mat);
```

Resumable projections

Sometimes you may need to implement “resumable” projections, that will not start from the beginning of time each time when run. In this case you will need to store the sequence number (or `offset`) of the processed event and use it the next time this projection is started. This pattern is not built-in, however is rather simple to implement yourself.

The example below additionally highlights how you would use Actors to implement the write side, in case you need to do some complex logic that would be best handled inside an Actor before persisting the event into the other datastore:

```
final Timeout timeout = Timeout.apply(3, TimeUnit.SECONDS);

final MyResumableProjection bidProjection = new MyResumableProjection("bid");

final Props writerProps = Props.create(TheOneWhoWritesToQueryJournal.class, "bid");
final ActorRef writer = system.actorOf(writerProps, "bid-projection-writer");

long startFromOffset = Await.result(bidProjection.latestOffset(), timeout.duration());

readJournal
  .eventsByTag("bid", startFromOffset)
  .<Long>mapAsync(8, envelope -> {
    final Future<Object> f = ask(writer, envelope.event(), timeout);
    return f.<Long>map(new Mapper<Object, Long>() {
      @Override public Long apply(Object in) {
        return envelope.offset();
      }
    }, system.dispatcher());
  })
  .mapAsync(1, offset -> bidProjection.saveProgress(offset))
  .runWith(Sink.ignore(), mat);
}

final class TheOneWhoWritesToQueryJournal extends AbstractActor {
  private final ExampleStore store;

  private ComplexState state = new ComplexState();

  public TheOneWhoWritesToQueryJournal() {
    store = new ExampleStore();
  }

  receive(ReceiveBuilder.matchAny(message -> {
    state = updateState(state, message);

    // example saving logic that requires state to become ready:
```

```

    if (state.readyToSave())
        store.save(Record.of(state));

    }).build());
}

ComplexState updateState(ComplexState state, Object msg) {
    // some complicated aggregation logic here ...
    return state;
}
}

```

3.10.5 Query plugins

Query plugins are various (mostly community driven) `ReadJournal` implementations for all kinds of available datastores. The complete list of available plugins is maintained on the Akka Persistence Query [Community Plugins](#) page.

The plugin for LevelDB is described in *Persistence Query for LevelDB*.

This section aims to provide tips and guide plugin developers through implementing a custom query plugin. Most users will not need to implement journals themselves, except if targeting a not yet supported datastore.

Note: Since different data stores provide different query capabilities journal plugins **must extensively document** their exposed semantics as well as handled query scenarios.

ReadJournal plugin API

A read journal plugin must implement `akka.persistence.query.ReadJournalProvider` which creates instances of `akka.persistence.query.scaladsl.ReadJournal` and `akka.persistence.query.javaadsl.ReadJournal`. The plugin must implement both the `scaladsl` and the `javaadsl` interfaces because the `akka.stream.scaladsl.Source` and `akka.stream.javaadsl.Source` are different types and even though those types can easily be converted to each other it is most convenient for the end user to get access to the Java or Scala `Source` directly. As illustrated below one of the implementations can delegate to the other.

Below is a simple journal implementation:

```

public class MyReadJournalProvider implements ReadJournalProvider {
    private final MyJavaadslReadJournal javaadslReadJournal;

    public MyReadJournalProvider(ExtendedActorSystem system, Config config) {
        this.javaadslReadJournal = new MyJavaadslReadJournal(system, config);
    }

    @Override
    public MyScaladslReadJournal scaladslReadJournal() {
        return new MyScaladslReadJournal(javaadslReadJournal);
    }

    @Override
    public MyJavaadslReadJournal javaadslReadJournal() {
        return this.javaadslReadJournal;
    }
}

public class MyJavaadslReadJournal implements
akka.persistence.query.javaadsl.ReadJournal,
akka.persistence.query.javaadsl.EventsByTagQuery,

```

```

akka.persistence.query.javadsl.EventsByPersistenceIdQuery,
akka.persistence.query.javadsl.AllPersistenceIdsQuery,
akka.persistence.query.javadsl.CurrentPersistenceIdsQuery {

    private final FiniteDuration refreshInterval;

    public MyJavadslReadJournal(ExtendedActorSystem system, Config config) {
        refreshInterval =
            FiniteDuration.create(config.getDuration("refresh-interval",
                TimeUnit.MILLISECONDS), TimeUnit.MILLISECONDS);
    }

    @Override
    public Source<EventEnvelope, BoxedUnit> eventsByTag(String tag, long offset) {
        final Props props = MyEventsByTagPublisher.props(tag, offset, refreshInterval);
        return Source.<EventEnvelope>actorPublisher(props).
            mapMaterializedValue(m -> BoxedUnit.UNIT);
    }

    @Override
    public Source<EventEnvelope, BoxedUnit> eventsByPersistenceId(String persistenceId,
        long fromSequenceNr, long toSequenceNr) {
        // implement in a similar way as eventsByTag
        throw new UnsupportedOperationException("Not implemented yet");
    }

    @Override
    public Source<String, BoxedUnit> allPersistenceIds() {
        // implement in a similar way as eventsByTag
        throw new UnsupportedOperationException("Not implemented yet");
    }

    @Override
    public Source<String, BoxedUnit> currentPersistenceIds() {
        // implement in a similar way as eventsByTag
        throw new UnsupportedOperationException("Not implemented yet");
    }

    // possibility to add more plugin specific queries

    public Source<RichEvent, QueryMetadata> byTagsWithMeta(Set<String> tags) {
        // implement in a similar way as eventsByTag
        throw new UnsupportedOperationException("Not implemented yet");
    }
}

public class MyScaladslReadJournal implements
akka.persistence.query.scaladsl.ReadJournal,
akka.persistence.query.scaladsl.EventsByTagQuery,
akka.persistence.query.scaladsl.EventsByPersistenceIdQuery,
akka.persistence.query.scaladsl.AllPersistenceIdsQuery,
akka.persistence.query.scaladsl.CurrentPersistenceIdsQuery {

    private final MyJavadslReadJournal javadslReadJournal;

    public MyScaladslReadJournal(MyJavadslReadJournal javadslReadJournal) {
        this.javadslReadJournal = javadslReadJournal;
    }

    @Override
    public akka.stream.scaladsl.Source<EventEnvelope, BoxedUnit> eventsByTag(
        String tag, long offset) {
        return javadslReadJournal.eventsByTag(tag, offset).asScala();
    }
}

```

```

}

@Override
public akka.stream.scaladsl.Source<EventEnvelope, BoxedUnit> eventsByPersistenceId(
    String persistenceId, long fromSequenceNr, long toSequenceNr) {
    return javadslReadJournal.eventsByPersistenceId(persistenceId, fromSequenceNr,
        toSequenceNr).asScala();
}

@Override
public akka.stream.scaladsl.Source<String, BoxedUnit> allPersistenceIds() {
    return javadslReadJournal.allPersistenceIds().asScala();
}

@Override
public akka.stream.scaladsl.Source<String, BoxedUnit> currentPersistenceIds() {
    return javadslReadJournal.currentPersistenceIds().asScala();
}

// possibility to add more plugin specific queries

public akka.stream.scaladsl.Source<RichEvent, QueryMetadata> byTagsWithMeta(
    scala.collection.Set<String> tags) {
    Set<String> jTags = scala.collection.JavaConversions.setAsJavaSet(tags);
    return javadslReadJournal.byTagsWithMeta(jTags).asScala();
}
}

```

And the `EventsByTag` could be backed by such an Actor for example:

```

class MyEventsByTagJavaPublisher extends AbstractActorPublisher<EventEnvelope> {
    private final Serialization serialization =
        SerializationExtension.get(context().system());

    private final Connection connection;

    private final String tag;

    private final String CONTINUE = "CONTINUE";
    private final int LIMIT = 1000;
    private long currentOffset;
    private List<EventEnvelope> buf = new LinkedList<>();

    private Cancellable continueTask;

    public MyEventsByTagJavaPublisher(Connection connection,
        String tag,
        Long offset,
        FiniteDuration refreshInterval) {
        this.connection = connection;
        this.tag = tag;
        this.currentOffset = offset;

        final Scheduler scheduler = context().system().scheduler();
        this.continueTask = scheduler
            .schedule(refreshInterval, refreshInterval, self(), CONTINUE,
                context().dispatcher(), self());

        receive(ReceiveBuilder
            .matchEquals(CONTINUE, (in) -> {
                query();
                deliverBuf();
            })
        );
    }
}

```

```

        })
        .match(Cancel.class, (in) -> {
            context().stop(self());
        })
        .build());
    }

    public static Props props(Connection conn, String tag, Long offset,
        FiniteDuration refreshInterval) {
        return Props.create(() ->
            new MyEventsByTagJavaPublisher(conn, tag, offset, refreshInterval));
    }

    @Override
    public void postStop() {
        continueTask.cancel();
    }

    private void query() {
        if (buf.isEmpty()) {
            final String query = "SELECT id, persistent_repr " +
                "FROM journal WHERE tag = ? AND id >= ? " +
                "ORDER BY id LIMIT ?";

            try (PreparedStatement s = connection.prepareStatement(query)) {
                s.setString(1, tag);
                s.setLong(2, currentOffset);
                s.setLong(3, LIMIT);
                try (ResultSet rs = s.executeQuery()) {

                    final List<Pair<Long, byte[]>> res = new ArrayList<>(LIMIT);
                    while (rs.next())
                        res.add(Pair.create(rs.getLong(1), rs.getBytes(2)));

                    if (!res.isEmpty()) {
                        currentOffset = res.get(res.size() - 1).first();
                    }

                    buf = res.stream().map(in -> {
                        final Long id = in.first();
                        final byte[] bytes = in.second();

                        final PersistentRepr p =
                            serialization.deserialize(bytes, PersistentRepr.class).get();

                        return new EventEnvelope(id, p.persistenceId(), p.sequenceNr(), p.payload());
                    }).collect(toList());
                }
            } catch (Exception e) {
                onErrorThenStop(e);
            }
        }
    }

    private void deliverBuf() {
        while (totalDemand() > 0 && !buf.isEmpty())
            onNext(buf.remove(0));
    }
}

```

If the underlying datastore only supports queries that are completed when they reach the end of the “result set”, the journal has to submit new queries after a while in order to support “infinite” event streams that include events stored after the initial query has completed. It is recommended that the plugin use a configuration property named

`refresh-interval` for defining such a refresh interval.

Plugin TCK

TODO, not available yet.

3.11 Persistence Query for LevelDB

This is documentation for the LevelDB implementation of the *Persistence Query* API. Note that implementations for other journals may have different semantics.

Warning: This module is marked as “**experimental**” as of its introduction in Akka 2.4.0. We will continue to improve this API based on our users’ feedback, which implies that while we try to keep incompatible changes to a minimum the binary compatibility guarantee for maintenance releases does not apply to the contents of the `akka.persistence.query` package.

3.11.1 Dependencies

Akka persistence LevelDB query implementation is bundled in the `akka-persistence-query-experimental` artifact. Make sure that you have the following dependency in your project:

```
<dependency>
  <groupId>com.typesafe.akka</groupId>
  <artifactId>akka-persistence-query-experimental_2.11</artifactId>
  <version>2.4.0</version>
</dependency>
```

3.11.2 How to get the ReadJournal

The `ReadJournal` is retrieved via the `akka.persistence.query.PersistenceQuery` extension:

```
final ActorMaterializer mat = ActorMaterializer.create(system);

LeveldbReadJournal queries =
  PersistenceQuery.get(system).getReadJournalFor(LeveldbReadJournal.class,
    LeveldbReadJournal.Identifier());
```

3.11.3 Supported Queries

EventsByPersistenceIdQuery and CurrentEventsByPersistenceIdQuery

`eventsByPersistenceId` is used for retrieving events for a specific `PersistentActor` identified by `persistenceId`.

```
LeveldbReadJournal queries =
  PersistenceQuery.get(system).getReadJournalFor(LeveldbReadJournal.class,
    LeveldbReadJournal.Identifier());

Source<EventEnvelope, BoxedUnit> source =
  queries.eventsByPersistenceId("some-persistence-id", 0, Long.MAX_VALUE);
```

You can retrieve a subset of all events by specifying `fromSequenceNr` and `toSequenceNr` or use `0L` and `Long.MAX_VALUE` respectively to retrieve all events. Note that the corresponding sequence number of each

event is provided in the `EventEnvelope`, which makes it possible to resume the stream at a later point from a given sequence number.

The returned event stream is ordered by sequence number, i.e. the same order as the `PersistentActor` persisted the events. The same prefix of stream elements (in same order) are returned for multiple executions of the query, except for when events have been deleted.

The stream is not completed when it reaches the end of the currently stored events, but it continues to push new events when new events are persisted. Corresponding query that is completed when it reaches the end of the currently stored events is provided by `currentEventsByPersistenceId`.

The LevelDB write journal is notifying the query side as soon as events are persisted, but for efficiency reasons the query side retrieves the events in batches that sometimes can be delayed up to the configured `refresh-interval` or given `RefreshInterval` hint.

The stream is completed with failure if there is a failure in executing the query in the backend journal.

AllPersistenceIdsQuery and CurrentPersistenceIdsQuery

`allPersistenceIds` is used for retrieving all `persistenceIds` of all persistent actors.

```
LeveldbReadJournal queries =
    PersistenceQuery.get(system).getReadJournalFor(LeveldbReadJournal.class,
        LeveldbReadJournal.Identifier());

Source<String, BoxedUnit> source = queries.allPersistenceIds();
```

The returned event stream is unordered and you can expect different order for multiple executions of the query.

The stream is not completed when it reaches the end of the currently used `persistenceIds`, but it continues to push new `persistenceIds` when new persistent actors are created. Corresponding query that is completed when it reaches the end of the currently used `persistenceIds` is provided by `currentPersistenceIds`.

The LevelDB write journal is notifying the query side as soon as new `persistenceIds` are created and there is no periodic polling or batching involved in this query.

The stream is completed with failure if there is a failure in executing the query in the backend journal.

EventsByTag and CurrentEventsByTag

`eventsByTag` is used for retrieving events that were marked with a given tag, e.g. all domain events of an `Aggregate Root` type.

```
LeveldbReadJournal queries =
    PersistenceQuery.get(system).getReadJournalFor(LeveldbReadJournal.class,
        LeveldbReadJournal.Identifier());

Source<EventEnvelope, BoxedUnit> source =
    queries.eventsByTag("green", 0);
```

To tag events you create an *Event Adapters* that wraps the events in a `akka.persistence.journal.Tagged` with the given tags.

```
public class MyTaggingEventAdapter implements WriteEventAdapter {

    @Override
    public Object toJournal(Object event) {
        if (event instanceof String) {
            String s = (String) event;
            Set<String> tags = new HashSet<String>();
            if (s.contains("green")) tags.add("green");
            if (s.contains("black")) tags.add("black");
            if (s.contains("blue")) tags.add("blue");
        }
    }
}
```

```

    if (tags.isEmpty())
        return event;
    else
        return new Tagged(event, tags);
} else {
    return event;
}
}

@Override
public String manifest(Object event) {
    return "";
}
}

```

You can retrieve a subset of all events by specifying `offset`, or use `0L` to retrieve all events with a given tag. The `offset` corresponds to an ordered sequence number for the specific tag. Note that the corresponding offset of each event is provided in the `EventEnvelope`, which makes it possible to resume the stream at a later point from a given offset.

In addition to the `offset` the `EventEnvelope` also provides `persistenceId` and `sequenceNr` for each event. The `sequenceNr` is the sequence number for the persistent actor with the `persistenceId` that persisted the event. The `persistenceId + sequenceNr` is a unique identifier for the event.

The returned event stream is ordered by the `offset` (tag sequence number), which corresponds to the same order as the write journal stored the events. The same stream elements (in same order) are returned for multiple executions of the query. Deleted events are not deleted from the tagged event stream.

Note: Events deleted using `deleteMessages(toSequenceNr)` are not deleted from the “tagged stream”.

The stream is not completed when it reaches the end of the currently stored events, but it continues to push new events when new events are persisted. Corresponding query that is completed when it reaches the end of the currently stored events is provided by `currentEventsByTag`.

The LevelDB write journal is notifying the query side as soon as tagged events are persisted, but for efficiency reasons the query side retrieves the events in batches that sometimes can be delayed up to the configured `refresh-interval` or given `RefreshInterval` hint.

The stream is completed with failure if there is a failure in executing the query in the backend journal.

3.11.4 Configuration

Configuration settings can be defined in the configuration section with the absolute path corresponding to the identifier, which is `"akka.persistence.query.journal.leveldb"` for the default `LeveldbReadJournal.Identifier`.

It can be configured with the following properties:

```

# Configuration for the LeveldbReadJournal
akka.persistence.query.journal.leveldb {
  # Implementation class of the LevelDB ReadJournalProvider
  class = "akka.persistence.query.journal.leveldb.LeveldbReadJournalProvider"

  # Absolute path to the write journal plugin configuration entry that this
  # query journal will connect to. That must be a LeveldbJournal or SharedLeveldbJournal.
  # If undefined (or "") it will connect to the default journal as specified by the
  # akka.persistence.journal.plugin property.
  write-plugin = ""

  # The LevelDB write journal is notifying the query side as soon as things
  # are persisted, but for efficiency reasons the query side retrieves the events

```

```
# in batches that sometimes can be delayed up to the configured 'refresh-interval'.
refresh-interval = 3s

# How many events to fetch in one query (replay) and keep buffered until they
# are delivered downstream.
max-buffer-size = 100
}
```

3.12 Testing Actor Systems

As with any piece of software, automated tests are a very important part of the development cycle. The actor model presents a different view on how units of code are delimited and how they interact, which has an influence on how to perform tests.

Note: Due to the conciseness of test DSLs available for Scala (`ScalaTest`, `Specs2`, `ScalaCheck`), it may be a good idea to write the test suite in that language even if the main project is written in Java. If that is not desirable, you can also use `TestKit` and friends from Java, albeit with more verbose syntax which is covered below. Munish Gupta has [published a nice post](#) showing several patterns you may find useful.

Akka comes with a dedicated module `akka-testkit` for supporting tests at different levels, which fall into two clearly distinct categories:

- Testing isolated pieces of code without involving the actor model, meaning without multiple threads; this implies completely deterministic behavior concerning the ordering of events and no concurrency concerns and will be called **Unit Testing** in the following.
- Testing (multiple) encapsulated actors including multi-threaded scheduling; this implies non-deterministic order of events but shielding from concurrency concerns by the actor model and will be called **Integration Testing** in the following.

There are of course variations on the granularity of tests in both categories, where unit testing reaches down to white-box tests and integration testing can encompass functional tests of complete actor networks. The important distinction lies in whether concurrency concerns are part of the test or not. The tools offered are described in detail in the following sections.

Note: Be sure to add the module `akka-testkit` to your dependencies.

3.12.1 Synchronous Unit Testing with `TestActorRef`

Testing the business logic inside `Actor` classes can be divided into two parts: first, each atomic operation must work in isolation, then sequences of incoming events must be processed correctly, even in the presence of some possible variability in the ordering of events. The former is the primary use case for single-threaded unit testing, while the latter can only be verified in integration tests.

Normally, the `ActorRef` shields the underlying `Actor` instance from the outside, the only communications channel is the actor's mailbox. This restriction is an impediment to unit testing, which led to the inception of the `TestActorRef`. This special type of reference is designed specifically for test purposes and allows access to the actor in two ways: either by obtaining a reference to the underlying actor instance, or by invoking or querying the actor's behaviour (`receive`). Each one warrants its own section below.

Note: It is highly recommended to stick to traditional behavioural testing (using messaging to ask the Actor to reply with the state you want to run assertions against), instead of using `TestActorRef` whenever possible.

Warning: Due to the synchronous nature of `TestActorRef` it will **not** work with some support traits that Akka provides as they require asynchronous behaviours to function properly. Examples of traits that do not mix well with test actor refs are *PersistentActor* and *AtLeastOnceDelivery* provided by *Akka Persistence*.

Obtaining a Reference to an Actor

Having access to the actual `Actor` object allows application of all traditional unit testing techniques on the contained methods. Obtaining a reference is done like this:

```
static class MyActor extends UntypedActor {
  public void onReceive(Object o) throws Exception {
    if (o.equals("say42")) {
      getSender().tell(42, getSelf());
    } else if (o instanceof Exception) {
      throw (Exception) o;
    }
  }
  public boolean testMe() { return true; }
}

@Test
public void demonstrateTestActorRef() {
  final Props props = Props.create(MyActor.class);
  final TestActorRef<MyActor> ref = TestActorRef.create(system, props, "testA");
  final MyActor actor = ref.underlyingActor();
  assertTrue(actor.testMe());
}
```

Since `TestActorRef` is generic in the actor type it returns the underlying actor with its proper static type. From this point on you may bring any unit testing tool to bear on your actor as usual.

Testing the Actor's Behavior

When the dispatcher invokes the processing behavior of an actor on a message, it actually calls `apply` on the current behavior registered for the actor. This starts out with the return value of the declared `receive` method, but it may also be changed using `become` and `unbecome` in response to external messages. All of this contributes to the overall actor behavior and it does not lend itself to easy testing on the `Actor` itself. Therefore the `TestActorRef` offers a different mode of operation to complement the `Actor` testing: it supports all operations also valid on normal `ActorRef`. Messages sent to the actor are processed synchronously on the current thread and answers may be sent back as usual. This trick is made possible by the `CallingThreadDispatcher` described below (see [CallingThreadDispatcher](#)); this dispatcher is set implicitly for any actor instantiated into a `TestActorRef`.

```
final Props props = Props.create(MyActor.class);
final TestActorRef<MyActor> ref = TestActorRef.create(system, props, "testB");
final Future<Object> future = akka.pattern.Patterns.ask(ref, "say42", 3000);
assertTrue(future.isCompleted());
assertEquals(42, Await.result(future, Duration.Zero()));
```

As the `TestActorRef` is a subclass of `LocalActorRef` with a few special extras, also aspects like supervision and restarting work properly, but beware that execution is only strictly synchronous as long as all actors involved use the `CallingThreadDispatcher`. As soon as you add elements which include more sophisticated scheduling you leave the realm of unit testing as you then need to think about asynchronicity again (in most cases the problem will be to wait until the desired effect had a chance to happen).

One more special aspect which is overridden for single-threaded tests is the `receiveTimeout`, as including that would entail asynchronous queuing of `ReceiveTimeout` messages, violating the synchronous contract.

Note: To summarize: `TestActorRef` overwrites two fields: it sets the dispatcher to `CallingThreadDispatcher.global` and it sets the `receiveTimeout` to `None`.

The Way In-Between: Expecting Exceptions

If you want to test the actor behavior, including hotswapping, but without involving a dispatcher and without having the `TestActorRef` swallow any thrown exceptions, then there is another mode available for you: just use the `receive` method on `TestActorRef`, which will be forwarded to the underlying actor:

```
final Props props = Props.create(MyActor.class);
final TestActorRef<MyActor> ref = TestActorRef.create(system, props, "myActor");
try {
  ref.receive(new Exception("expected"));
  fail("expected an exception to be thrown");
} catch (Exception e) {
  assertEquals("expected", e.getMessage());
}
```

Use Cases

You may of course mix and match both modi operandi of `TestActorRef` as suits your test needs:

- one common use case is setting up the actor into a specific internal state before sending the test message
- another is to verify correct internal state transitions after having sent the test message

Feel free to experiment with the possibilities, and if you find useful patterns, don't hesitate to let the Akka forums know about them! Who knows, common operations might even be worked into nice DSLs.

3.12.2 Asynchronous Integration Testing with `JavaTestKit`

When you are reasonably sure that your actor's business logic is correct, the next step is verifying that it works correctly within its intended environment. The definition of the environment depends of course very much on the problem at hand and the level at which you intend to test, ranging for functional/integration tests to full system tests. The minimal setup consists of the test procedure, which provides the desired stimuli, the actor under test, and an actor receiving replies. Bigger systems replace the actor under test with a network of actors, apply stimuli at varying injection points and arrange results to be sent from different emission points, but the basic principle stays the same in that a single procedure drives the test.

The `JavaTestKit` class contains a collection of tools which makes this common task easy.

```
import org.junit.AfterClass;
import org.junit.Assert;
import org.junit.BeforeClass;
import org.junit.Test;

import akka.actor.ActorRef;
import akka.actor.ActorSystem;
import akka.actor.Props;
import akka.actor.UntypedActor;
import akka.testkit.JavaTestKit;
import scala.concurrent.duration.Duration;

public class TestKitSampleTest {

  public static class SomeActor extends UntypedActor {
    ActorRef target = null;

    public void onReceive(Object msg) {
```

```

    if (msg.equals("hello")) {
        getSender().tell("world", getSelf());
        if (target != null) target.forward(msg, getContext());
    } else if (msg instanceof ActorRef) {
        target = (ActorRef) msg;
        getSender().tell("done", getSelf());
    }
}

static ActorSystem system;

@BeforeClass
public static void setup() {
    system = ActorSystem.create();
}

@AfterClass
public static void teardown() {
    JavaTestKit.shutdownActorSystem(system);
    system = null;
}

@Test
public void testIt() {
    /*
     * Wrap the whole test procedure within a testkit constructor
     * if you want to receive actor replies or use Within(), etc.
     */
    new JavaTestKit(system) {{
        final Props props = Props.create(SomeActor.class);
        final ActorRef subject = system.actorOf(props);

        // can also use JavaTestKit "from the outside"
        final JavaTestKit probe = new JavaTestKit(system);
        // "inject" the probe by passing it to the test subject
        // like a real resource would be passed in production
        subject.tell(probe.getRef(), getRef());
        // await the correct response
        expectMsgEquals(duration("1 second"), "done");

        // the run() method needs to finish within 3 seconds
        new Within(duration("3 seconds")) {
            protected void run() {

                subject.tell("hello", getRef());

                // This is a demo: would normally use expectMsgEquals().
                // Wait time is bounded by 3-second deadline above.
                new AwaitCond() {
                    protected boolean cond() {
                        return probe.msgAvailable();
                    }
                };
            };

            // response must have been enqueued to us before probe
            expectMsgEquals(Duration.Zero(), "world");
            // check that the probe we injected earlier got the msg
            probe.expectMsgEquals(Duration.Zero(), "hello");
            Assert.assertEquals(getRef(), probe.getLastSender());
        }
    }
}

```

```

        // Will wait for the rest of the 3 seconds
        expectNoMsg();
    }
};
}};
}
}

```

The `JavaTestKit` contains an actor named `testActor` which is the entry point for messages to be examined with the various `expectMsg...` assertions detailed below. The test actor's reference is obtained using the `getRef` method as demonstrated above. The `testActor` may also be passed to other actors as usual, usually subscribing it as notification listener. There is a whole set of examination methods, e.g. receiving all consecutive messages matching certain criteria, receiving a whole sequence of fixed messages or classes, receiving nothing for some time, etc.

The `ActorSystem` passed in to the constructor of `JavaTestKit` is accessible via the `getSystem` method.

Note: Remember to shut down the actor system after the test is finished (also in case of failure) so that all actors—including the test actor—are stopped.

Built-In Assertions

The above mentioned `expectMsgEquals` is not the only method for formulating assertions concerning received messages, the full set is this:

```

final String hello = expectMsgEquals("hello");
final Object any = expectMsgAnyOf("hello", "world");
final Object[] all = expectMsgAllOf("hello", "world");
final int i = expectMsgClass(Integer.class);
final Number j = expectMsgAnyClassOf(Integer.class, Long.class);
expectNoMsg();
final Object[] two = receiveN(2);

```

In these examples, the maximum durations you will find mentioned below are left out, in which case they use the default value from configuration item `akka.test.single-expect-default` which itself defaults to 3 seconds (or they obey the innermost enclosing `Within` as detailed *below*). The full signatures are:

- `public <T> T expectMsgEquals(Duration max, T msg)`
The given message object must be received within the specified time; the object will be returned.
- `public Object expectMsgAnyOf(Duration max, Object... msg)`
An object must be received within the given time, and it must be equal (compared with `equals()`) to at least one of the passed reference objects; the received object will be returned.
- `public Object[] expectMsgAllOf(Duration max, Object... msg)`
A number of objects matching the size of the supplied object array must be received within the given time, and for each of the given objects there must exist at least one among the received ones which equals it (compared with `equals()`). The full sequence of received objects is returned in the order received.
- `public <T> T expectMsgClass(Duration max, Class<T> c)`
An object which is an instance of the given `Class` must be received within the allotted time frame; the object will be returned. Note that this does a conformance check, if you need the class to be equal you need to verify that afterwards.
- `public <T> T expectMsgAnyClassOf(Duration max, Class<? extends T>... c)`

An object must be received within the given time, and it must be an instance of at least one of the supplied Class objects; the received object will be returned. Note that this does a conformance check, if you need the class to be equal you need to verify that afterwards.

Note: Because of a limitation in Java's type system it may be necessary to add `@SuppressWarnings("unchecked")` when using this method.

- `public void expectNoMsg(Duration max)`

No message must be received within the given time. This also fails if a message has been received before calling this method which has not been removed from the queue using one of the other methods.

- `Object[] receiveN(int n, Duration max)`

`n` messages must be received within the given time; the received messages are returned.

For cases which require more refined conditions there are constructs which take code blocks:

- **ExpectMsg<T>**

```
new JavaTestKit(system) {{
  getRef().tell(42, ActorRef.noSender());
  final String out = new ExpectMsg<String>("match hint") {
    // do not put code outside this method, will run afterwards
    protected String match(Object in) {
      if (in instanceof Integer) {
        return "match";
      } else {
        throw noMatch();
      }
    }
  }.get(); // this extracts the received message
  assertEquals("match", out);
}};
```

The `match(Object in)` method will be evaluated once a message has been received within the allotted time (which may be given as constructor argument). If it throws `noMatch()` (where it is sufficient to call that method; the `throw` keyword is only needed in cases where the compiler would otherwise complain about wrong return types—Java is lacking Scala's notion of a type which signifies “will not ever return normally”), then the expectation fails with an `AssertionError`, otherwise the matched and possibly transformed object is stored for retrieval using the `get` method.

- **ReceiveWhile<T>**

```
new JavaTestKit(system) {{
  getRef().tell(42, ActorRef.noSender());
  getRef().tell(43, ActorRef.noSender());
  getRef().tell("hello", ActorRef.noSender());
  final String[] out =
    new ReceiveWhile<String>(String.class, duration("1 second")) {
      // do not put code outside this method, will run afterwards
      protected String match(Object in) {
        if (in instanceof Integer) {
          return in.toString();
        } else {
          throw noMatch();
        }
      }
    }.get(); // this extracts the received messages
  assertEquals(new String[] {"42", "43"}, out);
  expectMsgEquals("hello");
}};
```

This construct works like `ExpectMsg`, but it continually collects messages as long as they match the criteria,

and it does not fail when a non-matching one is encountered. Collecting messages also ends when the time is up, when too much time passes between messages or when enough messages have been received.

```
new ReceiveWhile<String>( // type of array to be created must match ...
    String.class, // ... this class which is needed to that end
    duration("100 millis"), // maximum collect time
    duration("50 millis"), // maximum time between messages
    12 // maximum number of messages to collect
) {
    // match elided ...
};
```

The need to specify the `String` result type twice results from the need to create a correctly typed array and Java's inability to infer the class's type argument.

- **AwaitCond**

```
new JavaTestKit(system) {{
    getRef().tell(42, ActorRef.noSender());
    new AwaitCond(
        duration("1 second"), // maximum wait time
        duration("100 millis") // interval at which to check the condition
    ) {
        // do not put code outside this method, will run afterwards
    protected boolean cond() {
        // typically used to wait for something to start up
        return msgAvailable();
    }
};
}};
```

This general construct is not connected with the test kit's message reception, the embedded condition can compute the boolean result from anything in scope.

- **AwaitAssert**

```
new JavaTestKit(system) {{
    getRef().tell(42, ActorRef.noSender());
    new AwaitAssert(
        duration("1 second"), // maximum wait time
        duration("100 millis") // interval at which to check the condition
    ) {
        // do not put code outside this method, will run afterwards
    protected void check() {
        assertEquals(msgAvailable(), true);
    }
};
}};
```

This general construct is not connected with the test kit's message reception, the embedded assert can check anything in scope.

There are also cases where not all messages sent to the test kit are actually relevant to the test, but removing them would mean altering the actors under test. For this purpose it is possible to ignore certain messages:

- **IgnoreMsg**

```
new JavaTestKit(system) {{
    // ignore all Strings
    new IgnoreMsg() {
        protected boolean ignore(Object msg) {
            return msg instanceof String;
        }
};
    getRef().tell("hello", ActorRef.noSender());
```

```

getRef().tell(42, ActorRef.noSender());
expectMsgEquals(42);
// remove message filter
ignoreNoMsg();
getRef().tell("hello", ActorRef.noSender());
expectMsgEquals("hello");
});

```

Expecting Log Messages

Since an integration test does not allow to the internal processing of the participating actors, verifying expected exceptions cannot be done directly. Instead, use the logging system for this purpose: replacing the normal event handler with the `TestEventListener` and using an `EventFilter` allows assertions on log messages, including those which are generated by exceptions:

```

new JavaTestKit(system) {{
  assertEquals("TestKitDocTest", system.name());
  final ActorRef victim = system.actorOf(Props.empty(), "victim");

  final int result = new EventFilter<Integer>(ActorKilledException.class) {
    protected Integer run() {
      victim.tell(Kill.getInstance(), ActorRef.noSender());
      return 42;
    }
  }.from("akka://TestKitDocTest/user/victim").occurrences(1).exec();

  assertEquals(42, result);
}};

```

If a number of occurrences is specific—as demonstrated above—then `exec()` will block until that number of matching messages have been received or the timeout configured in `akka.test.filter-leeway` is used up (time starts counting after the `run()` method returns). In case of a timeout the test fails.

Note: Be sure to exchange the default logger with the `TestEventListener` in your `application.conf` to enable this function:

```
akka.loggers = [akka.testkit.TestEventListener]
```

Timing Assertions

Another important part of functional testing concerns timing: certain events must not happen immediately (like a timer), others need to happen before a deadline. Therefore, all examination methods accept an upper time limit within the positive or negative result must be obtained. Lower time limits need to be checked external to the examination, which is facilitated by a new construct for managing time constraints:

```

new JavaTestKit(system) {{
  getRef().tell(42, ActorRef.noSender());
  new Within(Duration.Zero(), Duration.create(1, "second")) {
    // do not put code outside this method, will run afterwards
    public void run() {
      assertEquals((Integer) 42, expectMsgClass(Integer.class));
    }
  };
}};

```

The block in `Within.run` must complete after a *Duration* which is between `min` and `max`, where the former defaults to zero. The deadline calculated by adding the `max` parameter to the block's start time is implicitly available within the block to all examination methods, if you do not specify it, it is inherited from the innermost enclosing `within` block.

It should be noted that if the last message-receiving assertion of the block is `expectNoMsg` or `receiveWhile`, the final check of the `within` is skipped in order to avoid false positives due to wake-up latencies. This means that while individual contained assertions still use the maximum time bound, the overall block may take arbitrarily longer in this case.

Note: All times are measured using `System.nanoTime`, meaning that they describe wall time, not CPU time or system time.

Accounting for Slow Test Systems

The tight timeouts you use during testing on your lightning-fast notebook will invariably lead to spurious test failures on the heavily loaded Jenkins server (or similar). To account for this situation, all maximum durations are internally scaled by a factor taken from the *Configuration*, `akka.test.timefactor`, which defaults to 1.

You can scale other durations with the same factor by using `dilated` method in `JavaTestKit`.

```
new JavaTestKit(system) {{
  final Duration original = duration("1 second");
  final Duration stretched = dilated(original);
  assertTrue("dilated", stretched.gteq(original));
}};
```

Using Multiple Probe Actors

When the actors under test are supposed to send various messages to different destinations, it may be difficult distinguishing the message streams arriving at the `testActor` when using the `JavaTestKit` as shown until now. Another approach is to use it for creation of simple probe actors to be inserted in the message flows. The functionality is best explained using a small example:

```
new JavaTestKit(system) {{
  // simple actor which just forwards messages
  class Forwarder extends UntypedActor {
    final ActorRef target;
    @SuppressWarnings("unused")
    public Forwarder(ActorRef target) {
      this.target = target;
    }
    public void onReceive(Object msg) {
      target.forward(msg, getContext());
    }
  }

  // create a test probe
  final JavaTestKit probe = new JavaTestKit(system);

  // create a forwarder, injecting the probe's testActor
  final Props props = Props.create(Forwarder.class, this, probe.getRef());
  final ActorRef forwarder = system.actorOf(props, "forwarder");

  // verify correct forwarding
  forwarder.tell(42, getRef());
  probe.expectMsgEquals(42);
  assertEquals(getRef(), probe.getLastSender());
}};
```

This simple test verifies an equally simple `Forwarder` actor by injecting a probe as the forwarder's target. Another example would be two actors A and B which collaborate by A sending messages to B. In order to verify this message flow, a `TestProbe` could be inserted as target of A, using the forwarding capabilities or auto-pilot described below to include a real B in the test setup.

If you have many test probes, you can name them to get meaningful actor names in test logs and assertions:

```
new JavaTestKit(system) {{
  final TestProbe worker = new TestProbe(system, "worker");
  final TestProbe aggregator = new TestProbe(system, "aggregator");

  assertTrue(worker.ref().path().name().startsWith("worker"));
  assertTrue(aggregator.ref().path().name().startsWith("aggregator"));
}};
```

Probes may also be equipped with custom assertions to make your test code even more concise and clear:

```
new JavaTestKit(system) {{
  class MyProbe extends JavaTestKit {
    public MyProbe() {
      super(system);
    }
    public void assertHello() {
      expectMsgEquals("hello");
    }
  }

  final MyProbe probe = new MyProbe();
  probe.getRef().tell("hello", ActorRef.noSender());
  probe.assertHello();
}};
```

You have complete flexibility here in mixing and matching the `JavaTestKit` facilities with your own checks and choosing an intuitive name for it. In real life your code will probably be a bit more complicated than the example given above; just use the power!

Warning: Any message send from a `TestProbe` to another actor which runs on the `CallingThreadDispatcher` runs the risk of dead-lock, if that other actor might also send to this probe. The implementation of `TestProbe.watch` and `TestProbe.unwatch` will also send a message to the watchee, which means that it is dangerous to try watching e.g. `TestActorRef` from a `TestProbe`.

Watching Other Actors from Probes

A `JavaTestKit` can register itself for `DeathWatch` of any other actor:

```
new JavaTestKit(system) {{
  final JavaTestKit probe = new JavaTestKit(system);
  probe.watch(target);
  target.tell(PoisonPill.getInstance(), ActorRef.noSender());
  final Terminated msg = probe.expectMsgClass(Terminated.class);
  assertEquals(msg.getActor(), target);
}};
```

Replying to Messages Received by Probes

The probe stores the sender of the last dequeued message (i.e. after its `expectMsg*` reception), which may be retrieved using the `getLastSender` method. This information can also implicitly be used for having the probe reply to the last received message:

```
new JavaTestKit(system) {{
  final JavaTestKit probe = new JavaTestKit(system);
  probe.getRef().tell("hello", getRef());
  probe.expectMsgEquals("hello");
  probe.reply("world");
  expectMsgEquals("world");
}};
```

```
    assertEquals(probe.getRef(), getLastSender());
  });
```

Forwarding Messages Received by Probes

The probe can also forward a received message (i.e. after its `expectMsg*` reception), retaining the original sender:

```
new JavaTestKit(system) {{
  final JavaTestKit probe = new JavaTestKit(system);
  probe.getRef().tell("hello", getRef());
  probe.expectMsgEquals("hello");
  probe.forward(getRef());
  expectMsgEquals("hello");
  assertEquals(getRef(), getLastSender());
}};
```

Auto-Pilot

Receiving messages in a queue for later inspection is nice, but in order to keep a test running and verify traces later you can also install an `AutoPilot` in the participating test probes (actually in any `TestKit`) which is invoked before enqueueing to the inspection queue. This code can be used to forward messages, e.g. in a chain `A --> Probe --> B`, as long as a certain protocol is obeyed.

```
new JavaTestKit(system) {{
  final JavaTestKit probe = new JavaTestKit(system);
  // install auto-pilot
  probe.setAutoPilot(new TestActor.AutoPilot() {
    public AutoPilot run(ActorRef sender, Object msg) {
      sender.tell(msg, ActorRef.noSender());
      return noAutoPilot();
    }
  });
  // first one is replied to directly ...
  probe.getRef().tell("hello", getRef());
  expectMsgEquals("hello");
  // ... but then the auto-pilot switched itself off
  probe.getRef().tell("world", getRef());
  expectNoMsg();
}};
```

The `run` method must return the auto-pilot for the next message, wrapped in an `Option`; setting it to `None` terminates the auto-pilot.

Caution about Timing Assertions

The behavior of `within` blocks when using test probes might be perceived as counter-intuitive: you need to remember that the nicely scoped deadline as described *above* is local to each probe. Hence, probes do not react to each other's deadlines or to the deadline set in an enclosing `JavaTestKit` instance:

```
new JavaTestKit(system) {{
  final JavaTestKit probe = new JavaTestKit(system);
  new Within(duration("1 second")) {
    public void run() {
      probe.expectMsgEquals("hello");
    }
  };
}};
```

Here, the `expectMsgEquals` call will use the default timeout.

Testing parent-child relationships

The parent of an actor is always the actor that created it. At times this leads to a coupling between the two that may not be straightforward to test. Broadly, there are three approaches to improve testability of parent-child relationships:

1. when creating a child, pass an explicit reference to its parent
2. when creating a parent, tell the parent how to create its child
3. create a fabricated parent when testing

For example, the structure of the code you want to test may follow this pattern:

```
static class Parent extends UntypedActor {
    final ActorRef child = context().actorOf(Props.create(Child.class), "child");
    boolean ponged = false;

    @Override public void onReceive(Object message) throws Exception {
        if ("pingit".equals(message)) {
            child.tell("ping", self());
        } else if ("pong".equals(message)) {
            ponged = true;
        } else {
            unhandled(message);
        }
    }
}

static class Child extends UntypedActor {
    @Override public void onReceive(Object message) throws Exception {
        if ("ping".equals(message)) {
            context().parent().tell("pong", self());
        } else {
            unhandled(message);
        }
    }
}
```

Using dependency-injection

The first option is to avoid use of the `context.parent` function and create a child with a custom parent by passing an explicit reference to its parent instead.

```
class DependentChild extends UntypedActor {
    private final ActorRef parent;

    public DependentChild(ActorRef parent) {
        this.parent = parent;
    }

    @Override public void onReceive(Object message) throws Exception {
        if ("ping".equals(message)) {
            parent.tell("pong", self());
        } else {
            unhandled(message);
        }
    }
}
```

Alternatively, you can tell the parent how to create its child. There are two ways to do this: by giving it a Props object or by giving it a function which takes care of creating the child actor:

```
class DependentParent extends UntypedActor {
  final ActorRef child;
  boolean ponged = false;

  public DependentParent(Props childProps) {
    child = context().actorOf(childProps, "child");
  }

  @Override public void onReceive(Object message) throws Exception {
    if ("pingit".equals(message)) {
      child.tell("ping", self());
    } else if ("pong".equals(message)) {
      ponged = true;
    } else {
      unhandled(message);
    }
  }
}
```

```
class GenericDependentParent extends UntypedActor {
  final ActorRef child;
  boolean ponged = false;

  public GenericDependentParent(Function<ActorRefFactory, ActorRef> childMaker)
    throws Exception {
    child = childMaker.apply(context());
  }

  @Override public void onReceive(Object message) throws Exception {
    if ("pingit".equals(message)) {
      child.tell("ping", self());
    } else if ("pong".equals(message)) {
      ponged = true;
    } else {
      unhandled(message);
    }
  }
}
```

Creating the Actor is straightforward and the function may look like this in your test code:

```
Function<ActorRefFactory, ActorRef> maker = new Function<ActorRefFactory, ActorRef>() {
  @Override public ActorRef apply(ActorRefFactory param) throws Exception {
    return probe.ref();
  }
};
ActorRef parent = system.actorOf(Props.create(GenericDependentParent.class, maker));
```

And like this in your application code:

```
Function<ActorRefFactory, ActorRef> maker = new Function<ActorRefFactory, ActorRef>() {
  @Override public ActorRef apply(ActorRefFactory f) throws Exception {
    return f.actorOf(Props.create(Child.class));
  }
};
ActorRef parent = system.actorOf(Props.create(GenericDependentParent.class, maker));
```

Using a fabricated parent

If you prefer to avoid modifying the parent or child constructor you can create a fabricated parent in your test. This, however, does not enable you to test the parent actor in isolation.

```
class FabricatedParentCreator implements Creator<Actor> {
    private final TestProbe proxy;

    public FabricatedParentCreator(TestProbe proxy) {
        this.proxy = proxy;
    }

    @Override public Actor create() throws Exception {
        return new UntypedActor() {
            final ActorRef child = context().actorOf(Props.create(Child.class), "child");

            @Override public void onReceive(Object x) throws Exception {
                if (sender().equals(child)) {
                    proxy.ref().forward(x, context());
                } else {
                    child.forward(x, context());
                }
            }
        };
    }
}
```

```
TestProbe proxy = new TestProbe(system);
ActorRef parent = system.actorOf(Props.create(new FabricatedParentCreator(proxy)));

proxy.send(parent, "ping");
proxy.expectMsg("pong");
```

Which of these methods is the best depends on what is most important to test. The most generic option is to create the parent actor by passing it a function that is responsible for the Actor creation, but the fabricated parent is often sufficient.

3.12.3 CallingThreadDispatcher

The `CallingThreadDispatcher` serves good purposes in unit testing, as described above, but originally it was conceived in order to allow contiguous stack traces to be generated in case of an error. As this special dispatcher runs everything which would normally be queued directly on the current thread, the full history of a message's processing chain is recorded on the call stack, so long as all intervening actors run on this dispatcher.

How to use it

Just set the dispatcher as you normally would:

```
system.actorOf(
    Props.create(MyActor.class)
        .withDispatcher(CallingThreadDispatcher.Id()));
```

How it works

When receiving an invocation, the `CallingThreadDispatcher` checks whether the receiving actor is already active on the current thread. The simplest example for this situation is an actor which sends a message to itself. In this case, processing cannot continue immediately as that would violate the actor model, so the invocation is queued and will be processed when the active invocation on that actor finishes its processing; thus, it will be processed on the calling thread, but simply after the actor finishes its previous work. In the other case, the

invocation is simply processed immediately on the current thread. Futures scheduled via this dispatcher are also executed immediately.

This scheme makes the `CallingThreadDispatcher` work like a general purpose dispatcher for any actors which never block on external events.

In the presence of multiple threads it may happen that two invocations of an actor running on this dispatcher happen on two different threads at the same time. In this case, both will be processed directly on their respective threads, where both compete for the actor's lock and the loser has to wait. Thus, the actor model is left intact, but the price is loss of concurrency due to limited scheduling. In a sense this is equivalent to traditional mutex style concurrency.

The other remaining difficulty is correct handling of suspend and resume: when an actor is suspended, subsequent invocations will be queued in thread-local queues (the same ones used for queuing in the normal case). The call to `resume`, however, is done by one specific thread, and all other threads in the system will probably not be executing this specific actor, which leads to the problem that the thread-local queues cannot be emptied by their native threads. Hence, the thread calling `resume` will collect all currently queued invocations from all threads into its own queue and process them.

Limitations

Warning: In case the `CallingThreadDispatcher` is used for top-level actors, but without going through `TestActorRef`, then there is a time window during which the actor is awaiting construction by the user guardian actor. Sending messages to the actor during this time period will result in them being enqueued and then executed on the guardian's thread instead of the caller's thread. To avoid this, use `TestActorRef`.

If an actor's behavior blocks on something which would normally be affected by the calling actor after having sent the message, this will obviously dead-lock when using this dispatcher. This is a common scenario in actor tests based on `CountDownLatch` for synchronization:

```
val latch = new CountDownLatch(1)
actor ! startWorkAfter(latch) // actor will call latch.await() before proceeding
doSomeSetupStuff()
latch.countDown()
```

The example would hang indefinitely within the message processing initiated on the second line and never reach the fourth line, which would unblock it on a normal dispatcher.

Thus, keep in mind that the `CallingThreadDispatcher` is not a general-purpose replacement for the normal dispatchers. On the other hand it may be quite useful to run your actor network on it for testing, because if it runs without dead-locking chances are very high that it will not dead-lock in production.

Warning: The above sentence is unfortunately not a strong guarantee, because your code might directly or indirectly change its behavior when running on a different dispatcher. If you are looking for a tool to help you debug dead-locks, the `CallingThreadDispatcher` may help with certain error scenarios, but keep in mind that it has may give false negatives as well as false positives.

Thread Interruptions

If the `CallingThreadDispatcher` sees that the current thread has its `isInterrupted()` flag set when message processing returns, it will throw an `InterruptedException` after finishing all its processing (i.e. all messages which need processing as described above are processed before this happens). As `tell` cannot throw exceptions due to its contract, this exception will then be caught and logged, and the thread's interrupted status will be set again.

If during message processing an `InterruptedException` is thrown then it will be caught inside the `CallingThreadDispatcher`'s message handling loop, the thread's interrupted flag will be set and processing continues normally.

Note: The summary of these two paragraphs is that if the current thread is interrupted while doing work under the `CallingThreadDispatcher`, then that will result in the `isInterrupted` flag to be `true` when the message send returns and no `InterruptedException` will be thrown.

Benefits

To summarize, these are the features with the `CallingThreadDispatcher` has to offer:

- Deterministic execution of single-threaded tests while retaining nearly full actor semantics
- Full message processing history leading up to the point of failure in exception stack traces
- Exclusion of certain classes of dead-lock scenarios

3.12.4 Tracing Actor Invocations

The testing facilities described up to this point were aiming at formulating assertions about a system's behavior. If a test fails, it is usually your job to find the cause, fix it and verify the test again. This process is supported by debuggers as well as logging, where the Akka toolkit offers the following options:

- *Logging of exceptions thrown within Actor instances*

This is always on; in contrast to the other logging mechanisms, this logs at `ERROR` level.

- *Logging of special messages*

Actors handle certain special messages automatically, e.g. `Kill`, `PoisonPill`, etc. Tracing of these message invocations is enabled by the setting `akka.actor.debug.autoreceive`, which enables this on all actors.

- *Logging of the actor lifecycle*

Actor creation, start, restart, monitor start, monitor stop and stop may be traced by enabling the setting `akka.actor.debug.lifecycle`; this, too, is enabled uniformly on all actors.

All these messages are logged at `DEBUG` level. To summarize, you can enable full logging of actor activities using this configuration fragment:

```
akka {
  loglevel = "DEBUG"
  actor {
    debug {
      autoreceive = on
      lifecycle = on
    }
  }
}
```

3.12.5 Configuration

There are several configuration properties for the `TestKit` module, please refer to the [reference configuration](#).

ACTORS (JAVA WITH LAMBDA SUPPORT)

4.1 Actors (Java with Lambda Support)

The **Actor Model** provides a higher level of abstraction for writing concurrent and distributed systems. It alleviates the developer from having to deal with explicit locking and thread management, making it easier to write correct concurrent and parallel systems. Actors were defined in the 1973 paper by Carl Hewitt but have been popularized by the Erlang language, and used for example at Ericsson with great success to build highly concurrent and reliable telecom systems.

The API of Akka's Actors is similar to Scala Actors which has borrowed some of its syntax from Erlang.

Warning: The Java with lambda support part of Akka is marked as “**experimental**” as of its introduction in Akka 2.3.0. We will continue to improve this API based on our users' feedback, which implies that while we try to keep incompatible changes to a minimum, but the binary compatibility guarantee for maintenance releases does not apply to the `akka.actor.AbstractActor`, related classes and the `akka.japi.pf` package.

4.1.1 Creating Actors

Note: Since Akka enforces parental supervision every actor is supervised and (potentially) the supervisor of its children, it is advisable that you familiarize yourself with *Actor Systems* and *Supervision and Monitoring* and it may also help to read *Actor References, Paths and Addresses*.

Defining an Actor class

Actor classes are implemented by extending the `AbstractActor` class and setting the “initial behavior” in the constructor by calling the `receive` method in the `AbstractActor`.

The argument to the `receive` method is a `PartialFunction<Object, BoxedUnit>` that defines which messages your Actor can handle, along with the implementation of how the messages should be processed.

Don't let the type signature scare you. To allow you to easily build up a partial function there is a builder named `ReceiveBuilder` that you can use.

Here is an example:

```
import akka.actor.AbstractActor;
import akka.event.Logging;
import akka.event.LoggingAdapter;
import akka.japi.pf.ReceiveBuilder;
```

```
public class MyActor extends AbstractActor {
    private final LoggingAdapter log = Logging.getLogger(context().system(), this);

    public MyActor() {
        receive(ReceiveBuilder.
            match(String.class, s -> {
                log.info("Received String message: {}", s);
            }).
            matchAny(o -> log.info("received unknown message")).build()
        );
    }
}
```

Please note that the Akka Actor `receive` message loop is exhaustive, which is different compared to Erlang and the late Scala Actors. This means that you need to provide a pattern match for all messages that it can accept and if you want to be able to handle unknown messages then you need to have a default case as in the example above. Otherwise an `akka.actor.UnhandledMessage(message, sender, recipient)` will be published to the ActorSystem's EventStream.

Note further that the return type of the behavior defined above is `Unit`; if the actor shall reply to the received message then this must be done explicitly as explained below.

The argument to the `receive` method is a partial function object, which is stored within the actor as its “initial behavior”, see [Become/Unbecome](#) for further information on changing the behavior of an actor after its construction.

Props

`Props` is a configuration class to specify options for the creation of actors, think of it as an immutable and thus freely shareable recipe for creating an actor including associated deployment information (e.g. which dispatcher to use, see more below). Here are some examples of how to create a `Props` instance.

```
import akka.actor.Props;
```

```
Props props1 = Props.create(MyActor.class);
Props props2 = Props.create(ActorWithArgs.class,
    () -> new ActorWithArgs("arg")); // careful, see below
Props props3 = Props.create(ActorWithArgs.class, "arg");
```

The second variant shows how to pass constructor arguments to the Actor being created, but it should only be used outside of actors as explained below.

The last line shows a possibility to pass constructor arguments regardless of the context it is being used in. The presence of a matching constructor is verified during construction of the `Props` object, resulting in an `IllegalArgumentException` if no or multiple matching constructors are found.

Dangerous Variants

```
// NOT RECOMMENDED within another actor:
// encourages to close over enclosing class
Props props7 = Props.create(ActorWithArgs.class,
    () -> new ActorWithArgs("arg"));
```

This method is not recommended to be used within another actor because it encourages to close over the enclosing scope, resulting in non-serializable `Props` and possibly race conditions (breaking the actor encapsulation). On the other hand using this variant in a `Props` factory in the actor's companion object as documented under “Recommended Practices” below is completely fine.

There were two use-cases for these methods: passing constructor arguments to the actor—which is solved by the newly introduced `Props.create(clazz, args)` method above or the recommended practice below—and

creating actors “on the spot” as anonymous classes. The latter should be solved by making these actors named classes instead (if they are not declared within a top-level `object` then the enclosing instance’s `this` reference needs to be passed as the first argument).

Warning: Declaring one actor within another is very dangerous and breaks actor encapsulation. Never pass an actor’s `this` reference into `Props`!

Recommended Practices

It is a good idea to provide factory methods on the companion object of each `Actor` which help keeping the creation of suitable `Props` as close to the actor definition as possible. This also avoids the pitfalls associated with using the `Props.create(...)` method which takes a by-name argument, since within a companion object the given code block will not retain a reference to its enclosing scope:

```
public class DemoActor extends AbstractActor {
    /**
     * Create Props for an actor of this type.
     * @param magicNumber The magic number to be passed to this actor's constructor.
     * @return a Props for creating this actor, which can then be further configured
     *         (e.g. calling .withDispatcher() on it)
     */
    static Props props(Integer magicNumber) {
        // You need to specify the actual type of the returned actor
        // since Java 8 lambdas have some runtime type information erased
        return Props.create(DemoActor.class, () -> new DemoActor(magicNumber));
    }

    private final Integer magicNumber;

    DemoActor(Integer magicNumber) {
        this.magicNumber = magicNumber;
        receive(ReceiveBuilder.
            match(Integer.class, i -> {
                sender().tell(i + magicNumber, self());
            }).build()
        );
    }
}

public class SomeOtherActor extends AbstractActor {
    // Props(new DemoActor(42)) would not be safe
    ActorRef demoActor = context().actorOf(DemoActor.props(42), "demo");
    // ...
}
```

Another good practice is to declare what messages an Actor can receive as close to the actor definition as possible (e.g. as static classes inside the Actor or using other suitable class), which makes it easier to know what it can receive.

```
public class DemoMessagesActor extends AbstractLoggingActor {

    static public class Greeting {
        private final String from;

        public Greeting(String from) {
            this.from = from;
        }

        public String getGreeter() {
            return from;
        }
    }
}
```

```

}

DemoMessagesActor() {
  receive(ReceiveBuilder.
    match(Greeting.class, g -> {
      log().info("I was greeted by {}", g.getGreeter());
    }).build()
  );
};
}

```

Creating Actors with Props

Actors are created by passing a `Props` instance into the `actorOf` factory method which is available on `ActorSystem` and `ActorContext`.

```

import akka.actor.ActorRef;
import akka.actor.ActorSystem;

// ActorSystem is a heavy object: create only one per application
final ActorSystem system = ActorSystem.create("MySystem", config);
final ActorRef myActor = system.actorOf(Props.create(MyActor.class), "myactor");

```

Using the `ActorSystem` will create top-level actors, supervised by the actor system's provided guardian actor, while using an actor's context will create a child actor.

```

public class FirstActor extends AbstractActor {
  final ActorRef child = context().actorOf(Props.create(MyActor.class), "myChild");
  // plus some behavior ...
}

```

It is recommended to create a hierarchy of children, grand-children and so on such that it fits the logical failure-handling structure of the application, see [Actor Systems](#).

The call to `actorOf` returns an instance of `ActorRef`. This is a handle to the actor instance and the only way to interact with it. The `ActorRef` is immutable and has a one to one relationship with the Actor it represents. The `ActorRef` is also serializable and network-aware. This means that you can serialize it, send it over the wire and use it on a remote host and it will still be representing the same Actor on the original node, across the network.

The name parameter is optional, but you should preferably name your actors, since that is used in log messages and for identifying actors. The name must not be empty or start with `$`, but it may contain URL encoded characters (eg. `%20` for a blank space). If the given name is already in use by another child to the same parent an `InvalidActorNameException` is thrown.

Actors are automatically started asynchronously when created.

Dependency Injection

If your `UntypedActor` has a constructor that takes parameters then those need to be part of the `Props` as well, as described above. But there are cases when a factory method must be used, for example when the actual constructor arguments are determined by a dependency injection framework.

```

import akka.actor.Actor;
import akka.actor.IndirectActorProducer;

class DependencyInjector implements IndirectActorProducer {
  final Object applicationContext;
  final String beanName;

  public DependencyInjector(Object applicationContext, String beanName) {
    this.applicationContext = applicationContext;
  }
}

```

```

    this.beanName = beanName;
  }

  @Override
  public Class<? extends Actor> actorClass() {
    return MyActor.class;
  }

  @Override
  public MyActor produce() {
    MyActor result;
    // obtain fresh Actor instance from DI framework ...
    return result;
  }
}

final ActorRef myActor = getContext().actorOf(
    Props.create(DependencyInjector.class, applicationContext, "MyActor"),
    "myactor3");

```

Warning: You might be tempted at times to offer an `IndirectActorProducer` which always returns the same instance, e.g. by using a static field. This is not supported, as it goes against the meaning of an actor restart, which is described here: [What Restarting Means](#).
When using a dependency injection framework, actor beans *MUST NOT* have singleton scope.

Techniques for dependency injection and integration with dependency injection frameworks are described in more depth in the [Using Akka with Dependency Injection](#) guideline and the [Akka Java Spring](#) tutorial in `Typesafe Activator`.

The Inbox

When writing code outside of actors which shall communicate with actors, the `ask` pattern can be a solution (see below), but there are two things it cannot do: receiving multiple replies (e.g. by subscribing an `ActorRef` to a notification service) and watching other actors' lifecycle. For these purposes there is the `Inbox` class:

```

final Inbox inbox = Inbox.create(system);
inbox.send(target, "hello");
try {
  assert inbox.receive(Duration.create(1, TimeUnit.SECONDS)).equals("world");
} catch (java.util.concurrent.TimeoutException e) {
  // timeout
}

```

The `send` method wraps a normal `tell` and supplies the internal actor's reference as the sender. This allows the reply to be received on the last line. Watching an actor is quite simple as well:

```

final Inbox inbox = Inbox.create(system);
inbox.watch(target);
target.tell(PoisonPill.getInstance(), ActorRef.noSender());
try {
  assert inbox.receive(Duration.create(1, TimeUnit.SECONDS)) instanceof Terminated;
} catch (java.util.concurrent.TimeoutException e) {
  // timeout
}

```

4.1.2 Actor API

The `AbstractActor` class defines a method called `receive`, that is used to set the "initial behavior" of the actor.

If the current actor behavior does not match a received message, `unhandled` is called, which by default publishes an `akka.actor.UnhandledMessage(message, sender, recipient)` on the actor system's event stream (set configuration item `akka.actor.debug.unhandled` to `on` to have them converted into actual Debug messages).

In addition, it offers:

- `self` reference to the `ActorRef` of the actor
- `sender` reference `sender Actor` of the last received message, typically used as described in [Reply to messages](#)
- `supervisorStrategy` user overridable definition the strategy to use for supervising child actors

This strategy is typically declared inside the actor in order to have access to the actor's internal state within the decider function: since failure is communicated as a message sent to the supervisor and processed like other messages (albeit outside of the normal behavior), all values and variables within the actor are available, as is the `sender` reference (which will be the immediate child reporting the failure; if the original failure occurred within a distant descendant it is still reported one level up at a time).

- `context` exposes contextual information for the actor and the current message, such as:
 - factory methods to create child actors (`actorOf`)
 - system that the actor belongs to
 - parent supervisor
 - supervised children
 - lifecycle monitoring
 - hotswap behavior stack as described in [Become/Unbecome](#)

The remaining visible methods are user-overridable life-cycle hooks which are described in the following:

```
public void preStart() {
}


public void preRestart(Throwable reason, scala.Option<Object> message) {
    for (ActorRef each : getContext().getChildren()) {
        getContext().unwatch(each);
        getContext().stop(each);
    }
    postStop();
}

public void postRestart(Throwable reason) {
    preStart();
}

public void postStop() {
}
```

The implementations shown above are the defaults provided by the `AbstractActor` class.

Actor Lifecycle



A path in an actor system represents a “place” which might be occupied by a living actor. Initially (apart from system initialized actors) a path is empty. When `actorOf()` is called it assigns an *incarnation* of the actor described by the passed `Props` to the given path. An actor incarnation is identified by the path *and a UID*. A restart only swaps the `Actor` instance defined by the `Props` but the incarnation and hence the UID remains the same.

The lifecycle of an incarnation ends when the actor is stopped. At that point the appropriate lifecycle events are called and watching actors are notified of the termination. After the incarnation is stopped, the path can be reused again by creating an actor with `actorOf()`. In this case the name of the new incarnation will be the same as the previous one but the UIDs will differ.

An `ActorRef` always represents an incarnation (path and UID) not just a given path. Therefore if an actor is stopped and a new one with the same name is created an `ActorRef` of the old incarnation will not point to the new one.

`ActorSelection` on the other hand points to the path (or multiple paths if wildcards are used) and is completely oblivious to which incarnation is currently occupying it. `ActorSelection` cannot be watched for this reason. It is possible to resolve the current incarnation’s `ActorRef` living under the path by sending an `Identify` message to the `ActorSelection` which will be replied to with an `ActorIdentity` containing the correct reference (see *Identifying Actors via Actor Selection*). This can also be done with the `resolveOne` method of the `ActorSelection`, which returns a `Future` of the matching `ActorRef`.

Lifecycle Monitoring aka DeathWatch

In order to be notified when another actor terminates (i.e. stops permanently, not temporary failure and restart), an actor may register itself for reception of the `Terminated` message dispatched by the other actor upon termination (see *Stopping Actors*). This service is provided by the `DeathWatch` component of the actor system.

Registering a monitor is easy:

```
public class WatchActor extends AbstractActor {
    private final ActorRef child = context().actorOf(Props.empty(), "target");
    private ActorRef lastSender = system.deadLetters();

    public WatchActor() {
        context().watch(child); // <-- this is the only call needed for registration

        receive(ReceiveBuilder.
            matchEquals("kill", s -> {
                context().stop(child);
                lastSender = sender();
            }).
            match(Terminated.class, t -> t.actor().equals(child), t -> {
                lastSender.tell("finished", self());
            }).build()
        );
    }
}
```

It should be noted that the `Terminated` message is generated independent of the order in which registration and termination occur. In particular, the watching actor will receive a `Terminated` message even if the watched actor has already been terminated at the time of registration.

Registering multiple times does not necessarily lead to multiple messages being generated, but there is no guarantee that only exactly one such message is received: if termination of the watched actor has generated and queued the message, and another registration is done before this message has been processed, then a second message will be queued, because registering for monitoring of an already terminated actor leads to the immediate generation of the `Terminated` message.

It is also possible to deregister from watching another actor's liveliness using `context.unwatch(target)`. This works even if the `Terminated` message has already been enqueued in the mailbox; after calling `unwatch` no `Terminated` message for that actor will be processed anymore.

Start Hook

Right after starting the actor, its `preStart` method is invoked.

```
@Override
public void preStart() {
    target = context().actorOf(Props.create(MyActor.class, "target"));
}
```

This method is called when the actor is first created. During restarts it is called by the default implementation of `postRestart`, which means that by overriding that method you can choose whether the initialization code in this method is called only exactly once for this actor or for every restart. Initialization code which is part of the actor's constructor will always be called when an instance of the actor class is created, which happens at every restart.

Restart Hooks

All actors are supervised, i.e. linked to another actor with a fault handling strategy. Actors may be restarted in case an exception is thrown while processing a message (see *Supervision and Monitoring*). This restart involves the hooks mentioned above:

1. The old actor is informed by calling `preRestart` with the exception which caused the restart and the message which triggered that exception; the latter may be `None` if the restart was not caused by processing a message, e.g. when a supervisor does not trap the exception and is restarted in turn by its supervisor, or if an actor is restarted due to a sibling's failure. If the message is available, then that message's sender is also accessible in the usual way (i.e. by calling `sender`).

This method is the best place for cleaning up, preparing hand-over to the fresh actor instance, etc. By default it stops all children and calls `postStop`.

2. The initial factory from the `actorOf` call is used to produce the fresh instance.
3. The new actor's `postRestart` method is invoked with the exception which caused the restart. By default the `preStart` is called, just as in the normal start-up case.

An actor restart replaces only the actual actor object; the contents of the mailbox is unaffected by the restart, so processing of messages will resume after the `postRestart` hook returns. The message that triggered the exception will not be received again. Any message sent to an actor while it is being restarted will be queued to its mailbox as usual.

Warning: Be aware that the ordering of failure notifications relative to user messages is not deterministic. In particular, a parent might restart its child before it has processed the last messages sent by the child before the failure. See *Discussion: Message Ordering* for details.

Stop Hook

After stopping an actor, its `postStop` hook is called, which may be used e.g. for deregistering this actor from other services. This hook is guaranteed to run after message queuing has been disabled for this actor, i.e. messages sent to a stopped actor will be redirected to the `deadLetters` of the `ActorSystem`.

4.1.3 Identifying Actors via Actor Selection

As described in *Actor References, Paths and Addresses*, each actor has a unique logical path, which is obtained by following the chain of actors from child to parent until reaching the root of the actor system, and it has a physical path, which may differ if the supervision chain includes any remote supervisors. These paths are used by the system to look up actors, e.g. when a remote message is received and the recipient is searched, but they are also useful more directly: actors may look up other actors by specifying absolute or relative paths—logical or physical—and receive back an `ActorSelection` with the result:

```
// will look up this absolute path
context().actorSelection("/user/serviceA/actor");
// will look up sibling beneath same supervisor
context().actorSelection("../joe");
```

Note: It is always preferable to communicate with other Actors using their `ActorRef` instead of relying upon `ActorSelection`. Exceptions are

- sending messages using the *At-Least-Once Delivery* facility
- initiating first contact with a remote system

In all other cases `ActorRefs` can be provided during Actor creation or initialization, passing them from parent to child or introducing Actors by sending their `ActorRefs` to other Actors within messages.

The supplied path is parsed as a `java.net.URI`, which basically means that it is split on `/` into path elements. If the path starts with `/`, it is absolute and the look-up starts at the root guardian (which is the parent of `/user`); otherwise it starts at the current actor. If a path element equals `..`, the look-up will take a step “up” towards the supervisor of the currently traversed actor, otherwise it will step “down” to the named child. It should be noted that the `..` in actor paths here always means the logical structure, i.e. the supervisor.

The path elements of an actor selection may contain wildcard patterns allowing for broadcasting of messages to that section:

```
// will look all children to serviceB with names starting with worker
context().actorSelection("/user/serviceB/worker*");
// will look up all siblings beneath same supervisor
context().actorSelection("../*");
```

Messages can be sent via the `ActorSelection` and the path of the `ActorSelection` is looked up when delivering each message. If the selection does not match any actors the message will be dropped.

To acquire an `ActorRef` for an `ActorSelection` you need to send a message to the selection and use the `sender()` reference of the reply from the actor. There is a built-in `Identify` message that all Actors will understand and automatically reply to with a `ActorIdentity` message containing the `ActorRef`. This message is handled specially by the actors which are traversed in the sense that if a concrete name lookup fails (i.e. a non-wildcard path element does not correspond to a live actor) then a negative result is generated. Please note that this does not mean that delivery of that reply is guaranteed, it still is a normal message.

```
import akka.actor.ActorIdentity;
import akka.actor.ActorSelection;
import akka.actor.Identify;

public class Follower extends AbstractActor {
    final Integer identifyId = 1;

    public Follower() {
        ActorSelection selection = context().actorSelection("/user/another");
        selection.tell(new Identify(identifyId), self());

        receive(ReceiveBuilder.
            match(ActorIdentity.class, id -> id.getRef() != null, id -> {
                ActorRef ref = id.getRef();
                context().watch(ref);
                context().become(active(ref));
            }).
            match(ActorIdentity.class, id -> id.getRef() == null, id -> {
                context().stop(self());
            }).build()
        );
    }

    final PartialFunction<Object, BoxedUnit> active(final ActorRef another) {
        return ReceiveBuilder.
            match(Terminated.class, t -> t.actor().equals(another), t -> {
                context().stop(self());
            }).build();
    }
}
```

You can also acquire an `ActorRef` for an `ActorSelection` with the `resolveOne` method of the `ActorSelection`. It returns a `Future` of the matching `ActorRef` if such an actor exists. It is completed with failure `[[akka.actor.ActorNotFound]]` if no such actor exists or the identification didn't complete within the supplied *timeout*.

Remote actor addresses may also be looked up, if *remoting* is enabled:

```
context().actorSelection("akka.tcp://app@otherhost:1234/user/serviceB");
```

An example demonstrating actor look-up is given in *Remoting Sample*.

4.1.4 Messages and immutability

IMPORTANT: Messages can be any kind of object but have to be immutable. Akka can't enforce immutability (yet) so this has to be by convention.

Here is an example of an immutable message:

```
public class ImmutableMessage {
    private final int sequenceNumber;
    private final List<String> values;
```

```

public ImmutableMessage(int sequenceNumber, List<String> values) {
    this.sequenceNumber = sequenceNumber;
    this.values = Collections.unmodifiableList(new ArrayList<String>(values));
}

public int getSequenceNumber() {
    return sequenceNumber;
}

public List<String> getValues() {
    return values;
}
}

```

4.1.5 Send messages

Messages are sent to an Actor through one of the following methods.

- `tell` means “fire-and-forget”, e.g. send a message asynchronously and return immediately.
- `ask` sends a message asynchronously and returns a `Future` representing a possible reply.

Message ordering is guaranteed on a per-sender basis.

Note: There are performance implications of using `ask` since something needs to keep track of when it times out, there needs to be something that bridges a `Promise` into an `ActorRef` and it also needs to be reachable through remoting. So always prefer `tell` for performance, and only `ask` if you must.

In all these methods you have the option of passing along your own `ActorRef`. Make it a practice of doing so because it will allow the receiver actors to be able to respond to your message, since the sender reference is sent along with the message.

Tell: Fire-forget

This is the preferred way of sending messages. No blocking waiting for a message. This gives the best concurrency and scalability characteristics.

```

// don't forget to think about who is the sender (2nd argument)
target.tell(message, self());

```

The sender reference is passed along with the message and available within the receiving actor via its `sender` method while processing this message. Inside of an actor it is usually `self` who shall be the sender, but there can be cases where replies shall be routed to some other actor—e.g. the parent—in which the second argument to `tell` would be a different one. Outside of an actor and if no reply is needed the second argument can be `null`; if a reply is needed outside of an actor you can use the `ask`-pattern described next..

Ask: Send-And-Receive-Future

The `ask` pattern involves actors as well as futures, hence it is offered as a use pattern rather than a method on `ActorRef`:

```

import static akka.pattern.Patterns.ask;
import static akka.pattern.Patterns.pipe;
import scala.concurrent.Future;
import scala.concurrent.duration.Duration;
import akka.dispatch.Futures;
import akka.dispatch.Mapper;
import akka.util.Timeout;

```

```

final Timeout t = new Timeout(Duration.create(5, TimeUnit.SECONDS));

final ArrayList<Future<Object>> futures = new ArrayList<Future<Object>>();
futures.add(ask(actorA, "request", 1000)); // using 1000ms timeout
futures.add(ask(actorB, "another request", t)); // using timeout from
// above

final Future<Iterable<Object>> aggregate = Futures.sequence(futures,
    system.dispatcher());

final Future<Result> transformed = aggregate.map(
    new Mapper<Iterable<Object>, Result>() {
        public Result apply(Iterable<Object> coll) {
            final Iterator<Object> it = coll.iterator();
            final String x = (String) it.next();
            final String s = (String) it.next();
            return new Result(x, s);
        }
    }, system.dispatcher());

pipe(transformed, system.dispatcher()).to(actorC);

```

This example demonstrates `ask` together with the `pipe` pattern on futures, because this is likely to be a common combination. Please note that all of the above is completely non-blocking and asynchronous: `ask` produces a `Future`, two of which are composed into a new future using the `Futures.sequence` and `map` methods and then `pipe` installs an `onComplete`-handler on the future to effect the submission of the aggregated `Result` to another actor.

Using `ask` will send a message to the receiving Actor as with `tell`, and the receiving actor must reply with `sender().tell(reply, self())` in order to complete the returned `Future` with a value. The `ask` operation involves creating an internal actor for handling this reply, which needs to have a timeout after which it is destroyed in order not to leak resources; see more below.

Warning: To complete the future with an exception you need send a `Failure` message to the sender. This is *not done automatically* when an actor throws an exception while processing a message.

```

try {
    String result = operation();
    sender().tell(result, self());
} catch (Exception e) {
    sender().tell(new akka.actor.Status.Failure(e), self());
    throw e;
}

```

If the actor does not complete the future, it will expire after the timeout period, specified as parameter to the `ask` method; this will complete the `Future` with an `AskTimeoutException`.

See [Futures](#) for more information on how to await or query a future.

The `onComplete`, `onSuccess`, or `onFailure` methods of the `Future` can be used to register a callback to get a notification when the `Future` completes. Gives you a way to avoid blocking.

Warning: When using future callbacks, inside actors you need to carefully avoid closing over the containing actor's reference, i.e. do not call methods or access mutable state on the enclosing actor from within the callback. This would break the actor encapsulation and may introduce synchronization bugs and race conditions because the callback will be scheduled concurrently to the enclosing actor. Unfortunately there is not yet a way to detect these illegal accesses at compile time. See also: [Actors and shared mutable state](#)

Forward message

You can forward a message from one actor to another. This means that the original sender address/reference is maintained even though the message is going through a ‘mediator’. This can be useful when writing actors that work as routers, load-balancers, replicators etc.

```
target.forward(result, context());
```

4.1.6 Receive messages

An Actor either has to set its initial receive behavior in the constructor by calling the `receive` method in the `AbstractActor`:

```
public SomeActor() {
    receive(ReceiveBuilder.
        // and some behavior ...
        build());
}
```

or by implementing the `receive` method in the `Actor` interface:

```
public abstract PartialFunction<Object, BoxedUnit> receive();
```

Both the argument to the `AbstractActor` `receive` method and the return type of the `Actor` `receive` method is a `PartialFunction<Object, BoxedUnit>` that defines which messages your Actor can handle, along with the implementation of how the messages should be processed.

Don't let the type signature scare you. To allow you to easily build up a partial function there is a builder named `ReceiveBuilder` that you can use.

Here is an example:

```
import akka.actor.AbstractActor;
import akka.event.Logging;
import akka.event.LoggingAdapter;
import akka.japi.pf.ReceiveBuilder;

public class MyActor extends AbstractActor {
    private final LoggingAdapter log = Logging.getLogger(context().system(), this);

    public MyActor() {
        receive(ReceiveBuilder.
            match(String.class, s -> {
                log.info("Received String message: {}", s);
            }).
            matchAny(o -> log.info("received unknown message")).build()
        );
    }
}
```

4.1.7 Reply to messages

If you want to have a handle for replying to a message, you can use `sender()`, which gives you an `ActorRef`. You can reply by sending to that `ActorRef` with `sender().tell(replyMsg, self())`. You can also store the `ActorRef` for replying later, or passing on to other actors. If there is no sender (a message was sent without an actor or future context) then the sender defaults to a ‘dead-letter’ actor ref.

```
sender().tell(s, self());
```

4.1.8 Receive timeout

The `ActorContext` `setReceiveTimeout` defines the inactivity timeout after which the sending of a `ReceiveTimeout` message is triggered. When specified, the receive function should be able to handle an `akka.actor.ReceiveTimeout` message. 1 millisecond is the minimum supported timeout.

Please note that the receive timeout might fire and enqueue the `ReceiveTimeout` message right after another message was enqueued; hence it is **not guaranteed** that upon reception of the receive timeout there must have been an idle period beforehand as configured via this method.

Once set, the receive timeout stays in effect (i.e. continues firing repeatedly after inactivity periods). Pass in `Duration.Undefined` to switch off this feature.

```
public class ReceiveTimeoutActor extends AbstractActor {
    public ReceiveTimeoutActor() {
        // To set an initial delay
        context().setReceiveTimeout(Duration.create("10 seconds"));

        receive(ReceiveBuilder.
            matchEquals("Hello", s -> {
                // To set in a response to a message
                context().setReceiveTimeout(Duration.create("1 second"));
            }).
            match(ReceiveTimeout.class, r -> {
                // To turn it off
                context().setReceiveTimeout(Duration.Undefined());
            }).build()
        );
    }
}
```

Messages marked with `NotInfluenceReceiveTimeout` will not reset the timer. This can be useful when `ReceiveTimeout` should be fired by external inactivity but not influenced by internal activity, e.g. scheduled tick messages.

4.1.9 Stopping actors

Actors are stopped by invoking the `stop` method of a `ActorRefFactory`, i.e. `ActorContext` or `ActorSystem`. Typically the context is used for stopping child actors and the system for stopping top level actors. The actual termination of the actor is performed asynchronously, i.e. `stop` may return before the actor is stopped.

Processing of the current message, if any, will continue before the actor is stopped, but additional messages in the mailbox will not be processed. By default these messages are sent to the `deadLetters` of the `ActorSystem`, but that depends on the mailbox implementation.

Termination of an actor proceeds in two steps: first the actor suspends its mailbox processing and sends a stop command to all its children, then it keeps processing the internal termination notifications from its children until the last one is gone, finally terminating itself (invoking `postStop`, dumping mailbox, publishing `Terminated` on the `DeathWatch`, telling its supervisor). This procedure ensures that actor system sub-trees terminate in an orderly fashion, propagating the stop command to the leaves and collecting their confirmation back to the stopped supervisor. If one of the actors does not respond (i.e. processing a message for extended periods of time and therefore not receiving the stop command), this whole process will be stuck.

Upon `ActorSystem.terminate`, the system guardian actors will be stopped, and the aforementioned process will ensure proper termination of the whole system.

The `postStop` hook is invoked after an actor is fully stopped. This enables cleaning up of resources:

```
@Override
public void postStop() {
    // clean up some resources ...
}
```

Note: Since stopping an actor is asynchronous, you cannot immediately reuse the name of the child you just stopped; this will result in an `InvalidActorNameException`. Instead, watch the terminating actor and create its replacement in response to the `Terminated` message which will eventually arrive.

PoisonPill

You can also send an actor the `akka.actor.PoisonPill` message, which will stop the actor when the message is processed. `PoisonPill` is enqueued as ordinary messages and will be handled after messages that were already queued in the mailbox.

Graceful Stop

`gracefulStop` is useful if you need to wait for termination or compose ordered termination of several actors:

```
import static akka.pattern.Patterns.gracefulStop;
import scala.concurrent.Await;
import scala.concurrent.Future;
import scala.concurrent.duration.Duration;
import akka.pattern.AskTimeoutException;
```

```
try {
    Future<Boolean> stopped =
        gracefulStop(actorRef, Duration.create(5, TimeUnit.SECONDS), Manager.SHUTDOWN);
    Await.result(stopped, Duration.create(6, TimeUnit.SECONDS));
    // the actor has been stopped
} catch (AskTimeoutException e) {
    // the actor wasn't stopped within 5 seconds
}
```

```
public class Manager extends AbstractActor {
    private static enum Shutdown {
        Shutdown
    }
    public static final Shutdown SHUTDOWN = Shutdown.Shutdown;

    private ActorRef worker =
        context().watch(context().actorOf(Props.create(Cruncher.class), "worker"));

    public Manager() {
        receive(ReceiveBuilder.
            matchEquals("job", s -> {
                worker.tell("crunch", self());
            }).
            matchEquals(SHUTDOWN, x -> {
                worker.tell(PoisonPill.getInstance(), self());
                context().become(shuttingDown);
            }).build()
        );
    }

    public PartialFunction<Object, BoxedUnit> shuttingDown =
        ReceiveBuilder.
            matchEquals("job", s -> {
                sender().tell("service unavailable, shutting down", self());
            }).
            match(Terminated.class, t -> t.actor().equals(worker), t -> {
                context().stop(self());
            }).build();
}
```

When `gracefulStop()` returns successfully, the actor's `postStop()` hook will have been executed: there exists a happens-before edge between the end of `postStop()` and the return of `gracefulStop()`.

In the above example a custom `Manager.Shutdown` message is sent to the target actor to initiate the process of stopping the actor. You can use `PoisonPill` for this, but then you have limited possibilities to perform interactions with other actors before stopping the target actor. Simple cleanup tasks can be handled in `postStop`.

Warning: Keep in mind that an actor stopping and its name being deregistered are separate events which happen asynchronously from each other. Therefore it may be that you will find the name still in use after `gracefulStop()` returned. In order to guarantee proper deregistration, only reuse names from within a supervisor you control and only in response to a `Terminated` message, i.e. not for top-level actors.

4.1.10 Become/Unbecome

Upgrade

Akka supports hotswapping the Actor's message loop (e.g. its implementation) at runtime: invoke the `context.become` method from within the Actor. `become` takes a `PartialFunction<Object, BoxedUnit>` that implements the new message handler. The hotswapped code is kept in a `Stack` which can be pushed and popped.

Warning: Please note that the actor will revert to its original behavior when restarted by its Supervisor.

To hotswap the Actor behavior using `become`:

```
public class HotSwapActor extends AbstractActor {
    private PartialFunction<Object, BoxedUnit> angry;
    private PartialFunction<Object, BoxedUnit> happy;

    public HotSwapActor() {
        angry =
            ReceiveBuilder.
                matchEquals("foo", s -> {
                    sender().tell("I am already angry?", self());
                }).
                matchEquals("bar", s -> {
                    context().become(happy);
                }).build();

        happy = ReceiveBuilder.
            matchEquals("bar", s -> {
                sender().tell("I am already happy :-)", self());
            }).
            matchEquals("foo", s -> {
                context().become(angry);
            }).build();

        receive(ReceiveBuilder.
            matchEquals("foo", s -> {
                context().become(angry);
            }).
            matchEquals("bar", s -> {
                context().become(happy);
            }).build()
        );
    }
}
```

This variant of the `become` method is useful for many different things, such as to implement a Finite State Machine (FSM, for an example see [Dining Hakkers](#)). It will replace the current behavior (i.e. the top of the behavior stack), which means that you do not use `unbecome`, instead always the next behavior is explicitly installed.

The other way of using `become` does not replace but add to the top of the behavior stack. In this case care must be taken to ensure that the number of “pop” operations (i.e. `unbecome`) matches the number of “push” ones in the long run, otherwise this amounts to a memory leak (which is why this behavior is not the default).

```
public class Swapper extends AbstractLoggingActor {
  public Swapper() {
    receive(ReceiveBuilder.
      matchEquals(Swap, s -> {
        log().info("Hi");
        context().become(ReceiveBuilder.
          matchEquals(Swap, x -> {
            log().info("Ho");
            context().unbecome(); // resets the latest 'become' (just for fun)
          }).build(), false); // push on top instead of replace
        }).build()
    );
  }
}

public class SwapperApp {
  public static void main(String[] args) {
    ActorSystem system = ActorSystem.create("SwapperSystem");
    ActorRef swapper = system.actorOf(Props.create(Swapper.class), "swapper");
    swapper.tell(Swap, ActorRef.noSender()); // logs Hi
    swapper.tell(Swap, ActorRef.noSender()); // logs Ho
    swapper.tell(Swap, ActorRef.noSender()); // logs Hi
    swapper.tell(Swap, ActorRef.noSender()); // logs Ho
    swapper.tell(Swap, ActorRef.noSender()); // logs Hi
    swapper.tell(Swap, ActorRef.noSender()); // logs Ho
    system.terminate();
  }
}
```

4.1.11 Stash

The `AbstractActorWithStash` class enables an actor to temporarily stash away messages that can not or should not be handled using the actor’s current behavior. Upon changing the actor’s message handler, i.e., right before invoking `context().become()` or `context().unbecome()`, all stashed messages can be “unstashed”, thereby prepending them to the actor’s mailbox. This way, the stashed messages can be processed in the same order as they have been received originally. An actor that extends `AbstractActorWithStash` will automatically get a deque-based mailbox.

Note: The abstract class `AbstractActorWithStash` implements the marker interface `RequiresMessageQueue<DequeBasedMessageQueueSemantics>` which requests the system to automatically choose a deque based mailbox implementation for the actor. If you want more control over the mailbox, see the documentation on mailboxes: [Mailboxes](#).

Here is an example of the `AbstractActorWithStash` class in action:

```
public class ActorWithProtocol extends AbstractActorWithStash {
  public ActorWithProtocol() {
    receive(ReceiveBuilder.
      matchEquals("open", s -> {
        context().become(ReceiveBuilder.
          matchEquals("write", ws -> { /* do writing */ }).

```

```

    matchEquals("close", cs -> {
        unstashAll();
        context().unbecome();
    }).
    matchAny(msg -> stash().build(), false);
}).
matchAny(msg -> stash().build()
);
}
}

```

Invoking `stash()` adds the current message (the message that the actor received last) to the actor's stash. It is typically invoked when handling the default case in the actor's message handler to stash messages that aren't handled by the other cases. It is illegal to stash the same message twice; to do so results in an `IllegalStateException` being thrown. The stash may also be bounded in which case invoking `stash()` may lead to a capacity violation, which results in a `StashOverflowException`. The capacity of the stash can be configured using the `stash-capacity` setting (an `Int`) of the mailbox's configuration.

Invoking `unstashAll()` enqueues messages from the stash to the actor's mailbox until the capacity of the mailbox (if any) has been reached (note that messages from the stash are prepended to the mailbox). In case a bounded mailbox overflows, a `MessageQueueAppendFailedException` is thrown. The stash is guaranteed to be empty after calling `unstashAll()`.

The stash is backed by a `scala.collection.immutable.Vector`. As a result, even a very large number of messages may be stashed without a major impact on performance.

Note that the stash is part of the ephemeral actor state, unlike the mailbox. Therefore, it should be managed like other parts of the actor's state which have the same property. The `AbstractActorWithStash` implementation of `preRestart` will call `unstashAll()`, which is usually the desired behavior.

Note: If you want to enforce that your actor can only work with an unbounded stash, then you should use the `AbstractActorWithUnboundedStash` class instead.

4.1.12 Killing an Actor

You can kill an actor by sending a `Kill` message. This will cause the actor to throw a `ActorKilledException`, triggering a failure. The actor will suspend operation and its supervisor will be asked how to handle the failure, which may mean resuming the actor, restarting it or terminating it completely. See *What Supervision Means* for more information.

Use `Kill` like this:

```
victim.tell(akka.actor.Kill.getInstance(), ActorRef.noSender());
```

4.1.13 Actors and exceptions

It can happen that while a message is being processed by an actor, that some kind of exception is thrown, e.g. a database exception.

What happens to the Message

If an exception is thrown while a message is being processed (i.e. taken out of its mailbox and handed over to the current behavior), then this message will be lost. It is important to understand that it is not put back on the mailbox. So if you want to retry processing of a message, you need to deal with it yourself by catching the exception and retry your flow. Make sure that you put a bound on the number of retries since you don't want a system to livelock (so consuming a lot of cpu cycles without making progress). Another possibility would be to have a look at the *PeekMailbox pattern*.

What happens to the mailbox

If an exception is thrown while a message is being processed, nothing happens to the mailbox. If the actor is restarted, the same mailbox will be there. So all messages on that mailbox will be there as well.

What happens to the actor

If code within an actor throws an exception, that actor is suspended and the supervision process is started (see *Supervision and Monitoring*). Depending on the supervisor's decision the actor is resumed (as if nothing happened), restarted (wiping out its internal state and starting from scratch) or terminated.

4.1.14 Initialization patterns

The rich lifecycle hooks of Actors provide a useful toolkit to implement various initialization patterns. During the lifetime of an `ActorRef`, an actor can potentially go through several restarts, where the old instance is replaced by a fresh one, invisibly to the outside observer who only sees the `ActorRef`.

One may think about the new instances as “incarnations”. Initialization might be necessary for every incarnation of an actor, but sometimes one needs initialization to happen only at the birth of the first instance when the `ActorRef` is created. The following sections provide patterns for different initialization needs.

Initialization via constructor

Using the constructor for initialization has various benefits. First of all, it makes it possible to use `val` fields to store any state that does not change during the life of the actor instance, making the implementation of the actor more robust. The constructor is invoked for every incarnation of the actor, therefore the internals of the actor can always assume that proper initialization happened. This is also the drawback of this approach, as there are cases when one would like to avoid reinitializing internals on restart. For example, it is often useful to preserve child actors across restarts. The following section provides a pattern for this case.

Initialization via `preStart`

The method `preStart()` of an actor is only called once directly during the initialization of the first instance, that is, at creation of its `ActorRef`. In the case of restarts, `preStart()` is called from `postRestart()`, therefore if not overridden, `preStart()` is called on every incarnation. However, overriding `postRestart()` one can disable this behavior, and ensure that there is only one call to `preStart()`.

One useful usage of this pattern is to disable creation of new `ActorRefs` for children during restarts. This can be achieved by overriding `preRestart()`:

```
@Override
public void preStart() {
    // Initialize children here
}

// Overriding postRestart to disable the call to preStart()
// after restarts
@Override
public void postRestart(Throwable reason) {
}

// The default implementation of preRestart() stops all the children
// of the actor. To opt-out from stopping the children, we
// have to override preRestart()
@Override
public void preRestart(Throwable reason, Option<Object> message)
    throws Exception {
    // Keep the call to postStop(), but no stopping of children
}
```

```
postStop();
}
```

Please note, that the child actors are *still restarted*, but no new `ActorRef` is created. One can recursively apply the same principles for the children, ensuring that their `preStart()` method is called only at the creation of their refs.

For more information see [What Restarting Means](#).

Initialization via message passing

There are cases when it is impossible to pass all the information needed for actor initialization in the constructor, for example in the presence of circular dependencies. In this case the actor should listen for an initialization message, and use `become()` or a finite state-machine state transition to encode the initialized and uninitialized states of the actor.

```
receive(ReceiveBuilder.
  matchEquals("init", m1 -> {
    initializeMe = "Up and running";
    context().become(ReceiveBuilder.
      matchEquals("U OK?", m2 -> {
        sender().tell(initializeMe, self());
      }).build());
  }).build())
```

If the actor may receive messages before it has been initialized, a useful tool can be the `Stash` to save messages until the initialization finishes, and replaying them after the actor became initialized.

Warning: This pattern should be used with care, and applied only when none of the patterns above are applicable. One of the potential issues is that messages might be lost when sent to remote actors. Also, publishing an `ActorRef` in an uninitialized state might lead to the condition that it receives a user message before the initialization has been done.

4.1.15 Lambdas and Performance

There is one big difference between the optimized partial functions created by the Scala compiler and the ones created by the `ReceiveBuilder`. The partial functions created by the `ReceiveBuilder` consist of multiple lambda expressions for every match statement, where each lambda is an object referencing the code to be run. This is something that the JVM can have problems optimizing and the resulting code might not be as performant as the Scala equivalent or the corresponding *untyped actor* version.

4.2 Fault Tolerance (Java with Lambda Support)


As explained in [Actor Systems](#) each actor is the supervisor of its children, and as such each actor defines fault handling supervisor strategy. This strategy cannot be changed afterwards as it is an integral part of the actor system's structure.

4.2.1 Fault Handling in Practice

First, let us look at a sample that illustrates one way to handle data store errors, which is a typical source of failure in real world applications. Of course it depends on the actual application what is possible to do when the data store is unavailable, but in this sample we use a best effort re-connect approach.

Read the following source code. The inlined comments explain the different pieces of the fault handling and why they are added. It is also highly recommended to run this sample as it is easy to follow the log output to understand what is happening in runtime.


Diagrams of the Fault Tolerance Sample



The above diagram illustrates the normal message flow.

Normal flow:

Step	Description
1	The progress Listener starts the work.
2	The Worker schedules work by sending Do messages periodically to itself
3, 4, 5	When receiving Do the Worker tells the CounterService to increment the counter, three times. The Increment message is forwarded to the Counter, which updates its counter variable and sends current value to the Storage.
6, 7	The Worker asks the CounterService of current value of the counter and pipes the result back to the Listener.



The above diagram illustrates what happens in case of storage failure.

Failure flow:

Step	Description
1	The Storage throws StorageException.
2	The CounterService is supervisor of the Storage and restarts the Storage when StorageException is thrown.
3, 4, 5, 6	The Storage continues to fail and is restarted.
7	After 3 failures and restarts within 5 seconds the Storage is stopped by its supervisor, i.e. the CounterService.
8	The CounterService is also watching the Storage for termination and receives the Terminated message when the Storage has been stopped ...
9, 10, 11	and tells the Counter that there is no Storage.
12	The CounterService schedules a Reconnect message to itself.
13, 14	When it receives the Reconnect message it creates a new Storage ...
15, 16	and tells the Counter to use the new Storage

Full Source Code of the Fault Tolerance Sample

```
import java.util.ArrayList;
import java.util.HashMap;
import java.util.List;
import java.util.Map;

import akka.actor.*;
import akka.dispatch.Mapper;
import akka.event.LoggingReceive;
import akka.japi.pf.DeciderBuilder;
import akka.japi.pf.ReceiveBuilder;
import akka.util.Timeout;
import com.typesafe.config.Config;
import com.typesafe.config.ConfigFactory;
import scala.concurrent.duration.Duration;

import static akka.japi.Util.classTag;
import static akka.actor.SupervisorStrategy.restart;
import static akka.actor.SupervisorStrategy.stop;
import static akka.actor.SupervisorStrategy.escalate;

import static akka.pattern.Patterns.ask;
import static akka.pattern.Patterns.pipe;

import static docs.actorlambda.japi.FaultHandlingDocSample.WorkerApi.*;
import static docs.actorlambda.japi.FaultHandlingDocSample.CounterServiceApi.*;
import static docs.actorlambda.japi.FaultHandlingDocSample.CounterApi.*;
import static docs.actorlambda.japi.FaultHandlingDocSample.StorageApi.*;

public class FaultHandlingDocSample {

    /**
     * Runs the sample
     */
    public static void main(String[] args) {
        Config config = ConfigFactory.parseString(
            "akka.loglevel = \"DEBUG\"\\n" +
            "akka.actor.debug {\\n" +
            "  receive = on\\n" +
            "  lifecycle = on\\n" +

```

```

    ")\n");

    ActorSystem system = ActorSystem.create("FaultToleranceSample", config);
    ActorRef worker = system.actorOf(Props.create(Worker.class), "worker");
    ActorRef listener = system.actorOf(Props.create(Listener.class), "listener");
    // start the work and listen on progress
    // note that the listener is used as sender of the tell,
    // i.e. it will receive replies from the worker
    worker.tell(Start, listener);
}

/**
 * Listens on progress from the worker and shuts down the system when enough
 * work has been done.
 */
public static class Listener extends AbstractLoggingActor {

    @Override
    public void preStart() {
        // If we don't get any progress within 15 seconds then the service
        // is unavailable
        context().setReceiveTimeout(Duration.create("15 seconds"));
    }

    public Listener() {
        receive(LoggingReceive.create(ReceiveBuilder.
            match(Progress.class, progress -> {
                log().info("Current progress: {} %", progress.percent);
                if (progress.percent >= 100.0) {
                    log().info("That's all, shutting down");
                    context().system().terminate();
                }
            })
            .matchEquals(ReceiveTimeout.getInstance(), x -> {
                // No progress within 15 seconds, ServiceUnavailable
                log().error("Shutting down due to unavailable service");
                context().system().terminate();
            }).build(), context())
        );
    }
}

public interface WorkerApi {
    public static final Object Start = "Start";
    public static final Object Do = "Do";

    public static class Progress {
        public final double percent;

        public Progress(double percent) {
            this.percent = percent;
        }

        public String toString() {
            return String.format("%s(%s)", getClass().getSimpleName(), percent);
        }
    }
}

/**
 * Worker performs some work when it receives the Start message. It will
 * continuously notify the sender of the Start message of current Progress.

```

```

* The Worker supervise the CounterService.
*/
public static class Worker extends AbstractLoggingActor {
    final Timeout askTimeout = new Timeout(Duration.create(5, "seconds"));

    // The sender of the initial Start message will continuously be notified
    // about progress
    ActorRef progressListener;
    final ActorRef counterService = context().actorOf(
        Props.create(CounterService.class), "counter");
    final int totalCount = 51;

    // Stop the CounterService child if it throws ServiceUnavailable
    private static final SupervisorStrategy strategy =
        new OneForOneStrategy(DeciderBuilder.
            match(ServiceUnavailable.class, e -> stop()).
            matchAny(o -> escalate()).build());

    @Override
    public SupervisorStrategy supervisorStrategy() {
        return strategy;
    }

    public Worker() {
        receive(LoggingReceive.create(ReceiveBuilder.
            matchEquals(Start, x -> progressListener == null, x -> {
                progressListener = sender();
                context().system().scheduler().schedule(
                    Duration.Zero(), Duration.create(1, "second"), self(), Do,
                    context().dispatcher(), null
                );
            })
            matchEquals(Do, x -> {
                counterService.tell(new Increment(1), self());
                counterService.tell(new Increment(1), self());
                counterService.tell(new Increment(1), self());
                // Send current progress to the initial sender
                pipe(ask(counterService, GetCurrentCount, askTimeout)
                    .mapTo(classTag(CurrentCount.class))
                    .map(new Mapper<CurrentCount, Progress>() {
                        public Progress apply(CurrentCount c) {
                            return new Progress(100.0 * c.count / totalCount);
                        }
                    }, context().dispatcher()), context().dispatcher())
                    .to(progressListener);
            })
            .build(), context())
        );
    }
}

public interface CounterServiceApi {

    public static final Object GetCurrentCount = "GetCurrentCount";

    public static class CurrentCount {
        public final String key;
        public final long count;

        public CurrentCount(String key, long count) {
            this.key = key;
            this.count = count;
        }
    }
}

```

```

public String toString() {
    return String.format("%s(%s, %s)", getClass().getSimpleName(), key, count);
}
}

public static class Increment {
    public final long n;

    public Increment(long n) {
        this.n = n;
    }

    public String toString() {
        return String.format("%s(%s)", getClass().getSimpleName(), n);
    }
}

public static class ServiceUnavailable extends RuntimeException {
    private static final long serialVersionUID = 1L;
    public ServiceUnavailable(String msg) {
        super(msg);
    }
}
}

/**
 * Adds the value received in Increment message to a persistent counter.
 * Replies with CurrentCount when it is asked for CurrentCount. CounterService
 * supervise Storage and Counter.
 */
public static class CounterService extends AbstractLoggingActor {

    // Reconnect message
    static final Object Reconnect = "Reconnect";

    private static class SenderMsgPair {
        final ActorRef sender;
        final Object msg;

        SenderMsgPair(ActorRef sender, Object msg) {
            this.msg = msg;
            this.sender = sender;
        }
    }

    final String key = self().path().name();
    ActorRef storage;
    ActorRef counter;
    final List<SenderMsgPair> backlog = new ArrayList<>();
    final int MAX_BACKLOG = 10000;

    // Restart the storage child when StorageException is thrown.
    // After 3 restarts within 5 seconds it will be stopped.
    private static final SupervisorStrategy strategy =
        new OneForOneStrategy(3, Duration.create("5 seconds"), DeciderBuilder.
            match(StorageException.class, e -> restart()).
            matchAny(o -> escalate()).build());

    @Override
    public SupervisorStrategy supervisorStrategy() {
        return strategy;
    }
}

```

```

}

@Override
public void preStart() {
    initStorage();
}

/**
 * The child storage is restarted in case of failure, but after 3 restarts,
 * and still failing it will be stopped. Better to back-off than
 * continuously failing. When it has been stopped we will schedule a
 * Reconnect after a delay. Watch the child so we receive Terminated message
 * when it has been terminated.
 */
void initStorage() {
    storage = context().watch(context().actorOf(
        Props.create(Storage.class), "storage"));
    // Tell the counter, if any, to use the new storage
    if (counter != null)
        counter.tell(new UseStorage(storage), self());
    // We need the initial value to be able to operate
    storage.tell(new Get(key), self());
}

public CounterService() {
    receive(LoggingReceive.create(ReceiveBuilder.
        match(Entry.class, entry -> entry.key.equals(key) && counter == null, entry -> {
            // Reply from Storage of the initial value, now we can create the Counter
            final long value = entry.value;
            counter = context().actorOf(Props.create(Counter.class, key, value));
            // Tell the counter to use current storage
            counter.tell(new UseStorage(storage), self());
            // and send the buffered backlog to the counter
            for (SenderMsgPair each : backlog) {
                counter.tell(each.msg, each.sender);
            }
            backlog.clear();
        })).
        match(Increment.class, increment -> {
            forwardOrPlaceInBacklog(increment);
        })).
        matchEquals(GetCurrentCount, gcc -> {
            forwardOrPlaceInBacklog(gcc);
        })).
        match(Terminated.class, o -> {
            // After 3 restarts the storage child is stopped.
            // We receive Terminated because we watch the child, see initStorage.
            storage = null;
            // Tell the counter that there is no storage for the moment
            counter.tell(new UseStorage(null), self());
            // Try to re-establish storage after while
            context().system().scheduler().scheduleOnce(
                Duration.create(10, "seconds"), self(), Reconnect,
                context().dispatcher(), null);
        })).
        matchEquals(Reconnect, o -> {
            // Re-establish storage after the scheduled delay
            initStorage();
        })).build(), context());
};
}

void forwardOrPlaceInBacklog(Object msg) {

```

```

// We need the initial value from storage before we can start delegate to
// the counter. Before that we place the messages in a backlog, to be sent
// to the counter when it is initialized.
if (counter == null) {
    if (backlog.size() >= MAX_BACKLOG)
        throw new ServiceUnavailable("CounterService not available," +
            " lack of initial value");
    backlog.add(new SenderMsgPair(sender(), msg));
} else {
    counter.forward(msg, context());
}
}
}

public interface CounterApi {
    public static class UseStorage {
        public final ActorRef storage;

        public UseStorage(ActorRef storage) {
            this.storage = storage;
        }

        public String toString() {
            return String.format("%s(%s)", getClass().getSimpleName(), storage);
        }
    }
}

/**
 * The in memory count variable that will send current value to the Storage,
 * if there is any storage available at the moment.
 */
public static class Counter extends AbstractLoggingActor {
    final String key;
    long count;
    ActorRef storage;

    public Counter(String key, long initialValue) {
        this.key = key;
        this.count = initialValue;

        receive(LoggingReceive.create(ReceiveBuilder.
            match(UseStorage.class, useStorage -> {
                storage = useStorage.storage;
                storeCount();
            }).
            match(Increment.class, increment -> {
                count += increment.n;
                storeCount();
            }).
            matchEquals(GetCurrentCount, gcc -> {
                sender().tell(new CurrentCount(key, count), self());
            }).build(), context());
    }

    void storeCount() {
        // Delegate dangerous work, to protect our valuable state.
        // We can continue without storage.
        if (storage != null) {
            storage.tell(new Store(new Entry(key, count)), self());
        }
    }
}

```

```

    }
}

public interface StorageApi {

    public static class Store {
        public final Entry entry;

        public Store(Entry entry) {
            this.entry = entry;
        }

        public String toString() {
            return String.format("%s(%s)", getClass().getSimpleName(), entry);
        }
    }

    public static class Entry {
        public final String key;
        public final long value;

        public Entry(String key, long value) {
            this.key = key;
            this.value = value;
        }

        public String toString() {
            return String.format("%s(%s, %s)", getClass().getSimpleName(), key, value);
        }
    }

    public static class Get {
        public final String key;

        public Get(String key) {
            this.key = key;
        }

        public String toString() {
            return String.format("%s(%s)", getClass().getSimpleName(), key);
        }
    }

    public static class StorageException extends RuntimeException {
        private static final long serialVersionUID = 1L;
        public StorageException(String msg) {
            super(msg);
        }
    }
}

/**
 * Saves key/value pairs to persistent storage when receiving Store message.
 * Replies with current value when receiving Get message. Will throw
 * StorageException if the underlying data store is out of order.
 */
public static class Storage extends AbstractLoggingActor {

    final DummyDB db = DummyDB.instance;

    public Storage() {
        receive(LoggingReceive.create(ReceiveBuilder.

```

```

    match(Store.class, store -> {
      db.save(store.entry.key, store.entry.value);
    }).
    match(Get.class, get -> {
      Long value = db.load(get.key);
      sender().tell(new Entry(get.key, value == null ?
        Long.valueOf(0L) : value), self());
    }).build(), context())
  );
}
}

public static class DummyDB {
  public static final DummyDB instance = new DummyDB();
  private final Map<String, Long> db = new HashMap<String, Long>();

  private DummyDB() {
  }

  public synchronized void save(String key, Long value) throws StorageException {
    if (11 <= value && value <= 14)
      throw new StorageException("Simulated store failure " + value);
    db.put(key, value);
  }

  public synchronized Long load(String key) throws StorageException {
    return db.get(key);
  }
}
}

```

4.2.2 Creating a Supervisor Strategy

The following sections explain the fault handling mechanism and alternatives in more depth.

For the sake of demonstration let us consider the following strategy:

```

private static SupervisorStrategy strategy =
  new OneForOneStrategy(10, Duration.create("1 minute"), DeciderBuilder.
    match(ArithmeticException.class, e -> resume()).
    match(NullPointerException.class, e -> restart()).
    match(IllegalArgumentException.class, e -> stop()).
    matchAny(o -> escalate()).build());

@Override
public SupervisorStrategy supervisorStrategy() {
  return strategy;
}

```

I have chosen a few well-known exception types in order to demonstrate the application of the fault handling directives described in *Supervision and Monitoring*. First off, it is a one-for-one strategy, meaning that each child is treated separately (an all-for-one strategy works very similarly, the only difference is that any decision is applied to all children of the supervisor, not only the failing one). There are limits set on the restart frequency, namely maximum 10 restarts per minute. `-1` and `Duration.Inf()` means that the respective limit does not apply, leaving the possibility to specify an absolute upper limit on the restarts or to make the restarts work infinitely. The child actor is stopped if the limit is exceeded.

Note: If the strategy is declared inside the supervising actor (as opposed to a separate class) its decider has access to all internal state of the actor in a thread-safe fashion, including obtaining a reference to the currently failed child (available as the `getSender` of the failure message).

Default Supervisor Strategy

`Escalate` is used if the defined strategy doesn't cover the exception that was thrown.

When the supervisor strategy is not defined for an actor the following exceptions are handled by default:

- `ActorInitializationException` will stop the failing child actor
- `ActorKilledException` will stop the failing child actor
- `Exception` will restart the failing child actor
- Other types of `Throwable` will be escalated to parent actor

If the exception escalate all the way up to the root guardian it will handle it in the same way as the default strategy defined above.

Stopping Supervisor Strategy

Closer to the Erlang way is the strategy to just stop children when they fail and then take corrective action in the supervisor when `DeathWatch` signals the loss of the child. This strategy is also provided pre-packaged as `SupervisorStrategy.stoppingStrategy` with an accompanying `StoppingSupervisorStrategy` configurator to be used when you want the `"/user"` guardian to apply it.

Logging of Actor Failures

By default the `SupervisorStrategy` logs failures unless they are escalated. Escalated failures are supposed to be handled, and potentially logged, at a level higher in the hierarchy.

You can mute the default logging of a `SupervisorStrategy` by setting `loggingEnabled` to `false` when instantiating it. Customized logging can be done inside the `Decider`. Note that the reference to the currently failed child is available as the `getSender` when the `SupervisorStrategy` is declared inside the supervising actor.

You may also customize the logging in your own `SupervisorStrategy` implementation by overriding the `logFailure` method.

4.2.3 Supervision of Top-Level Actors

Toplevel actors means those which are created using `system.actorOf()`, and they are children of the *User Guardian*. There are no special rules applied in this case, the guardian simply applies the configured strategy.

4.2.4 Test Application

The following section shows the effects of the different directives in practice, where a test setup is needed. First off, we need a suitable supervisor:

```
public class Supervisor extends AbstractActor {

    private static SupervisorStrategy strategy =
        new OneForOneStrategy(10, Duration.create("1 minute"), DeciderBuilder.
            match(ArithmeticException.class, e -> resume()).
            match(NullPointerException.class, e -> restart()).
            match(IllegalArgumentException.class, e -> stop()).
            matchAny(o -> escalate()).build());

    @Override
    public SupervisorStrategy supervisorStrategy() {
        return strategy;
    }
}
```

```

public Supervisor() {
    receive(ReceiveBuilder.
        match(Props.class, props -> {
            sender().tell(context().actorOf(props), self());
        }).build()
    );
}
}

```

This supervisor will be used to create a child, with which we can experiment:

```

public class Child extends AbstractActor {
    int state = 0;

    public Child() {
        receive(ReceiveBuilder.
            match(Exception.class, exception -> { throw exception; }).
            match(Integer.class, i -> state = i).
            matchEquals("get", s -> sender().tell(state, self())).build()
        );
    }
}

```

The test is easier by using the utilities described in *akka-testkit*, where `TestProbe` provides an actor ref useful for receiving and inspecting replies.

```

import akka.actor.*;

import static akka.actor.SupervisorStrategy.resume;
import static akka.actor.SupervisorStrategy.restart;
import static akka.actor.SupervisorStrategy.stop;
import static akka.actor.SupervisorStrategy.escalate;
import akka.japi.pf.DeciderBuilder;
import akka.japi.pf.ReceiveBuilder;
import com.typesafe.config.Config;
import com.typesafe.config.ConfigFactory;
import scala.PartialFunction;
import scala.concurrent.Await;
import static akka.pattern.Patterns.ask;
import scala.concurrent.duration.Duration;
import akka.testkit.TestProbe;

public class FaultHandlingTest {
    static ActorSystem system;
    Duration timeout = Duration.create(5, SECONDS);

    @BeforeClass
    public static void start() {
        system = ActorSystem.create("FaultHandlingTest", config);
    }

    @AfterClass
    public static void cleanup() {
        JavaTestKit.shutdownActorSystem(system);
        system = null;
    }

    @Test
    public void mustEmploySupervisorStrategy() throws Exception {
        // code here
    }
}

```

```
}

```

Let us create actors:

```
Props superprops = Props.create(Supervisor.class);
ActorRef supervisor = system.actorOf(superprops, "supervisor");
ActorRef child = (ActorRef) Await.result(ask(supervisor,
    Props.create(Child.class), 5000), timeout);

```

The first test shall demonstrate the Resume directive, so we try it out by setting some non-initial state in the actor and have it fail:

```
child.tell(42, ActorRef.noSender());
assert Await.result(ask(child, "get", 5000), timeout).equals(42);
child.tell(new ArithmeticException(), ActorRef.noSender());
assert Await.result(ask(child, "get", 5000), timeout).equals(42);

```

As you can see the value 42 survives the fault handling directive. Now, if we change the failure to a more serious `NullPointerException`, that will no longer be the case:

```
child.tell(new NullPointerException(), ActorRef.noSender());
assert Await.result(ask(child, "get", 5000), timeout).equals(0);

```

And finally in case of the fatal `IllegalArgumentException` the child will be terminated by the supervisor:

```
final TestProbe probe = new TestProbe(system);
probe.watch(child);
child.tell(new IllegalArgumentException(), ActorRef.noSender());
probe.expectMsgClass(Terminated.class);

```

Up to now the supervisor was completely unaffected by the child's failure, because the directives set did handle it. In case of an Exception, this is not true anymore and the supervisor escalates the failure.

```
child = (ActorRef) Await.result(ask(supervisor,
    Props.create(Child.class), 5000), timeout);
probe.watch(child);
assert Await.result(ask(child, "get", 5000), timeout).equals(0);
child.tell(new Exception(), ActorRef.noSender());
probe.expectMsgClass(Terminated.class);

```

The supervisor itself is supervised by the top-level actor provided by the `ActorSystem`, which has the default policy to restart in case of all `Exception` cases (with the notable exceptions of `ActorInitializationException` and `ActorKilledException`). Since the default directive in case of a restart is to kill all children, we expected our poor child not to survive this failure.

In case this is not desired (which depends on the use case), we need to use a different supervisor which overrides this behavior.

```
public class Supervisor2 extends AbstractActor {

    private static SupervisorStrategy strategy =
        new OneForOneStrategy(10, Duration.create("1 minute"), DeciderBuilder.
            match(ArithmeticException.class, e -> resume()).
            match(NullPointerException.class, e -> restart()).
            match(IllegalArgumentException.class, e -> stop()).
            matchAny(o -> escalate()).build());

    @Override
    public SupervisorStrategy supervisorStrategy() {
        return strategy;
    }

    public Supervisor2() {

```

```

receive(ReceiveBuilder.
  match(Props.class, props -> {
    sender().tell(context().actorOf(props), self());
  }).build()
);

@Override
public void preRestart(Throwable cause, Option<Object> msg) {
  // do not kill all children, which is the default here
}
}

```

With this parent, the child survives the escalated restart, as demonstrated in the last test:

```

superprops = Props.create(Supervisor2.class);
supervisor = system.actorOf(superprops);
child = (ActorRef) Await.result(ask(supervisor,
  Props.create(Child.class), 5000), timeout);
child.tell(23, ActorRef.noSender());
assert Await.result(ask(child, "get", 5000), timeout).equals(23);
child.tell(new Exception(), ActorRef.noSender());
assert Await.result(ask(child, "get", 5000), timeout).equals(0);

```

4.3 FSM (Java with Lambda Support)

4.3.1 Overview

The FSM (Finite State Machine) is available as an abstract base class that implements an akka Actor and is best described in the [Erlang design principles](#)

A FSM can be described as a set of relations of the form:

State(S) x Event(E) -> Actions (A), State(S')

These relations are interpreted as meaning:

If we are in state S and the event E occurs, we should perform the actions A and make a transition to the state S'.

Warning: The Java with lambda support part of Akka is marked as “**experimental**” as of its introduction in Akka 2.3.0. We will continue to improve this API based on our users’ feedback, which implies that while we try to keep incompatible changes to a minimum, but the binary compatibility guarantee for maintenance releases does not apply to the akka.actor.AbstractFSM, related classes and the akka.japi.pf package.

4.3.2 A Simple Example

To demonstrate most of the features of the AbstractFSM class, consider an actor which shall receive and queue messages while they arrive in a burst and send them on after the burst ended or a flush request is received.

First, consider all of the below to use these import statements:

```

import akka.actor.AbstractFSM;
import akka.actor.ActorRef;
import akka.japi.pf.UnitMatch;
import java.util.Arrays;
import java.util.LinkedList;
import java.util.List;
import scala.concurrent.duration.Duration;

```

The contract of our “Buncher” actor is that it accepts or produces the following messages:

```
public final class SetTarget {
    private final ActorRef ref;

    public SetTarget(ActorRef ref) {
        this.ref = ref;
    }

    public ActorRef getRef() {
        return ref;
    }
    // boilerplate ...
}

public final class Queue {
    private final Object obj;

    public Queue(Object obj) {
        this.obj = obj;
    }

    public Object getObj() {
        return obj;
    }
    // boilerplate ...
}

public final class Batch {
    private final List<Object> list;

    public Batch(List<Object> list) {
        this.list = list;
    }

    public List<Object> getList() {
        return list;
    }
    // boilerplate ...
}

public enum Flush {
    Flush
}
```

`SetTarget` is needed for starting it up, setting the destination for the `Batches` to be passed on; `Queue` will add to the internal queue while `Flush` will mark the end of a burst.

The actor can be in two states: no message queued (aka `Idle`) or some message queued (aka `Active`). The states and the state data is defined like this:

```
// states
enum State {
    Idle, Active
}

// state data
interface Data {
}

enum Uninitialized implements Data {
    Uninitialized
}
```

```
final class Todo implements Data {
    private final ActorRef target;
    private final List<Object> queue;

    public Todo(ActorRef target, List<Object> queue) {
        this.target = target;
        this.queue = queue;
    }

    public ActorRef getTarget() {
        return target;
    }

    public List<Object> getQueue() {
        return queue;
    }
    // boilerplate ...
}
```

The actor starts out in the idle state. Once a message arrives it will go to the active state and stay there as long as messages keep arriving and no flush is requested. The internal state data of the actor is made up of the target actor reference to send the batches to and the actual queue of messages.

Now let's take a look at the skeleton for our FSM actor:

```
public class Buncher extends AbstractFSM<State, Data> {
    {
        startWith(Idle, Uninitialized);

        when(Idle,
            matchEvent(SetTarget.class, Uninitialized.class,
                (setTarget, uninitialized) ->
                    stay().using(new Todo(setTarget.getRef(), new LinkedList<>()))));

        // transition elided ...

        when(Active, Duration.create(1, "second"),
            matchEvent(Arrays.asList(Flush.class, StateTimeout()), Todo.class,
                (event, todo) -> goTo(Idle).using(todo.copy(new LinkedList<>()))));

        // unhandled elided ...

        initialize();
    }
}
```

The basic strategy is to declare the actor, by inheriting the `AbstractFSM` class and specifying the possible states and data values as type parameters. Within the body of the actor a DSL is used for declaring the state machine:

- `startWith` defines the initial state and initial data
- then there is one `when(<state>) { ... }` declaration per state to be handled (could potentially be multiple ones, the passed `PartialFunction` will be concatenated using `orElse`)
- finally starting it up using `initialize`, which performs the transition into the initial state and sets up timers (if required).

In this case, we start out in the `Idle` and `Uninitialized` state, where only the `SetTarget()` message is handled; `stay` prepares to end this event's processing for not leaving the current state, while the `using` modifier makes the FSM replace the internal state (which is `Uninitialized` at this point) with a fresh `Todo()` object containing the target actor reference. The `Active` state has a state timeout declared, which means that if no message is received for 1 second, a `FSM.StateTimeout` message will be generated. This has the same effect as receiving the `Flush` command in this case, namely to transition back into the `Idle` state and resetting the internal queue to the empty vector. But how do messages get queued? Since this shall work identically in

both states, we make use of the fact that any event which is not handled by the `when()` block is passed to the `whenUnhandled()` block:

```
whenUnhandled(
  matchEvent(Queue.class, Todo.class,
    (queue, todo) -> goTo(Active).using(todo.addElement(queue.getObj()))).
  anyEvent((event, state) -> {
    log().warning("received unhandled request {} in state {}/{}",
      event, stateName(), state);
    return stay();
  }));
```

The first case handled here is adding `Queue()` requests to the internal queue and going to the `Active` state (this does the obvious thing of staying in the `Active` state if already there), but only if the FSM data are not `Uninitialized` when the `Queue()` event is received. Otherwise—and in all other non-handled cases—the second case just logs a warning and does not change the internal state.

The only missing piece is where the `Batches` are actually sent to the target, for which we use the `onTransition` mechanism: you can declare multiple such blocks and all of them will be tried for matching behavior in case a state transition occurs (i.e. only when the state actually changes).

```
onTransition(
  matchState(Active, Idle, () -> {
    // reuse this matcher
    final UnitMatch<Data> m = UnitMatch.create(
      matchData(Todo.class,
        todo -> todo.getTarget().tell(new Batch(todo.getQueue()), self()));
    m.match(stateData());
  }).
  state(Idle, Active, () -> { /* Do something here */ }));
```

The transition callback is a partial function which takes as input a pair of states—the current and the next state. During the state change, the old state data is available via `stateData` as shown, and the new state data would be available as `nextStateData`.

To verify that this buncher actually works, it is quite easy to write a test using the *akka-testkit*, here using JUnit as an example:

```
public class BuncherTest {

  static ActorSystem system;

  @BeforeClass
  public static void setup() {
    system = ActorSystem.create("BuncherTest");
  }

  @AfterClass
  public static void tearDown() {
    JavaTestKit.shutdownActorSystem(system);
    system = null;
  }

  @Test
  public void testBuncherActorBatchesCorrectly() {
    new JavaTestKit(system) {{
      final ActorRef buncher =
        system.actorOf(Props.create(Buncher.class));
      final ActorRef probe = getRef();

      buncher.tell(new SetTarget(probe), probe);
      buncher.tell(new Queue(42), probe);
      buncher.tell(new Queue(43), probe);
      LinkedList<Object> list1 = new LinkedList<>();
```

```

list1.add(42);
list1.add(43);
expectMsgEquals(new Batch(list1));
buncher.tell(new Queue(44), probe);
buncher.tell(Flush, probe);
buncher.tell(new Queue(45), probe);
LinkedList<Object> list2 = new LinkedList<>();
list2.add(44);
expectMsgEquals(new Batch(list2));
LinkedList<Object> list3 = new LinkedList<>();
list3.add(45);
expectMsgEquals(new Batch(list3));
system.stop(buncher);
});
}

@Test
public void testBuncherActorDoesntBatchUninitialized() {
    new JavaTestKit(system) {{
        final ActorRef buncher =
            system.actorOf(Props.create(Buncher.class));
        final ActorRef probe = getRef();

        buncher.tell(new Queue(42), probe);
        expectNoMsg();
        system.stop(buncher);
    }};
}
}

```

4.3.3 Reference

The AbstractFSM Class

The `AbstractFSM` abstract class is the base class used to implement an FSM. It implements `Actor` since an `Actor` is created to drive the FSM.

```

public class Buncher extends AbstractFSM<State, Data> {
    {
        // fsm body ...
    }
}

```

Note: The `AbstractFSM` class defines a `receive` method which handles internal messages and passes everything else through to the FSM logic (according to the current state). When overriding the `receive` method, keep in mind that e.g. state timeout handling depends on actually passing the messages through the FSM logic.

The `AbstractFSM` class takes two type parameters:

1. the supertype of all state names, usually an enum,
 2. the type of the state data which are tracked by the `AbstractFSM` module itself.
-

Note: The state data together with the state name describe the internal state of the state machine; if you stick to this scheme and do not add mutable fields to the FSM class you have the advantage of making all changes of the internal state explicit in a few well-known places.

Defining States

A state is defined by one or more invocations of the method

```
when(<name>[, stateTimeout = <timeout>])(stateFunction).
```

The given name must be an object which is type-compatible with the first type parameter given to the `AbstractFSM` class. This object is used as a hash key, so you must ensure that it properly implements `equals` and `hashCode`; in particular it must not be mutable. The easiest fit for these requirements are case objects.

If the `stateTimeout` parameter is given, then all transitions into this state, including staying, receive this timeout by default. Initiating the transition with an explicit timeout may be used to override this default, see [Initiating Transitions](#) for more information. The state timeout of any state may be changed during action processing with `setStateTimeout(state, duration)`. This enables runtime configuration e.g. via external message.

The `stateFunction` argument is a `PartialFunction[Event, State]`, which is conveniently given using the state function builder syntax as demonstrated below:

```
when(Idle,
  matchEvent(SetTarget.class, Uninitialized.class,
    (setTarget, uninitialized) ->
      stay().using(new Todo(setTarget.getRef(), new LinkedList<>())));
```

Warning: It is required that you define handlers for each of the possible FSM states, otherwise there will be failures when trying to switch to undeclared states.

It is recommended practice to declare the states as an enum and then verify that there is a `when` clause for each of the states. If you want to leave the handling of a state “unhandled” (more below), it still needs to be declared like this:

```
when(SomeState, AbstractFSM.NullFunction());
```

Defining the Initial State

Each FSM needs a starting point, which is declared using

```
startWith(state, data[, timeout])
```

The optionally given timeout argument overrides any specification given for the desired initial state. If you want to cancel a default timeout, use `Duration.Inf`.

Unhandled Events

If a state doesn’t handle a received event a warning is logged. If you want to do something else in this case you can specify that with `whenUnhandled(stateFunction)`:

```
whenUnhandled(
  matchEvent(X.class, (x, data) -> {
    log().info("Received unhandled event: " + x);
    return stay();
  }).
  anyEvent((event, data) -> {
    log().warning("Received unknown event: " + event);
    return goTo(Error);
  }));
}
```

Within this handler the state of the FSM may be queried using the `stateName` method.

IMPORTANT: This handler is not stacked, meaning that each invocation of `whenUnhandled` replaces the previously installed handler.

Initiating Transitions

The result of any `stateFunction` must be a definition of the next state unless terminating the FSM, which is described in [Termination from Inside](#). The state definition can either be the current state, as described by the `stay` directive, or it is a different state as given by `goto(state)`. The resulting object allows further qualification by way of the modifiers described in the following:

- `forMax(duration)`

This modifier sets a state timeout on the next state. This means that a timer is started which upon expiry sends a `StateTimeout` message to the FSM. This timer is canceled upon reception of any other message in the meantime; you can rely on the fact that the `StateTimeout` message will not be processed after an intervening message.

This modifier can also be used to override any default timeout which is specified for the target state. If you want to cancel the default timeout, use `Duration.Inf`.

- `using(data)`

This modifier replaces the old state data with the new data given. If you follow the advice *above*, this is the only place where internal state data are ever modified.

- `replying(msg)`

This modifier sends a reply to the currently processed message and otherwise does not modify the state transition.

All modifiers can be chained to achieve a nice and concise description:

```
when(SomeState, matchAnyEvent((msg, data) -> {
  return goto(Processing).using(newData).
    forMax(Duration.create(5, SECONDS)).replying(WillDo);
}));
```

The parentheses are not actually needed in all cases, but they visually distinguish between modifiers and their arguments and therefore make the code even more pleasant to read for foreigners.

Note: Please note that the `return` statement may not be used in `when` blocks or similar; this is a Scala restriction. Either refactor your code using `if () ... else ...` or move it into a method definition.

Monitoring Transitions

Transitions occur “between states” conceptually, which means after any actions you have put into the event handling block; this is obvious since the next state is only defined by the value returned by the event handling logic. You do not need to worry about the exact order with respect to setting the internal state variable, as everything within the FSM actor is running single-threaded anyway.

Internal Monitoring

Up to this point, the FSM DSL has been centered on states and events. The dual view is to describe it as a series of transitions. This is enabled by the method

```
onTransition(handler)
```

which associates actions with a transition instead of with a state and event. The handler is a partial function which takes a pair of states as input; no resulting state is needed as it is not possible to modify the transition in progress.

```
onTransition(
  matchState(Active, Idle, () -> setTimer("timeout",
    Tick, Duration.create(1, SECONDS), true)).
  state(Active, null, () -> cancelTimer("timeout")).
  state(null, Idle, (f, t) -> log().info("entering Idle from " + f));
```

It is also possible to pass a function object accepting two states to `onTransition`, in case your transition handling logic is implemented as a method:

```
public void handler(StateType from, StateType to) {
    // handle transition here
}

onTransition(this::handler);
```

The handlers registered with this method are stacked, so you can intersperse `onTransition` blocks with `when` blocks as suits your design. It should be noted, however, that *all handlers will be invoked for each transition*, not only the first matching one. This is designed specifically so you can put all transition handling for a certain aspect into one place without having to worry about earlier declarations shadowing later ones; the actions are still executed in declaration order, though.

Note: This kind of internal monitoring may be used to structure your FSM according to transitions, so that for example the cancellation of a timer upon leaving a certain state cannot be forgot when adding new target states.

External Monitoring

External actors may be registered to be notified of state transitions by sending a message `SubscribeTransitionCallback(actorRef)`. The named actor will be sent a `CurrentState(self, stateName)` message immediately and will receive `Transition(actorRef, oldState, newState)` messages whenever a new state is reached. External monitors may be unregistered by sending `UnsubscribeTransitionCallback(actorRef)` to the FSM actor.

Stopping a listener without unregistering will not remove the listener from the subscription list; use `UnsubscribeTransitionCallback` before stopping the listener.

Timers

Besides state timeouts, FSM manages timers identified by `String` names. You may set a timer using

```
setTimer(name, msg, interval, repeat)
```

where `msg` is the message object which will be sent after the duration `interval` has elapsed. If `repeat` is `true`, then the timer is scheduled at fixed rate given by the `interval` parameter. Any existing timer with the same name will automatically be canceled before adding the new timer.

Timers may be canceled using

```
cancelTimer(name)
```

which is guaranteed to work immediately, meaning that the scheduled message will not be processed after this call even if the timer already fired and queued it. The status of any timer may be inquired with

```
isTimerActive(name)
```

These named timers complement state timeouts because they are not affected by intervening reception of other messages.

Termination from Inside

The FSM is stopped by specifying the result state as

```
stop([reason[, data]])
```

The reason must be one of `Normal` (which is the default), `Shutdown` or `Failure(reason)`, and the second argument may be given to change the state data which is available during termination handling.

Note: It should be noted that `stop` does not abort the actions and stop the FSM immediately. The stop action must be returned from the event handler in the same way as a state transition (but note that the `return` statement may not be used within a `when` block).

```
when(Error, matchEventEquals("stop", (event, data) -> {
  // do cleanup ...
  return stop();
}));
```

You can use `onTermination(handler)` to specify custom code that is executed when the FSM is stopped. The handler is a partial function which takes a `StopEvent(reason, stateName, stateData)` as argument:

```
onTermination(
  matchStop(Normal(),
    (state, data) -> { /* Do something here */}).
  stop(Shutdown(),
    (state, data) -> { /* Do something here */}).
  stop(Failure.class,
    (reason, state, data) -> { /* Do something here */}));
```

As for the `whenUnhandled` case, this handler is not stacked, so each invocation of `onTermination` replaces the previously installed handler.

Termination from Outside

When an `ActorRef` associated to a FSM is stopped using the `stop` method, its `postStop` hook will be executed. The default implementation by the `AbstractFSM` class is to execute the `onTermination` handler if that is prepared to handle a `StopEvent(Shutdown, ...)`.

Warning: In case you override `postStop` and want to have your `onTermination` handler called, do not forget to call `super.postStop`.

4.3.4 Testing and Debugging Finite State Machines

During development and for trouble shooting FSMs need care just as any other actor. There are specialized tools available as described in *TestFSMRef* and in the following.

Event Tracing

The setting `akka.actor.debug.fsm` in *Configuration* enables logging of an event trace by `LoggingFSM` instances:

```
public class MyFSM extends AbstractLoggingFSM<StateType, Data> {
  // body elided ...
}
```

This FSM will log at `DEBUG` level:

- all processed events, including `StateTimeout` and scheduled timer messages
- every setting and cancellation of named timers
- all state transitions

Life cycle changes and special messages can be logged as described for *Actors*.

Rolling Event Log

The `AbstractLoggingFSM` class adds one more feature to the FSM: a rolling event log which may be used during debugging (for tracing how the FSM entered a certain failure state) or for other creative uses:

```
public class MyFSM extends AbstractLoggingFSM<StateType, Data> {
  @Override
  public int logDepth() { return 12; }
  {
    onTermination(
      matchStop(Failure.class, (reason, state, data) -> {
        String lastEvents = getLog().mkString("\n\t");
        log().warning("Failure in state " + state + " with data " + data + "\n" +
          "Events leading up to this point:\n\t" + lastEvents);
      })
    );
    //...
  }
}
```

The `logDepth` defaults to zero, which turns off the event log.

Warning: The log buffer is allocated during actor creation, which is why the configuration is done using a virtual method call. If you want to override with a `val`, make sure that its initialization happens before the initializer of `LoggingFSM` runs, and do not change the value returned by `logDepth` after the buffer has been allocated.

The contents of the event log are available using method `getLog`, which returns an `IndexedSeq[LogEntry]` where the oldest entry is at index zero.

4.3.5 Examples

A bigger FSM example contrasted with Actor's `become/unbecome` can be found in the [Typesafe Activator](#) template named [Akka FSM in Scala](#)

4.4 Persistence (Java with Lambda Support)

Akka persistence enables stateful actors to persist their internal state so that it can be recovered when an actor is started, restarted after a JVM crash or by a supervisor, or migrated in a cluster. The key concept behind Akka persistence is that only changes to an actor's internal state are persisted but never its current state directly (except for optional snapshots). These changes are only ever appended to storage, nothing is ever mutated, which allows for very high transaction rates and efficient replication. Stateful actors are recovered by replaying stored changes to these actors from which they can rebuild internal state. This can be either the full history of changes or starting from a snapshot which can dramatically reduce recovery times. Akka persistence also provides point-to-point communication with at-least-once message delivery semantics.

Akka persistence is inspired by the [eventsourced](#) library. It follows the same concepts and architecture of [eventsourced](#) but significantly differs on API and implementation level.

4.4.1 Dependencies

Akka persistence is a separate jar file. Make sure that you have the following dependency in your project:

```
<dependency>
  <groupId>com.typesafe.akka</groupId>
  <artifactId>akka-persistence_2.11</artifactId>
```

```
<version>2.4.0</version>
</dependency>
```

Akka persistence extension comes with few built-in persistence plugins, including in-memory heap based journal, local file-system based snapshot-store and LevelDB based journal.

LevelDB based plugins will require the following additional dependency declaration:

```
<dependency>
  <groupId>org.iq80.leveldb</groupId>
  <artifactId>leveldb</artifactId>
  <version>0.7</version>
</dependency>
<dependency>
  <groupId>org.fusesource.leveldbjni</groupId>
  <artifactId>leveldbjni-all</artifactId>
  <version>1.8</version>
</dependency>
```

4.4.2 Architecture

- *AbstractPersistentActor*: Is a persistent, stateful actor. It is able to persist events to a journal and can react to them in a thread-safe manner. It can be used to implement both *command* as well as *event sourced* actors. When a persistent actor is started or restarted, journaled messages are replayed to that actor, so that it can recover internal state from these messages.
- *AbstractPersistentView*: A view is a persistent, stateful actor that receives journaled messages that have been written by another persistent actor. A view itself does not journal new messages, instead, it updates internal state only from a persistent actor's replicated message stream.
- *AbstractPersistentActorAtLeastOnceDelivery*: To send messages with at-least-once delivery semantics to destinations, also in case of sender and receiver JVM crashes.
- *AsyncWriteJournal*: A journal stores the sequence of messages sent to a persistent actor. An application can control which messages are journaled and which are received by the persistent actor without being journaled. The storage backend of a journal is pluggable. Persistence extension comes with a "leveldb" journal plugin, which writes to the local filesystem, and replicated journals are available as [Community plugins](#).
- *Snapshot store*: A snapshot store persists snapshots of a persistent actor's or a view's internal state. Snapshots are used for optimizing recovery times. The storage backend of a snapshot store is pluggable. Persistence extension comes with a "local" snapshot storage plugin, which writes to the local filesystem, and replicated snapshot stores are available as [Community plugins](#).
- *Event sourcing*. Based on the building blocks described above, Akka persistence provides abstractions for the development of event sourced applications (see section [Event sourcing](#))

4.4.3 Event sourcing

The basic idea behind [Event Sourcing](#) is quite simple. A persistent actor receives a (non-persistent) command which is first validated if it can be applied to the current state. Here, validation can mean anything, from simple inspection of a command message's fields up to a conversation with several external services, for example. If validation succeeds, events are generated from the command, representing the effect of the command. These events are then persisted and, after successful persistence, used to change the actor's state. When the persistent actor needs to be recovered, only the persisted events are replayed of which we know that they can be successfully applied. In other words, events cannot fail when being replayed to a persistent actor, in contrast to commands. Event sourced actors may of course also process commands that do not change application state, such as query commands, for example.

Akka persistence supports event sourcing with the `AbstractPersistentActor` abstract class. An actor that extends this class uses the `persist` method to persist and handle events. The behavior of an

`AbstractPersistentActor` is defined by implementing `receiveRecover` and `receiveCommand`. This is demonstrated in the following example.

```
import akka.actor.ActorRef;
import akka.actor.ActorSystem;
import akka.actor.Props;
import akka.japi.pf.ReceiveBuilder;
import akka.persistence.AbstractPersistentActor;
import akka.persistence.SnapshotOffer;
import scala.Function;
import scala.runtime.BoxedUnit;

import java.io.Serializable;
import java.util.ArrayList;

import static java.util.Arrays.asList;

class Cmd implements Serializable {
    private static final long serialVersionUID = 1L;
    private final String data;

    public Cmd(String data) {
        this.data = data;
    }

    public String getData() {
        return data;
    }
}

class Evt implements Serializable {
    private static final long serialVersionUID = 1L;
    private final String data;

    public Evt(String data) {
        this.data = data;
    }

    public String getData() {
        return data;
    }
}

class ExampleState implements Serializable {
    private static final long serialVersionUID = 1L;
    private final ArrayList<String> events;

    public ExampleState() {
        this(new ArrayList<>());
    }

    public ExampleState(ArrayList<String> events) {
        this.events = events;
    }

    public ExampleState copy() {
        return new ExampleState(new ArrayList<>(events));
    }

    public void update(Evt evt) {
        events.add(evt.getData());
    }
}
```

```

public int size() {
    return events.size();
}

@Override
public String toString() {
    return events.toString();
}
}

class ExamplePersistentActor extends AbstractPersistentActor {

    private ExampleState state = new ExampleState();

    public int getNumEvents() {
        return state.size();
    }

    @Override
    public String persistenceId() { return "sample-id-1"; }

    @Override
    public PartialFunction<Object, BoxedUnit> receiveRecover() {
        return ReceiveBuilder.
            match(Evt.class, state::update).
            match(SnapshotOffer.class, ss -> state = (ExampleState) ss.snapshot()).build();
    }

    @Override
    public PartialFunction<Object, BoxedUnit> receiveCommand() {
        return ReceiveBuilder.
            match(Cmd.class, c -> {
                final String data = c.getData();
                final Evt evt1 = new Evt(data + "-" + getNumEvents());
                final Evt evt2 = new Evt(data + "-" + (getNumEvents() + 1));
                persistAll(asList(evt1, evt2), (Evt evt) -> {
                    state.update(evt);
                    if (evt.equals(evt2)) {
                        context().system().eventStream().publish(evt);
                    }
                });
            }).
            match(String.class, s -> s.equals("snap"), s -> saveSnapshot(state.copy())).
            match(String.class, s -> s.equals("print"), s -> System.out.println(state)).
            build();
    }
}

```

The example defines two data types, `Cmd` and `Evt` to represent commands and events, respectively. The state of the `ExamplePersistentActor` is a list of persisted event data contained in `ExampleState`.

The persistent actor's `receiveRecover` method defines how state is updated during recovery by handling `Evt` and `SnapshotOffer` messages. The persistent actor's `receiveCommand` method is a command handler. In this example, a command is handled by generating two events which are then persisted and handled. Events are persisted by calling `persist` with an event (or a sequence of events) as first argument and an event handler as second argument.

The `persist` method persists events asynchronously and the event handler is executed for successfully persisted events. Successfully persisted events are internally sent back to the persistent actor as individual messages that trigger event handler executions. An event handler may close over persistent actor state and mutate it. The sender of a persisted event is the sender of the corresponding command. This allows event handlers to reply to the sender of a command (not shown).

The main responsibility of an event handler is changing persistent actor state using event data and notifying others about successful state changes by publishing events.

When persisting events with `persist` it is guaranteed that the persistent actor will not receive further commands between the `persist` call and the execution(s) of the associated event handler. This also holds for multiple `persist` calls in context of a single command. Incoming messages are *stashed* until the `persist` is completed. You should be careful to not send more messages to a persistent actor than it can keep up with, otherwise the number of stashed messages will grow. It can be wise to protect against *OutOfMemoryError* by defining a maximum stash capacity in the mailbox configuration:

```
akka.actor.default-mailbox.stash-capacity=10000
```

If the stash capacity is exceeded for an actor the stashed messages are discarded and a `MessageQueueAppendFailedException` is thrown, causing actor restart if default supervision strategy is used.

Note that the stash capacity is per actor. If you have many persistent actors, e.g. when using cluster sharding, you may need to define a small stash capacity to ensure that the total number of stashed messages in the system don't consume too much memory.

If persistence of an event fails, `onPersistFailure` will be invoked (logging the error by default) and the actor will unconditionally be stopped. If persistence of an event is rejected before it is stored, e.g. due to serialization error, `onPersistRejected` will be invoked (logging a warning by default) and the actor continues with next message.

The easiest way to run this example yourself is to download [Typesafe Activator](#) and open the tutorial named [Akka Persistence Samples in Java with Lambdas](#). It contains instructions on how to run the `PersistentActorExample`.

Note: It's also possible to switch between different command handlers during normal processing and recovery with `context().become()` and `context().unbecome()`. To get the actor into the same state after recovery you need to take special care to perform the same state transitions with `become` and `unbecome` in the `receiveRecover` method as you would have done in the command handler. Note that when using `become` from `receiveRecover` it will still only use the `receiveRecover` behavior when replaying the events. When replay is completed it will use the new behavior.

Identifiers

A persistent actor must have an identifier that doesn't change across different actor incarnations. The identifier must be defined with the `persistenceId` method.

```
@Override
public String persistenceId() {
    return "my-stable-persistence-id";
}
```

Recovery

By default, a persistent actor is automatically recovered on start and on restart by replaying journaled messages. New messages sent to a persistent actor during recovery do not interfere with replayed messages. New messages will only be received by a persistent actor after recovery completes.

Note: Accessing the `sender()` for replayed messages will always result in a `deadLetters` reference, as the original sender is presumed to be long gone. If you indeed have to notify an actor during recovery in the future, store its `ActorPath` explicitly in your persisted events.

Recovery customization

Applications may also customise how recovery is performed by returning a customised `Recovery` object in the recovery method of a `AbstractPersistentActor`, for example setting an upper bound to the replay, which allows the actor to be replayed to a certain point “in the past” instead to its most up to date state:

```
@Override
public Recovery recovery() {
    return Recovery.create(457L);
}
```

Recovery can be disabled by returning `Recovery.none` in the recovery method of a `PersistentActor`:

```
@Override
public Recovery recovery() {
    return Recovery.none();
}
```

Recovery status

A persistent actor can query its own recovery status via the methods

```
public boolean recoveryRunning();
public boolean recoveryFinished();
```

Sometimes there is a need for performing additional initialization when the recovery has completed, before processing any other message sent to the persistent actor. The persistent actor will receive a special `RecoveryCompleted` message right after recovery and before any other received messages.

```
class MyPersistentActor5 extends AbstractPersistentActor {

    @Override public String persistenceId() {
        return "my-stable-persistence-id";
    }

    @Override public PartialFunction<Object, BoxedUnit> receiveRecover() {
        return ReceiveBuilder.
            match(RecoveryCompleted.class, r -> {
                // perform init after recovery, before any other messages
                // ...
            }).
            match(String.class, this::handleEvent).build();
    }

    @Override public PartialFunction<Object, BoxedUnit> receiveCommand() {
        return ReceiveBuilder.
            match(String.class, s -> s.equals("cmd"),
                s -> persist("evt", this::handleEvent)).build();
    }

    private void handleEvent(String event) {
        // update state
        // ...
    }
}
```

If there is a problem with recovering the state of the actor from the journal, `onRecoveryFailure` is called (logging the error by default) and the actor will be stopped.

Relaxed local consistency requirements and high throughput use-cases

If faced with relaxed local consistency requirements and high throughput demands sometimes `PersistentActor` and its `persist` may not be enough in terms of consuming incoming `Commands` at a high rate, because it has to wait until all `Events` related to a given `Command` are processed in order to start processing the next `Command`. While this abstraction is very useful for most cases, sometimes you may be faced with relaxed requirements about consistency – for example you may want to process commands as fast as you can, assuming that `Event` will eventually be persisted and handled properly in the background and retroactively reacting to persistence failures if needed.

The `persistAsync` method provides a tool for implementing high-throughput persistent actors. It will *not* stash incoming `Commands` while the `Journal` is still working on persisting and/or user code is executing event callbacks.

In the below example, the event callbacks may be called “at any time”, even after the next `Command` has been processed. The ordering between events is still guaranteed (“`evt-b-1`” will be sent after “`evt-a-2`”, which will be sent after “`evt-a-1`” etc.).

```
class MyPersistentActor extends AbstractPersistentActor {

  @Override public String persistenceId() {
    return "my-stable-persistence-id";
  }

  private void handleCommand(String c) {
    sender().tell(c, self());

    persistAsync(String.format("evt-%s-1", c), e -> {
      sender().tell(e, self());
    });
    persistAsync(String.format("evt-%s-2", c), e -> {
      sender().tell(e, self());
    });
  }

  @Override public PartialFunction<Object, BoxedUnit> receiveRecover() {
    return ReceiveBuilder.
      match(String.class, this::handleCommand).build();
  }

  @Override public PartialFunction<Object, BoxedUnit> receiveCommand() {
    return ReceiveBuilder.
      match(String.class, this::handleCommand).build();
  }
}
```

Note: In order to implement the pattern known as “*command sourcing*” simply call `persistAsync` on all incoming messages right away, and handle them in the callback.

Warning: The callback will not be invoked if the actor is restarted (or stopped) in between the call to `persistAsync` and the journal has confirmed the write.

Deferring actions until preceding persist handlers have executed

Sometimes when working with `persistAsync` you may find that it would be nice to define some actions in terms of “happens-after the previous `persistAsync` handlers have been invoked”. `PersistentActor` provides an utility method called `deferAsync`, which works similarly to `persistAsync` yet does not persist the passed in event. It is recommended to use it for *read* operations, and actions which do not have corresponding events in your domain model.

Using this method is very similar to the `persist` family of methods, yet it does **not** persist the passed in event. It will be kept in memory and used when invoking the handler.

```
class MyPersistentActor extends AbstractPersistentActor {

  @Override public String persistenceId() {
    return "my-stable-persistence-id";
  }

  private void handleCommand(String c) {
    persistAsync(String.format("evt-%s-1", c), e -> {
      sender().tell(e, self());
    });
    persistAsync(String.format("evt-%s-2", c), e -> {
      sender().tell(e, self());
    });

    deferAsync(String.format("evt-%s-3", c), e -> {
      sender().tell(e, self());
    });
  }

  @Override public PartialFunction<Object, BoxedUnit> receiveRecover() {
    return ReceiveBuilder.
      match(String.class, this::handleCommand).build();
  }

  @Override public PartialFunction<Object, BoxedUnit> receiveCommand() {
    return ReceiveBuilder.
      match(String.class, this::handleCommand).build();
  }
}
```

Notice that the `sender()` is **safe** to access in the handler callback, and will be pointing to the original sender of the command for which this `deferAsync` handler was called.

```
final ActorRef persistentActor = system.actorOf(Props.create(MyPersistentActor.class));
persistentActor.tell("a", sender);
persistentActor.tell("b", sender);

// order of received messages:
// a
// b
// evt-a-1
// evt-a-2
// evt-a-3
// evt-b-1
// evt-b-2
// evt-b-3
```

Warning: The callback will not be invoked if the actor is restarted (or stopped) in between the call to `deferAsync` and the journal has processed and confirmed all preceding writes.

Nested persist calls

It is possible to call `persist` and `persistAsync` inside their respective callback blocks and they will properly retain both the thread safety (including the right value of `sender()`) as well as stashing guarantees.

In general it is encouraged to create command handlers which do not need to resort to nested event persisting, however there are situations where it may be useful. It is important to understand the ordering of callback execution in those situations, as well as their implication on the stashing behaviour (that `persist()` enforces). In the following example two `persist` calls are issued, and each of them issues another `persist` inside its callback:

```

@Override public PartialFunction<Object, BoxedUnit> receiveRecover() {
    final Procedure<String> replyToSender = event -> sender().tell(event, self());

    return ReceiveBuilder
        .match(String.class, msg -> {
            persist(String.format("%s-outer-1", msg), event -> {
                sender().tell(event, self());
                persist(String.format("%s-inner-1", event), replyToSender);
            });

            persist(String.format("%s-outer-2", msg), event -> {
                sender().tell(event, self());
                persist(String.format("%s-inner-2", event), replyToSender);
            });
        })
        .build();
}

```

When sending two commands to this `PersistentActor`, the `persist` handlers will be executed in the following order:

```

persistentActor.tell("a", ActorRef.noSender());
persistentActor.tell("b", ActorRef.noSender());

// order of received messages:
// a
// a-outer-1
// a-outer-2
// a-inner-1
// a-inner-2
// and only then process "b"
// b
// b-outer-1
// b-outer-2
// b-inner-1
// b-inner-2

```

First the “outer layer” of `persist` calls is issued and their callbacks applied, after these have successfully completed the inner callbacks will be invoked (once the events they are persisting have been confirmed to be persisted by the journal). And only after all these handlers have been successfully invoked, the next command will be delivered to the persistent Actor. In other words, the stashing of incoming commands that is guaranteed by initially calling `persist()` on the outer layer is extended until all nested `persist` callbacks have been handled.

It is also possible to nest `persistAsync` calls, using the same pattern:

```

@Override public PartialFunction<Object, BoxedUnit> receiveRecover() {
    final Procedure<String> replyToSender = event -> sender().tell(event, self());

    return ReceiveBuilder
        .match(String.class, msg -> {
            persistAsync(String.format("%s-outer-1", msg), event -> {
                sender().tell(event, self());
                persistAsync(String.format("%s-inner-1", event), replyToSender);
            });

            persistAsync(String.format("%s-outer-2", msg), event -> {
                sender().tell(event, self());
                persistAsync(String.format("%s-inner-1", event), replyToSender);
            });
        })
        .build();
}

```

In this case no stashing is happening, yet the events are still persisted and callbacks executed in the expected order:

```

persistentActor.tell("a", self());
persistentActor.tell("b", self());

// order of received messages:
// a
// b
// a-outer-1
// a-outer-2
// b-outer-1
// b-outer-2
// a-inner-1
// a-inner-2
// b-inner-1
// b-inner-2

// which can be seen as the following causal relationship:
// a -> a-outer-1 -> a-outer-2 -> a-inner-1 -> a-inner-2
// b -> b-outer-1 -> b-outer-2 -> b-inner-1 -> b-inner-2

```

While it is possible to nest mixed `persist` and `persistAsync` with keeping their respective semantics it is not a recommended practice as it may lead to overly complex nesting.

Failures

If persistence of an event fails, `onPersistFailure` will be invoked (logging the error by default) and the actor will unconditionally be stopped.

The reason that it cannot resume when `persist` fails is that it is unknown if the event was actually persisted or not, and therefore it is in an inconsistent state. Restarting on persistent failures will most likely fail anyway, since the journal is probably unavailable. It is better to stop the actor and after a back-off timeout start it again. The `akka.pattern.BackoffSupervisor` actor is provided to support such restarts.

```

@Override
public void preStart() throws Exception {
    final Props childProps = Props.create(MyPersistentActor1.class);
    final Props props = BackoffSupervisor.props(
        childProps,
        "myActor",
        Duration.create(3, TimeUnit.SECONDS),
        Duration.create(30, TimeUnit.SECONDS),
        0.2);
    context().actorOf(props, "mySupervisor");
    super.preStart();
}

```

If persistence of an event is rejected before it is stored, e.g. due to serialization error, `onPersistRejected` will be invoked (logging a warning by default) and the actor continues with next message.

If there is a problem with recovering the state of the actor from the journal when the actor is started, `onRecoveryFailure` is called (logging the error by default) and the actor will be stopped.

Atomic writes

Each event is of course stored atomically, but it is also possible to store several events atomically by using the `persistAll` or `persistAllAsync` method. That means that all events passed to that method are stored or none of them are stored if there is an error.

The recovery of a persistent actor will therefore never be done partially with only a subset of events persisted by `persistAll`.

Some journals may not support atomic writes of several events and they will then reject the `persistAll` command, i.e. `onPersistRejected` is called with an exception (typically `UnsupportedOperationException`).

Batch writes

To optimize throughput, a persistent actor internally batches events to be stored under high load before writing them to the journal (as a single batch). The batch size dynamically grows from 1 under low and moderate loads to a configurable maximum size (default is 200) under high load. When using `persistAsync` this increases the maximum throughput dramatically.

```
akka.persistence.journal.leveldb.max-message-batch-size = 200
```

A new batch write is triggered by a persistent actor as soon as a batch reaches the maximum size or if the journal completed writing the previous batch. Batch writes are never timer-based which keeps latencies at a minimum.

Message deletion

It is possible to delete all messages (journaled by a single persistent actor) up to a specified sequence number, persistent actors may call the `deleteMessages` method.

Deleting messages in event sourcing based applications is typically either not used at all, or used in conjunction with *snapshotting*, i.e. after a snapshot has been successfully stored, a `deleteMessages(toSequenceNr)` up until the sequence number of the data held by that snapshot can be issued, to safely delete the previous events, while still having access to the accumulated state during replays - by loading the snapshot.

The result of the `deleteMessages` request is signaled to the persistent actor with a `DeleteMessagesSuccess` message if the delete was successful or a `DeleteMessagesFailure` message if it failed.

Persistence status handling

Persisting, deleting and replaying messages can either succeed or fail.

Method	Success	Failure / Rejection	After failure handler invoked
<code>persist/persistAsync</code> <code>onPersistRejected</code>	persist handler invoked	<code>onPersistFailure</code> No automatic actions.	Actor is stopped.
<code>recovery</code>	<code>RecoverySuccess</code>	<code>onRecoveryFailure</code>	Actor is stopped.
<code>deleteMessages</code>	<code>DeleteMessagesSuccess</code>	<code>DeleteMessagesFailure</code>	No automatic actions.

The most important operations (`persist` and `recovery`) have failure handlers modelled as explicit callbacks which the user can override in the `PersistentActor`. The default implementations of these handlers emit a log message (`error` for `persist/recovery` failures, and `warning` for others), logging the failure cause and information about which message caused the failure.

For critical failures, such as `recovery` or `persisting` events failing, the persistent actor will be stopped after the failure handler is invoked. This is because if the underlying journal implementation is signalling persistence failures it is most likely either failing completely or overloaded and restarting right-away and trying to `persist` the event again will most likely not help the journal recover – as it would likely cause a [Thundering herd problem](#), as many persistent actors would restart and try to `persist` their events again. Instead, using a `BackoffSupervisor` (as described in [Failures](#)) which implements an exponential-backoff strategy which allows for more breathing room for the journal to recover between restarts of the persistent actor.

Note: Journal implementations may choose to implement a retry mechanisms, e.g. such that only after a write fails N number of times a persistence failure is signalled back to the user. In other words, once a journal returns a failure, it is considered *fatal* by Akka Persistence, and the persistent actor which caused the failure will be stopped.

Check the documentation of the journal implementation you are using for details if/how it is using this technique.

Safely shutting down persistent actors

Special care should be given when when shutting down persistent actors from the outside. With normal Actors it is often acceptable to use the special *PoisonPill* message to signal to an Actor that it should stop itself once it receives this message – in fact this message is handled automatically by Akka, leaving the target actor no way to refuse stopping itself when given a poison pill.

This can be dangerous when used with `PersistentActor` due to the fact that incoming commands are *stashed* while the persistent actor is awaiting confirmation from the Journal that events have been written when `persist()` was used. Since the incoming commands will be drained from the Actor's mailbox and put into it's internal stash while awaiting the confirmation (thus, before calling the `persist` handlers) the Actor **may receive and (auto)handle the PoisonPill before it processes the other messages which have been put into its stash**, causing a pre-mature shutdown of the Actor.

Warning: Consider using explicit shut-down messages instead of `PoisonPill` when working with persistent actors.

The example below highlights how messages arrive in the Actor's mailbox and how they interact with it's internal stashing mechanism when `persist()` is used, notice the early stop behaviour that occurs when `PoisonPill` is used:

```
final class Shutdown {
}

class MyPersistentActor extends AbstractPersistentActor {
  @Override
  public String persistenceId() {
    return "some-persistence-id";
  }

  @Override
  public PartialFunction<Object, BoxedUnit> receiveCommand() {
    return ReceiveBuilder
      .match(Shutdown.class, shutdown -> {
        context().stop(self());
      })
      .match(String.class, msg -> {
        System.out.println(msg);
        persist("handle-" + msg, e -> System.out.println(e));
      })
      .build();
  }

  @Override
  public PartialFunction<Object, BoxedUnit> receiveRecover() {
    return ReceiveBuilder.matchAny(any -> {}).build();
  }
}
```

```
// UN-SAFE, due to PersistentActor's command stashing:
persistentActor.tell("a", ActorRef.noSender());
persistentActor.tell("b", ActorRef.noSender());
persistentActor.tell(PoisonPill.getInstance(), ActorRef.noSender());
// order of received messages:
// a
// # b arrives at mailbox, stashing;          internal-stash = [b]
// # PoisonPill arrives at mailbox, stashing; internal-stash = [b, Shutdown]
// PoisonPill is an AutoReceivedMessage, is handled automatically
```

```
// !! stop !!
// Actor is stopped without handling `b` nor the `a` handler!

// SAFE:
persistentActor.tell("a", ActorRef.noSender());
persistentActor.tell("b", ActorRef.noSender());
persistentActor.tell(new Shutdown(), ActorRef.noSender());
// order of received messages:
// a
// # b arrives at mailbox, stashing;           internal-stash = [b]
// # Shutdown arrives at mailbox, stashing;    internal-stash = [b, Shutdown]
// handle-a
// # unstashing;                               internal-stash = [Shutdown]
// b
// handle-b
// # unstashing;                               internal-stash = []
// Shutdown
// -- stop --
```

4.4.4 Persistent Views

Warning: `AbstractPersistentView` is deprecated. Use *Persistence Query* instead. The corresponding query type is `EventsByPersistenceId`. There are several alternatives for connecting the Source to an actor corresponding to a previous `UntypedPersistentView` actor:

- `Sink.actorRef` is simple, but has the disadvantage that there is no back-pressure signal from the destination actor, i.e. if the actor is not consuming the messages fast enough the mailbox of the actor will grow
- `mapAsync` combined with *Ask: Send-And-Receive-Future* is almost as simple with the advantage of back-pressure being propagated all the way
- `ActorSubscriber` in case you need more fine grained control

The consuming actor may be a plain `AbstractActor` or an `AbstractPersistentActor` if it needs to store its own state (e.g. `fromSequenceNr` offset).

Persistent views can be implemented by extending the `AbstractView` abstract class, implement the `persistenceId` method and setting the “initial behavior” in the constructor by calling the `receive` method.

```
class MyView extends AbstractPersistentView {
  @Override public String persistenceId() { return "some-persistence-id"; }
  @Override public String viewId() { return "some-persistence-id-view"; }

  public MyView() {
    receive(ReceiveBuilder.
      match(Object.class, p -> isPersistent(), persistent -> {
        // ...
      }).build()
    );
  }
}
```

The `persistenceId` identifies the persistent actor from which the view receives journaled messages. It is not necessary the referenced persistent actor is actually running. Views read messages from a persistent actor’s journal directly. When a persistent actor is started later and begins to write new messages, the corresponding view is updated automatically, by default.

It is possible to determine if a message was sent from the Journal or from another actor in user-land by calling the `isPersistent` method. Having that said, very often you don’t need this information at all and can simply apply the same logic to both cases (skip the `if isPersistent` check).

Updates

The default update interval of all persistent views of an actor system is configurable:

```
akka.persistence.view.auto-update-interval = 5s
```

`AbstractPersistentView` implementation classes may also override the `autoUpdateInterval` method to return a custom update interval for a specific view class or view instance. Applications may also trigger additional updates at any time by sending a view an `Update` message.

```
final ActorRef view = system.actorOf(Props.create(MyView.class));
view.tell(Update.create(true), null);
```

If the `await` parameter is set to `true`, messages that follow the `Update` request are processed when the incremental message replay, triggered by that update request, completed. If set to `false` (default), messages following the update request may interleave with the replayed message stream. Automated updates always run with `await = false`.

Automated updates of all persistent views of an actor system can be turned off by configuration:

```
akka.persistence.view.auto-update = off
```

Implementation classes may override the configured default value by overriding the `autoUpdate` method. To limit the number of replayed messages per update request, applications can configure a custom `akka.persistence.view.auto-update-replay-max` value or override the `autoUpdateReplayMax` method. The number of replayed messages for manual updates can be limited with the `replayMax` parameter of the `Update` message.

Recovery

Initial recovery of persistent views works in the very same way as for a persistent actor (i.e. by sending a `Recover` message to self). The maximum number of replayed messages during initial recovery is determined by `autoUpdateReplayMax`. Further possibilities to customize initial recovery are explained in section [Recovery](#).

Identifiers

A persistent view must have an identifier that doesn't change across different actor incarnations. The identifier must be defined with the `viewId` method.

The `viewId` must differ from the referenced `persistenceId`, unless *Snapshots* of a view and its persistent actor shall be shared (which is what applications usually do not want).

4.4.5 Snapshots

Snapshots can dramatically reduce recovery times of persistent actors and views. The following discusses snapshots in context of persistent actors but this is also applicable to persistent views.

Persistent actor can save snapshots of internal state by calling the `saveSnapshot` method. If saving of a snapshot succeeds, the persistent actor receives a `SaveSnapshotSuccess` message, otherwise a `SaveSnapshotFailure` message

```
private Object state;

@Override public PartialFunction<Object, BoxedUnit> receiveCommand() {
    return ReceiveBuilder.
        match(String.class, s -> s.equals("snap"),
            s -> saveSnapshot(state)).
        match(SaveSnapshotSuccess.class, ss -> {
            SnapshotMetadata metadata = ss.metadata();
            // ...
        });
}
```

```

    }).
    match(SaveSnapshotFailure.class, sf -> {
        SnapshotMetadata metadata = sf.metadata();
        // ...
    }).build();
}

```

During recovery, the persistent actor is offered a previously saved snapshot via a `SnapshotOffer` message from which it can initialize internal state.

```

private Object state;

@Override public PartialFunction<Object, BoxedUnit> receiveRecover() {
    return ReceiveBuilder.
        match(SnapshotOffer.class, s -> {
            state = s.snapshot();
            // ...
        }).
        match(String.class, s -> { /* ... */ }).build();
}

```

The replayed messages that follow the `SnapshotOffer` message, if any, are younger than the offered snapshot. They finally recover the persistent actor to its current (i.e. latest) state.

In general, a persistent actor is only offered a snapshot if that persistent actor has previously saved one or more snapshots and at least one of these snapshots matches the `SnapshotSelectionCriteria` that can be specified for recovery.

```

@Override
public Recovery recovery() {
    return Recovery.create(
        SnapshotSelectionCriteria
            .create(457L, System.currentTimeMillis()));
}

```

If not specified, they default to `SnapshotSelectionCriteria.latest()` which selects the latest (= youngest) snapshot. To disable snapshot-based recovery, applications should use `SnapshotSelectionCriteria.none()`. A recovery where no saved snapshot matches the specified `SnapshotSelectionCriteria` will replay all journaled messages.

Note: In order to use snapshots a default snapshot-store (`akka.persistence.snapshot-store.plugin`) must be configured, or the persistent actor can pick a snapshot store explicitly by overriding `String snapshotPluginId()`.

Since it is acceptable for some applications to not use any snapshotting, it is legal to not configure a snapshot store, however Akka will log a warning message when this situation is detected and then continue to operate until an actor tries to store a snapshot, at which point the the operation will fail (by replying with an `SaveSnapshotFailure` for example).

Note that *Cluster Sharding* is using snapshots, so if you use Cluster Sharding you need to define a snapshot store plugin.

Snapshot deletion

A persistent actor can delete individual snapshots by calling the `deleteSnapshot` method with the sequence number of when the snapshot was taken.

To bulk-delete a range of snapshots matching `SnapshotSelectionCriteria`, persistent actors should use the `deleteSnapshots` method.

Snapshot status handling

Saving or deleting snapshots can either succeed or fail – this information is reported back to the persistent actor via status messages as illustrated in the following table.

Method	Success	Failure message
<code>saveSnapshot (Any)</code>	<code>SaveSnapshotSuccess</code>	<code>SaveSnapshotFailure</code>
<code>deleteSnapshot (Long)</code>	<code>DeleteSnapshotSuccess</code>	<code>DeleteSnapshotFailure</code>
<code>deleteSnapshots (SnapshotSelectionCriteria)</code>	<code>DeleteSnapshotsSuccess</code>	<code>DeleteSnapshotsFailure</code>

4.4.6 At-Least-Once Delivery

To send messages with at-least-once delivery semantics to destinations you can extend the `AbstractPersistentActorWithAtLeastOnceDelivery` class instead of `AbstractPersistentActor` on the sending side. It takes care of re-sending messages when they have not been confirmed within a configurable timeout.

The state of the sending actor, including which messages that have been sent and still not been confirmed by the recipient, must be persistent so that it can survive a crash of the sending actor or JVM. The `AbstractPersistentActorWithAtLeastOnceDelivery` class does not persist anything by itself. It is your responsibility to persist the intent that a message is sent and that a confirmation has been received.

Note: At-least-once delivery implies that original message send order is not always preserved and the destination may receive duplicate messages. That means that the semantics do not match those of a normal `ActorRef` send operation:

- it is not at-most-once delivery
- message order for the same sender–receiver pair is not preserved due to possible resends
- after a crash and restart of the destination messages are still delivered—to the new actor incarnation

These semantics is similar to what an `ActorPath` represents (see *actor-lifecycle-scala*), therefore you need to supply a path and not a reference when delivering messages. The messages are sent to the path with an actor selection.

Use the `deliver` method to send a message to a destination. Call the `confirmDelivery` method when the destination has replied with a confirmation message.

Relationship between `deliver` and `confirmDelivery`

To send messages to the destination path, use the `deliver` method after you have persisted the intent to send the message.

The destination actor must send back a confirmation message. When the sending actor receives this confirmation message you should persist the fact that the message was delivered successfully and then call the `confirmDelivery` method.

If the persistent actor is not currently recovering, the `deliver` method will send the message to the destination actor. When recovering, messages will be buffered until they have been confirmed using `confirmDelivery`. Once recovery has completed, if there are outstanding messages that have not been confirmed (during the message replay), the persistent actor will resend these before sending any other messages.

`Deliver` requires a `deliveryIdToMessage` function to pass the provided `deliveryId` into the message so that correlation between `deliver` and `confirmDelivery` is possible. The `deliveryId` must do the round trip. Upon receipt of the message, destination actor will send the same “`deliveryId`” wrapped in a confirmation message back to the sender. The sender will then use it to call `confirmDelivery` method to complete delivery routine.

```

class Msg implements Serializable {
    private static final long serialVersionUID = 1L;
    public final long deliveryId;
    public final String s;

    public Msg(long deliveryId, String s) {
        this.deliveryId = deliveryId;
        this.s = s;
    }
}

class Confirm implements Serializable {
    private static final long serialVersionUID = 1L;
    public final long deliveryId;

    public Confirm(long deliveryId) {
        this.deliveryId = deliveryId;
    }
}

class MsgSent implements Serializable {
    private static final long serialVersionUID = 1L;
    public final String s;

    public MsgSent(String s) {
        this.s = s;
    }
}

class MsgConfirmed implements Serializable {
    private static final long serialVersionUID = 1L;
    public final long deliveryId;

    public MsgConfirmed(long deliveryId) {
        this.deliveryId = deliveryId;
    }
}

class MyPersistentActor extends AbstractPersistentActorWithAtLeastOnceDelivery {
    private final ActorSelection destination;

    public MyPersistentActor(ActorSelection destination) {
        this.destination = destination;
    }

    @Override public String persistenceId() {
        return "persistence-id";
    }

    @Override
    public PartialFunction<Object, BoxedUnit> receiveCommand() {
        return ReceiveBuilder.
            match(String.class, s -> {
                persist(new MsgSent(s), evt -> updateState(evt));
            }).
            match(Confirm.class, confirm -> {
                persist(new MsgConfirmed(confirm.deliveryId), evt -> updateState(evt));
            }).
            build();
    }

    @Override
    public PartialFunction<Object, BoxedUnit> receiveRecover() {

```

```

return ReceiveBuilder.
    match(Object.class, evt -> updateState(evt)).build();
}

void updateState(Object event) {
    if (event instanceof MsgSent) {
        final MsgSent evt = (MsgSent) event;
        deliver(destination, deliveryId -> new Msg(deliveryId, evt.s));
    } else if (event instanceof MsgConfirmed) {
        final MsgConfirmed evt = (MsgConfirmed) event;
        confirmDelivery(evt.deliveryId);
    }
}
}

class MyDestination extends AbstractActor {
    public MyDestination() {
        receive(ReceiveBuilder.
            match(Msg.class, msg -> {
                // ...
                sender().tell(new Confirm(msg.deliveryId), self());
            }).build()
        );
    }
}
}

```

The `deliveryId` generated by the persistence module is a strictly monotonically increasing sequence number without gaps. The same sequence is used for all destinations of the actor, i.e. when sending to multiple destinations the destinations will see gaps in the sequence. It is not possible to use custom `deliveryId`. However, you can send a custom correlation identifier in the message to the destination. You must then retain a mapping between the internal `deliveryId` (passed into the `deliveryIdToMessage` function) and your custom correlation id (passed into the message). You can do this by storing such mapping in a `Map(correlationId -> deliveryId)` from which you can retrieve the `deliveryId` to be passed into the `confirmDelivery` method once the receiver of your message has replied with your custom correlation id.

The `AbstractPersistentActorWithAtLeastOnceDelivery` class has a state consisting of unconfirmed messages and a sequence number. It does not store this state itself. You must persist events corresponding to the `deliver` and `confirmDelivery` invocations from your `PersistentActor` so that the state can be restored by calling the same methods during the recovery phase of the `PersistentActor`. Sometimes these events can be derived from other business level events, and sometimes you must create separate events. During recovery calls to `deliver` will not send out the message, but it will be sent later if no matching `confirmDelivery` was performed.

Support for snapshots is provided by `getDeliverySnapshot` and `setDeliverySnapshot`. The `AtLeastOnceDeliverySnapshot` contains the full delivery state, including unconfirmed messages. If you need a custom snapshot for other parts of the actor state you must also include the `AtLeastOnceDeliverySnapshot`. It is serialized using protobuf with the ordinary Akka serialization mechanism. It is easiest to include the bytes of the `AtLeastOnceDeliverySnapshot` as a blob in your custom snapshot.

The interval between redelivery attempts is defined by the `redeliverInterval` method. The default value can be configured with the `akka.persistence.at-least-once-delivery.redeliver-interval` configuration key. The method can be overridden by implementation classes to return non-default values.

After a number of delivery attempts a `AtLeastOnceDelivery.UnconfirmedWarning` message will be sent to `self`. The re-sending will still continue, but you can choose to call `confirmDelivery` to cancel the re-sending. The number of delivery attempts before emitting the warning is defined by the `warnAfterNumberOfUnconfirmedAttempts` method. The default value can be configured with the `akka.persistence.at-least-once-delivery.warn-after-number-of-unconfirmed-attempts` configuration key. The method can be overridden by implementation classes to return non-default values.

The `AbstractPersistentActorWithAtLeastOnceDelivery` class holds messages in mem-

ory until their successful delivery has been confirmed. The limit of maximum number of unconfirmed messages that the actor is allowed to hold in memory is defined by the `maxUnconfirmedMessages` method. If this limit is exceeded the `deliver` method will not accept more messages and it will throw `AtLeastOnceDelivery.MaxUnconfirmedMessagesExceededException`. The default value can be configured with the `akka.persistence.at-least-once-delivery.max-unconfirmed-messages` configuration key. The method can be overridden by implementation classes to return non-default values.

4.4.7 Event Adapters

In long running projects using event sourcing sometimes the need arises to detach the data model from the domain model completely.

Event Adapters help in situations where:

- **Version Migrations** – existing events stored in *Version 1* should be “upcasted” to a new *Version 2* representation, and the process of doing so involves actual code, not just changes on the serialization layer. For these scenarios the `toJournal` function is usually an identity function, however the `fromJournal` is implemented as `v1.Event=>v2.Event`, performing the necessary mapping inside the `fromJournal` method. This technique is sometimes referred to as “upcasting” in other CQRS libraries.
- **Separating Domain and Data models** – thanks to `EventAdapters` it is possible to completely separate the domain model from the model used to persist data in the Journals. For example one may want to use case classes in the domain model, however persist their protocol-buffer (or any other binary serialization format) counter-parts to the Journal. A simple `toJournal:MyModel=>MyDataModel` and `fromJournal:MyDataModel=>MyModel` adapter can be used to implement this feature.
- **Journal Specialized Data Types** – exposing data types understood by the underlying Journal, for example for data stores which understand JSON it is possible to write an `EventAdapter toJournal:Any=>JSON` such that the Journal can *directly* store the json instead of serializing the object to its binary representation.

Implementing an `EventAdapter` is rather straightforward:

```
class MyEventAdapter implements EventAdapter {
  @Override
  public String manifest(Object event) {
    return ""; // if no manifest needed, return ""
  }

  @Override
  public Object toJournal(Object event) {
    return event; // identity
  }

  @Override
  public EventSeq fromJournal(Object event, String manifest) {
    return EventSeq.single(event); // identity
  }
}
```

Then in order for it to be used on events coming to and from the journal you must bind it using the below configuration syntax:

```
akka.persistence.journal {
  inmem {
    event-adapters {
      tagging          = "docs.persistence.MyTaggingEventAdapter"
      user-upcasting   = "docs.persistence.UserUpcastingEventAdapter"
      item-upcasting   = "docs.persistence.ItemUpcastingEventAdapter"
    }

    event-adapter-bindings {
      "docs.persistence.Item"           = tagging
      "docs.persistence.TaggedEvent"    = tagging
    }
  }
}
```

```

    "docs.persistence.v1.Event" = [user-upcasting, item-upcasting]
  }
}
}

```

It is possible to bind multiple adapters to one class *for recovery*, in which case the `fromJournal` methods of all bound adapters will be applied to a given matching event (in order of definition in the configuration). Since each adapter may return from 0 to *n* adapted events (called as `EventSeq`), each adapter can investigate the event and if it should indeed adapt it return the adapted event(s) for it, other adapters which do not have anything to contribute during this adaptation simply return `EventSeq.empty`. The adapted events are then delivered in-order to the `PersistentActor` during replay.

Note: For more advanced schema evolution techniques refer to the *persistence-schema-evolution-scala* documentation.

4.4.8 Persistent FSM

`AbstractPersistentFSM` handles the incoming messages in an FSM like fashion. Its internal state is persisted as a sequence of changes, later referred to as domain events. Relationship between incoming messages, FSM's states and transitions, persistence of domain events is defined by a DSL.

Warning: `AbstractPersistentFSM` is marked as “**experimental**” as of its introduction in Akka 2.4.0. We will continue to improve this API based on our users’ feedback, which implies that while we try to keep incompatible changes to a minimum the binary compatibility guarantee for maintenance releases does not apply to the contents of the *classes related to ‘AbstractPersistentFSM’*.

A Simple Example

To demonstrate the features of the `AbstractPersistentFSM`, consider an actor which represents a Web store customer. The contract of our “`WebStoreCustomerFSMACTOR`” is that it accepts the following commands:

```

public static final class AddItem implements Command {
    private final Item item;

    public AddItem(Item item) {
        this.item = item;
    }

    public Item getItem() {
        return item;
    }
}

public enum Buy implements Command {INSTANCE}

public enum Leave implements Command {INSTANCE}

public enum GetCurrentCart implements Command {INSTANCE}

```

`AddItem` sent when the customer adds an item to a shopping cart `Buy` - when the customer finishes the purchase `Leave` - when the customer leaves the store without purchasing anything `GetCurrentCart` allows to query the current state of customer’s shopping cart

The customer can be in one of the following states:

```

enum UserState implements PersistentFSM.FSMState {
    LOOKING_AROUND("Looking Around"),
    SHOPPING("Shopping"),

```

```

    INACTIVE("Inactive"),
    PAID("Paid");

    private final String stateIdentifier;

    UserState(String stateIdentifier) {
        this.stateIdentifier = stateIdentifier;
    }

    @Override
    public String identifier() {
        return stateIdentifier;
    }
}

```

LookingAround customer is browsing the site, but hasn't added anything to the shopping cart Shopping customer has recently added items to the shopping cart Inactive customer has items in the shopping cart, but hasn't added anything recently, Paid customer has purchased the items

Note: AbstractPersistentFSM states must inherit from PersistentFSM.FSMState and implement the String identifier() method. This is required in order to simplify the serialization of FSM states. String identifiers should be unique!

Customer's actions are "recorded" as a sequence of "domain events", which are persisted. Those events are replayed on actor's start in order to restore the latest customer's state:

```

public static final class ItemAdded implements DomainEvent {
    private final Item item;

    public ItemAdded(Item item) {
        this.item = item;
    }

    public Item getItem() {
        return item;
    }
}

public enum OrderExecuted implements DomainEvent {INSTANCE}

public enum OrderDiscarded implements DomainEvent {INSTANCE}

```

Customer state data represents the items in customer's shopping cart:

```

static class ShoppingCart {
    private final List<Item> items = new ArrayList<>();

    public List<Item> getItems() {
        return Collections.unmodifiableList(items);
    }

    void addItem(Item item) {
        items.add(item);
    }

    void empty() {
        items.clear();
    }
}

static class Item implements Serializable {
    private final String id;
}

```

```

private final String name;
private final float price;

Item(String id, String name, float price) {
    this.id = id;
    this.name = name;
    this.price = price;
}

public String getId() {
    return id;
}

public float getPrice() {
    return price;
}

public String getName() {
    return name;
}

@Override
public String toString() {
    return String.format("Item{id=%s, name=%s, price=%s}", id, price, name);
}

@Override
public boolean equals(Object o) {
    if (this == o) return true;
    if (o == null || getClass() != o.getClass()) return false;

    Item item = (Item) o;

    return item.price == price && id.equals(item.id) && name.equals(item.name);
}
}

```

Here is how everything is wired together:

```

startWith(UserState.LOOKING_AROUND, new ShoppingCart());

when(UserState.LOOKING_AROUND,
    matchEvent(AddItem.class,
        (event, data) ->
            goTo(UserState.SHOPPING).applying(new ItemAdded(event.getItem()))
                .forMax(Duration.create(1, TimeUnit.SECONDS))
    )
    .event(GetCurrentCart.class, (event, data) -> stay().replying(data))
);

when(UserState.SHOPPING,
    matchEvent(AddItem.class,
        (event, data) ->
            stay().applying(new ItemAdded(event.getItem()))
                .forMax(Duration.create(1, TimeUnit.SECONDS))
    )
    .event(Buy.class,
        (event, data) ->
            goTo(UserState.PAID).applying(OrderExecuted.INSTANCE)
                .andThen(exec(cart ->
                    reportActor.tell(new PurchaseWasMade(cart.getItems()), self()))
            )
    )
    .event(Leave.class,
        (event, data) ->

```

```

        stop().applying(OrderDiscarded.INSTANCE)
            .andThen(exec(cart ->
                reportActor.tell(ShoppingCardDiscarded.INSTANCE, self())
            ))
        .event(GetCurrentCart.class, (event, data) -> stay().replying(data))
        .event(StateTimeout$.class,
            (event, data) ->
                goTo(UserState.INACTIVE).forMax(Duration.create(2, TimeUnit.SECONDS)))
    );

when(UserState.INACTIVE,
    matchEvent(AddItem.class,
        (event, data) ->
            goTo(UserState.SHOPPING).applying(new ItemAdded(event.getItem()))
                .forMax(Duration.create(1, TimeUnit.SECONDS)))
        .event(GetCurrentCart.class, (event, data) -> stay().replying(data))
        .event(StateTimeout$.class,
            (event, data) ->
                stop().applying(OrderDiscarded.INSTANCE)
                    .andThen(exec(cart ->
                        reportActor.tell(ShoppingCardDiscarded.INSTANCE, self())
                    ))
    ))
);

when(UserState.PAID,
    matchEvent(Leave.class, (event, data) -> stop())
    .event(GetCurrentCart.class, (event, data) -> stay().replying(data))
);

initialize();

```

Note: State data can only be modified directly on initialization. Later it's modified only as a result of applying domain events. Override the `applyEvent` method to define how state data is affected by domain events, see the example below

```

@Override
public ShoppingCart applyEvent(DomainEvent event, ShoppingCart currentData) {
    if (event instanceof ItemAdded) {
        currentData.addItem(((ItemAdded) event).getItem());
        return currentData;
    } else if (event instanceof OrderExecuted) {
        return currentData;
    } else if (event instanceof OrderDiscarded) {
        currentData.empty();
        return currentData;
    }
    throw new RuntimeException("Unhandled");
}

```

4.4.9 Storage plugins

Storage backends for journals and snapshot stores are pluggable in the Akka persistence extension.

Directory of persistence journal and snapshot store plugins is available at the Akka Community Projects page, see [Community plugins](#)

Plugins can be selected either by “default”, for all persistent actors and views, or “individually”, when persistent actor or view defines it's own set of plugins.

When persistent actor or view does NOT override `journalPluginId` and `snapshotPluginId` methods, persistence extension will use “default” journal and snapshot-store plugins configured in the `reference.conf`:

```
akka.persistence.journal.plugin = ""
akka.persistence.snapshot-store.plugin = ""
```

However, these entries are provided as empty “”, and require explicit user configuration via override in the user `application.conf`. For an example of journal plugin which writes messages to LevelDB see [Local LevelDB journal](#). For an example of snapshot store plugin which writes snapshots as individual files to the local filesystem see [Local snapshot store](#).

Applications can provide their own plugins by implementing a plugin API and activate them by configuration. Plugin development requires the following imports:

```
import akka.dispatch.Futures;
import akka.persistence.*;
import akka.persistence.journal.japi.*;
import akka.persistence.snapshot.japi.*;
```

Journal plugin API

A journal plugin extends `AsyncWriteJournal`.

`AsyncWriteJournal` is an actor and the methods to be implemented are:

```
/**
 * Java API, Plugin API: asynchronously writes a batch ('Iterable') of
 * persistent messages to the journal.
 *
 * The batch is only for performance reasons, i.e. all messages don't have to
 * be written atomically. Higher throughput can typically be achieved by using
 * batch inserts of many records compared to inserting records one-by-one, but
 * this aspect depends on the underlying data store and a journal
 * implementation can implement it as efficient as possible. Journals should
 * aim to persist events in-order for a given 'persistenceId' as otherwise in
 * case of a failure, the persistent state may be end up being inconsistent.
 *
 * Each 'AtomicWrite' message contains the single 'PersistentRepr' that
 * corresponds to the event that was passed to the 'persist' method of the
 * 'PersistentActor', or it contains several 'PersistentRepr' that corresponds
 * to the events that were passed to the 'persistAll' method of the
 * 'PersistentActor'. All 'PersistentRepr' of the 'AtomicWrite' must be
 * written to the data store atomically, i.e. all or none must be stored. If
 * the journal (data store) cannot support atomic writes of multiple events it
 * should reject such writes with an 'Optional' with an
 * 'UnsupportedOperationException' describing the issue. This limitation
 * should also be documented by the journal plugin.
 *
 * If there are failures when storing any of the messages in the batch the
 * returned 'Future' must be completed with failure. The 'Future' must only be
 * completed with success when all messages in the batch have been confirmed
 * to be stored successfully, i.e. they will be readable, and visible, in a
 * subsequent replay. If there is uncertainty about if the messages were
 * stored or not the 'Future' must be completed with failure.
 *
 * Data store connection problems must be signaled by completing the 'Future'
 * with failure.
 *
 * The journal can also signal that it rejects individual messages
 * ('AtomicWrite') by the returned
 * 'Iterable<Optional<Exception>>'. The returned 'Iterable' must
 * have as many elements as the input 'messages' 'Iterable'. Each 'Optional'
 * element signals if the corresponding 'AtomicWrite' is rejected or not, with
```

```

* an exception describing the problem. Rejecting a message means it was not
* stored, i.e. it must not be included in a later replay. Rejecting a message
* is typically done before attempting to store it, e.g. because of
* serialization error.
*
* Data store connection problems must not be signaled as rejections.
*
* Note that it is possible to reduce number of allocations by caching some
* result `Iterable` for the happy path, i.e. when no messages are rejected.
*/
Future<Iterable<Optional<Exception>>> doAsyncWriteMessages(Iterable<AtomicWrite> messages);

/**
 * Java API, Plugin API: synchronously deletes all persistent messages up to
 * `toSequenceNr`.
 *
 * @see AsyncRecoveryPlugin
 */
Future<Void> doAsyncDeleteMessagesTo(String persistenceId, long toSequenceNr);

```

If the storage backend API only supports synchronous, blocking writes, the methods should be implemented as:

```

@Override
public Future<Iterable<Optional<Exception>>> doAsyncWriteMessages(
    Iterable<AtomicWrite> messages) {
    try {
        Iterable<Optional<Exception>> result = new ArrayList<Optional<Exception>>();
        // blocking call here...
        // result.add(..)
        return Futures.successful(result);
    } catch (Exception e) {
        return Futures.failed(e);
    }
}

```

A journal plugin must also implement the methods defined in `AsyncRecovery` for replays and sequence number recovery:

```

/**
 * Java API, Plugin API: asynchronously replays persistent messages.
 * Implementations replay a message by calling `replayCallback`. The returned
 * future must be completed when all messages (matching the sequence number
 * bounds) have been replayed. The future must be completed with a failure if
 * any of the persistent messages could not be replayed.
 *
 * The `replayCallback` must also be called with messages that have been
 * marked as deleted. In this case a replayed message's `deleted` method must
 * return `true`.
 *
 * The `toSequenceNr` is the lowest of what was returned by
 * {@link #doAsyncReadHighestSequenceNr} and what the user specified as
 * recovery {@link akka.persistence.Recovery} parameter.
 *
 * @param persistenceId
 *     id of the persistent actor.
 * @param fromSequenceNr
 *     sequence number where replay should start (inclusive).
 * @param toSequenceNr
 *     sequence number where replay should end (inclusive).
 * @param max
 *     maximum number of messages to be replayed.
 * @param replayCallback
 *     called to replay a single message. Can be called from any thread.

```

```

*/
Future<Void> doAsyncReplayMessages(String persistenceId, long fromSequenceNr,
    long toSequenceNr, long max, Consumer<PersistentRepr> replayCallback);

/**
 * Java API, Plugin API: asynchronously reads the highest stored sequence
 * number for the given `persistenceId`. The persistent actor will use the
 * highest sequence number after recovery as the starting point when
 * persisting new events. This sequence number is also used as `toSequenceNr`
 * in subsequent call to [[#asyncReplayMessages]] unless the user has
 * specified a lower `toSequenceNr`.
 *
 * @param persistenceId
 *       id of the persistent actor.
 * @param fromSequenceNr
 *       hint where to start searching for the highest sequence number.
 */
Future<Long> doAsyncReadHighestSequenceNr(String persistenceId, long fromSequenceNr);

```

A journal plugin can be activated with the following minimal configuration:

```

# Path to the journal plugin to be used
akka.persistence.journal.plugin = "my-journal"

# My custom journal plugin
my-journal {
  # Class name of the plugin.
  class = "docs.persistence.MyJournal"
  # Dispatcher for the plugin actor.
  plugin-dispatcher = "akka.actor.default-dispatcher"
}

```

The specified plugin class must have a no-arg constructor. The plugin-dispatcher is the dispatcher used for the plugin actor. If not specified, it defaults to akka.persistence.dispatchers.default-plugin-dispatcher.

The journal plugin instance is an actor so the methods corresponding to requests from persistent actors are executed sequentially. It may delegate to asynchronous libraries, spawn futures, or delegate to other actors to achieve parallelism.

The journal plugin class must have a constructor without parameters or constructor with one com.typesafe.config.Config parameter. The plugin section of the actor system's config will be passed in the config constructor parameter.

Don't run journal tasks/futures on the system default dispatcher, since that might starve other tasks.

Snapshot store plugin API

A snapshot store plugin must extend the SnapshotStore actor and implement the following methods:

```

/**
 * Java API, Plugin API: asynchronously loads a snapshot.
 *
 * @param persistenceId
 *       id of the persistent actor.
 * @param criteria
 *       selection criteria for loading.
 */
Future<Optional<SelectedSnapshot>> doLoadAsync(String persistenceId,
    SnapshotSelectionCriteria criteria);

/**
 * Java API, Plugin API: asynchronously saves a snapshot.

```

```

*
* @param metadata
*       snapshot metadata.
* @param snapshot
*       snapshot.
*/
Future<Void> doSaveAsync(SnapshotMetadata metadata, Object snapshot);

/**
 * Java API, Plugin API: deletes the snapshot identified by 'metadata'.
 *
 * @param metadata
 *       snapshot metadata.
 */
Future<Void> doDeleteAsync(SnapshotMetadata metadata);

/**
 * Java API, Plugin API: deletes all snapshots matching 'criteria'.
 *
 * @param persistenceId
 *       id of the persistent actor.
 * @param criteria
 *       selection criteria for deleting.
 */
Future<Void> doDeleteAsync(String persistenceId, SnapshotSelectionCriteria criteria);

```

A snapshot store plugin can be activated with the following minimal configuration:

```

# Path to the snapshot store plugin to be used
akka.persistence.snapshot-store.plugin = "my-snapshot-store"

# My custom snapshot store plugin
my-snapshot-store {
  # Class name of the plugin.
  class = "docs.persistence.MySnapshotStore"
  # Dispatcher for the plugin actor.
  plugin-dispatcher = "akka.persistence.dispatchers.default-plugin-dispatcher"
}

```

The specified plugin class must have a no-arg constructor. The plugin-dispatcher is the dispatcher used for the plugin actor. If not specified, it defaults to `akka.persistence.dispatchers.default-plugin-dispatcher`.

The snapshot store instance is an actor so the methods corresponding to requests from persistent actors are executed sequentially. It may delegate to asynchronous libraries, spawn futures, or delegate to other actors to achieve parallelism.

The snapshot store plugin class must have a constructor without parameters or constructor with one `com.typesafe.config.Config` parameter. The plugin section of the actor system's config will be passed in the config constructor parameter.

Don't run snapshot store tasks/futures on the system default dispatcher, since that might starve other tasks.

4.4.10 Pre-packaged plugins

Local LevelDB journal

LevelDB journal plugin config entry is `akka.persistence.journal.leveldb` and it writes messages to a local LevelDB instance. Enable this plugin by defining config property:

```

# Path to the journal plugin to be used
akka.persistence.journal.plugin = "akka.persistence.journal.leveldb"

```

LevelDB based plugins will also require the following additional dependency declaration:

```
<dependency>
  <groupId>org.iq80.leveldb</groupId>
  <artifactId>leveldb</artifactId>
  <version>0.7</version>
</dependency>
<dependency>
  <groupId>org.fusesource.leveldbjni</groupId>
  <artifactId>leveldbjni-all</artifactId>
  <version>1.8</version>
</dependency>
```

The default location of the LevelDB files is a directory named `journal` in the current working directory. This location can be changed by configuration where the specified path can be relative or absolute:

```
akka.persistence.journal.leveldb.dir = "target/journal"
```

With this plugin, each actor system runs its own private LevelDB instance.

Shared LevelDB journal

A LevelDB instance can also be shared by multiple actor systems (on the same or on different nodes). This, for example, allows persistent actors to failover to a backup node and continue using the shared journal instance from the backup node.

Warning: A shared LevelDB instance is a single point of failure and should therefore only be used for testing purposes. Highly-available, replicated journal are available as [Community plugins](#).

A shared LevelDB instance is started by instantiating the `SharedLeveldbStore` actor.

```
final ActorRef store = system.actorOf(Props.create(SharedLeveldbStore.class), "store");
```

By default, the shared instance writes journaled messages to a local directory named `journal` in the current working directory. The storage location can be changed by configuration:

```
akka.persistence.journal.leveldb-shared.store.dir = "target/shared"
```

Actor systems that use a shared LevelDB store must activate the `akka.persistence.journal.leveldb-shared` plugin.

```
akka.persistence.journal.plugin = "akka.persistence.journal.leveldb-shared"
```

This plugin must be initialized by injecting the (remote) `SharedLeveldbStore` actor reference. Injection is done by calling the `SharedLeveldbJournal.setStore` method with the actor reference as argument.

```
class SharedStorageUsage extends UntypedActor {
  @Override
  public void preStart() throws Exception {
    String path = "akka.tcp://example@127.0.0.1:2552/user/store";
    ActorSelection selection = getContext().actorSelection(path);
    selection.tell(new Identify(1), getSelf());
  }

  @Override
  public void onReceive(Object message) throws Exception {
    if (message instanceof ActorIdentity) {
      ActorIdentity identity = (ActorIdentity) message;
      if (identity.correlationId().equals(1)) {
        ActorRef store = identity.getRef();
        if (store != null) {
          SharedLeveldbJournal.setStore(store, getContext().system());
        }
      }
    }
  }
}
```

```

    }
  }
}

```

Internal journal commands (sent by persistent actors) are buffered until injection completes. Injection is idempotent i.e. only the first injection is used.

Local snapshot store

Local snapshot store plugin config entry is `akka.persistence.snapshot-store.local` and it writes snapshot files to the local filesystem. Enable this plugin by defining config property:

```

# Path to the snapshot store plugin to be used
akka.persistence.snapshot-store.plugin = "akka.persistence.snapshot-store.local"

```

The default storage location is a directory named `snapshots` in the current working directory. This can be changed by configuration where the specified path can be relative or absolute:

```

akka.persistence.snapshot-store.local.dir = "target/snapshots"

```

Note that it is not mandatory to specify a snapshot store plugin. If you don't use snapshots you don't have to configure it.

4.4.11 Custom serialization

Serialization of snapshots and payloads of `Persistent` messages is configurable with Akka's *Serialization* infrastructure. For example, if an application wants to serialize

- payloads of type `MyPayload` with a custom `MyPayloadSerializer` and
- snapshots of type `MySnapshot` with a custom `MySnapshotSerializer`

it must add

```

akka.actor {
  serializers {
    my-payload = "docs.persistence.MyPayloadSerializer"
    my-snapshot = "docs.persistence.MySnapshotSerializer"
  }
  serialization-bindings {
    "docs.persistence.MyPayload" = my-payload
    "docs.persistence.MySnapshot" = my-snapshot
  }
}

```

to the application configuration. If not specified, a default serializer is used.

For more advanced schema evolution techniques refer to the *persistence-schema-evolution-scala* documentation.

4.4.12 Testing

When running tests with LevelDB default settings in `sbt`, make sure to set `fork := true` in your `sbt` project otherwise, you'll see an `UnsatisfiedLinkError`. Alternatively, you can switch to a LevelDB Java port by setting

```

akka.persistence.journal.leveldb.native = off

```

or

```
akka.persistence.journal.leveldb-shared.store.native = off
```

in your Akka configuration. The LevelDB Java port is for testing purposes only.

Warning: It is not possible to test persistence provided classes (i.e. *PersistentActor* and *AtLeastOnceDelivery*) using *TestActorRef* due to its *synchronous* nature. These traits need to be able to perform asynchronous tasks in the background in order to handle internal persistence related events.
When testing Persistence based projects always rely on *asynchronous messaging using the TestKit*.

4.4.13 Multiple persistence plugin configurations

By default, persistent actor or view will use “default” journal and snapshot store plugins configured in the following sections of the `reference.conf` configuration resource:

```
# Absolute path to the default journal plugin configuration entry.
akka.persistence.journal.plugin = "akka.persistence.journal.inmem"
# Absolute path to the default snapshot store plugin configuration entry.
akka.persistence.snapshot-store.plugin = "akka.persistence.snapshot-store.local"
```

Note that in this case actor or view overrides only `persistenceId` method:

```
abstract class ActorWithDefaultPlugins extends UntypedPersistentActor {
  @Override
  public String persistenceId() { return "123"; }
}
```

When persistent actor or view overrides `journalPluginId` and `snapshotPluginId` methods, the actor or view will be serviced by these specific persistence plugins instead of the defaults:

```
abstract class ActorWithOverridePlugins extends UntypedPersistentActor {
  @Override
  public String persistenceId() { return "123"; }
  // Absolute path to the journal plugin configuration entry in the 'reference.conf'
  @Override
  public String journalPluginId() { return "akka.persistence.chronicle.journal"; }
  // Absolute path to the snapshot store plugin configuration entry in the 'reference.conf'
  @Override
  public String snapshotPluginId() { return "akka.persistence.chronicle.snapshot-store"; }
}
```

Note that `journalPluginId` and `snapshotPluginId` must refer to properly configured `reference.conf` plugin entries with standard class property as well as settings which are specific for those plugins, i.e.:

```
# Configuration entry for the custom journal plugin, see 'journalPluginId'.
akka.persistence.chronicle.journal {
  # Standard persistence extension property: provider FQCN.
  class = "akka.persistence.chronicle.ChronicleSyncJournal"
  # Custom setting specific for the journal 'ChronicleSyncJournal'.
  folder = ${user.dir}/store/journal
}
# Configuration entry for the custom snapshot store plugin, see 'snapshotPluginId'.
akka.persistence.chronicle.snapshot-store {
  # Standard persistence extension property: provider FQCN.
  class = "akka.persistence.chronicle.ChronicleSnapshotStore"
  # Custom setting specific for the snapshot store 'ChronicleSnapshotStore'.
  folder = ${user.dir}/store/snapshot
}
```

FUTURES AND AGENTS

5.1 Futures

5.1.1 Introduction

In the Scala Standard Library, a `Future` is a data structure used to retrieve the result of some concurrent operation. This result can be accessed synchronously (blocking) or asynchronously (non-blocking). To be able to use this from Java, Akka provides a java friendly interface in `akka.dispatch.Futures`.

5.1.2 Execution Contexts

In order to execute callbacks and operations, Futures need something called an `ExecutionContext`, which is very similar to a `java.util.concurrent.Executor`. If you have an `ActorSystem` in scope, it will use its default dispatcher as the `ExecutionContext`, or you can use the factory methods provided by the `ExecutionContexts` class to wrap `Executors` and `ExecutorServices`, or even create your own.

```
import akka.dispatch.*;
import scala.concurrent.ExecutionContext;
import scala.concurrent.Future;
import scala.concurrent.Await;
import scala.concurrent.Promise;
import akka.util.Timeout;
import scala.concurrent.ExecutionContext;
import scala.concurrent.ExecutionContext$;
```

```
ExecutionContext ec =
    ExecutionContexts.fromExecutorService(yourExecutorServiceGoesHere);

//Use ec with your Futures
Future<String> f1 = Futures.successful("foo");

// Then you shut down the ExecutorService at the end of your application.
yourExecutorServiceGoesHere.shutdown();
```

5.1.3 Use with Actors

There are generally two ways of getting a reply from an `UntypedActor`: the first is by a sent message (`actorRef.tell(msg, sender)`), which only works if the original sender was an `UntypedActor` and the second is through a `Future`.

Using the `ActorRef`'s `ask` method to send a message will return a `Future`. To wait for and retrieve the actual result the simplest method is:

```
import akka.dispatch.*;
import scala.concurrent.ExecutionContext;
import scala.concurrent.Future;
import scala.concurrent.Await;
import scala.concurrent.Promise;
import akka.util.Timeout;
```

```
Timeout timeout = new Timeout(Duration.create(5, "seconds"));
Future<Object> future = Patterns.ask(actor, msg, timeout);
String result = (String) Await.result(future, timeout.duration());
```

This will cause the current thread to block and wait for the `UntypedActor` to ‘complete’ the `Future` with its reply. Blocking is discouraged though as it can cause performance problem. The blocking operations are located in `Await.result` and `Await.ready` to make it easy to spot where blocking occurs. Alternatives to blocking are discussed further within this documentation. Also note that the `Future` returned by an `UntypedActor` is a `Future<Object>` since an `UntypedActor` is dynamic. That is why the cast to `String` is used in the above sample.

Warning: `Await.result` and `Await.ready` are provided for exceptional situations where you **must** block, a good rule of thumb is to only use them if you know why you **must** block. For all other cases, use asynchronous composition as described below.

To send the result of a `Future` to an `Actor`, you can use the pipe construct:

```
akka.pattern.Patterns.pipe(future, system.dispatcher()).to(actor);
```

5.1.4 Use Directly

A common use case within Akka is to have some computation performed concurrently without needing the extra utility of an `UntypedActor`. If you find yourself creating a pool of `UntypedActors` for the sole reason of performing a calculation in parallel, there is an easier (and faster) way:

```
import scala.concurrent.duration.Duration;
import akka.japi.Function;
import java.util.concurrent.Callable;
import static akka.dispatch.Futures.future;
import static java.util.concurrent.TimeUnit.SECONDS;
```

```
Future<String> f = future(new Callable<String>() {
    public String call() {
        return "Hello" + "World";
    }
}, system.dispatcher());

f onSuccess(new PrintResult<String>(), system.dispatcher());
```

In the above code the block passed to `future` will be executed by the default `Dispatcher`, with the return value of the block used to complete the `Future` (in this case, the result would be the string: “HelloWorld”). Unlike a `Future` that is returned from an `UntypedActor`, this `Future` is properly typed, and we also avoid the overhead of managing an `UntypedActor`.

You can also create already completed `Futures` using the `Futures` class, which can be either successes:

```
Future<String> future = Futures.successful("Yay!");
```

Or failures:

```
Future<String> otherFuture = Futures.failed(
    new IllegalArgumentException("Bang!"));
```

It is also possible to create an empty `Promise`, to be filled later, and obtain the corresponding `Future`:

```
Promise<String> promise = Futures.promise();
Future<String> theFuture = promise.future();
promise.success("hello");
```

For these examples `PrintResult` is defined as follows:

```
public final static class PrintResult<T> extends OnSuccess<T> {
    @Override public final void onSuccess(T t) {
        System.out.println(t);
    }
}
```

5.1.5 Functional Futures

Scala's `Future` has several monadic methods that are very similar to the ones used by Scala's collections. These allow you to create 'pipelines' or 'streams' that the result will travel through.

Future is a Monad

The first method for working with `Future` functionally is `map`. This method takes a `Mapper` which performs some operation on the result of the `Future`, and returning a new result. The return value of the `map` method is another `Future` that will contain the new result:

```
import scala.concurrent.duration.Duration;
import akka.japi.Function;
import java.util.concurrent.Callable;
import static akka.dispatch.Futures.future;
import static java.util.concurrent.TimeUnit.SECONDS;

final ExecutionContext ec = system.dispatcher();

Future<String> f1 = future(new Callable<String>() {
    public String call() {
        return "Hello" + "World";
    }
}, ec);

Future<Integer> f2 = f1.map(new Mapper<String, Integer>() {
    public Integer apply(String s) {
        return s.length();
    }
}, ec);

f2.onSuccess(new PrintResult<Integer>(), system.dispatcher());
```

In this example we are joining two strings together within a `Future`. Instead of waiting for `f1` to complete, we apply our function that calculates the length of the string using the `map` method. Now we have a second `Future`, `f2`, that will eventually contain an `Integer`. When our original `Future`, `f1`, completes, it will also apply our function and complete the second `Future` with its result. When we finally get the result, it will contain the number 10. Our original `Future` still contains the string "HelloWorld" and is unaffected by the `map`.

Something to note when using these methods: if the `Future` is still being processed when one of these methods are called, it will be the completing thread that actually does the work. If the `Future` is already complete though, it will be run in our current thread. For example:

```
final ExecutionContext ec = system.dispatcher();

Future<String> f1 = future(new Callable<String>() {
    public String call() throws Exception {
        Thread.sleep(100);
    }
}, ec);
```

```

    return "Hello" + "World";
  }
}, ec);

Future<Integer> f2 = f1.map(new Mapper<String, Integer>() {
  public Integer apply(String s) {
    return s.length();
  }
}, ec);

f2.onSuccess(new PrintResult<Integer>(), system.dispatcher());

```

The original `Future` will take at least 0.1 second to execute now, which means it is still being processed at the time we call `map`. The function we provide gets stored within the `Future` and later executed automatically by the dispatcher when the result is ready.

If we do the opposite:

```

final ExecutionContext ec = system.dispatcher();

Future<String> f1 = future(new Callable<String>() {
  public String call() {
    return "Hello" + "World";
  }
}, ec);

// Thread.sleep is only here to prove a point
Thread.sleep(100); // Do not use this in your code

Future<Integer> f2 = f1.map(new Mapper<String, Integer>() {
  public Integer apply(String s) {
    return s.length();
  }
}, ec);

f2.onSuccess(new PrintResult<Integer>(), system.dispatcher());

```

Our little string has been processed long before our 0.1 second sleep has finished. Because of this, the dispatcher has moved onto other messages that need processing and can no longer calculate the length of the string for us, instead it gets calculated in the current thread just as if we weren't using a `Future`.

Normally this works quite well as it means there is very little overhead to running a quick function. If there is a possibility of the function taking a non-trivial amount of time to process it might be better to have this done concurrently, and for that we use `flatMap`:

```

final ExecutionContext ec = system.dispatcher();

Future<String> f1 = future(new Callable<String>() {
  public String call() {
    return "Hello" + "World";
  }
}, ec);

Future<Integer> f2 = f1.flatMap(new Mapper<String, Future<Integer>>() {
  public Future<Integer> apply(final String s) {
    return future(new Callable<Integer>() {
      public Integer call() {
        return s.length();
      }
    }, ec);
  }
}, ec);

f2.onSuccess(new PrintResult<Integer>(), system.dispatcher());

```

Now our second `Future` is executed concurrently as well. This technique can also be used to combine the results of several `Futures` into a single calculation, which will be better explained in the following sections.

If you need to do conditional propagation, you can use `filter`:

```
final ExecutionContext ec = system.dispatcher();
Future<Integer> future1 = Futures.successful(4);
Future<Integer> successfulFilter = future1.filter(Filter.filterOf(
    new Function<Integer, Boolean>() {
        public Boolean apply(Integer i) {
            return i % 2 == 0;
        }
    }
), ec);

Future<Integer> failedFilter = future1.filter(Filter.filterOf(
    new Function<Integer, Boolean>() {
        public Boolean apply(Integer i) {
            return i % 2 != 0;
        }
    }
), ec);
//When filter fails, the returned Future will be failed with a scala.MatchError
```

Composing Futures

It is very often desirable to be able to combine different `Futures` with each other, below are some examples on how that can be done in a non-blocking fashion.

```
import static akka.dispatch.Futures.sequence;

final ExecutionContext ec = system.dispatcher();
//Some source generating a sequence of Future<Integer>:s
Iterable<Future<Integer>> listOfFutureInts = source;

// now we have a Future[Iterable[Integer]]
Future<Iterable<Integer>> futureListOfInts = sequence(listOfFutureInts, ec);

// Find the sum of the odd numbers
Future<Long> futureSum = futureListOfInts.map(
    new Mapper<Iterable<Integer>, Long>() {
        public Long apply(Iterable<Integer> ints) {
            long sum = 0;
            for (Integer i : ints)
                sum += i;
            return sum;
        }
    }, ec);

futureSum.onSuccess(new PrintResult<Long>(), system.dispatcher());
```

To better explain what happened in the example, `Future.sequence` is taking the `Iterable<Future<Integer>>` and turning it into a `Future<Iterable<Integer>>`. We can then use `map` to work with the `Iterable<Integer>` directly, and we aggregate the sum of the `Iterable`.

The `traverse` method is similar to `sequence`, but it takes a sequence of `A` and applies a function from `A` to `Future` and returns a `Future<Iterable>`, enabling parallel `map` over the sequence, if you use `Futures.future` to create the `Future`.

```
import static akka.dispatch.Futures.traverse;
```

```

final ExecutionContext ec = system.dispatcher();
//Just a sequence of Strings
Iterable<String> listStrings = Arrays.asList("a", "b", "c");

Future<Iterable<String>> futureResult = traverse(listStrings,
    new Function<String, Future<String>>() {
        public Future<String> apply(final String r) {
            return future(new Callable<String>() {
                public String call() {
                    return r.toUpperCase();
                }
            }, ec);
        }
    }, ec);

//Returns the sequence of strings as upper case
futureResult onSuccess(new PrintResult<Iterable<String>>(), system.dispatcher());

```

It's as simple as that!

Then there's a method that's called `fold` that takes a start-value, a sequence of `Future`s and a function from the type of the start-value, a timeout, and the type of the futures and returns something with the same type as the start-value, and then applies the function to all elements in the sequence of futures, non-blockingly, the execution will be started when the last of the `Futures` is completed.

```

import akka.japi.Function2;
import static akka.dispatch.Futures.fold;

```

```

final ExecutionContext ec = system.dispatcher();

//A sequence of Futures, in this case Strings
Iterable<Future<String>> futures = source;

//Start value is the empty string
Future<String> resultFuture = fold("", futures,
    new Function2<String, String, String>() {
        public String apply(String r, String t) {
            return r + t; //Just concatenate
        }
    }, ec);

resultFuture onSuccess(new PrintResult<String>(), system.dispatcher());

```

That's all it takes!

If the sequence passed to `fold` is empty, it will return the start-value, in the case above, that will be empty `String`. In some cases you don't have a start-value and you're able to use the value of the first completing `Future` in the sequence as the start-value, you can use `reduce`, it works like this:

```

import static akka.dispatch.Futures.reduce;

```

```

final ExecutionContext ec = system.dispatcher();
//A sequence of Futures, in this case Strings
Iterable<Future<String>> futures = source;

Future<Object> resultFuture = reduce(futures,
    new Function2<Object, String, Object>() {
        public Object apply(Object r, String t) {
            return r + t; //Just concatenate
        }
    }, ec);

resultFuture onSuccess(new PrintResult<Object>(), system.dispatcher());

```

Same as with `fold`, the execution will be started when the last of the `Futures` is completed, you can also parallelize it by chunking your futures into sub-sequences and reduce them, and then reduce the reduced results again.

This is just a sample of what can be done.

5.1.6 Callbacks

Sometimes you just want to listen to a `Future` being completed, and react to that not by creating a new `Future`, but by side-effecting. For this Scala supports `onComplete`, `onSuccess` and `onFailure`, of which the latter two are specializations of the first.

```
final ExecutionContext ec = system.dispatcher();

future.onSuccess(new OnSuccess<String>() {
  public void onSuccess(String result) {
    if ("bar" == result) {
      //Do something if it resulted in "bar"
    } else {
      //Do something if it was some other String
    }
  }
}, ec);
```

```
final ExecutionContext ec = system.dispatcher();

future.onFailure(new OnFailure() {
  public void onFailure(Throwable failure) {
    if (failure instanceof IllegalStateException) {
      //Do something if it was this particular failure
    } else {
      //Do something if it was some other failure
    }
  }
}, ec);
```

```
final ExecutionContext ec = system.dispatcher();

future.onComplete(new OnComplete<String>() {
  public void onComplete(Throwable failure, String result) {
    if (failure != null) {
      //We got a failure, handle it here
    } else {
      // We got a result, do something with it
    }
  }
}, ec);
```

5.1.7 Ordering

Since callbacks are executed in any order and potentially in parallel, it can be tricky at the times when you need sequential ordering of operations. But there's a solution! And it's name is `andThen`, and it creates a new `Future` with the specified callback, a `Future` that will have the same result as the `Future` it's called on, which allows for ordering like in the following sample:

```
final ExecutionContext ec = system.dispatcher();
Future<String> future1 = Futures.successful("value").andThen(
  new OnComplete<String>() {
    public void onComplete(Throwable failure, String result) {
      if (failure != null)
        sendToIssueTracker(failure);
    }
  }
);
```

```

    }, ec).andThen(new OnComplete<String>() {
    public void onComplete(Throwable failure, String result) {
        if (result != null)
            sendToTheInternetz(result);
    }
    }, ec);

```

5.1.8 Auxiliary methods

Future `fallbackTo` combines 2 Futures into a new Future, and will hold the successful value of the second Future if the first Future fails.

```

Future<String> future1 = Futures.failed(new IllegalStateException("OHNOES1"));
Future<String> future2 = Futures.failed(new IllegalStateException("OHNOES2"));
Future<String> future3 = Futures.successful("bar");
// Will have "bar" in this case
Future<String> future4 = future1.fallbackTo(future2).fallbackTo(future3);
future4.onSuccess(new PrintResult<String>(), system.dispatcher());

```

You can also combine two Futures into a new Future that will hold a tuple of the two Futures successful results, using the `zip` operation.

```

final ExecutionContext ec = system.dispatcher();
Future<String> future1 = Futures.successful("foo");
Future<String> future2 = Futures.successful("bar");
Future<String> future3 = future1.zip(future2).map(
    new Mapper<scala.Tuple2<String, String>, String>() {
        public String apply(scala.Tuple2<String, String> zipped) {
            return zipped._1() + " " + zipped._2();
        }
    }, ec);

future3.onSuccess(new PrintResult<String>(), system.dispatcher());

```

5.1.9 Exceptions

Since the result of a Future is created concurrently to the rest of the program, exceptions must be handled differently. It doesn't matter if an `UntypedActor` or the dispatcher is completing the Future, if an Exception is caught the Future will contain it instead of a valid result. If a Future does contain an Exception, calling `Await.result` will cause it to be thrown again so it can be handled properly.

It is also possible to handle an Exception by returning a different result. This is done with the `recover` method. For example:

```

final ExecutionContext ec = system.dispatcher();

Future<Integer> future = future(new Callable<Integer>() {
    public Integer call() {
        return 1 / 0;
    }
}, ec).recover(new Recover<Integer>() {
    public Integer recover(Throwable problem) throws Throwable {
        if (problem instanceof ArithmeticException)
            return 0;
        else
            throw problem;
    }
}, ec);

future.onSuccess(new PrintResult<Integer>(), system.dispatcher());

```

In this example, if the actor replied with a `akka.actor.Status.Failure` containing the `ArithmeticException`, our `Future` would have a result of 0. The `recover` method works very similarly to the standard `try/catch` blocks, so multiple `Exceptions` can be handled in this manner, and if an `Exception` is not handled this way it will behave as if we hadn't used the `recover` method.

You can also use the `recoverWith` method, which has the same relationship to `recover` as `flatMap` has to `map`, and is use like this:

```
final ExecutionContext ec = system.dispatcher();

Future<Integer> future = future(new Callable<Integer>() {
    public Integer call() {
        return 1 / 0;
    }
}, ec).recoverWith(new Recover<Future<Integer>>() {
    public Future<Integer> recover(Throwable problem) throws Throwable {
        if (problem instanceof ArithmeticException) {
            return future(new Callable<Integer>() {
                public Integer call() {
                    return 0;
                }
            }, ec);
        } else
            throw problem;
    }
}, ec);

future.onSuccess(new PrintResult<Integer>(), system.dispatcher());
```

5.1.10 After

`akka.pattern.Patterns.after` makes it easy to complete a `Future` with a value or exception after a timeout.

```
import static akka.pattern.Patterns.after;
import java.util.Arrays;

final ExecutionContext ec = system.dispatcher();
Future<String> failExc = Futures.failed(new IllegalStateException("OHNOES1"));
Future<String> delayed = Patterns.after(Duration.create(200, "millis"),
    system.scheduler(), ec, failExc);
Future<String> future = future(new Callable<String>() {
    public String call() throws InterruptedException {
        Thread.sleep(1000);
        return "foo";
    }
}, ec);
Future<String> result = Futures.firstCompletedOf(
    Arrays.<Future<String>>asList(future, delayed), ec);
```

5.2 Agents

Agents in Akka are inspired by [agents in Clojure](#).

Agents provide asynchronous change of individual locations. Agents are bound to a single storage location for their lifetime, and only allow mutation of that location (to a new state) to occur as a result of an action. Update actions are functions that are asynchronously applied to the Agent's state and whose return value becomes the Agent's new state. The state of an Agent should be immutable.

While updates to Agents are asynchronous, the state of an Agent is always immediately available for reading by any thread (using `get`) without any messages.

Agents are reactive. The update actions of all Agents get interleaved amongst threads in an `ExecutionContext`. At any point in time, at most one `send` action for each Agent is being executed. Actions dispatched to an agent from another thread will occur in the order they were sent, potentially interleaved with actions dispatched to the same agent from other threads.

5.2.1 Creating Agents

Agents are created by invoking `new Agent<ValueType>(value, executionContext)` – passing in the Agent's initial value and providing an `ExecutionContext` to be used for it:

```
import scala.concurrent.ExecutionContext;
import akka.agent.Agent;
import akka.dispatch.ExecutionContexts;
ExecutionContext ec = ExecutionContexts.global();
Agent<Integer> agent = Agent.create(5, ec);
```

5.2.2 Reading an Agent's value

Agents can be dereferenced (you can get an Agent's value) by invoking the Agent with `get()` like this:

```
Integer result = agent.get();
```

Reading an Agent's current value does not involve any message passing and happens immediately. So while updates to an Agent are asynchronous, reading the state of an Agent is synchronous.

You can also get a `Future` to the Agent's value, that will be completed after the currently queued updates have completed:

```
import scala.concurrent.Future;
Future<Integer> future = agent.future();
```

See *Futures* for more information on Futures.

5.2.3 Updating Agents (send & alter)

You update an Agent by sending a function (`akka.dispatch.Mapper`) that transforms the current value or by sending just a new value. The Agent will apply the new value or function atomically and asynchronously. The update is done in a fire-forget manner and you are only guaranteed that it will be applied. There is no guarantee of when the update will be applied but dispatches to an Agent from a single thread will occur in order. You apply a value or a function by invoking the `send` function.

```
import akka.dispatch.Mapper;
// send a value, enqueues this change
// of the value of the Agent
agent.send(7);

// send a Mapper, enqueues this change
// to the value of the Agent
agent.send(new Mapper<Integer, Integer>() {
    public Integer apply(Integer i) {
        return i * 2;
    }
});
```

You can also dispatch a function to update the internal state but on its own thread. This does not use the reactive thread pool and can be used for long-running or blocking operations. You do this with the `sendOff` method. Dispatches using either `sendOff` or `send` will still be executed in order.

```
import akka.dispatch.Mapper;
// sendOff a function
agent.sendOff(longRunningOrBlockingFunction,
              theExecutionContextToExecuteItIn);
```

All send methods also have a corresponding alter method that returns a Future. See [Futures](#) for more information on Futures.

```
import akka.dispatch.Mapper;
import scala.concurrent.Future;
// alter a value
Future<Integer> f1 = agent.alter(7);

// alter a function (Mapper)
Future<Integer> f2 = agent.alter(new Mapper<Integer, Integer>() {
    public Integer apply(Integer i) {
        return i * 2;
    }
});
```

```
import akka.dispatch.Mapper;
import scala.concurrent.Future;
// alterOff a function (Mapper)
Future<Integer> f3 = agent.alterOff(longRunningOrBlockingFunction,
                                  theExecutionContextToExecuteItIn);
```

5.2.4 Configuration

There are several configuration properties for the agents module, please refer to the [reference configuration](#).

5.2.5 Deprecated Transactional Agents

Agents participating in enclosing STM transaction is a deprecated feature in 2.3.

If an Agent is used within an enclosing Scala STM transaction, then it will participate in that transaction. If you send to an Agent within a transaction then the dispatch to the Agent will be held until that transaction commits, and discarded if the transaction is aborted.

NETWORKING

6.1 Cluster Specification

Note: This document describes the design concepts of the clustering.

6.1.1 Intro

Akka Cluster provides a fault-tolerant decentralized peer-to-peer based cluster `membership` service with no single point of failure or single point of bottleneck. It does this using `gossip` protocols and an automatic `failure detector`.

6.1.2 Terms

node A logical member of a cluster. There could be multiple nodes on a physical machine. Defined by a `hostname:port:uid` tuple.

cluster A set of nodes joined together through the `membership` service.

leader A single node in the cluster that acts as the leader. Managing cluster convergence and membership state transitions.

6.1.3 Membership

A cluster is made up of a set of member nodes. The identifier for each node is a `hostname:port:uid` tuple. An Akka application can be distributed over a cluster with each node hosting some part of the application. Cluster membership and the actors running on that node of the application are decoupled. A node could be a member of a cluster without hosting any actors. Joining a cluster is initiated by issuing a `Join` command to one of the nodes in the cluster to join.

The node identifier internally also contains a `UID` that uniquely identifies this actor system instance at that `hostname:port`. Akka uses the `UID` to be able to reliably trigger remote death watch. This means that the same actor system can never join a cluster again once it's been removed from that cluster. To re-join an actor system with the same `hostname:port` to a cluster you have to stop the actor system and start a new one with the same `hostname:port` which will then receive a different `UID`.

The cluster membership state is a specialized `CRDT`, which means that it has a monotonic merge function. When concurrent changes occur on different nodes the updates can always be merged and converge to the same end result.

Gossip

The cluster membership used in Akka is based on Amazon's [Dynamo](#) system and particularly the approach taken in Basho's [Riak](#) distributed database. Cluster membership is communicated using a [Gossip Protocol](#), where the current state of the cluster is gossiped randomly through the cluster, with preference to members that have not seen the latest version.

Vector Clocks

[Vector clocks](#) are a type of data structure and algorithm for generating a partial ordering of events in a distributed system and detecting causality violations.

We use vector clocks to reconcile and merge differences in cluster state during gossiping. A vector clock is a set of (node, counter) pairs. Each update to the cluster state has an accompanying update to the vector clock.

Gossip Convergence

Information about the cluster converges locally at a node at certain points in time. This is when a node can prove that the cluster state he is observing has been observed by all other nodes in the cluster. Convergence is implemented by passing a set of nodes that have seen current state version during gossip. This information is referred to as the seen set in the gossip overview. When all nodes are included in the seen set there is convergence.

Gossip convergence cannot occur while any nodes are `unreachable`. The nodes need to become `reachable` again, or moved to the `down` and `removed` states (see the [Membership Lifecycle](#) section below). This only blocks the leader from performing its cluster membership management and does not influence the application running on top of the cluster. For example this means that during a network partition it is not possible to add more nodes to the cluster. The nodes can join, but they will not be moved to the `up` state until the partition has healed or the unreachable nodes have been downed.

Failure Detector

The failure detector is responsible for trying to detect if a node is `unreachable` from the rest of the cluster. For this we are using an implementation of [The Phi Accrual Failure Detector](#) by Hayashibara et al.

An accrual failure detector decouple monitoring and interpretation. That makes them applicable to a wider area of scenarios and more adequate to build generic failure detection services. The idea is that it is keeping a history of failure statistics, calculated from heartbeats received from other nodes, and is trying to do educated guesses by taking multiple factors, and how they accumulate over time, into account in order to come up with a better guess if a specific node is up or down. Rather than just answering "yes" or "no" to the question "is the node down?" it returns a `phi` value representing the likelihood that the node is down.

The `threshold` that is the basis for the calculation is configurable by the user. A low `threshold` is prone to generate many wrong suspicions but ensures a quick detection in the event of a real crash. Conversely, a high `threshold` generates fewer mistakes but needs more time to detect actual crashes. The default `threshold` is 8 and is appropriate for most situations. However in cloud environments, such as Amazon EC2, the value could be increased to 12 in order to account for network issues that sometimes occur on such platforms.

In a cluster each node is monitored by a few (default maximum 5) other nodes, and when any of these detects the node as `unreachable` that information will spread to the rest of the cluster through the gossip. In other words, only one node needs to mark a node `unreachable` to have the rest of the cluster mark that node `unreachable`.

The nodes to monitor are picked out of neighbors in a hashed ordered node ring. This is to increase the likelihood to monitor across racks and data centers, but the order is the same on all nodes, which ensures full coverage.

Heartbeats are sent out every second and every heartbeat is performed in a request/reply handshake with the replies used as input to the failure detector.

The failure detector will also detect if the node becomes `reachable` again. When all nodes that monitored the `unreachable` node detects it as `reachable` again the cluster, after gossip dissemination, will consider it as `reachable`.

If system messages cannot be delivered to a node it will be quarantined and then it cannot come back from `unreachable`. This can happen if there are too many unacknowledged system messages (e.g. `watch`, `Terminated`, remote actor deployment, failures of actors supervised by remote parent). Then the node needs to be moved to the `down` or `removed` states (see the [Membership Lifecycle](#) section below) and the actor system must be restarted before it can join the cluster again.

Leader

After gossip convergence a `leader` for the cluster can be determined. There is no `leader` election process, the `leader` can always be recognised deterministically by any node whenever there is gossip convergence. The `leader` is just a role, any node can be the `leader` and it can change between convergence rounds. The `leader` is simply the first node in sorted order that is able to take the leadership role, where the preferred member states for a `leader` are `up` and `leaving` (see the [Membership Lifecycle](#) section below for more information about member states).

The role of the `leader` is to shift members in and out of the cluster, changing `joining` members to the `up` state or `exiting` members to the `removed` state. Currently `leader` actions are only triggered by receiving a new cluster state with gossip convergence.

The `leader` also has the power, if configured so, to “auto-down” a node that according to the [Failure Detector](#) is considered `unreachable`. This means setting the `unreachable` node status to `down` automatically after a configured time of unreachability.

Seed Nodes

The seed nodes are configured contact points for new nodes joining the cluster. When a new node is started it sends a message to all seed nodes and then sends a `join` command to the seed node that answers first.

The seed nodes configuration value does not have any influence on the running cluster itself, it is only relevant for new nodes joining the cluster as it helps them to find contact points to send the `join` command to; a new member can send this command to any current member of the cluster, not only to the seed nodes.

Gossip Protocol

A variation of *push-pull gossip* is used to reduce the amount of gossip information sent around the cluster. In push-pull gossip a digest is sent representing current versions but not actual values; the recipient of the gossip can then send back any values for which it has newer versions and also request values for which it has outdated versions. Akka uses a single shared state with a vector clock for versioning, so the variant of push-pull gossip used in Akka makes use of this version to only push the actual state as needed.

Periodically, the default is every 1 second, each node chooses another random node to initiate a round of gossip with. If less than $\frac{1}{2}$ of the nodes resides in the seen set (have seen the new state) then the cluster gossips 3 times instead of once every second. This adjusted gossip interval is a way to speed up the convergence process in the early dissemination phase after a state change.

The choice of node to gossip with is random but it is biased to towards nodes that might not have seen the current state version. During each round of gossip exchange when no convergence it uses a probability of 0.8 (configurable) to gossip to a node not part of the seen set, i.e. that probably has an older version of the state. Otherwise gossip to any random live node.

This biased selection is a way to speed up the convergence process in the late dissemination phase after a state change.

For clusters larger than 400 nodes (configurable, and suggested by empirical evidence) the 0.8 probability is gradually reduced to avoid overwhelming single stragglers with too many concurrent gossip requests. The gossip

receiver also has a mechanism to protect itself from too many simultaneous gossip messages by dropping messages that have been enqueued in the mailbox for too long time.

While the cluster is in a converged state the gossip only sends a small gossip status message containing the gossip version to the chosen node. As soon as there is a change to the cluster (meaning non-convergence) then it goes back to biased gossip again.

The recipient of the gossip state or the gossip status can use the gossip version (vector clock) to determine whether:

1. it has a newer version of the gossip state, in which case it sends that back to the gossip
2. it has an outdated version of the state, in which case the recipient requests the current state from the gossip by sending back its version of the gossip state
3. it has conflicting gossip versions, in which case the different versions are merged and sent back

If the recipient and the gossip have the same version then the gossip state is not sent or requested.

The periodic nature of the gossip has a nice batching effect of state changes, e.g. joining several nodes quickly after each other to one node will result in only one state change to be spread to other members in the cluster.

The gossip messages are serialized with `protobuf` and also gzipped to reduce payload size.

Membership Lifecycle

A node begins in the `joining` state. Once all nodes have seen that the new node is joining (through gossip convergence) the `leader` will set the member state to `up`.



If a node is leaving the cluster in a safe, expected manner then it switches to the `leaving` state. Once the leader sees the convergence on the node in the `leaving` state, the leader will then move it to `exiting`. Once all nodes have seen the `exiting` state (convergence) the `leader` will remove the node from the cluster, marking it as `removed`.

If a node is `unreachable` then gossip convergence is not possible and therefore any `leader` actions are also not possible (for instance, allowing a node to become a part of the cluster). To be able to move forward the state of the `unreachable` nodes must be changed. It must become `reachable` again or marked as `down`. If the node is to join the cluster again the actor system must be restarted and go through the joining process again. The cluster can, through the leader, also *auto-down* a node after a configured time of unreachability..

Note: If you have *auto-down* enabled and the failure detector triggers, you can over time end up with a lot of single node clusters if you don't put measures in place to shut down nodes that have become `unreachable`. This follows from the fact that the `unreachable` node will likely see the rest of the cluster as `unreachable`, become its own leader and form its own cluster.

As mentioned before, if a node is `unreachable` then gossip convergence is not possible and therefore any `leader` actions are also not possible. By enabling `akka.cluster.allow-weakly-up-members` it is possible to let new joining nodes be promoted while convergence is not yet reached. These `Joining` nodes will be promoted as `WeaklyUp`. Once gossip convergence is reached, the leader will move `WeaklyUp` members to `Up`.

Note that members on the other side of a network partition have no knowledge about the existence of the new members. You should for example not count `WeaklyUp` members in quorum decisions.

State Diagram for the Member States (`akka.cluster.allow-weakly-up-members=off`)State Diagram for the Member States (`akka.cluster.allow-weakly-up-members=on`)

Member States

- **joining** transient state when joining a cluster

- **weakly up** transient state while network split (only if `akka.cluster.allow-weakly-up-members=on`)
- **up** normal operating state
- **leaving / exiting** states during graceful removal
- **down** marked as down (no longer part of cluster decisions)
- **removed** tombstone state (no longer a member)

User Actions

- **join** join a single node to a cluster - can be explicit or automatic on startup if a node to join have been specified in the configuration
- **leave** tell a node to leave the cluster gracefully
- **down** mark a node as down

Leader Actions

The leader has the following duties:

- shifting members in and out of the cluster
 - joining -> up
 - exiting -> removed

Failure Detection and Unreachability

- **fd*** the failure detector of one of the monitoring nodes has triggered causing the monitored node to be marked as unreachable
- **unreachable*** unreachable is not a real member states but more of a flag in addition to the state signaling that the cluster is unable to talk to this node, after being unreachable the failure detector may detect it as reachable again and thereby remove the flag

6.2 Cluster Usage

For introduction to the Akka Cluster concepts please see *Cluster Specification*.

6.2.1 Preparing Your Project for Clustering

The Akka cluster is a separate jar file. Make sure that you have the following dependency in your project:

```
<dependency>
  <groupId>com.typesafe.akka</groupId>
  <artifactId>akka-cluster_2.11</artifactId>
  <version>2.4.0</version>
</dependency>
```

6.2.2 A Simple Cluster Example

The following configuration enables the `Cluster` extension to be used. It joins the cluster and an actor subscribes to cluster membership events and logs them.

The `application.conf` configuration looks like this:

```
akka {
  actor {
    provider = "akka.cluster.ClusterActorRefProvider"
  }
  remote {
    log-remote-lifecycle-events = off
    netty.tcp {
      hostname = "127.0.0.1"
      port = 0
    }
  }

  cluster {
    seed-nodes = [
      "akka.tcp://ClusterSystem@127.0.0.1:2551",
      "akka.tcp://ClusterSystem@127.0.0.1:2552"]

    auto-down-unreachable-after = 10s
  }
}

# Disable legacy metrics in akka-cluster.
akka.cluster.metrics.enabled=off

# Enable metrics extension in akka-cluster-metrics.
akka.extensions=["akka.cluster.metrics.ClusterMetricsExtension"]

# Sigar native library extract location during tests.
# Note: use per-jvm-instance folder when running multiple jvm on one host.
akka.cluster.metrics.native-library-extract-folder=${user.dir}/target/native
```

To enable cluster capabilities in your Akka project you should, at a minimum, add the *Remoting* settings, but with `akka.cluster.ClusterActorRefProvider`. The `akka.cluster.seed-nodes` should normally also be added to your `application.conf` file.

Note: If you are using Docker or the nodes for some other reason have separate internal and external ip addresses you must configure remoting according to *Remote configuration for NAT and Docker*

The seed nodes are configured contact points for initial, automatic, join of the cluster.

Note that if you are going to start the nodes on different machines you need to specify the ip-addresses or host names of the machines in `application.conf` instead of `127.0.0.1`

An actor that uses the cluster extension may look like this:

```
package sample.cluster.simple;

import akka.actor.UntypedActor;
import akka.cluster.Cluster;
import akka.cluster.ClusterEvent;
import akka.cluster.ClusterEvent.MemberEvent;
import akka.cluster.ClusterEvent.MemberUp;
import akka.cluster.ClusterEvent.MemberRemoved;
import akka.cluster.ClusterEvent.UnreachableMember;
import akka.event.Logging;
import akka.event.LoggingAdapter;

public class SimpleClusterListener extends UntypedActor {
  LoggingAdapter log = Logging.getLogger(getContext().system(), this);
  Cluster cluster = Cluster.get(getContext().system());

  //subscribe to cluster changes
```

```

@Override
public void preStart() {
    // #subscribe
    cluster.subscribe(getSelf(), ClusterEvent.initialStateAsEvents(),
        MemberEvent.class, UnreachableMember.class);
    // #subscribe
}

// re-subscribe when restart
@Override
public void postStop() {
    cluster.unsubscribe(getSelf());
}

@Override
public void onReceive(Object message) {
    if (message instanceof MemberUp) {
        MemberUp mUp = (MemberUp) message;
        log.info("Member is Up: {}", mUp.member());
    } else if (message instanceof UnreachableMember) {
        UnreachableMember mUnreachable = (UnreachableMember) message;
        log.info("Member detected as unreachable: {}", mUnreachable.member());
    } else if (message instanceof MemberRemoved) {
        MemberRemoved mRemoved = (MemberRemoved) message;
        log.info("Member is Removed: {}", mRemoved.member());
    } else if (message instanceof MemberEvent) {
        // ignore
    } else {
        unhandled(message);
    }
}
}

```

The actor registers itself as subscriber of certain cluster events. It receives events corresponding to the current state of the cluster when the subscription starts and then it receives events for changes that happen in the cluster.

The easiest way to run this example yourself is to download [Typesafe Activator](#) and open the tutorial named [Akka Cluster Samples with Java](#). It contains instructions of how to run the `SimpleClusterApp`.

6.2.3 Joining to Seed Nodes

You may decide if joining to the cluster should be done manually or automatically to configured initial contact points, so-called seed nodes. When a new node is started it sends a message to all seed nodes and then sends join command to the one that answers first. If no one of the seed nodes replied (might not be started yet) it retries this procedure until successful or shutdown.

You define the seed nodes in the *Configuration* file (`application.conf`):

```

akka.cluster.seed-nodes = [
    "akka.tcp://ClusterSystem@host1:2552",
    "akka.tcp://ClusterSystem@host2:2552"]

```

This can also be defined as Java system properties when starting the JVM using the following syntax:

```

-Dakka.cluster.seed-nodes.0=akka.tcp://ClusterSystem@host1:2552
-Dakka.cluster.seed-nodes.1=akka.tcp://ClusterSystem@host2:2552

```

The seed nodes can be started in any order and it is not necessary to have all seed nodes running, but the node configured as the first element in the `seed-nodes` configuration list must be started when initially starting a cluster, otherwise the other seed-nodes will not become initialized and no other node can join the cluster. The reason for the special first seed node is to avoid forming separated islands when starting from an empty cluster. It is quickest to start all configured seed nodes at the same time (order doesn't matter), otherwise it can take up to the configured `seed-node-timeout` until the nodes can join.

Once more than two seed nodes have been started it is no problem to shut down the first seed node. If the first seed node is restarted, it will first try to join the other seed nodes in the existing cluster.

If you don't configure seed nodes you need to join the cluster programmatically or manually.

Manual joining can be performed by using `ref:cluster_jmx_java` or *Command Line Management*. Joining programmatically can be performed with `Cluster.get(system).join`. Unsuccessful join attempts are automatically retried after the time period defined in configuration property `retry-unsuccessful-join-after`. Retries can be disabled by setting the property to `off`.

You can join to any node in the cluster. It does not have to be configured as a seed node. Note that you can only join to an existing cluster member, which means that for bootstrapping some node must join itself, and then the following nodes could join them to make up a cluster.

You may also use `Cluster.get(system).joinSeedNodes` to join programmatically, which is attractive when dynamically discovering other nodes at startup by using some external tool or API. When using `joinSeedNodes` you should not include the node itself except for the node that is supposed to be the first seed node, and that should be placed first in parameter to `joinSeedNodes`.

Unsuccessful attempts to contact seed nodes are automatically retried after the time period defined in configuration property `seed-node-timeout`. Unsuccessful attempt to join a specific seed node is automatically retried after the configured `retry-unsuccessful-join-after`. Retrying means that it tries to contact all seed nodes and then joins the node that answers first. The first node in the list of seed nodes will join itself if it cannot contact any of the other seed nodes within the configured `seed-node-timeout`.

An actor system can only join a cluster once. Additional attempts will be ignored. When it has successfully joined it must be restarted to be able to join another cluster or to join the same cluster again. It can use the same host name and port after the restart, when it come up as new incarnation of existing member in the cluster, trying to join in, then the existing one will be removed from the cluster and then it will be allowed to join.

6.2.4 Automatic vs. Manual Downing

When a member is considered by the failure detector to be unreachable the leader is not allowed to perform its duties, such as changing status of new joining members to 'Up'. The node must first become reachable again, or the status of the unreachable member must be changed to 'Down'. Changing status to 'Down' can be performed automatically or manually. By default it must be done manually, using *JMX* or *Command Line Management*.

It can also be performed programmatically with `Cluster.get(system).down(address)`.

You can enable automatic downing with configuration:

```
akka.cluster.auto-down-unreachable-after = 120s
```

This means that the cluster leader member will change the `unreachable` node status to `down` automatically after the configured time of unreachability.

Be aware of that using `auto-down` implies that two separate clusters will automatically be formed in case of network partition. That might be desired by some applications but not by others.

Note: If you have *auto-down* enabled and the failure detector triggers, you can over time end up with a lot of single node clusters if you don't put measures in place to shut down nodes that have become `unreachable`. This follows from the fact that the `unreachable` node will likely see the rest of the cluster as `unreachable`, become its own leader and form its own cluster.

6.2.5 Leaving

There are two ways to remove a member from the cluster.

You can just stop the actor system (or the JVM process). It will be detected as unreachable and removed after the automatic or manual downing as described above.

A more graceful exit can be performed if you tell the cluster that a node shall leave. This can be performed using *JMX* or *Command Line Management*. It can also be performed programmatically with:

```
final Cluster cluster = Cluster.get(system);
cluster.leave(cluster.selfAddress());
```

Note that this command can be issued to any member in the cluster, not necessarily the one that is leaving. The cluster extension, but not the actor system or JVM, of the leaving member will be shutdown after the leader has changed status of the member to *Exiting*. Thereafter the member will be removed from the cluster. Normally this is handled automatically, but in case of network failures during this process it might still be necessary to set the node's status to *Down* in order to complete the removal.

6.2.6 WeaklyUp Members

If a node is `unreachable` then gossip convergence is not possible and therefore any `leader` actions are also not possible. However, we still might want new nodes to join the cluster in this scenario.

Warning: The `WeaklyUp` feature is marked as “**experimental**” as of its introduction in Akka 2.4.0. We will continue to improve this feature based on our users' feedback, which implies that while we try to keep incompatible changes to a minimum the binary compatibility guarantee for maintenance releases does not apply this feature.

This feature is disabled by default. With a configuration option you can allow this behavior:

```
akka.cluster.allow-weakly-up-members = on
```

When `allow-weakly-up-members` is enabled and there is no gossip convergence, `Joining` members will be promoted to `WeaklyUp` and they will become part of the cluster. Once gossip convergence is reached, the leader will move `WeaklyUp` members to `Up`.

You can subscribe to the `WeaklyUp` membership event to make use of the members that are in this state, but you should be aware of that members on the other side of a network partition have no knowledge about the existence of the new members. You should for example not count `WeaklyUp` members in quorum decisions.

Warning: This feature is only available from Akka 2.4.0 and cannot be used if some of your cluster members are running an older version of Akka.

6.2.7 Subscribe to Cluster Events

You can subscribe to change notifications of the cluster membership by using `Cluster.get(system).subscribe`.

```
cluster.subscribe(getSelf(), MemberEvent.class, UnreachableMember.class);
```

A snapshot of the full state, `akka.cluster.ClusterEvent.CurrentClusterState`, is sent to the subscriber as the first message, followed by events for incremental updates.

Note that you may receive an empty `CurrentClusterState`, containing no members, if you start the subscription before the initial join procedure has completed. This is expected behavior. When the node has been accepted in the cluster you will receive `MemberUp` for that node, and other nodes.

If you find it inconvenient to handle the `CurrentClusterState` you can use `ClusterEvent.initialStateAsEvents()` as parameter to subscribe. That means that instead of receiving `CurrentClusterState` as the first message you will receive the events corresponding to the current state to mimic what you would have seen if you were listening to the events when they occurred in the past. Note that those initial events only correspond to the current state and it is not the full history of all changes that actually has occurred in the cluster.

```
cluster.subscribe(getSelf(), ClusterEvent.initialStateAsEvents(),
    MemberEvent.class, UnreachableMember.class);
```

The events to track the life-cycle of members are:

- `ClusterEvent.MemberUp` - A new member has joined the cluster and its status has been changed to Up.
- `ClusterEvent.MemberExited` - A member is leaving the cluster and its status has been changed to Exiting Note that the node might already have been shutdown when this event is published on another node.
- `ClusterEvent.MemberRemoved` - Member completely removed from the cluster.
- `ClusterEvent.UnreachableMember` - A member is considered as unreachable, detected by the failure detector of at least one other node.
- `ClusterEvent.ReachableMember` - A member is considered as reachable again, after having been unreachable. All nodes that previously detected it as unreachable has detected it as reachable again.

There are more types of change events, consult the API documentation of classes that extends `akka.cluster.ClusterEvent.ClusterDomainEvent` for details about the events.

Instead of subscribing to cluster events it can sometimes be convenient to only get the full membership state with `Cluster.get(system).state()`. Note that this state is not necessarily in sync with the events published to a cluster subscription.

Worker Dial-in Example

Let's take a look at an example that illustrates how workers, here named *backend*, can detect and register to new master nodes, here named *frontend*.

The example application provides a service to transform text. When some text is sent to one of the frontend services, it will be delegated to one of the backend workers, which performs the transformation job, and sends the result back to the original client. New backend nodes, as well as new frontend nodes, can be added or removed to the cluster dynamically.

Messages:

```
public interface TransformationMessages {

    public static class TransformationJob implements Serializable {
        private final String text;

        public TransformationJob(String text) {
            this.text = text;
        }

        public String getText() {
            return text;
        }
    }

    public static class TransformationResult implements Serializable {
        private final String text;

        public TransformationResult(String text) {
```

```

    this.text = text;
  }

  public String getText() {
    return text;
  }

  @Override
  public String toString() {
    return "TransformationResult(" + text + ")";
  }
}

public static class JobFailed implements Serializable {
  private final String reason;
  private final TransformationJob job;

  public JobFailed(String reason, TransformationJob job) {
    this.reason = reason;
    this.job = job;
  }

  public String getReason() {
    return reason;
  }

  public TransformationJob getJob() {
    return job;
  }

  @Override
  public String toString() {
    return "JobFailed(" + reason + ")";
  }
}

public static final String BACKEND_REGISTRATION = "BackendRegistration";
}

```

The backend worker that performs the transformation job:

```

public class TransformationBackend extends UntypedActor {

  Cluster cluster = Cluster.get(getContext().system());

  //subscribe to cluster changes, MemberUp
  @Override
  public void preStart() {
    cluster.subscribe(getSelf(), MemberUp.class);
  }

  //re-subscribe when restart
  @Override
  public void postStop() {
    cluster.unsubscribe(getSelf());
  }

  @Override
  public void onReceive(Object message) {
    if (message instanceof TransformationJob) {
      TransformationJob job = (TransformationJob) message;
      getSender().tell(new TransformationResult(job.getText().toUpperCase()),

```

```

        getSelf());

    } else if (message instanceof CurrentClusterState) {
        CurrentClusterState state = (CurrentClusterState) message;
        for (Member member : state.getMembers()) {
            if (member.status().equals(MemberStatus.up())) {
                register(member);
            }
        }
    }

    } else if (message instanceof MemberUp) {
        MemberUp mUp = (MemberUp) message;
        register(mUp.member());
    }

    } else {
        unhandled(message);
    }
}

void register(Member member) {
    if (member.hasRole("frontend"))
        getContext().actorSelection(member.address() + "/user/frontend").tell(
            BACKEND_REGISTRATION, getSelf());
}
}

```

Note that the `TransformationBackend` actor subscribes to cluster events to detect new, potential, frontend nodes, and send them a registration message so that they know that they can use the backend worker.

The frontend that receives user jobs and delegates to one of the registered backend workers:

```

public class TransformationFrontend extends UntypedActor {

    List<ActorRef> backends = new ArrayList<ActorRef>();
    int jobCounter = 0;

    @Override
    public void onReceive(Object message) {
        if ((message instanceof TransformationJob) && backends.isEmpty()) {
            TransformationJob job = (TransformationJob) message;
            getSender().tell(
                new JobFailed("Service unavailable, try again later", job),
                getSender());
        }

        } else if (message instanceof TransformationJob) {
            TransformationJob job = (TransformationJob) message;
            jobCounter++;
            backends.get(jobCounter % backends.size())
                .forward(job, getContext());
        }

        } else if (message.equals(BACKEND_REGISTRATION)) {
            getContext().watch(getSender());
            backends.add(getSender());
        }

        } else if (message instanceof Terminated) {
            Terminated terminated = (Terminated) message;
            backends.remove(terminated.getActor());
        }

        } else {
            unhandled(message);
        }
    }
}

```

```
}

```

Note that the `TransformationFrontend` actor watch the registered backend to be able to remove it from its list of available backend workers. Death watch uses the cluster failure detector for nodes in the cluster, i.e. it detects network failures and JVM crashes, in addition to graceful termination of watched actor. Death watch generates the `Terminated` message to the watching actor when the unreachable cluster node has been downed and removed.

The [Typesafe Activator](#) tutorial named [Akka Cluster Samples with Java](#). contains the full source code and instructions of how to run the **Worker Dial-in Example**.

6.2.8 Node Roles

Not all nodes of a cluster need to perform the same function: there might be one sub-set which runs the web front-end, one which runs the data access layer and one for the number-crunching. Deployment of actors—for example by cluster-aware routers—can take node roles into account to achieve this distribution of responsibilities.

The roles of a node is defined in the configuration property named `akka.cluster.roles` and it is typically defined in the start script as a system property or environment variable.

The roles of the nodes is part of the membership information in `MemberEvent` that you can subscribe to.

6.2.9 How To Startup when Cluster Size Reached

A common use case is to start actors after the cluster has been initialized, members have joined, and the cluster has reached a certain size.

With a configuration option you can define required number of members before the leader changes member status of ‘Joining’ members to ‘Up’.

```
akka.cluster.min-nr-of-members = 3

```

In a similar way you can define required number of members of a certain role before the leader changes member status of ‘Joining’ members to ‘Up’.

```
akka.cluster.role {
  frontend.min-nr-of-members = 1
  backend.min-nr-of-members = 2
}
```

You can start the actors in a `registerOnMemberUp` callback, which will be invoked when the current member status is changed to ‘Up’, i.e. the cluster has at least the defined number of members.

```
Cluster.get(system).registerOnMemberUp(new Runnable() {
  @Override
  public void run() {
    system.actorOf(Props.create(FactorialFrontend.class, upToN, true),
      "factorialFrontend");
  }
});
```

This callback can be used for other things than starting actors.

6.2.10 How To Cleanup when Member is Removed

You can do some clean up in a `registerOnMemberRemoved` callback, which will be invoked when the current member status is changed to ‘Removed’ or the cluster have been shutdown.

For example, this is how to shut down the `ActorSystem` and thereafter exit the JVM:

```

Cluster.get(system).registerOnMemberRemoved(new Runnable() {
    @Override
    public void run() {
        // exit JVM when ActorSystem has been terminated
        final Runnable exit = new Runnable() {
            @Override public void run() {
                System.exit(0);
            }
        };
        system.registerOnTermination(exit);

        // shut down ActorSystem
        system.terminate();

        // In case ActorSystem shutdown takes longer than 10 seconds,
        // exit the JVM forcefully anyway.
        // We must spawn a separate thread to not block current thread,
        // since that would have blocked the shutdown of the ActorSystem.
        new Thread() {
            @Override public void run(){
                try {
                    Await.ready(system.whenTerminated(), Duration.create(10, TimeUnit.SECONDS));
                } catch (Exception e) {
                    System.exit(-1);
                }
            }
        }.start();
    }
});

```

Note: Register a `OnMemberRemoved` callback on a cluster that have been shutdown, the callback will be invoked immediately on the caller thread, otherwise it will be invoked later when the current member status changed to 'Removed'. You may want to install some cleanup handling after the cluster was started up, but the cluster might already be shutting down when you installing, and depending on the race is not healthy.

6.2.11 Cluster Singleton

For some use cases it is convenient and sometimes also mandatory to ensure that you have exactly one actor of a certain type running somewhere in the cluster.

This can be implemented by subscribing to member events, but there are several corner cases to consider. Therefore, this specific use case is made easily accessible by the *Cluster Singleton*.

6.2.12 Cluster Sharding

Distributes actors across several nodes in the cluster and supports interaction with the actors using their logical identifier, but without having to care about their physical location in the cluster.

See *Cluster Sharding*.

6.2.13 Distributed Publish Subscribe

Publish-subscribe messaging between actors in the cluster, and point-to-point messaging using the logical path of the actors, i.e. the sender does not have to know on which node the destination actor is running.

See *distributed-pub-sub-scala*.

6.2.14 Cluster Client

Communication from an actor system that is not part of the cluster to actors running somewhere in the cluster. The client does not have to know on which node the destination actor is running.

See *Cluster Client*.

6.2.15 Distributed Data

Akka Distributed Data is useful when you need to share data between nodes in an Akka Cluster. The data is accessed with an actor providing a key-value store like API.

See *Distributed Data*.

6.2.16 Failure Detector

In a cluster each node is monitored by a few (default maximum 5) other nodes, and when any of these detects the node as `unreachable` that information will spread to the rest of the cluster through the gossip. In other words, only one node needs to mark a node `unreachable` to have the rest of the cluster mark that node `unreachable`.

The failure detector will also detect if the node becomes `reachable` again. When all nodes that monitored the `unreachable` node detects it as `reachable` again the cluster, after gossip dissemination, will consider it as `reachable`.

If system messages cannot be delivered to a node it will be quarantined and then it cannot come back from `unreachable`. This can happen if there are too many unacknowledged system messages (e.g. `watch`, `Terminated`, remote actor deployment, failures of actors supervised by remote parent). Then the node needs to be moved to the `down` or `removed` states and the actor system must be restarted before it can join the cluster again.

The nodes in the cluster monitor each other by sending heartbeats to detect if a node is unreachable from the rest of the cluster. The heartbeat arrival times is interpreted by an implementation of [The Phi Accrual Failure Detector](#).

The suspicion level of failure is given by a value called *phi*. The basic idea of the phi failure detector is to express the value of *phi* on a scale that is dynamically adjusted to reflect current network conditions.

The value of *phi* is calculated as:


```
phi = -log10(1 - F(timeSinceLastHeartbeat))
```

where *F* is the cumulative distribution function of a normal distribution with mean and standard deviation estimated from historical heartbeat inter-arrival times.


In the *Configuration* you can adjust the `akka.cluster.failure-detector.threshold` to define when a *phi* value is considered to be a failure.

A low `threshold` is prone to generate many false positives but ensures a quick detection in the event of a real crash. Conversely, a high `threshold` generates fewer mistakes but needs more time to detect actual crashes. The default `threshold` is 8 and is appropriate for most situations. However in cloud environments, such as Amazon EC2, the value could be increased to 12 in order to account for network issues that sometimes occur on such platforms.


The following chart illustrates how *phi* increase with increasing time since the previous heartbeat.



Phi is calculated from the mean and standard deviation of historical inter arrival times. The previous chart is an example for standard deviation of 200 ms. If the heartbeats arrive with less deviation the curve becomes steeper, i.e. it is possible to determine failure more quickly. The curve looks like this for a standard deviation of 100 ms.



To be able to survive sudden abnormalities, such as garbage collection pauses and transient network failures the failure detector is configured with a margin, `akka.cluster.failure-detector.acceptable-heartbeat-pause`. You may want to adjust the *Configuration* of this depending on you environment. This is how the curve looks like for `acceptable-heartbeat-pause` configured to 3 seconds.



Death watch uses the cluster failure detector for nodes in the cluster, i.e. it detects network failures and JVM crashes, in addition to graceful termination of watched actor. Death watch generates the `Terminated` message to the watching actor when the unreachable cluster node has been downed and removed.

If you encounter suspicious false positives when the system is under load you should define a separate dispatcher for the cluster actors as described in [Cluster Dispatcher](#).

6.2.17 Cluster Aware Routers

All *routers* can be made aware of member nodes in the cluster, i.e. deploying new routees or looking up routees on nodes in the cluster. When a node becomes unreachable or leaves the cluster the routees of that node are automatically unregistered from the router. When new nodes join the cluster additional routees are added to the router, according to the configuration. Routees are also added when a node becomes reachable again, after having been unreachable.

Cluster aware routers make use of members with status *WeaklyUp* if that feature is enabled.

There are two distinct types of routers.

- **Group - router that sends messages to the specified path using actor selection** The routees can be shared between routers running on different nodes in the cluster. One example of a use case for this type of router is a service running on some backend nodes in the cluster and used by routers running on front-end nodes in the cluster.
- **Pool - router that creates routees as child actors and deploys them on remote nodes.** Each router will have its own routee instances. For example, if you start a router on 3 nodes in a 10 nodes cluster you will have 30 routee actors in total if the router is configured to use one instance per node. The routees created by the different routers will not be shared between the routers. One example of a use case for this type of router is a single master that coordinate jobs and delegates the actual work to routees running on other nodes in the cluster.

Router with Group of Routees

When using a `Group` you must start the routee actors on the cluster member nodes. That is not done by the router. The configuration for a group looks like this:

```
akka.actor.deployment {
  /statsService/workerRouter {
    router = consistent-hashing-group
    routees.paths = ["/user/statsWorker"]
    cluster {
      enabled = on
      allow-local-routees = on
      use-role = compute
    }
  }
}
```

Note: The routee actors should be started as early as possible when starting the actor system, because the router will try to use them as soon as the member status is changed to 'Up'.

The actor paths without address information that are defined in `routees.paths` are used for selecting the actors to which the messages will be forwarded to by the router. Messages will be forwarded to the routees using [ActorSelection](#), so the same delivery semantics should be expected. It is possible to limit the lookup of routees to member nodes tagged with a certain role by specifying `use-role`.

`max-total-nr-of-instances` defines total number of routees in the cluster. By default `max-total-nr-of-instances` is set to a high value (10000) that will result in new routees added to the router when nodes join the cluster. Set it to a lower value if you want to limit total number of routees.

The same type of router could also have been defined in code:

```
int totalInstances = 100;
Iterable<String> routeesPaths = Collections
    .singletonList("/user/statsWorker");
boolean allowLocalRoutees = true;
String useRole = "compute";
ActorRef workerRouter = getContext().actorOf(
    new ClusterRouterGroup(new ConsistentHashingGroup(routeesPaths),
        new ClusterRouterGroupSettings(totalInstances, routeesPaths,
            allowLocalRoutees, useRole)).props(), "workerRouter2");
```

See [Configuration](#) section for further descriptions of the settings.

Router Example with Group of Routees

Let's take a look at how to use a cluster aware router with a group of routees, i.e. router sending to the paths of the routees.

The example application provides a service to calculate statistics for a text. When some text is sent to the service it splits it into words, and delegates the task to count number of characters in each word to a separate worker, a routee of a router. The character count for each word is sent back to an aggregator that calculates the average number of characters per word when all results have been collected.

Messages:

```
public interface StatsMessages {

    public static class StatsJob implements Serializable {
        private final String text;

        public StatsJob(String text) {
            this.text = text;
        }

        public String getText() {
            return text;
        }
    }
}
```

```

    }
}

public static class StatsResult implements Serializable {
    private final double meanWordLength;

    public StatsResult(double meanWordLength) {
        this.meanWordLength = meanWordLength;
    }

    public double getMeanWordLength() {
        return meanWordLength;
    }

    @Override
    public String toString() {
        return "meanWordLength: " + meanWordLength;
    }
}

public static class JobFailed implements Serializable {
    private final String reason;

    public JobFailed(String reason) {
        this.reason = reason;
    }

    public String getReason() {
        return reason;
    }

    @Override
    public String toString() {
        return "JobFailed(" + reason + ")";
    }
}
}

```

The worker that counts number of characters in each word:

```

public class StatsWorker extends UntypedActor {

    Map<String, Integer> cache = new HashMap<String, Integer>();

    @Override
    public void onReceive(Object message) {
        if (message instanceof String) {
            String word = (String) message;
            Integer length = cache.get(word);
            if (length == null) {
                length = word.length();
                cache.put(word, length);
            }
            getSender().tell(length, getSelf());
        } else {
            unhandled(message);
        }
    }
}

```

The service that receives text from users and splits it up into words, delegates to workers and aggregates:

```
public class StatsService extends UntypedActor {

    // This router is used both with lookup and deploy of routees. If you
    // have a router with only lookup of routees you can use Props.empty()
    // instead of Props.create(StatsWorker.class).
    ActorRef workerRouter = getContext().actorOf(
        FromConfig.getInstance().props(Props.create(StatsWorker.class)),
        "workerRouter");

    @Override
    public void onReceive(Object message) {
        if (message instanceof StatsJob) {
            StatsJob job = (StatsJob) message;
            if (job.getText().equals("")) {
                unhandled(message);
            } else {
                final String[] words = job.getText().split(" ");
                final ActorRef replyTo = getSender();

                // create actor that collects replies from workers
                ActorRef aggregator = getContext().actorOf(
                    Props.create(StatsAggregator.class, words.length, replyTo));

                // send each word to a worker
                for (String word : words) {
                    workerRouter.tell(new ConsistentHashableEnvelope(word, word),
                        aggregator);
                }
            }
        } else {
            unhandled(message);
        }
    }
}
```

```
public class StatsAggregator extends UntypedActor {

    final int expectedResults;
    final ActorRef replyTo;
    final List<Integer> results = new ArrayList<Integer>();

    public StatsAggregator(int expectedResults, ActorRef replyTo) {
        this.expectedResults = expectedResults;
        this.replyTo = replyTo;
    }

    @Override
    public void preStart() {
        getContext().setReceiveTimeout(Duration.create(3, TimeUnit.SECONDS));
    }

    @Override
    public void onReceive(Object message) {
        if (message instanceof Integer) {
            Integer wordCount = (Integer) message;
            results.add(wordCount);
            if (results.size() == expectedResults) {
                int sum = 0;
                for (int c : results)
                    sum += c;
            }
        }
    }
}
```

```

    double meanWordLength = ((double) sum) / results.size();
    replyTo.tell(new StatsResult(meanWordLength), getSelf());
    getContext().stop(getSelf());
  }

  } else if (message == ReceiveTimeout.getInstance()) {
    replyTo.tell(new JobFailed("Service unavailable, try again later"),
      getSelf());
    getContext().stop(getSelf());
  } else {
    unhandled(message);
  }
}
}
}

```

Note, nothing cluster specific so far, just plain actors.

All nodes start `StatsService` and `StatsWorker` actors. Remember, routees are the workers in this case. The router is configured with `routees.paths`:

```

akka.actor.deployment {
  /statsService/workerRouter {
    router = consistent-hashing-group
    routees.paths = ["/user/statsWorker"]
    cluster {
      enabled = on
      allow-local-routees = on
      use-role = compute
    }
  }
}

```

This means that user requests can be sent to `StatsService` on any node and it will use `StatsWorker` on all nodes.

The [Typesafe Activator tutorial](#) named [Akka Cluster Samples with Java](#). contains the full source code and instructions of how to run the **Router Example with Group of Routees**.

Router with Pool of Remote Deployed Routees

When using a `Pool` with routees created and deployed on the cluster member nodes the configuration for a router looks like this:

```

akka.actor.deployment {
  /statsService/singleton/workerRouter {
    router = consistent-hashing-pool
    cluster {
      enabled = on
      max-nr-of-instances-per-node = 3
      allow-local-routees = on
      use-role = compute
    }
  }
}

```

It is possible to limit the deployment of routees to member nodes tagged with a certain role by specifying `use-role`.

`max-total-nr-of-instances` defines total number of routees in the cluster, but the number of routees per node, `max-nr-of-instances-per-node`, will not be exceeded. By default

`max-total-nr-of-instances` is set to a high value (10000) that will result in new routees added to the router when nodes join the cluster. Set it to a lower value if you want to limit total number of routees.

The same type of router could also have been defined in code:

```
int totalInstances = 100;
int maxInstancesPerNode = 3;
boolean allowLocalRoutees = false;
String useRole = "compute";
ActorRef workerRouter = getContext().actorOf(
    new ClusterRouterPool(new ConsistentHashingPool(0),
        new ClusterRouterPoolSettings(totalInstances, maxInstancesPerNode,
            allowLocalRoutees, useRole)).props(Props
                .create(StatsWorker.class), "workerRouter3");
```

See [Configuration](#) section for further descriptions of the settings.

Router Example with Pool of Remote Deployed Routees

Let's take a look at how to use a cluster aware router on single master node that creates and deploys workers. To keep track of a single master we use the [Cluster Singleton](#) in the `contrib` module. The `ClusterSingletonManager` is started on each node.

```
ClusterSingletonManagerSettings settings = ClusterSingletonManagerSettings.create(system)
    .withRole("compute");
system.actorOf(ClusterSingletonManager.props(
    Props.create(StatsService.class), PoisonPill.getInstance(), settings),
    "statsService");
```

We also need an actor on each node that keeps track of where current single master exists and delegates jobs to the `StatsService`. That is provided by the `ClusterSingletonProxy`.

```
ClusterSingletonProxySettings proxySettings =
    ClusterSingletonProxySettings.create(system).withRole("compute");
system.actorOf(ClusterSingletonProxy.props("/user/statsService",
    proxySettings), "statsServiceProxy");
```

The `ClusterSingletonProxy` receives text from users and delegates to the current `StatsService`, the single master. It listens to cluster events to lookup the `StatsService` on the oldest node.

All nodes start `ClusterSingletonProxy` and the `ClusterSingletonManager`. The router is now configured like this:

```
akka.actor.deployment {
  /statsService singleton/workerRouter {
    router = consistent-hashing-pool
    cluster {
      enabled = on
      max-nr-of-instances-per-node = 3
      allow-local-routees = on
      use-role = compute
    }
  }
}
```

The [Typesafe Activator](#) tutorial named [Akka Cluster Samples with Java](#), contains the full source code and instructions of how to run the **Router Example with Pool of Remote Deployed Routees**.

6.2.18 Cluster Metrics

The member nodes of the cluster can collect system health metrics and publish that to other cluster nodes and to the registered subscribers on the system event bus with the help of [Cluster Metrics Extension](#).

6.2.19 JMX

Information and management of the cluster is available as JMX MBeans with the root name `akka.Cluster`. The JMX information can be displayed with an ordinary JMX console such as JConsole or JVisualVM.

From JMX you can:

- see what members that are part of the cluster
- see status of this node
- see roles of each member
- join this node to another node in cluster
- mark any node in the cluster as down
- tell any node in the cluster to leave

Member nodes are identified by their address, in format `akka.<protocol>://<actor-system-name>@<hostname>:<port>`.

6.2.20 Command Line Management

The cluster can be managed with the script `bin/akka-cluster` provided in the Akka distribution.

Run it without parameters to see instructions about how to use the script:

```
Usage: bin/akka-cluster <node-hostname> <jmx-port> <command> ...

Supported commands are:
  join <node-url> - Sends request a JOIN node with the specified URL
  leave <node-url> - Sends a request for node with URL to LEAVE the cluster
  down <node-url> - Sends a request for marking node with URL as DOWN
  member-status - Asks the member node for its current status
  members        - Asks the cluster for addresses of current members
  unreachable    - Asks the cluster for addresses of unreachable members
  cluster-status - Asks the cluster for its current status (member ring,
                  unavailable nodes, meta data etc.)
  leader         - Asks the cluster who the current leader is
  is-singleton   - Checks if the cluster is a singleton cluster (single
                  node cluster)
  is-available   - Checks if the member node is available

Where the <node-url> should be on the format of
  'akka.<protocol>://<actor-system-name>@<hostname>:<port>'

Examples: bin/akka-cluster localhost 9999 is-available
          bin/akka-cluster localhost 9999 join akka.tcp://MySystem@darkstar:2552
          bin/akka-cluster localhost 9999 cluster-status
```

To be able to use the script you must enable remote monitoring and management when starting the JVMs of the cluster nodes, as described in [Monitoring and Management Using JMX Technology](#)

Example of system properties to enable remote monitoring and management:

```
java -Dcom.sun.management.jmxremote.port=9999 \
-Dcom.sun.management.jmxremote.authenticate=false \
-Dcom.sun.management.jmxremote.ssl=false
```

6.2.21 Configuration

There are several configuration properties for the cluster. We refer to the [reference configuration](#) for more information.

Cluster Info Logging

You can silence the logging of cluster events at info level with configuration property:

```
akka.cluster.log-info = off
```

Cluster Dispatcher

Under the hood the cluster extension is implemented with actors and it can be necessary to create a bulkhead for those actors to avoid disturbance from other actors. Especially the heartbeating actors that is used for failure detection can generate false positives if they are not given a chance to run at regular intervals. For this purpose you can define a separate dispatcher to be used for the cluster actors:

```
akka.cluster.use-dispatcher = cluster-dispatcher

cluster-dispatcher {
  type = "Dispatcher"
  executor = "fork-join-executor"
  fork-join-executor {
    parallelism-min = 2
    parallelism-max = 4
  }
}
```

6.3 Cluster Singleton

For some use cases it is convenient and sometimes also mandatory to ensure that you have exactly one actor of a certain type running somewhere in the cluster.

Some examples:

- single point of responsibility for certain cluster-wide consistent decisions, or coordination of actions across the cluster system
- single entry point to an external system
- single master, many workers
- centralized naming service, or routing logic

Using a singleton should not be the first design choice. It has several drawbacks, such as single-point of bottleneck. Single-point of failure is also a relevant concern, but for some cases this feature takes care of that by making sure that another singleton instance will eventually be started.

The cluster singleton pattern is implemented by `akka.cluster.singleton.ClusterSingletonManager`. It manages one singleton actor instance among all cluster nodes or a group of nodes tagged with a specific role. `ClusterSingletonManager` is an actor that is supposed to be started on all nodes, or all nodes with specified role, in the cluster. The actual singleton actor is started by the `ClusterSingletonManager` on the oldest node by creating a child actor from supplied `Props`. `ClusterSingletonManager` makes sure that at most one singleton instance is running at any point in time.

The singleton actor is always running on the oldest member with specified role. The oldest member is determined by `akka.cluster.Member#isOlderThan`. This can change when removing that member from the cluster. Be aware that there is a short time period when there is no active singleton during the hand-over process.

The cluster failure detector will notice when oldest node becomes unreachable due to things like JVM crash, hard shut down, or network failure. Then a new oldest node will take over and a new singleton actor is created. For these failure scenarios there will not be a graceful hand-over, but more than one active singletons is prevented by all reasonable means. Some corner cases are eventually resolved by configurable timeouts.

You can access the singleton actor by using the provided `akka.cluster.singleton.ClusterSingletonProxy`, which will route all messages to the current instance of the singleton. The proxy will keep track of the oldest node

in the cluster and resolve the singleton's `ActorRef` by explicitly sending the singleton's `actorSelection` the `akka.actor.Identify` message and waiting for it to reply. This is performed periodically if the singleton doesn't reply within a certain (configurable) time. Given the implementation, there might be periods of time during which the `ActorRef` is unavailable, e.g., when a node leaves the cluster. In these cases, the proxy will buffer the messages sent to the singleton and then deliver them when the singleton is finally available. If the buffer is full the `ClusterSingletonProxy` will drop old messages when new messages are sent via the proxy. The size of the buffer is configurable and it can be disabled by using a buffer size of 0.

It's worth noting that messages can always be lost because of the distributed nature of these actors. As always, additional logic should be implemented in the singleton (acknowledgement) and in the client (retry) actors to ensure at-least-once message delivery.

The singleton instance will not run on members with status *WeaklyUp* if that feature is enabled.

6.3.1 Potential problems to be aware of

This pattern may seem to be very tempting to use at first, but it has several drawbacks, some of them are listed below:

- the cluster singleton may quickly become a *performance bottleneck*,
- you can not rely on the cluster singleton to be *non-stop* available — e.g. when the node on which the singleton has been running dies, it will take a few seconds for this to be noticed and the singleton be migrated to another node,
- in the case of a *network partition* appearing in a Cluster that is using Automatic Downing (see docs for *Automatic vs. Manual Downing*), it may happen that the isolated clusters each decide to spin up their own singleton, meaning that there might be multiple singletons running in the system, yet the Clusters have no way of finding out about them (because of the partition).

Especially the last point is something you should be aware of — in general when using the Cluster Singleton pattern you should take care of downing nodes yourself and not rely on the timing based auto-down feature.

Warning: Be very careful when using Cluster Singleton together with Automatic Downing, since it allows the cluster to split up into two separate clusters, which in turn will result in *multiple Singletons* being started, one in each separate cluster!

6.3.2 An Example

Assume that we need one single entry point to an external system. An actor that receives messages from a JMS queue with the strict requirement that only one JMS consumer must exist to be make sure that the messages are processed in order. That is perhaps not how one would like to design things, but a typical real-world scenario when integrating with external systems.

On each node in the cluster you need to start the `ClusterSingletonManager` and supply the `Props` of the singleton actor, in this case the JMS queue consumer.

```
final ClusterSingletonManagerSettings settings =
    ClusterSingletonManagerSettings.create(system).withRole("worker");
system.actorOf(ClusterSingletonManager.props(
    Props.create(Consumer.class, queue, testActor),
    new End(), settings), "consumer");
```

Here we limit the singleton to nodes tagged with the "worker" role, but all nodes, independent of role, can be used by not specifying `withRole`.

Here we use an application specific `terminationMessage` to be able to close the resources before actually stopping the singleton actor. Note that `PoisonPill` is a perfectly fine `terminationMessage` if you only need to stop the actor.

With the names given above, access to the singleton can be obtained from any cluster node using a properly configured proxy.

```
ClusterSingletonProxySettings proxySettings =
    ClusterSingletonProxySettings.create(system).withRole("worker");
system.actorOf(ClusterSingletonProxy.props("/user/consumer", proxySettings),
    "consumerProxy");
```

A more comprehensive sample is available in the [Typesafe Activator tutorial](#) named [Distributed workers with Akka and Java!](#).

6.3.3 Dependencies

To use the Cluster Singleton you must add the following dependency in your project.

sbt:

```
"com.typesafe.akka" %% "akka-cluster-tools" % "2.4.0"
```

maven:

```
<dependency>
  <groupId>com.typesafe.akka</groupId>
  <artifactId>akka-cluster-tools_2.11</artifactId>
  <version>2.4.0</version>
</dependency>
```

6.3.4 Configuration

The following configuration properties are read by the `ClusterSingletonManagerSettings` when created with a `ActorSystem` parameter. It is also possible to amend the `ClusterSingletonManagerSettings` or create it from another config section with the same layout as below. `ClusterSingletonManagerSettings` is a parameter to the `ClusterSingletonManager.props` factory method, i.e. each singleton can be configured with different settings if needed.

```
akka.cluster.singleton {
  # The actor name of the child singleton actor.
  singleton-name = "singleton"

  # Singleton among the nodes tagged with specified role.
  # If the role is not specified it's a singleton among all nodes in the cluster.
  role = ""

  # When a node is becoming oldest it sends hand-over request to previous oldest,
  # that might be leaving the cluster. This is retried with this interval until
  # the previous oldest confirms that the hand over has started or the previous
  # oldest member is removed from the cluster (+ akka.cluster.down-removal-margin).
  hand-over-retry-interval = 1s

  # The number of retries are derived from hand-over-retry-interval and
  # akka.cluster.down-removal-margin (or ClusterSingletonManagerSettings.removalMargin),
  # but it will never be less than this property.
  min-number-of-hand-over-retries = 10
}
```

The following configuration properties are read by the `ClusterSingletonProxySettings` when created with a `ActorSystem` parameter. It is also possible to amend the `ClusterSingletonProxySettings` or create it from another config section with the same layout as below. `ClusterSingletonProxySettings` is a parameter to the `ClusterSingletonProxy.props` factory method, i.e. each singleton proxy can be configured with different settings if needed.

```

akka.cluster.singleton-proxy {
  # The actor name of the singleton actor that is started by the ClusterSingletonManager
  singleton-name = ${akka.cluster.singleton.singleton-name}

  # The role of the cluster nodes where the singleton can be deployed.
  # If the role is not specified then any node will do.
  role = ""

  # Interval at which the proxy will try to resolve the singleton instance.
  singleton-identification-interval = 1s

  # If the location of the singleton is unknown the proxy will buffer this
  # number of messages and deliver them when the singleton is identified.
  # When the buffer is full old messages will be dropped when new messages are
  # sent via the proxy.
  # Use 0 to disable buffering, i.e. messages will be dropped immediately if
  # the location of the singleton is unknown.
  # Maximum allowed buffer size is 10000.
  buffer-size = 1000
}

```

6.4 Distributed Publish Subscribe in Cluster

How do I send a message to an actor without knowing which node it is running on?

How do I send messages to all actors in the cluster that have registered interest in a named topic?

This pattern provides a mediator actor, `akka.cluster.pubsub.DistributedPubSubMediator`, that manages a registry of actor references and replicates the entries to peer actors among all cluster nodes or a group of nodes tagged with a specific role.

The `DistributedPubSubMediator` actor is supposed to be started on all nodes, or all nodes with specified role, in the cluster. The mediator can be started with the `DistributedPubSub` extension or as an ordinary actor.

The registry is eventually consistent, i.e. changes are not immediately visible at other nodes, but typically they will be fully replicated to all other nodes after a few seconds. Changes are only performed in the own part of the registry and those changes are versioned. Deltas are disseminated in a scalable way to other nodes with a gossip protocol.

Cluster members with status *WeaklyUp*, if that feature is enabled, will participate in Distributed Publish Subscribe, i.e. subscribers on nodes with *WeaklyUp* status will receive published messages if the publisher and subscriber are on same side of a network partition.

You can send messages via the mediator on any node to registered actors on any other node.

There are two different modes of message delivery, explained in the sections *Publish* and *Send* below.

6.4.1 Publish

This is the true pub/sub mode. A typical usage of this mode is a chat room in an instant messaging application.

Actors are registered to a named topic. This enables many subscribers on each node. The message will be delivered to all subscribers of the topic.

For efficiency the message is sent over the wire only once per node (that has a matching topic), and then delivered to all subscribers of the local topic representation.

You register actors to the local mediator with `DistributedPubSubMediator.Subscribe`. Successful `Subscribe` and `Unsubscribe` is acknowledged with `DistributedPubSubMediator.SubscribeAck` and `DistributedPubSubMediator.UnsubscribeAck` replies. The acknowledgment means that the subscription is registered, but it can still take some time until it is replicated to other nodes.

You publish messages by sending `DistributedPubSubMediator.Publish` message to the local mediator.

Actors are automatically removed from the registry when they are terminated, or you can explicitly remove entries with `DistributedPubSubMediator.Unsubscribe`.

An example of a subscriber actor:

```
public class Subscriber extends UntypedActor {
    LoggingAdapter log = Logging.getLogger(getContext().system(), this);

    public Subscriber() {
        ActorRef mediator =
            DistributedPubSub.get(getContext().system()).mediator();
        // subscribe to the topic named "content"
        mediator.tell(new DistributedPubSubMediator.Subscribe("content", getSelf()),
            getSelf());
    }

    public void onReceive(Object msg) {
        if (msg instanceof String)
            log.info("Got: {}", msg);
        else if (msg instanceof DistributedPubSubMediator.SubscribeAck)
            log.info("subscribing");
        else
            unhandled(msg);
    }
}
```

Subscriber actors can be started on several nodes in the cluster, and all will receive messages published to the “content” topic.

```
system.actorOf(Props.create(Subscriber.class), "subscriber1");
//another node
system.actorOf(Props.create(Subscriber.class), "subscriber2");
system.actorOf(Props.create(Subscriber.class), "subscriber3");
```

A simple actor that publishes to this “content” topic:

```
public class Publisher extends UntypedActor {

    // activate the extension
    ActorRef mediator =
        DistributedPubSub.get(getContext().system()).mediator();

    public void onReceive(Object msg) {
        if (msg instanceof String) {
            String in = (String) msg;
            String out = in.toUpperCase();
            mediator.tell(new DistributedPubSubMediator.Publish("content", out),
                getSelf());
        } else {
            unhandled(msg);
        }
    }
}
```

It can publish messages to the topic from anywhere in the cluster:

```
//somewhere else
ActorRef publisher = system.actorOf(Props.create(Publisher.class), "publisher");
// after a while the subscriptions are replicated
publisher.tell("hello", null);
```

Actors may also be subscribed to a named topic with an optional property (`group`). If subscribing with a group name, each message published to a topic with the (`sendOneMessageToEachGroup`) flag is delivered via the

supplied `RoutingLogic` (default `random`) to one actor within each subscribing group. If all the subscribed actors have the same group name, then this works just like `Send` and all messages are delivered to one subscriber. If all the subscribed actors have different group names, then this works like normal `Publish` and all messages are broadcast to all subscribers.

6.4.2 Send

This is a point-to-point mode where each message is delivered to one destination, but you still does not have to know where the destination is located. A typical usage of this mode is private chat to one other user in an instant messaging application. It can also be used for distributing tasks to registered workers, like a cluster aware router where the routees dynamically can register themselves.

The message will be delivered to one recipient with a matching path, if any such exists in the registry. If several entries match the path because it has been registered on several nodes the message will be sent via the supplied `RoutingLogic` (default `random`) to one destination. The `sender()` of the message can specify that local affinity is preferred, i.e. the message is sent to an actor in the same local actor system as the used mediator actor, if any such exists, otherwise route to any other matching entry.

You register actors to the local mediator with `DistributedPubSubMediator.Put`. The `ActorRef` in `Put` must belong to the same local actor system as the mediator. The path without address information is the key to which you send messages. On each node there can only be one actor for a given path, since the path is unique within one local actor system.

You send messages by sending `DistributedPubSubMediator.Send` message to the local mediator with the path (without address information) of the destination actors.

Actors are automatically removed from the registry when they are terminated, or you can explicitly remove entries with `DistributedPubSubMediator.Remove`.

An example of a destination actor:

```
public class Destination extends UntypedActor {
    LoggingAdapter log = Logging.getLogger(getContext().system(), this);

    public Destination() {
        ActorRef mediator =
            DistributedPubSub.get(getContext().system()).mediator();
        // register to the path
        mediator.tell(new DistributedPubSubMediator.Put(getSelf()), getSelf());
    }

    public void onReceive(Object msg) {
        if (msg instanceof String)
            log.info("Got: {}", msg);
        else if (msg instanceof DistributedPubSubMediator.SubscribeAck)
            log.info("subscribing");
        else
            unhandled(msg);
    }
}
```

Subscriber actors can be started on several nodes in the cluster, and all will receive messages published to the “content” topic.

```
system.actorOf(Props.create(Destination.class), "destination");
//another node
system.actorOf(Props.create(Destination.class), "destination");
```

A simple actor that publishes to this “content” topic:

```
public class Sender extends UntypedActor {

    // activate the extension
```

```

ActorRef mediator =
    DistributedPubSub.get(getContext().system()).mediator();

public void onReceive(Object msg) {
    if (msg instanceof String) {
        String in = (String) msg;
        String out = in.toUpperCase();
        boolean localAffinity = true;
        mediator.tell(new DistributedPubSubMediator.Send("/user/destination", out,
            localAffinity), getSelf());
    } else {
        unhandled(msg);
    }
}
}

```

It can publish messages to the topic from anywhere in the cluster:

```

//somewhere else
ActorRef sender = system.actorOf(Props.create(Publisher.class), "sender");
// after a while the destinations are replicated
sender.tell("hello", null);

```

It is also possible to broadcast messages to the actors that have been registered with `Put`. Send `DistributedPubSubMediator.SendToAll` message to the local mediator and the wrapped message will then be delivered to all recipients with a matching path. Actors with the same path, without address information, can be registered on different nodes. On each node there can only be one such actor, since the path is unique within one local actor system.

Typical usage of this mode is to broadcast messages to all replicas with the same path, e.g. 3 actors on different nodes that all perform the same actions, for redundancy. You can also optionally specify a property (`allButSelf`) deciding if the message should be sent to a matching path on the self node or not.

6.4.3 DistributedPubSub Extension

In the example above the mediator is started and accessed with the `akka.cluster.pubsub.DistributedPubSub` extension. That is convenient and perfectly fine in most cases, but it can be good to know that it is possible to start the mediator actor as an ordinary actor and you can have several different mediators at the same time to be able to divide a large number of actors/topics to different mediators. For example you might want to use different cluster roles for different mediators.

The `DistributedPubSub` extension can be configured with the following properties:

```

# Settings for the DistributedPubSub extension
akka.cluster.pub-sub {
    # Actor name of the mediator actor, /system/distributedPubSubMediator
    name = distributedPubSubMediator

    # Start the mediator on members tagged with this role.
    # All members are used if undefined or empty.
    role = ""

    # The routing logic to use for 'Send'
    # Possible values: random, round-robin, broadcast
    routing-logic = random

    # How often the DistributedPubSubMediator should send out gossip information
    gossip-interval = 1s

    # Removed entries are pruned after this duration
    removed-time-to-live = 120s
}

```

```

# Maximum number of elements to transfer in one message when synchronizing the registries.
# Next chunk will be transferred in next round of gossip.
max-delta-elements = 3000

# The id of the dispatcher to use for DistributedPubSubMediator actors.
# If not specified default dispatcher is used.
# If specified you need to define the settings of the actual dispatcher.
use-dispatcher = ""

}

```

It is recommended to load the extension when the actor system is started by defining it in `akka.extensions` configuration property. Otherwise it will be activated when first used and then it takes a while for it to be populated.

```
akka.extensions = ["akka.cluster.pubsub.DistributedPubSub"]
```

6.4.4 Dependencies

To use the Cluster Singleton you must add the following dependency in your project.

sbt:

```
"com.typesafe.akka" %% "akka-cluster-tools" % "2.4.0"
```

maven:

```

<dependency>
  <groupId>com.typesafe.akka</groupId>
  <artifactId>akka-cluster-tools_2.11</artifactId>
  <version>2.4.0</version>
</dependency>

```

6.5 Cluster Client

An actor system that is not part of the cluster can communicate with actors somewhere in the cluster via this `ClusterClient`. The client can of course be part of another cluster. It only needs to know the location of one (or more) nodes to use as initial contact points. It will establish a connection to a `ClusterReceptionist` somewhere in the cluster. It will monitor the connection to the receptionist and establish a new connection if the link goes down. When looking for a new receptionist it uses fresh contact points retrieved from previous establishment, or periodically refreshed contacts, i.e. not necessarily the initial contact points.

Note: `ClusterClient` should not be used when sending messages to actors that run within the same cluster. Similar functionality as the `ClusterClient` is provided in a more efficient way by *Distributed Publish Subscribe in Cluster* for actors that belong to the same cluster.

Also, note it's necessary to change `akka.actor.provider` from `akka.actor.LocalActorRefProvider` to `akka.remote.RemoteActorRefProvider` or `akka.cluster.ClusterActorRefProvider` when using the cluster client.

The receptionist is supposed to be started on all nodes, or all nodes with specified role, in the cluster. The receptionist can be started with the `ClusterClientReceptionist` extension or as an ordinary actor.

You can send messages via the `ClusterClient` to any actor in the cluster that is registered in the `DistributedPubSubMediator` used by the `ClusterReceptionist`. The `ClusterClientReceptionist` provides methods for registration of actors that should be reachable from the client. Messages are wrapped in `ClusterClient.Send`, `ClusterClient.SendToAll` or `ClusterClient.Publish`.

1. ClusterClient.Send

The message will be delivered to one recipient with a matching path, if any such exists. If several entries match the path the message will be delivered to one random destination. The `sender()` of the message can specify that local affinity is preferred, i.e. the message is sent to an actor in the same local actor system as the used receptionist actor, if any such exists, otherwise random to any other matching entry.

2. `ClusterClient.SendToAll`

The message will be delivered to all recipients with a matching path.

3. `ClusterClient.Publish`

The message will be delivered to all recipients Actors that have been registered as subscribers to the named topic.

Response messages from the destination actor are tunneled via the receptionist to avoid inbound connections from other cluster nodes to the client, i.e. the `sender()`, as seen by the destination actor, is not the client itself. The `sender()` of the response messages, as seen by the client, is `deadLetters` since the client should normally send subsequent messages via the `ClusterClient`. It is possible to pass the the original sender inside the reply messages if the client is supposed to communicate directly to the actor in the cluster.

While establishing a connection to a receptionist the `ClusterClient` will buffer messages and send them when the connection is established. If the buffer is full the `ClusterClient` will drop old messages when new messages are sent via the client. The size of the buffer is configurable and it can be disabled by using a buffer size of 0.

It's worth noting that messages can always be lost because of the distributed nature of these actors. As always, additional logic should be implemented in the destination (acknowledgement) and in the client (retry) actors to ensure at-least-once message delivery.

6.5.1 An Example

On the cluster nodes first start the receptionist. Note, it is recommended to load the extension when the actor system is started by defining it in the `akka.extensions` configuration property:

```
akka.extensions = ["akka.cluster.client.ClusterClientReceptionist"]
```

Next, register the actors that should be available for the client.

```
ActorRef serviceA = system.actorOf(Props.create(Service.class), "serviceA");
ClusterClientReceptionist.get(system).registerService(serviceA);

ActorRef serviceB = system.actorOf(Props.create(Service.class), "serviceB");
ClusterClientReceptionist.get(system).registerService(serviceB);
```

On the client you create the `ClusterClient` actor and use it as a gateway for sending messages to the actors identified by their path (without address information) somewhere in the cluster.

```
final ActorRef c = system.actorOf(ClusterClient.props(
    ClusterClientSettings.create(system).withInitialContacts(initialContacts()),
    "client");
c.tell(new ClusterClient.Send("/user/serviceA", "hello", true), ActorRef.noSender());
c.tell(new ClusterClient.SendToAll("/user/serviceB", "hi"), ActorRef.noSender());
```

The `initialContacts` parameter is a `Set<ActorPath>`, which can be created like this:

```
Set<ActorPath> initialContacts() {
    return new HashSet<ActorPath>(Arrays.asList(
        ActorPaths.fromString("akka.tcp://OtherSys@host1:2552/system/receptionist"),
        ActorPaths.fromString("akka.tcp://OtherSys@host2:2552/system/receptionist")));
}
```

You will probably define the address information of the initial contact points in configuration or system property. See also [Configuration](#).

A more comprehensive sample is available in the [Typesafe Activator tutorial](#) named [Distributed workers with Akka and Java!](#)

6.5.2 ClusterClientReceptionist Extension

In the example above the receptionist is started and accessed with the `akka.cluster.client.ClusterClientReceptionist` extension. That is convenient and perfectly fine in most cases, but it can be good to know that it is possible to start the `akka.cluster.client.ClusterReceptionist` actor as an ordinary actor and you can have several different receptionists at the same time, serving different types of clients.

Note that the `ClusterClientReceptionist` uses the `DistributedPubSub` extension, which is described in *Distributed Publish Subscribe in Cluster*.

It is recommended to load the extension when the actor system is started by defining it in the `akka.extensions` configuration property:

```
akka.extensions = ["akka.cluster.client.ClusterClientReceptionist"]
```

6.5.3 Dependencies

To use the Cluster Client you must add the following dependency in your project.

sbt:

```
"com.typesafe.akka" %% "akka-cluster-tools" % "2.4.0"
```

maven:

```
<dependency>
  <groupId>com.typesafe.akka</groupId>
  <artifactId>akka-cluster-tools_2.11</artifactId>
  <version>2.4.0</version>
</dependency>
```

6.5.4 Configuration

The `ClusterClientReceptionist` extension (or `ClusterReceptionistSettings`) can be configured with the following properties:

```
# Settings for the ClusterClientReceptionist extension
akka.cluster.client.receptionist {
  # Actor name of the ClusterReceptionist actor, /system/receptionist
  name = receptionist

  # Start the receptionist on members tagged with this role.
  # All members are used if undefined or empty.
  role = ""

  # The receptionist will send this number of contact points to the client
  number-of-contacts = 3

  # The actor that tunnel response messages to the client will be stopped
  # after this time of inactivity.
  response-tunnel-receive-timeout = 30s

  # The id of the dispatcher to use for ClusterReceptionist actors.
  # If not specified default dispatcher is used.
  # If specified you need to define the settings of the actual dispatcher.
  use-dispatcher = ""
}
```

The following configuration properties are read by the `ClusterClientSettings` when created with a `ActorSystem` parameter. It is also possible to amend the `ClusterClientSettings` or create it from

another config section with the same layout as below. `ClusterClientSettings` is a parameter to the `ClusterClient.props` factory method, i.e. each client can be configured with different settings if needed.

```
# Settings for the ClusterClient
akka.cluster.client {
  # Actor paths of the ClusterReceptionist actors on the servers (cluster nodes)
  # that the client will try to contact initially. It is mandatory to specify
  # at least one initial contact.
  # Comma separated full actor paths defined by a string on the form of
  # "akka.tcp://system@hostname:port/system/receptionist"
  initial-contacts = []

  # Interval at which the client retries to establish contact with one of
  # ClusterReceptionist on the servers (cluster nodes)
  establishing-get-contacts-interval = 3s

  # Interval at which the client will ask the ClusterReceptionist for
  # new contact points to be used for next reconnect.
  refresh-contacts-interval = 60s

  # How often failure detection heartbeat messages should be sent
  heartbeat-interval = 2s

  # Number of potentially lost/delayed heartbeats that will be
  # accepted before considering it to be an anomaly.
  # The ClusterClient is using the akka.remote.DeadlineFailureDetector, which
  # will trigger if there are no heartbeats within the duration
  # heartbeat-interval + acceptable-heartbeat-pause, i.e. 15 seconds with
  # the default settings.
  acceptable-heartbeat-pause = 13s

  # If connection to the receptionist is not established the client will buffer
  # this number of messages and deliver them the connection is established.
  # When the buffer is full old messages will be dropped when new messages are sent
  # via the client. Use 0 to disable buffering, i.e. messages will be dropped
  # immediately if the location of the singleton is unknown.
  # Maximum allowed buffer size is 10000.
  buffer-size = 1000
}
```

6.6 Cluster Sharding

Cluster sharding is useful when you need to distribute actors across several nodes in the cluster and want to be able to interact with them using their logical identifier, but without having to care about their physical location in the cluster, which might also change over time.

It could for example be actors representing Aggregate Roots in Domain-Driven Design terminology. Here we call these actors “entities”. These actors typically have persistent (durable) state, but this feature is not limited to actors with persistent state.

Cluster sharding is typically used when you have many stateful actors that together consume more resources (e.g. memory) than fit on one machine. If you only have a few stateful actors it might be easier to run them on a *Cluster Singleton* node.

In this context sharding means that actors with an identifier, so called entities, can be automatically distributed across multiple nodes in the cluster. Each entity actor runs only at one place, and messages can be sent to the entity without requiring the sender to know the location of the destination actor. This is achieved by sending the messages via a `ShardRegion` actor provided by this extension, which knows how to route the message with the entity id to the final destination.

Cluster sharding will not be active on members with status *WeaklyUp* if that feature is enabled.

6.6.1 An Example

This is how an entity actor may look like:

```
public class Counter extends UntypedPersistentActor {

    public static enum CounterOp {
        INCREMENT, DECREMENT
    }

    public static class Get {
        final public long counterId;

        public Get(long counterId) {
            this.counterId = counterId;
        }
    }

    public static class EntityEnvelope {
        final public long id;
        final public Object payload;

        public EntityEnvelope(long id, Object payload) {
            this.id = id;
            this.payload = payload;
        }
    }

    public static class CounterChanged {
        final public int delta;

        public CounterChanged(int delta) {
            this.delta = delta;
        }
    }

    int count = 0;

    // getSelf().path().name() is the entity identifier (utf-8 URL-encoded)
    @Override
    public String persistenceId() {
        return "Counter-" + getSelf().path().name();
    }

    @Override
    public void preStart() throws Exception {
        super.preStart();
        context().setReceiveTimeout(Duration.create(120, TimeUnit.SECONDS));
    }

    void updateState(CounterChanged event) {
        count += event.delta;
    }

    @Override
    public void onReceiveRecover(Object msg) {
        if (msg instanceof CounterChanged)
            updateState((CounterChanged) msg);
        else
            unhandled(msg);
    }

    @Override
```

```

public void onReceiveCommand(Object msg) {
    if (msg instanceof Get)
        getSender().tell(count, getSelf());

    else if (msg == CounterOp.INCREMENT)
        persist(new CounterChanged(+1), new Procedure<CounterChanged>() {
            public void apply(CounterChanged evt) {
                updateState(evt);
            }
        });

    else if (msg == CounterOp.DECREMENT)
        persist(new CounterChanged(-1), new Procedure<CounterChanged>() {
            public void apply(CounterChanged evt) {
                updateState(evt);
            }
        });

    else if (msg.equals(ReceiveTimeout.getInstance()))
        getContext().parent().tell(
            new ShardRegion.Passivate(PoisonPill.getInstance()), getSelf());

    else
        unhandled(msg);
}
}

```

The above actor uses event sourcing and the support provided in `UntypedPersistentActor` to store its state. It does not have to be a persistent actor, but in case of failure or migration of entities between nodes it must be able to recover its state if it is valuable.

Note how the `persistenceId` is defined. The name of the actor is the entity identifier (utf-8 URL-encoded). You may define it another way, but it must be unique.

When using the sharding extension you are first, typically at system startup on each node in the cluster, supposed to register the supported entity types with the `ClusterSharding.start` method. `ClusterSharding.start` gives you the reference which you can pass along.

```

Option<String> roleOption = Option.none();
ClusterShardingSettings settings = ClusterShardingSettings.create(system);
ActorRef startedCounterRegion = ClusterSharding.get(system).start("Counter",
    Props.create(Counter.class), settings, messageExtractor);

```

The `messageExtractor` defines application specific methods to extract the entity identifier and the shard identifier from incoming messages.

```

ShardRegion.MessageExtractor messageExtractor = new ShardRegion.MessageExtractor() {

    @Override
    public String entityId(Object message) {
        if (message instanceof Counter.EntityEnvelope)
            return String.valueOf(((Counter.EntityEnvelope) message).id);
        else if (message instanceof Counter.Get)
            return String.valueOf(((Counter.Get) message).counterId);
        else
            return null;
    }

    @Override
    public Object entityMessage(Object message) {
        if (message instanceof Counter.EntityEnvelope)
            return ((Counter.EntityEnvelope) message).payload;
        else
            return message;
    }
}

```

```

}

@Override
public String shardId(Object message) {
    int numberOfShards = 100;
    if (message instanceof Counter.EntityEnvelope) {
        long id = ((Counter.EntityEnvelope) message).id;
        return String.valueOf(id % numberOfShards);
    } else if (message instanceof Counter.Get) {
        long id = ((Counter.Get) message).counterId;
        return String.valueOf(id % numberOfShards);
    } else {
        return null;
    }
}
};

```

This example illustrates two different ways to define the entity identifier in the messages:

- The `Get` message includes the identifier itself.
- The `EntityEnvelope` holds the identifier, and the actual message that is sent to the entity actor is wrapped in the envelope.

Note how these two messages types are handled in the `entityId` and `entityMessage` methods shown above. The message sent to the entity actor is what `entityMessage` returns and that makes it possible to unwrap envelopes if needed.

A shard is a group of entities that will be managed together. The grouping is defined by the `extractShardId` function shown above. For a specific entity identifier the shard identifier must always be the same. Otherwise the entity actor might accidentally be started in several places at the same time.

Creating a good sharding algorithm is an interesting challenge in itself. Try to produce a uniform distribution, i.e. same amount of entities in each shard. As a rule of thumb, the number of shards should be a factor ten greater than the planned maximum number of cluster nodes. Less shards than number of nodes will result in that some nodes will not host any shards. Too many shards will result in less efficient management of the shards, e.g. rebalancing overhead, and increased latency because the coordinator is involved in the routing of the first message for each shard. The sharding algorithm must be the same on all nodes in a running cluster. It can be changed after stopping all nodes in the cluster.

A simple sharding algorithm that works fine in most cases is to take the absolute value of the `hashCode` of the entity identifier modulo number of shards. As a convenience this is provided by the `ShardRegion.HashCodeMessageExtractor`.

Messages to the entities are always sent via the local `ShardRegion`. The `ShardRegion` actor reference for a named entity type is returned by `ClusterSharding.start` and it can also be retrieved with `ClusterSharding.shardRegion`. The `ShardRegion` will lookup the location of the shard for the entity if it does not already know its location. It will delegate the message to the right node and it will create the entity actor on demand, i.e. when the first message for a specific entity is delivered.

```

ActorRef counterRegion = ClusterSharding.get(system).shardRegion("Counter");
counterRegion.tell(new Counter.Get(123), getSelf());

counterRegion.tell(new Counter.EntityEnvelope(123,
    Counter.CounterOp.INCREMENT), getSelf());
counterRegion.tell(new Counter.Get(123), getSelf());

```

6.6.2 How it works

The `ShardRegion` actor is started on each node in the cluster, or group of nodes tagged with a specific role. The `ShardRegion` is created with two application specific functions to extract the entity identifier and the shard

identifier from incoming messages. A shard is a group of entities that will be managed together. For the first message in a specific shard the `ShardRegion` request the location of the shard from a central coordinator, the `ShardCoordinator`.

The `ShardCoordinator` decides which `ShardRegion` shall own the `Shard` and informs that `ShardRegion`. The region will confirm this request and create the `Shard` supervisor as a child actor. The individual `Entities` will then be created when needed by the `Shard` actor. Incoming messages thus travel via the `ShardRegion` and the `Shard` to the target `Entity`.

If the shard home is another `ShardRegion` instance messages will be forwarded to that `ShardRegion` instance instead. While resolving the location of a shard incoming messages for that shard are buffered and later delivered when the shard home is known. Subsequent messages to the resolved shard can be delivered to the target destination immediately without involving the `ShardCoordinator`.

Scenario 1:

1. Incoming message M1 to `ShardRegion` instance R1.
2. M1 is mapped to shard S1. R1 doesn't know about S1, so it asks the coordinator C for the location of S1.
3. C answers that the home of S1 is R1.
4. R1 creates child actor for the entity E1 and sends buffered messages for S1 to E1 child
5. All incoming messages for S1 which arrive at R1 can be handled by R1 without C. It creates entity children as needed, and forwards messages to them.

Scenario 2:

1. Incoming message M2 to R1.
2. M2 is mapped to S2. R1 doesn't know about S2, so it asks C for the location of S2.
3. C answers that the home of S2 is R2.
4. R1 sends buffered messages for S2 to R2
5. All incoming messages for S2 which arrive at R1 can be handled by R1 without C. It forwards messages to R2.
6. R2 receives message for S2, ask C, which answers that the home of S2 is R2, and we are in Scenario 1 (but for R2).

To make sure that at most one instance of a specific entity actor is running somewhere in the cluster it is important that all nodes have the same view of where the shards are located. Therefore the shard allocation decisions are taken by the central `ShardCoordinator`, which is running as a cluster singleton, i.e. one instance on the oldest member among all cluster nodes or a group of nodes tagged with a specific role.

The logic that decides where a shard is to be located is defined in a pluggable shard allocation strategy. The default implementation `ShardCoordinator.LeastShardAllocationStrategy` allocates new shards to the `ShardRegion` with least number of previously allocated shards. This strategy can be replaced by an application specific implementation.

To be able to use newly added members in the cluster the coordinator facilitates rebalancing of shards, i.e. migrate entities from one node to another. In the rebalance process the coordinator first notifies all `ShardRegion` actors that a handoff for a shard has started. That means they will start buffering incoming messages for that shard, in the same way as if the shard location is unknown. During the rebalance process the coordinator will not answer any requests for the location of shards that are being rebalanced, i.e. local buffering will continue until the handoff is completed. The `ShardRegion` responsible for the rebalanced shard will stop all entities in that shard by sending the specified `handOffStopMessage` (default `PoisonPill`) to them. When all entities have been terminated the `ShardRegion` owning the entities will acknowledge the handoff as completed to the coordinator. Thereafter the coordinator will reply to requests for the location of the shard and thereby allocate a new home for the shard and then buffered messages in the `ShardRegion` actors are delivered to the new location. This means that the state of the entities are not transferred or migrated. If the state of the entities are of importance it should be persistent (durable), e.g. with *Persistence*, so that it can be recovered at the new location.

The logic that decides which shards to rebalance is defined in a pluggable shard allocation strategy. The default implementation `ShardCoordinator.LeastShardAllocationStrategy` picks shards for hand-off from the `ShardRegion` with most number of previously allocated shards. They will then be allocated to the `ShardRegion` with least number of previously allocated shards, i.e. new members in the cluster. There is a configurable threshold of how large the difference must be to begin the rebalancing. This strategy can be replaced by an application specific implementation.

The state of shard locations in the `ShardCoordinator` is persistent (durable) with *Persistence* to survive failures. Since it is running in a cluster *Persistence* must be configured with a distributed journal. When a crashed or unreachable coordinator node has been removed (via `down`) from the cluster a new `ShardCoordinator` singleton actor will take over and the state is recovered. During such a failure period shards with known location are still available, while messages for new (unknown) shards are buffered until the new `ShardCoordinator` becomes available.

As long as a sender uses the same `ShardRegion` actor to deliver messages to an entity actor the order of the messages is preserved. As long as the buffer limit is not reached messages are delivered on a best effort basis, with at-most once delivery semantics, in the same way as ordinary message sending. Reliable end-to-end messaging, with at-least-once semantics can be added by using `AtLeastOnceDelivery` in *Persistence*.

Some additional latency is introduced for messages targeted to new or previously unused shards due to the round-trip to the coordinator. Rebalancing of shards may also add latency. This should be considered when designing the application specific shard resolution, e.g. to avoid too fine grained shards.

6.6.3 Distributed Data Mode

Instead of using *Persistence* it is possible to use the *Distributed Data* module as storage for the state of the sharding coordinator. In such case the state of the `ShardCoordinator` will be replicated inside a cluster by the *Distributed Data* module with `WriteMajority/ReadMajority` consistency.

This mode can be enabled by setting configuration property `akka.cluster.sharding.state-store-mode` as `ddata`.

You must explicitly add the `akka-distributed-data-experimental` dependency to your build if you use this mode. It is possible to remove `akka-persistence` dependency from a project if it is not used in user code and `remember-entities` is off.

Warning: The `ddata` mode is considered as “**experimental**” as of its introduction in Akka 2.4.0, since it depends on the experimental *Distributed Data* module.

6.6.4 Proxy Only Mode

The `ShardRegion` actor can also be started in proxy only mode, i.e. it will not host any entities itself, but knows how to delegate messages to the right location. A `ShardRegion` is started in proxy only mode with the method `ClusterSharding.startProxy` method.

6.6.5 Passivation

If the state of the entities are persistent you may stop entities that are not used to reduce memory consumption. This is done by the application specific implementation of the entity actors for example by defining `receiveTimeout` (`context.setReceiveTimeout`). If a message is already enqueued to the entity when it stops itself the enqueued message in the mailbox will be dropped. To support graceful passivation without losing such messages the entity actor can send `ShardRegion.Passivate` to its parent `Shard`. The specified wrapped message in `Passivate` will be sent back to the entity, which is then supposed to stop itself. Incoming messages will be buffered by the `Shard` between reception of `Passivate` and termination of the entity. Such buffered messages are thereafter delivered to a new incarnation of the entity.

6.6.6 Remembering Entities

The list of entities in each `Shard` can be made persistent (durable) by setting the `rememberEntities` flag to `true` in `ClusterShardingSettings` when calling `ClusterSharding.start`. When configured to remember entities, whenever a `Shard` is rebalanced onto another node or recovers after a crash it will recreate all the entities which were previously running in that `Shard`. To permanently stop entities, a `Passivate` message must be sent to the parent of the entity actor, otherwise the entity will be automatically restarted after the entity restart backoff specified in the configuration.

When `rememberEntities` is set to `false`, a `Shard` will not automatically restart any entities after a rebalance or recovering from a crash. Entities will only be started once the first message for that entity has been received in the `Shard`. Entities will not be restarted if they stop without using a `Passivate`.

Note that the state of the entities themselves will not be restored unless they have been made persistent, e.g. with *Persistence*.

6.6.7 Supervision

If you need to use another `supervisorStrategy` for the entity actors than the default (restarting) strategy you need to create an intermediate parent actor that defines the `supervisorStrategy` to the child entity actor.

```
public class CounterSupervisor extends UntypedActor {

    private final ActorRef counter = getContext().actorOf(
        Props.create(Counter.class), "theCounter");

    private static final SupervisorStrategy strategy =
        new OneForOneStrategy(DeciderBuilder.
            match(IllegalArgumentException.class, e -> SupervisorStrategy.resume()).
            match(ActorInitializationException.class, e -> SupervisorStrategy.stop()).
            match(Exception.class, e -> SupervisorStrategy.restart()).
            matchAny(o -> SupervisorStrategy.escalate()).build());

    @Override
    public SupervisorStrategy supervisorStrategy() {
        return strategy;
    }

    @Override
    public void onReceive(Object msg) {
        counter.forward(msg, getContext());
    }
}
```

You start such a supervisor in the same way as if it was the entity actor.

```
ClusterSharding.get(system).start("SupervisedCounter",
    Props.create(CounterSupervisor.class), settings, messageExtractor);
```

Note that stopped entities will be started again when a new message is targeted to the entity.

6.6.8 Graceful Shutdown

You can send the message `ClusterSharding.GracefulShutdown` message (`ClusterSharding.gracefulShutdownInstance` in Java) to the `ShardRegion` actor to hand-off all shards that are hosted by that `ShardRegion` and then the `ShardRegion` actor will be stopped. You can watch the `ShardRegion` actor to know when it is completed. During this period other regions will buffer messages for those shards in the same way as when a rebalance is triggered by the coordinator. When the shards have been stopped the coordinator will allocate these shards elsewhere.

When the `ShardRegion` has terminated you probably want to leave the cluster, and shut down the `ActorSystem`.

This is how to do that:

```
public class IllustrateGracefulShutdown extends AbstractActor {

    public IllustrateGracefulShutdown() {
        final ActorSystem system = context().system();
        final Cluster cluster = Cluster.get(system);
        final ActorRef region = ClusterSharding.get(system).shardRegion("Entity");

        receive(ReceiveBuilder.
            match(String.class, s -> s.equals("leave"), s -> {
                context().watch(region);
                region.tell(ShardRegion.gracefulShutdownInstance(), self());
            }).
            match(Terminated.class, t -> t.actor().equals(region), t -> {
                cluster.registerOnMemberRemoved(() -> system.terminate());
                cluster.leave(cluster.selfAddress());
            }).build());
    }
}
```

6.6.9 Removal of Internal Cluster Sharding Data

The Cluster Sharding coordinator stores the locations of the shards using Akka Persistence. This data can safely be removed when restarting the whole Akka Cluster. Note that this is not application data.

There is a utility program `akka.cluster.sharding.RemoveInternalClusterShardingData` that removes this data.

Warning: Never use this program while there are running Akka Cluster nodes that are using Cluster Sharding. Stop all Cluster nodes before using this program.

It can be needed to remove the data if the Cluster Sharding coordinator cannot startup because of corrupt data, which may happen if accidentally two clusters were running at the same time, e.g. caused by using auto-down and there was a network partition.

Use this program as a standalone Java main program:

```
java -classpath <jar files, including akka-cluster-sharding>
  akka.cluster.sharding.RemoveInternalClusterShardingData
  -2.3 entityType1 entityType2 entityType3
```

The program is included in the `akka-cluster-sharding` jar file. It is easiest to run it with same classpath and configuration as your ordinary application. It can be run from sbt or maven in similar way.

Specify the entity type names (same as you use in the `start` method of `ClusterSharding`) as program arguments.

If you specify `-2.3` as the first program argument it will also try to remove data that was stored by Cluster Sharding in Akka 2.3.x using different persistenceId.

6.6.10 Dependencies

To use the Cluster Sharding you must add the following dependency in your project.

sbt:

```
"com.typesafe.akka" %% "akka-cluster-sharding" % "2.4.0"
```

maven:

```
<dependency>
  <groupId>com.typesafe.akka</groupId>
  <artifactId>akka-cluster-sharding_2.11</artifactId>
  <version>2.4.0</version>
</dependency>
```

6.6.11 Configuration

The `ClusterSharding` extension can be configured with the following properties. These configuration properties are read by the `ClusterShardingSettings` when created with a `ActorSystem` parameter. It is also possible to amend the `ClusterShardingSettings` or create it from another config section with the same layout as below. `ClusterShardingSettings` is a parameter to the `start` method of the `ClusterSharding` extension, i.e. each entity type can be configured with different settings if needed.

```
# Settings for the ClusterShardingExtension
akka.cluster.sharding {

  # The extension creates a top level actor with this name in top level system scope,
  # e.g. '/system/sharding'
  guardian-name = sharding

  # Specifies that entities runs on cluster nodes with a specific role.
  # If the role is not specified (or empty) all nodes in the cluster are used.
  role = ""

  # When this is set to 'on' the active entity actors will automatically be restarted
  # upon Shard restart. i.e. if the Shard is started on a different ShardRegion
  # due to rebalance or crash.
  remember-entities = off

  # If the coordinator can't store state changes it will be stopped
  # and started again after this duration, with an exponential back-off
  # of up to 5 times this duration.
  coordinator-failure-backoff = 5 s

  # The ShardRegion retries registration and shard location requests to the
  # ShardCoordinator with this interval if it does not reply.
  retry-interval = 2 s

  # Maximum number of messages that are buffered by a ShardRegion actor.
  buffer-size = 100000

  # Timeout of the shard rebalancing process.
  handoff-timeout = 60 s

  # Time given to a region to acknowledge it's hosting a shard.
  shard-start-timeout = 10 s

  # If the shard is remembering entities and can't store state changes
  # will be stopped and then started again after this duration. Any messages
  # sent to an affected entity may be lost in this process.
  shard-failure-backoff = 10 s

  # If the shard is remembering entities and an entity stops itself without
  # using passivate. The entity will be restarted after this duration or when
  # the next message for it is received, which ever occurs first.
  entity-restart-backoff = 10 s

  # Rebalance check is performed periodically with this interval.
  rebalance-interval = 10 s
```

```

# Absolute path to the journal plugin configuration entity that is to be
# used for the internal persistence of ClusterSharding. If not defined
# the default journal plugin is used. Note that this is not related to
# persistence used by the entity actors.
journal-plugin-id = ""

# Absolute path to the snapshot plugin configuration entity that is to be
# used for the internal persistence of ClusterSharding. If not defined
# the default snapshot plugin is used. Note that this is not related to
# persistence used by the entity actors.
snapshot-plugin-id = ""

# Parameter which determines how the coordinator will be store a state
# valid values either "persistence" or "ddata"
# The "ddata" mode is experimental, since it depends on the experimental
# module akka-distributed-data-experimental.
state-store-mode = "persistence"

# The shard saves persistent snapshots after this number of persistent
# events. Snapshots are used to reduce recovery times.
snapshot-after = 1000

# Setting for the default shard allocation strategy
least-shard-allocation-strategy {
  # Threshold of how large the difference between most and least number of
  # allocated shards must be to begin the rebalancing.
  rebalance-threshold = 10

  # The number of ongoing rebalancing processes is limited to this number.
  max-simultaneous-rebalance = 3
}

# Timeout of waiting the initial distributed state (an initial state will be queried again if t
# works only for state-store-mode = "ddata"
waiting-for-state-timeout = 5 s

# Timeout of waiting for update the distributed state (update will be retried if the timeout hap
# works only for state-store-mode = "ddata"
updating-state-timeout = 5 s

# Settings for the coordinator singleton. Same layout as akka.cluster.singleton.
coordinator-singleton = ${akka.cluster.singleton}

# The id of the dispatcher to use for ClusterSharding actors.
# If not specified default dispatcher is used.
# If specified you need to define the settings of the actual dispatcher.
# This dispatcher for the entity actors is defined by the user provided
# Props, i.e. this dispatcher is not used for the entity actors.
use-dispatcher = ""
}

```

Custom shard allocation strategy can be defined in an optional parameter to `ClusterSharding.start`. See the API documentation of `AbstractShardAllocationStrategy` for details of how to implement a custom shard allocation strategy.

6.7 Cluster Metrics Extension

6.7.1 Introduction

The member nodes of the cluster can collect system health metrics and publish that to other cluster nodes and to the registered subscribers on the system event bus with the help of Cluster Metrics Extension.

Cluster metrics information is primarily used for load-balancing routers, and can also be used to implement advanced metrics-based node life cycles, such as “Node Let-it-crash” when CPU steal time becomes excessive.

Cluster Metrics Extension is a separate akka module delivered in `akka-cluster-metrics` jar.

To enable usage of the extension you need to add the following dependency to your project:

```
<dependency>
  <groupId>com.typesafe.akka</groupId>
  <artifactId>akka-cluster-metrics_2.11</artifactId>
  <version>2.4.0</version>
</dependency>
```

and add the following configuration stanza to your `application.conf`

```
akka.extensions = [ "akka.cluster.metrics.ClusterMetricsExtension" ]
```

Make sure to disable legacy metrics in akka-cluster: `akka.cluster.metrics.enabled=off`, since it is still enabled in akka-cluster by default (for compatibility with past releases).

Cluster members with status *WeaklyUp*, if that feature is enabled, will participate in Cluster Metrics collection and dissemination.

6.7.2 Metrics Collector

Metrics collection is delegated to an implementation of `akka.cluster.metrics.MetricsCollector`.

Different collector implementations provide different subsets of metrics published to the cluster. Certain message routing and let-it-crash functions may not work when Sigar is not provisioned.

Cluster metrics extension comes with two built-in collector implementations:

1. `akka.cluster.metrics.SigarMetricsCollector`, which requires Sigar provisioning, and is more rich/precise
2. `akka.cluster.metrics.JmxMetricsCollector`, which is used as fall back, and is less rich/precise

You can also plug-in your own metrics collector implementation.

By default, metrics extension will use collector provider fall back and will try to load them in this order:

1. configured user-provided collector
2. built-in `akka.cluster.metrics.SigarMetricsCollector`
3. and finally `akka.cluster.metrics.JmxMetricsCollector`

6.7.3 Metrics Events

Metrics extension periodically publishes current snapshot of the cluster metrics to the node system event bus.

The publication period is controlled by the `akka.cluster.metrics.collector.sample-period` setting.

The payload of the `akka.cluster.metrics.ClusterMetricsChanged` event will contain latest metrics of the node as well as other cluster member nodes metrics gossip which was received during the collector sample period.

You can subscribe your metrics listener actors to these events in order to implement custom node lifecycle

```
ClusterMetricsExtension.get(system).subscribe(metricsListenerActor);
```

6.7.4 Hyperic Sigar Provisioning

Both user-provided and built-in metrics collectors can optionally use [Hyperic Sigar](#) for a wider and more accurate range of metrics compared to what can be retrieved from ordinary JMX MBeans.

Sigar is using a native o/s library, and requires library provisioning, i.e. deployment, extraction and loading of the o/s native library into JVM at runtime.

User can provision Sigar classes and native library in one of the following ways:

1. Use [Kamon sigar-loader](#) as a project dependency for the user project. Metrics extension will extract and load sigar library on demand with help of Kamon sigar provisioner.
2. Use [Kamon sigar-loader](#) as java agent: `java -javaagent:/path/to/sigar-loader.jar`. Kamon sigar loader agent will extract and load sigar library during JVM start.
3. Place `sigar.jar` on the classpath and Sigar native library for the o/s on the `java.library.path`. User is required to manage both project dependency and library deployment manually.

To enable usage of Sigar you can add the following dependency to the user project

```
<dependency>
  <groupId>io.kamon</groupId>
  <artifactId>sigar-loader</artifactId>
  <version>1.6.6-rev002</version>
</dependency>
```

You can download Kamon sigar-loader from [Maven Central](#)

6.7.5 Adaptive Load Balancing

The `AdaptiveLoadBalancingPool` / `AdaptiveLoadBalancingGroup` performs load balancing of messages to cluster nodes based on the cluster metrics data. It uses random selection of routees with probabilities derived from the remaining capacity of the corresponding node. It can be configured to use a specific `MetricsSelector` to produce the probabilities, a.k.a. weights:

- `heap` / `HeapMetricsSelector` - Used and max JVM heap memory. Weights based on remaining heap capacity; $(\text{max} - \text{used}) / \text{max}$
- `load` / `SystemLoadAverageMetricsSelector` - System load average for the past 1 minute, corresponding value can be found in `top` of Linux systems. The system is possibly nearing a bottleneck if the system load average is nearing number of cpus/cores. Weights based on remaining load capacity; $1 - (\text{load} / \text{processors})$
- `cpu` / `CpuMetricsSelector` - CPU utilization in percentage, sum of User + Sys + Nice + Wait. Weights based on remaining cpu capacity; $1 - \text{utilization}$
- `mix` / `MixMetricsSelector` - Combines heap, cpu and load. Weights based on mean of remaining capacity of the combined selectors.
- Any custom implementation of `akka.cluster.metrics.MetricsSelector`

The collected metrics values are smoothed with [exponential weighted moving average](#). In the [Configuration](#) you can adjust how quickly past data is decayed compared to new data.

Let's take a look at this router in action. What can be more demanding than calculating factorials?

The backend worker that performs the factorial calculation:

```

public class FactorialBackend extends UntypedActor {

    @Override
    public void onReceive(Object message) {
        if (message instanceof Integer) {
            final Integer n = (Integer) message;
            Future<BigInteger> f = future(new Callable<BigInteger>() {
                public BigInteger call() {
                    return factorial(n);
                }
            }, getContext().dispatcher());

            Future<FactorialResult> result = f.map(
                new Mapper<BigInteger, FactorialResult>() {
                    public FactorialResult apply(BigInteger factorial) {
                        return new FactorialResult(n, factorial);
                    }
                }, getContext().dispatcher());

            pipe(result, getContext().dispatcher()).to(getSender());

        } else {
            unhandled(message);
        }
    }

    BigInteger factorial(int n) {
        BigInteger acc = BigInteger.ONE;
        for (int i = 1; i <= n; ++i) {
            acc = acc.multiply(BigInteger.valueOf(i));
        }
        return acc;
    }
}

```

The frontend that receives user jobs and delegates to the backends via the router:

```

public class FactorialFrontend extends UntypedActor {
    final int upToN;
    final boolean repeat;

    LoggingAdapter log = Logging.getLogger(getContext().system(), this);

    ActorRef backend = getContext().actorOf(FromConfig.getInstance().props(),
        "factorialBackendRouter");

    public FactorialFrontend(int upToN, boolean repeat) {
        this.upToN = upToN;
        this.repeat = repeat;
    }

    @Override
    public void preStart() {
        sendJobs();
        getContext().setReceiveTimeout(Duration.create(10, TimeUnit.SECONDS));
    }

    @Override
    public void onReceive(Object message) {
        if (message instanceof FactorialResult) {
            FactorialResult result = (FactorialResult) message;
            if (result.n == upToN) {
                log.debug("{}! = {}", result.n, result.factorial);
            }
        }
    }
}

```

```

        if (repeat)
            sendJobs();
        else
            getContext().stop(getSelf());
    }

} else if (message instanceof ReceiveTimeout) {
    log.info("Timeout");
    sendJobs();
} else {
    unhandled(message);
}
}

void sendJobs() {
    log.info("Starting batch of factorials up to [{}]", upToN);
    for (int n = 1; n <= upToN; n++) {
        backend.tell(n, getSelf());
    }
}
}
}

```

As you can see, the router is defined in the same way as other routers, and in this case it is configured as follows:

```

akka.actor.deployment {
  /factorialFrontend/factorialBackendRouter = {
    # Router type provided by metrics extension.
    router = cluster-metrics-adaptive-group
    # Router parameter specific for metrics extension.
    # metrics-selector = heap
    # metrics-selector = load
    # metrics-selector = cpu
    metrics-selector = mix
    #
    routees.paths = ["/user/factorialBackend"]
    cluster {
      enabled = on
      use-role = backend
      allow-local-routees = off
    }
  }
}

```

It is only router type and the metrics-selector parameter that is specific to this router, other things work in the same way as other routers.

The same type of router could also have been defined in code:

```

int totalInstances = 100;
Iterable<String> routeesPaths = Arrays.asList("/user/factorialBackend", "");
boolean allowLocalRoutees = true;
String useRole = "backend";
ActorRef backend = getContext().actorOf(
    new ClusterRouterGroup(new AdaptiveLoadBalancingGroup(
        HeapMetricsSelector.getInstance(), Collections.<String> emptyList()),
        new ClusterRouterGroupSettings(totalInstances, routeesPaths,
            allowLocalRoutees, useRole)).props(), "factorialBackendRouter2");

```

```

int totalInstances = 100;
int maxInstancesPerNode = 3;
boolean allowLocalRoutees = false;
String useRole = "backend";

```

```
ActorRef backend = getContext().actorOf(
    new ClusterRouterPool(new AdaptiveLoadBalancingPool(
        SystemLoadAverageMetricsSelector.getInstance(), 0),
        new ClusterRouterPoolSettings(totalInstances, maxInstancesPerNode,
            allowLocalRoutees, useRole)).props(Props
            .create(FactoryBackend.class)), "factorialBackendRouter3");
```

The [Typesafe Activator](#) tutorial named [Akka Cluster Samples with Java](#). contains the full source code and instructions of how to run the **Adaptive Load Balancing** sample.

6.7.6 Subscribe to Metrics Events

It is possible to subscribe to the metrics events directly to implement other functionality.

```
import akka.actor.UntypedActor;
import akka.cluster.Cluster;
import akka.cluster.ClusterEvent.CurrentClusterState;
import akka.cluster.metrics.ClusterMetricsChanged;
import akka.cluster.metrics.NodeMetrics;
import akka.cluster.metrics.StandardMetrics;
import akka.cluster.metrics.StandardMetrics.HeapMemory;
import akka.cluster.metrics.StandardMetrics.Cpu;
import akka.cluster.metrics.ClusterMetricsExtension;
import akka.event.Logging;
import akka.event.LoggingAdapter;

public class MetricsListener extends UntypedActor {
    LoggingAdapter log = Logging.getLogger(getContext().system(), this);

    Cluster cluster = Cluster.get(getContext().system());

    ClusterMetricsExtension extension = ClusterMetricsExtension.get(getContext().system());

    // Subscribe unto ClusterMetricsEvent events.
    @Override
    public void preStart() {
        extension.subscribe(getSelf());
    }

    // Unsubscribe from ClusterMetricsEvent events.
    @Override
    public void postStop() {
        extension.unsubscribe(getSelf());
    }

    @Override
    public void onReceive(Object message) {
        if (message instanceof ClusterMetricsChanged) {
            ClusterMetricsChanged clusterMetrics = (ClusterMetricsChanged) message;
            for (NodeMetrics nodeMetrics : clusterMetrics.getNodeMetrics()) {
                if (nodeMetrics.address().equals(cluster.selfAddress())) {
                    logHeap(nodeMetrics);
                    logCpu(nodeMetrics);
                }
            }
        }
        else if (message instanceof CurrentClusterState) {
            // Ignore.
        }
        else {
            unhandled(message);
        }
    }
}
```

```

    }
  }

  void logHeap(NodeMetrics nodeMetrics) {
    HeapMemory heap = StandardMetrics.extractHeapMemory(nodeMetrics);
    if (heap != null) {
      log.info("Used heap: {} MB", ((double) heap.used()) / 1024 / 1024);
    }
  }

  void logCpu(NodeMetrics nodeMetrics) {
    Cpu cpu = StandardMetrics.extractCpu(nodeMetrics);
    if (cpu != null && cpu.systemLoadAverage().isDefined()) {
      log.info("Load: {} ({} processors)", cpu.systemLoadAverage().get(),
        cpu.processors());
    }
  }
}
}

```

6.7.7 Custom Metrics Collector

Metrics collection is delegated to the implementation of `akka.cluster.metrics.MetricsCollector`

You can plug-in your own metrics collector instead of built-in `akka.cluster.metrics.SigarMetricsCollector` or `akka.cluster.metrics.JmxMetricsCollector`.

Look at those two implementations for inspiration.

Custom metrics collector implementation class must be specified in the `akka.cluster.metrics.collector.provider` configuration property.

6.7.8 Configuration

The Cluster metrics extension can be configured with the following properties:

```

#####
# Akka Cluster Metrics Reference Config File #
#####

# This is the reference config file that contains all the default settings.
# Make your edits in your application.conf in order to override these settings.

# Sigar provisioning:
#
# User can provision sigar classes and native library in one of the following ways:
#
# 1) Use https://github.com/kamon-io/sigar-loader Kamon sigar-loader as a project dependency for
# Metrics extension will extract and load sigar library on demand with help of Kamon sigar provi
#
# 2) Use https://github.com/kamon-io/sigar-loader Kamon sigar-loader as java agent: `java -javaagent:
# Kamon sigar loader agent will extract and load sigar library during JVM start.
#
# 3) Place `sigar.jar` on the `classpath` and sigar native library for the o/s on the `java.library
# User is required to manage both project dependency and library deployment manually.

# Cluster metrics extension.
# Provides periodic statistics collection and publication throughout the cluster.
akka.cluster.metrics {
  # Full path of dispatcher configuration key.
  # Use "" for default key `akka.actor.default-dispatcher`.

```

```

dispatcher = ""
# How long should any actor wait before starting the periodic tasks.
periodic-tasks-initial-delay = 1s
# Sigar native library extract location.
# Use per-application-instance scoped location, such as program working directory.
native-library-extract-folder = ${user.dir}/"native"
# Metrics supervisor actor.
supervisor {
  # Actor name. Example name space: /system/cluster-metrics
  name = "cluster-metrics"
  # Supervision strategy.
  strategy {
    #
    # FQCN of class providing `akka.actor.SupervisorStrategy`.
    # Must have a constructor with signature `(com.typesafe.config.Config)`.
    # Default metrics strategy provider is a configurable extension of `OneForOneStrategy`
    provider = "akka.cluster.metrics.ClusterMetricsStrategy"
    #
    # Configuration of the default strategy provider.
    # Replace with custom settings when overriding the provider.
    configuration = {
      # Log restart attempts.
      loggingEnabled = true
      # Child actor restart-on-failure window.
      withinTimeRange = 3s
      # Maximum number of restart attempts before child actor is stopped.
      maxNrOfRetries = 3
    }
  }
}
}
# Metrics collector actor.
collector {
  # Enable or disable metrics collector for load-balancing nodes.
  # Metrics collection can also be controlled at runtime by sending control messages
  # to /system/cluster-metrics actor: `akka.cluster.metrics.{CollectionStartMessage,Collectio
  enabled = on
  # FQCN of the metrics collector implementation.
  # It must implement `akka.cluster.metrics.MetricsCollector` and
  # have public constructor with akka.actor.ActorSystem parameter.
  # Will try to load in the following order of priority:
  # 1) configured custom collector 2) internal `SigarMetricsCollector` 3) internal `JmxMetric
  provider = ""
  # Try all 3 available collector providers, or else fail on the configured custom collector
  fallback = true
  # How often metrics are sampled on a node.
  # Shorter interval will collect the metrics more often.
  # Also controls frequency of the metrics publication to the node system event bus.
  sample-interval = 3s
  # How often a node publishes metrics information to the other nodes in the cluster.
  # Shorter interval will publish the metrics gossip more often.
  gossip-interval = 3s
  # How quickly the exponential weighting of past data is decayed compared to
  # new data. Set lower to increase the bias toward newer values.
  # The relevance of each data sample is halved for every passing half-life
  # duration, i.e. after 4 times the half-life, a data sample's relevance is
  # reduced to 6% of its original relevance. The initial relevance of a data
  # sample is given by 1 - 0.5 ^ (collect-interval / half-life).
  # See http://en.wikipedia.org/wiki/Moving\_average#Exponential\_moving\_average
  moving-average-half-life = 12s
}
}
# Cluster metrics extension serializers and routers.

```

```
akka.actor {
  # Protobuf serializer for remote cluster metrics messages.
  serializers {
    akka-cluster-metrics = "akka.cluster.metrics.protobuf.MessageSerializer"
  }
  # Interface binding for remote cluster metrics messages.
  serialization-bindings {
    "akka.cluster.metrics.ClusterMetricsMessage" = akka-cluster-metrics
  }
  # Globally unique metrics extension serializer identifier.
  serialization-identifiers {
    "akka.cluster.metrics.protobuf.MessageSerializer" = 10
  }
  # Provide routing of messages based on cluster metrics.
  router.type-mapping {
    cluster-metrics-adaptive-pool = "akka.cluster.metrics.AdaptiveLoadBalancingPool"
    cluster-metrics-adaptive-group = "akka.cluster.metrics.AdaptiveLoadBalancingGroup"
  }
}
```

6.8 Distributed Data

Akka Distributed Data is useful when you need to share data between nodes in an Akka Cluster. The data is accessed with an actor providing a key-value store like API. The keys are unique identifiers with type information of the data values. The values are *Conflict Free Replicated Data Types* (CRDTs).

All data entries are spread to all nodes, or nodes with a certain role, in the cluster via direct replication and gossip based dissemination. You have fine grained control of the consistency level for reads and writes.

The nature CRDTs makes it possible to perform updates from any node without coordination. Concurrent updates from different nodes will automatically be resolved by the monotonic merge function, which all data types must provide. The state changes always converge. Several useful data types for counters, sets, maps and registers are provided and you can also implement your own custom data types.

It is eventually consistent and geared toward providing high read and write availability (partition tolerance), with low latency. Note that in an eventually consistent system a read may return an out-of-date value.

Warning: This module is marked as “**experimental**” as of its introduction in Akka 2.4.0. We will continue to improve this API based on our users’ feedback, which implies that while we try to keep incompatible changes to a minimum the binary compatibility guarantee for maintenance releases does not apply to the contents of the `akka.persistence` package.

6.8.1 Using the Replicator

The `akka.cluster.ddata.Replicator` actor provides the API for interacting with the data. The `Replicator` actor must be started on each node in the cluster, or group of nodes tagged with a specific role. It communicates with other `Replicator` instances with the same path (without address) that are running on other nodes. For convenience it can be used with the `akka.cluster.ddata.DistributedData` extension.

Cluster members with status *WeaklyUp*, if that feature is enabled, will currently not participate in Distributed Data, but that is something that should be possible to add in a future release.

Below is an example of an actor that schedules tick messages to itself and for each tick adds or removes elements from a `ORSet` (observed-remove set). It also subscribes to changes of this.

```
import static java.util.concurrent.TimeUnit.SECONDS;
import scala.concurrent.duration.Duration;
import scala.concurrent.forkjoin.ThreadLocalRandom;
```

```

import akka.actor.AbstractActor;
import akka.actor.ActorRef;
import akka.actor.Cancellable;
import akka.cluster.Cluster;
import akka.cluster.ddata.DistributedData;
import akka.cluster.ddata.Key;
import akka.cluster.ddata.ORSet;
import akka.cluster.ddata.ORSetKey;
import akka.cluster.ddata.Replicator;
import akka.cluster.ddata.Replicator.Changed;
import akka.cluster.ddata.Replicator.Subscribe;
import akka.cluster.ddata.Replicator.Update;
import akka.cluster.ddata.Replicator.UpdateResponse;
import akka.event.Logging;
import akka.event.LoggingAdapter;
import akka.japi.pf.ReceiveBuilder;

public class DataBot extends AbstractActor {

    private static final String TICK = "tick";

    private final LoggingAdapter log = Logging.getLogger(context().system(), this);

    private final ActorRef replicator =
        DistributedData.get(context().system()).replicator();
    private final Cluster node = Cluster.get(context().system());

    private final Cancellable tickTask = context().system().scheduler().schedule(
        Duration.create(5, SECONDS), Duration.create(5, SECONDS), self(), TICK,
        context().dispatcher(), self());

    private final Key<ORSet<String>> dataKey = ORSetKey.create("key");

    @SuppressWarnings("unchecked")
    public DataBot() {
        receive(ReceiveBuilder
            .match(String.class, a -> a.equals(TICK), a -> receiveTick())
            .match(Changed.class, c -> c.key().equals(dataKey), c -> receiveChanged((Changed<ORSet<String>>) c))
            .match(UpdateResponse.class, r -> receiveUpdateResponse(r))
            .build());
    }

    private void receiveTick() {
        String s = String.valueOf((char) ThreadLocalRandom.current().nextInt(97, 123));
        if (ThreadLocalRandom.current().nextBoolean()) {
            // add
            log.info("Adding: {}", s);
            Update<ORSet<String>> update = new Update<>(
                dataKey,
                ORSet.create(),
                Replicator.writeLocal(),
                curr -> curr.add(node, s));
            replicator.tell(update, self());
        } else {
            // remove
            log.info("Removing: {}", s);
            Update<ORSet<String>> update = new Update<>(
                dataKey,
                ORSet.create(),
                Replicator.writeLocal(),
                curr -> curr.remove(node, s));
            replicator.tell(update, self());
        }
    }
}

```

```

    }
}

private void receiveChanged(Changed<ORSet<String>> c) {
    ORSet<String> data = c.dataValue();
    log.info("Current elements: {}", data.getElements());
}

private void receiveUpdateResoponse() {
    // ignore
}

@Override
public void preStart() {
    Subscribe<ORSet<String>> subscribe = new Subscribe<>(dataKey, self());
    replicator.tell(subscribe, ActorRef.noSender());
}

@Override
public void postStop(){
    tickTask.cancel();
}
}
}

```

Update

To modify and replicate a data value you send a `Replicator.Update` message to the local `Replicator`.

The current data value for the key of the `Update` is passed as parameter to the `modify` function of the `Update`. The function is supposed to return the new value of the data, which will then be replicated according to the given consistency level.

The `modify` function is called by the `Replicator` actor and must therefore be a pure function that only uses the data parameter and stable fields from enclosing scope. It must for example not access `sender()` reference of an enclosing actor.

Update is intended to only be sent from an actor running in same local `ActorSystem` as

- the *Replicator*, because the *modify* function is typically not serializable.

You supply a write consistency level which has the following meaning:

- `writeLocal` the value will immediately only be written to the local replica, and later disseminated with gossip
- `writeTo(n)` the value will immediately be written to at least `n` replicas, including the local replica
- `writeMajority` the value will immediately be written to a majority of replicas, i.e. at least $N/2 + 1$ replicas, where `N` is the number of nodes in the cluster (or cluster role group)
- `writeAll` the value will immediately be written to all nodes in the cluster (or all nodes in the cluster role group)

```

final Cluster node = Cluster.get(system);
final ActorRef replicator = DistributedData.get(system).replicator();

final Key<PNCounter> counter1Key = PNCounterKey.create("counter1");
final Key<GSet<String>> set1Key = GSetKey.create("set1");
final Key<ORSet<String>> set2Key = ORSetKey.create("set2");
final Key<Flag> activeFlagKey = FlagKey.create("active");

```

```

replicator.tell(new Replicator.Update<PNCCounter>(counter1Key, PNCCounter.create(),
    Replicator.writeLocal(), curr -> curr.increment(node, 1)), self());

final WriteConsistency writeTo3 = new WriteTo(3, Duration.create(1, SECONDS));
replicator.tell(new Replicator.Update<GSet<String>>(set1Key, GSet.create(),
    writeTo3, curr -> curr.add("hello")), self());

final WriteConsistency writeMajority =
    new WriteMajority(Duration.create(5, SECONDS));
replicator.tell(new Replicator.Update<ORSet<String>>(set2Key, ORSet.create(),
    writeMajority, curr -> curr.add(node, "hello")), self());

final WriteConsistency writeAll = new WriteAll(Duration.create(5, SECONDS));
replicator.tell(new Replicator.Update<Flag>(activeFlagKey, Flag.create(),
    writeAll, curr -> curr.switchOn()), self());

```

As reply of the Update a `Replicator.UpdateSuccess` is sent to the sender of the Update if the value was successfully replicated according to the supplied consistency level within the supplied timeout. Otherwise a `Replicator.UpdateFailure` subclass is sent back. Note that a `Replicator.UpdateTimeout` reply does not mean that the update completely failed or was rolled back. It may still have been replicated to some nodes, and will eventually be replicated to all nodes with the gossip protocol.

```

receive(ReceiveBuilder.
    match(UpdateSuccess.class, a -> a.key().equals(counter1Key), a -> {
        // ok
    }).build());

```

```

receive(ReceiveBuilder.
    match(UpdateSuccess.class, a -> a.key().equals(set1Key), a -> {
        // ok
    }).
    match(UpdateTimeout.class, a -> a.key().equals(set1Key), a -> {
        // write to 3 nodes failed within 1.second
    }).build());

```

You will always see your own writes. For example if you send two Update messages changing the value of the same key, the modify function of the second message will see the change that was performed by the first Update message.

In the Update message you can pass an optional request context, which the Replicator does not care about, but is included in the reply messages. This is a convenient way to pass contextual information (e.g. original sender) without having to use `ask` or maintain local correlation data structures.

```

final Cluster node = Cluster.get(system);
final ActorRef replicator = DistributedData.get(system).replicator();

final WriteConsistency writeTwo = new WriteTo(2, Duration.create(3, SECONDS));
final Key<PNCCounter> counter1Key = PNCCounterKey.create("counter1");

receive(ReceiveBuilder.
    match(String.class, a -> a.equals("increment"), a -> {
        // incoming command to increase the counter
        Optional<Object> reqContext = Optional.of(sender());
        Replicator.Update<PNCCounter> upd = new Replicator.Update<PNCCounter>(counter1Key,
            PNCCounter.create(), writeTwo, reqContext, curr -> curr.increment(node, 1));
        replicator.tell(upd, self());
    }).

    match(UpdateSuccess.class, a -> a.key().equals(counter1Key), a -> {
        ActorRef replyTo = (ActorRef) a.getRequest().get();
        replyTo.tell("ack", self());
    }).
);

```

```
match(UpdateTimeout.class, a -> a.key().equals(counter1Key), a -> {
    ActorRef replyTo = (ActorRef) a.getRequest().get();
    replyTo.tell("nack", self());
}).build();
```

Get

To retrieve the current value of a data you send `Replicator.Get` message to the `Replicator`. You supply a consistency level which has the following meaning:

- `readLocal` the value will only be read from the local replica
- `readFrom(n)` the value will be read and merged from `n` replicas, including the local replica
- `readMajority` the value will be read and merged from a majority of replicas, i.e. at least $N/2 + 1$ replicas, where `N` is the number of nodes in the cluster (or cluster role group)
- `readAll` the value will be read and merged from all nodes in the cluster (or all nodes in the cluster role group)

```
final ActorRef replicator = DistributedData.get(system).replicator();
final Key<PNCOUNTER> counter1Key = PNCOUNTERKey.create("counter1");
final Key<GSet<String>> set1Key = GSetKey.create("set1");
final Key<ORSet<String>> set2Key = ORSetKey.create("set2");
final Key<Flag> activeFlagKey = FlagKey.create("active");

replicator.tell(new Replicator.Get<PNCOUNTER>(counter1Key,
    Replicator.readLocal()), self());

final ReadConsistency readFrom3 = new ReadFrom(3, Duration.create(1, SECONDS));
replicator.tell(new Replicator.Get<GSet<String>>(set1Key,
    readFrom3), self());

final ReadConsistency readMajority = new ReadMajority(Duration.create(5, SECONDS));
replicator.tell(new Replicator.Get<ORSet<String>>(set2Key,
    readMajority), self());

final ReadConsistency readAll = new ReadAll(Duration.create(5, SECONDS));
replicator.tell(new Replicator.Get<Flag>(activeFlagKey,
    readAll), self());
```

As reply of the `Get` a `Replicator.GetSuccess` is sent to the sender of the `Get` if the value was successfully retrieved according to the supplied consistency level within the supplied timeout. Otherwise a `Replicator.GetFailure` is sent. If the key does not exist the reply will be `Replicator.NotFound`.

```
receive(ReceiveBuilder.
    match(GetSuccess.class, a -> a.key().equals(counter1Key), a -> {
        GetSuccess<PNCOUNTER> g = a;
        BigInteger value = g.dataValue().getValue();
    }).
    match(NotFound.class, a -> a.key().equals(counter1Key), a -> {
        // key counter1 does not exist
    }).build());
```

```
receive(ReceiveBuilder.
    match(GetSuccess.class, a -> a.key().equals(set1Key), a -> {
        GetSuccess<GSet<String>> g = a;
        Set<String> value = g.dataValue().getElements();
    }).
    match(GetFailure.class, a -> a.key().equals(set1Key), a -> {
        // read from 3 nodes failed within 1.second
    }).
    match(NotFound.class, a -> a.key().equals(set1Key), a -> {
```

```
// key set1 does not exist
}).build());
```

You will always read your own writes. For example if you send a `Update` message followed by a `Get` of the same key the `Get` will retrieve the change that was performed by the preceding `Update` message. However, the order of the reply messages are not defined, i.e. in the previous example you may receive the `GetSuccess` before the `UpdateSuccess`.

In the `Get` message you can pass an optional request context in the same way as for the `Update` message, described above. For example the original sender can be passed and replied to after receiving and transforming `GetSuccess`.

```
final ActorRef replicator = DistributedData.get(system).replicator();
final ReadConsistency readTwo = new ReadFrom(2, Duration.create(3, SECONDS));
final Key<PNCCounter> counter1Key = PNCCounterKey.create("counter1");

receive(ReceiveBuilder.
  match(String.class, a -> a.equals("get-count"), a -> {
    // incoming request to retrieve current value of the counter
    Optional<Object> reqContext = Optional.of(sender());
    replicator.tell(new Replicator.Get<PNCCounter>(counter1Key,
      readTwo), self());
  }).

  match(GetSuccess.class, a -> a.key().equals(counter1Key), a -> {
    ActorRef replyTo = (ActorRef) a.getRequest().get();
    GetSuccess<PNCCounter> g = a;
    long value = g.dataValue().getValue().longValue();
    replyTo.tell(value, self());
  }).

  match(GetFailure.class, a -> a.key().equals(counter1Key), a -> {
    ActorRef replyTo = (ActorRef) a.getRequest().get();
    replyTo.tell(-1L, self());
  }).

  match(NotFound.class, a -> a.key().equals(counter1Key), a -> {
    ActorRef replyTo = (ActorRef) a.getRequest().get();
    replyTo.tell(0L, self());
  }).build());
```

Consistency

The consistency level that is supplied in the `Update` and `Get` specifies per request how many replicas that must respond successfully to a write and read request.

For low latency reads you use `ReadLocal` with the risk of retrieving stale data, i.e. updates from other nodes might not be visible yet.

When using `writeLocal` the update is only written to the local replica and then disseminated in the background with the gossip protocol, which can take few seconds to spread to all nodes.

`writeAll` and `readAll` is the strongest consistency level, but also the slowest and with lowest availability. For example, it is enough that one node is unavailable for a `Get` request and you will not receive the value.

If consistency is important, you can ensure that a read always reflects the most recent write by using the following formula:

```
(nodes_written + nodes_read) > N
```

where `N` is the total number of nodes in the cluster, or the number of nodes with the role that is used for the `Replicator`.

For example, in a 7 node cluster these consistency properties are achieved by writing to 4 nodes and reading from 4 nodes, or writing to 5 nodes and reading from 3 nodes.

By combining `writeMajority` and `readMajority` levels a read always reflects the most recent write. The Replicator writes and reads to a majority of replicas, i.e. $N / 2 + 1$. For example, in a 5 node cluster it writes to 3 nodes and reads from 3 nodes. In a 6 node cluster it writes to 4 nodes and reads from 4 nodes.

Here is an example of using `writeMajority` and `readMajority`:

```
private final WriteConsistency writeMajority =
    new WriteMajority(Duration.create(3, SECONDS));
private final static ReadConsistency readMajority =
    new ReadMajority(Duration.create(3, SECONDS));

private PartialFunction<Object, BoxedUnit> matchGetCart() {
    return ReceiveBuilder
        .matchEquals((GET_CART),
            s -> receiveGetCart())
        .match(GetSuccess.class, g -> isResponseToGetCart(g),
            g -> receiveGetSuccess((GetSuccess<LWWMap<LineItem>>) g))
        .match(NotFound.class, n -> isResponseToGetCart(n),
            n -> receiveNotFound((NotFound<LWWMap<LineItem>>) n))
        .match(GetFailure.class, f -> isResponseToGetCart(f),
            f -> receiveGetFailure((GetFailure<LWWMap<LineItem>>) f))
        .build();
}

private void receiveGetCart() {
    Optional<Object> ctx = Optional.of(sender());
    replicator.tell(new Replicator.Get<LWWMap<LineItem>>(dataKey, readMajority, ctx),
        self());
}

private boolean isResponseToGetCart(GetResponse<?> response) {
    return response.key().equals(dataKey) &&
        (response.getRequest().orElse(null) instanceof ActorRef);
}

private void receiveGetSuccess(GetSuccess<LWWMap<LineItem>> g) {
    Set<LineItem> items = new HashSet<>(g.dataValue().getEntries().values());
    ActorRef replyTo = (ActorRef) g.getRequest().get();
    replyTo.tell(new Cart(items), self());
}

private void receiveNotFound(NotFound<LWWMap<LineItem>> n) {
    ActorRef replyTo = (ActorRef) n.getRequest().get();
    replyTo.tell(new Cart(new HashSet<>()), self());
}

private void receiveGetFailure(GetFailure<LWWMap<LineItem>> f) {
    // ReadMajority failure, try again with local read
    Optional<Object> ctx = Optional.of(sender());
    replicator.tell(new Replicator.Get<LWWMap<LineItem>>(dataKey, Replicator.readLocal(),
        ctx), self());
}

private PartialFunction<Object, BoxedUnit> matchAddItem() {
    return ReceiveBuilder
        .match(AddItem.class, r -> receiveAddItem(r))
        .build();
}

private void receiveAddItem(AddItem add) {
```

```
Update<LWWMap<LineItem>> update = new Update<>(dataKey, LWWMap.create(), writeMajority,
    cart -> updateCart(cart, add.item));
replicator.tell(update, self());
}
```

In some rare cases, when performing an Update it is needed to first try to fetch latest data from other nodes. That can be done by first sending a Get with ReadMajority and then continue with the Update when the GetSuccess, GetFailure or NotFound reply is received. This might be needed when you need to base a decision on latest information or when removing entries from ORSet or ORMap. If an entry is added to an ORSet or ORMap from one node and removed from another node the entry will only be removed if the added entry is visible on the node where the removal is performed (hence the name observed-removed set).

The following example illustrates how to do that:

```
private void receiveRemoveItem(RemoveItem rm) {
    // Try to fetch latest from a majority of nodes first, since ORMap
    // remove must have seen the item to be able to remove it.
    Optional<Object> ctx = Optional.of(rm);
    replicator.tell(new Replicator.Get<LWWMap<LineItem>>(dataKey, readMajority, ctx),
        self());
}

private void receiveRemoveItemGetSuccess(GetSuccess<LWWMap<LineItem>> g) {
    RemoveItem rm = (RemoveItem) g.getRequest().get();
    removeItem(rm.productId);
}

private void receiveRemoveItemGetFailure(GetFailure<LWWMap<LineItem>> f) {
    // ReadMajority failed, fall back to best effort local value
    RemoveItem rm = (RemoveItem) f.getRequest().get();
    removeItem(rm.productId);
}

private void removeItem(String productId) {
    Update<LWWMap<LineItem>> update = new Update<>(dataKey, LWWMap.create(), writeMajority,
        cart -> cart.remove(node, productId));
    replicator.tell(update, self());
}

private boolean isResponseToRemoveItem(GetResponse<?> response) {
    return response.key().equals(dataKey) &&
        (response.getRequest().orElse(null) instanceof RemoveItem);
}
```

Warning: *Caveat:* Even if you use writeMajority and readMajority there is small risk that you may read stale data if the cluster membership has changed between the Update and the Get. For example, in cluster of 5 nodes when you Update and that change is written to 3 nodes: n1, n2, n3. Then 2 more nodes are added and a Get request is reading from 4 nodes, which happens to be n4, n5, n6, n7, i.e. the value on n1, n2, n3 is not seen in the response of the Get request.

Subscribe

You may also register interest in change notifications by sending Replicator.Subscribe message to the Replicator. It will send Replicator.Changed messages to the registered subscriber when the data for the subscribed key is updated. Subscribers will be notified periodically with the configured notify-subscribers-interval, and it is also possible to send an explicit Replicator.FlushChanges message to the Replicator to notify the subscribers immediately.

The subscriber is automatically removed if the subscriber is terminated. A subscriber can also be deregistered with the Replicator.Unsubscribe message.

```

final ActorRef replicator = DistributedData.get(system).replicator();
final Key<PNCounter> counter1Key = PNCounterKey.create("counter1");

BigInteger currentValue = BigInteger.valueOf(0);

public MyActor() {
  receive(ReceiveBuilder.
    match(Changed.class, a -> a.key().equals(counter1Key), a -> {
      Changed<PNCounter> g = a;
      currentValue = g.dataValue().getValue();
    }).

    match(String.class, a -> a.equals("get-count"), a -> {
      // incoming request to retrieve current value of the counter
      sender().tell(currentValue, sender());
    }).build());
}

public void preStart() {
  // subscribe to changes of the Counter1Key value
  replicator.tell(new Subscribe<PNCounter>(counter1Key, self()), ActorRef.noSender());
}

```

Delete

A data entry can be deleted by sending a `Replicator.Delete` message to the local `Replicator`. As reply of the `Delete` a `Replicator.DeleteSuccess` is sent to the sender of the `Delete` if the value was successfully deleted according to the supplied consistency level within the supplied timeout. Otherwise a `Replicator.ReplicationDeleteFailure` is sent. Note that `ReplicationDeleteFailure` does not mean that the delete completely failed or was rolled back. It may still have been replicated to some nodes, and may eventually be replicated to all nodes.

A deleted key cannot be reused again, but it is still recommended to delete unused data entries because that reduces the replication overhead when new nodes join the cluster. Subsequent `Delete`, `Update` and `Get` requests will be replied with `Replicator.DataDeleted`. Subscribers will receive `Replicator.DataDeleted`.

```

final ActorRef replicator = DistributedData.get(system).replicator();
final Key<PNCounter> counter1Key = PNCounterKey.create("counter1");
final Key<ORSet<String>> set2Key = ORSetKey.create("set2");

replicator.tell(new Delete<PNCounter>(counter1Key,
  Replicator.writeLocal()), self());

final WriteConsistency writeMajority =
  new WriteMajority(Duration.create(5, SECONDS));
replicator.tell(new Delete<PNCounter>(counter1Key,
  writeMajority), self());

```

6.8.2 Data Types

The data types must be convergent (stateful) CRDTs and implement the `ReplicatedData` trait, i.e. they provide a monotonic merge function and the state changes always converge.

You can use your own custom `ReplicatedData` types, and several types are provided by this package, such as:

- Counters: `GCounter`, `PNCounter`
- Sets: `GSet`, `ORSet`
- Maps: `ORMap`, `ORMultiMap`, `LWWMap`, `PNCounterMap`

- Registers: `LWWRegister`, `Flag`

Counters

`GCounter` is a “grow only counter”. It only supports increments, no decrements.

It works in a similar way as a vector clock. It keeps track of one counter per node and the total value is the sum of these counters. The `merge` is implemented by taking the maximum count for each node.

If you need both increments and decrements you can use the `PNCCounter` (positive/negative counter).

It is tracking the increments (P) separate from the decrements (N). Both P and N are represented as two internal `GCounter`. Merge is handled by merging the internal P and N counters. The value of the counter is the value of the P counter minus the value of the N counter.

```
final Cluster node = Cluster.get(system);
final PNCCounter c0 = PNCCounter.create();
final PNCCounter c1 = c0.increment(node, 1);
final PNCCounter c2 = c1.increment(node, 7);
final PNCCounter c3 = c2.decrement(node, 2);
System.out.println(c3.value()); // 6
```

Several related counters can be managed in a map with the `PNCCounterMap` data type. When the counters are placed in a `PNCCounterMap` as opposed to placing them as separate top level values they are guaranteed to be replicated together as one unit, which is sometimes necessary for related data.

```
final Cluster node = Cluster.get(system);
final PNCCounterMap m0 = PNCCounterMap.create();
final PNCCounterMap m1 = m0.increment(node, "a", 7);
final PNCCounterMap m2 = m1.decrement(node, "a", 2);
final PNCCounterMap m3 = m2.increment(node, "b", 1);
System.out.println(m3.get("a")); // 5
System.out.println(m3.getEntries());
```

Sets

If you only need to add elements to a set and not remove elements the `GSet` (grow-only set) is the data type to use. The elements can be any type of values that can be serialized. Merge is simply the union of the two sets.

```
final GSet<String> s0 = GSet.create();
final GSet<String> s1 = s0.add("a");
final GSet<String> s2 = s1.add("b").add("c");
if (s2.contains("a"))
    System.out.println(s2.getElements()); // a, b, c
```

If you need add and remove operations you should use the `ORSet` (observed-remove set). Elements can be added and removed any number of times. If an element is concurrently added and removed, the add will win. You cannot remove an element that you have not seen.

The `ORSet` has a version vector that is incremented when an element is added to the set. The version for the node that added the element is also tracked for each element in a so called “birth dot”. The version vector and the dots are used by the `merge` function to track causality of the operations and resolve concurrent updates.

```
final Cluster node = Cluster.get(system);
final ORSet<String> s0 = ORSet.create();
final ORSet<String> s1 = s0.add(node, "a");
final ORSet<String> s2 = s1.add(node, "b");
final ORSet<String> s3 = s2.remove(node, "a");
System.out.println(s3.getElements()); // b
```

Maps

`ORMap` (observed-remove map) is a map with `String` keys and the values are `ReplicatedData` types themselves. It supports add, remove and delete any number of times for a map entry.

If an entry is concurrently added and removed, the add will win. You cannot remove an entry that you have not seen. This is the same semantics as for the `ORSet`.

If an entry is concurrently updated to different values the values will be merged, hence the requirement that the values must be `ReplicatedData` types.

It is rather inconvenient to use the `ORMap` directly since it does not expose specific types of the values. The `ORMap` is intended as a low level tool for building more specific maps, such as the following specialized maps.

`ORMultiMap` (observed-remove multi-map) is a multi-map implementation that wraps an `ORMap` with an `ORSet` for the map's value.

`PNCounterMap` (positive negative counter map) is a map of named counters. It is a specialized `ORMap` with `PNCounter` values.

`LWWMap` (last writer wins map) is a specialized `ORMap` with `LWWRegister` (last writer wins register) values.

```
final Cluster node = Cluster.get(system);
final ORMultiMap<Integer> m0 = ORMultiMap.create();
final ORMultiMap<Integer> m1 = m0.put(node, "a",
    new HashSet<Integer>(Arrays.asList(1, 2, 3)));
final ORMultiMap<Integer> m2 = m1.addBinding(node, "a", 4);
final ORMultiMap<Integer> m3 = m2.removeBinding(node, "a", 2);
final ORMultiMap<Integer> m4 = m3.addBinding(node, "b", 1);
System.out.println(m4.getEntries());
```

Note that `LWWRegister` and therefore `LWWMap` relies on synchronized clocks and should only be used when the choice of value is not important for concurrent updates occurring within the clock skew.

Instead of using timestamps based on `System.currentTimeMillis()` time it is possible to use a timestamp value based on something else, for example an increasing version number from a database record that is used for optimistic concurrency control.

When a data entry is changed the full state of that entry is replicated to other nodes, i.e. when you update a map the whole map is replicated. Therefore, instead of using one `ORMap` with 1000 elements it is more efficient to split that up in 10 top level `ORMap` entries with 100 elements each. Top level entries are replicated individually, which has the trade-off that different entries may not be replicated at the same time and you may see inconsistencies between related entries. Separate top level entries cannot be updated atomically together.

Flags and Registers

`Flag` is a data type for a boolean value that is initialized to `false` and can be switched to `true`. Thereafter it cannot be changed. `true` wins over `false` in merge.

```
final Flag f0 = Flag.create();
final Flag f1 = f0.switchOn();
System.out.println(f1.enabled());
```

`LWWRegister` (last writer wins register) can hold any (serializable) value.

Merge of a `LWWRegister` takes the register with highest timestamp. Note that this relies on synchronized clocks. `LWWRegister` should only be used when the choice of value is not important for concurrent updates occurring within the clock skew.

Merge takes the register updated by the node with lowest address (`UniqueAddress` is ordered) if the timestamps are exactly the same.

```
final Cluster node = Cluster.get(system);
final LWWRegister<String> r1 = LWWRegister.create(node, "Hello");
```

```
final LWWRegister<String> r2 = r1.withValue(node, "Hi");
System.out.println(r1.value() + " by " + r1.updatedBy() + " at " + r1.timestamp());
```

Instead of using timestamps based on `System.currentTimeMillis()` time it is possible to use a timestamp value based on something else, for example an increasing version number from a database record that is used for optimistic concurrency control.

```
class Record {
    public final int version;
    public final String name;
    public final String address;

    public Record(int version, String name, String address) {
        this.version = version;
        this.name = name;
        this.address = address;
    }
}

final Cluster node = Cluster.get(system);
final LWWRegister.Clock<Record> recordClock = new LWWRegister.Clock<Record>() {
    @Override
    public long apply(long currentTimestamp, Record value) {
        return value.version;
    }
};

final Record record1 = new Record(1, "Alice", "Union Square");
final LWWRegister<Record> r1 = LWWRegister.create(node, record1);

final Record record2 = new Record(2, "Alice", "Madison Square");
final LWWRegister<Record> r2 = LWWRegister.create(node, record2);

final LWWRegister<Record> r3 = r1.merge(r2);
System.out.println(r3.value());
```

For first-write-wins semantics you can use the `LWWRegister#reverseClock` instead of the `LWWRegister#defaultClock`.

Custom Data Type

You can rather easily implement your own data types. The only requirement is that it implements the `mergeData` function of the `AbstractReplicatedData` class.

A nice property of stateful CRDTs is that they typically compose nicely, i.e. you can combine several smaller data types to build richer data structures. For example, the `PNCCounter` is composed of two internal `GCounter` instances to keep track of increments and decrements separately.

Here is a simple implementation of a custom `TwoPhaseSet` that is using two internal `GSet` types to keep track of addition and removals. A `TwoPhaseSet` is a set where an element may be added and removed, but never added again thereafter.

```
public class TwoPhaseSet extends AbstractReplicatedData<TwoPhaseSet> {

    public final GSet<String> adds;
    public final GSet<String> removals;

    public TwoPhaseSet(GSet<String> adds, GSet<String> removals) {
        this.adds = adds;
        this.removals = removals;
    }
}
```

```

public static TwoPhaseSet create() {
    return new TwoPhaseSet(GSet.create(), GSet.create());
}

public TwoPhaseSet add(String element) {
    return new TwoPhaseSet(adds.add(element), removals);
}

public TwoPhaseSet remove(String element) {
    return new TwoPhaseSet(adds, removals.add(element));
}

public Set<String> getElements() {
    Set<String> result = new HashSet<>(adds.getElements());
    result.removeAll(removals.getElements());
    return result;
}

@Override
public TwoPhaseSet mergeData(TwoPhaseSet that) {
    return new TwoPhaseSet(this.adds.merge(that.adds),
        this.removals.merge(that.removals));
}
}

```

Data types should be immutable, i.e. “modifying” methods should return a new instance.

Serialization

The data types must be serializable with an *Akka Serializer*. It is highly recommended that you implement efficient serialization with Protobuf or similar for your custom data types. The built in data types are marked with `ReplicatedDataSerialization` and serialized with `akka.cluster.ddata.protobuf.ReplicatedDataSerializer`.

Serialization of the data types are used in remote messages and also for creating message digests (SHA-1) to detect changes. Therefore it is important that the serialization is efficient and produce the same bytes for the same content. For example sets and maps should be sorted deterministically in the serialization.

This is a protobuf representation of the above `TwoPhaseSet`:

```

option java_package = "docs.ddata.protobuf.msg";
option optimize_for = SPEED;

message TwoPhaseSet {
    repeated string adds = 1;
    repeated string removals = 2;
}

```

The serializer for the `TwoPhaseSet`:

```

import docs.ddata.japi.TwoPhaseSet;
import docs.ddata.protobuf.msg.TwoPhaseSetMessages;
import docs.ddata.protobuf.msg.TwoPhaseSetMessages.TwoPhaseSet.Builder;
import java.util.ArrayList;
import java.util.Collections;

import akka.actor.ExtendedActorSystem;
import akka.cluster.ddata.GSet;
import akka.cluster.ddata.protobuf.AbstractSerializationSupport;

public class TwoPhaseSetSerializer extends AbstractSerializationSupport {

```

```

private final ExtendedActorSystem system;

public TwoPhaseSetSerializer(ExtendedActorSystem system) {
    this.system = system;
}

@Override
public ExtendedActorSystem system() {
    return this.system;
}

@Override
public boolean includeManifest() {
    return false;
}

@Override
public int identifier() {
    return 99998;
}

@Override
public byte[] toBinary(Object obj) {
    if (obj instanceof TwoPhaseSet) {
        return twoPhaseSetToProto((TwoPhaseSet) obj).toByteArray();
    } else {
        throw new IllegalArgumentException(
            "Can't serialize object of type " + obj.getClass());
    }
}

@Override
public Object fromBinaryJava(byte[] bytes, Class<?> manifest) {
    return twoPhaseSetFromBinary(bytes);
}

protected TwoPhaseSetMessages.TwoPhaseSet twoPhaseSetToProto(TwoPhaseSet twoPhaseSet) {
    Builder b = TwoPhaseSetMessages.TwoPhaseSet.newBuilder();
    ArrayList<String> adds = new ArrayList<>(twoPhaseSet.adds.getElements());
    if (!adds.isEmpty()) {
        Collections.sort(adds);
        b.addAllAdds(adds);
    }
    ArrayList<String> removals = new ArrayList<>(twoPhaseSet.removals.getElements());
    if (!removals.isEmpty()) {
        Collections.sort(removals);
        b.addAllRemovals(removals);
    }
    return b.build();
}

protected TwoPhaseSet twoPhaseSetFromBinary(byte[] bytes) {
    try {
        TwoPhaseSetMessages.TwoPhaseSet msg =
            TwoPhaseSetMessages.TwoPhaseSet.parseFrom(bytes);
        GSet<String> adds = GSet.create();
        for (String elem : msg.getAddsList()) {
            adds = adds.add(elem);
        }
        GSet<String> removals = GSet.create();
        for (String elem : msg.getRemovalsList()) {
            removals = removals.add(elem);
        }
    }
}

```

```

    return new TwoPhaseSet(adds, removals);
  } catch (Exception e) {
    throw new RuntimeException(e.getMessage(), e);
  }
}
}

```

Note that the elements of the sets are sorted so the SHA-1 digests are the same for the same elements.

You register the serializer in configuration:

```

akka.actor {
  serializers {
    twophaseset = "docs.ddata.japi.protobuf.TwoPhaseSetSerializer"
  }
  serialization-bindings {
    "docs.ddata.japi.TwoPhaseSet" = twophaseset
  }
}

```

Using compression can sometimes be a good idea to reduce the data size. Gzip compression is provided by the `akka.cluster.ddata.protobuf.SerializationSupport` trait:

```

@Override
public byte[] toBinary(Object obj) {
  if (obj instanceof TwoPhaseSet) {
    return compress(twoPhaseSetToProto((TwoPhaseSet) obj));
  } else {
    throw new IllegalArgumentException(
      "Can't serialize object of type " + obj.getClass());
  }
}

@Override
public Object fromBinaryJava(byte[] bytes, Class<?> manifest) {
  return twoPhaseSetFromBinary(decompress(bytes));
}

```

The two embedded GSet can be serialized as illustrated above, but in general when composing new data types from the existing built in types it is better to make use of the existing serializer for those types. This can be done by declaring those as bytes fields in protobuf:

```

message TwoPhaseSet2 {
  optional bytes adds = 1;
  optional bytes removals = 2;
}

```

and use the methods `otherMessageToProto` and `otherMessageFromBinary` that are provided by the `SerializationSupport` trait to serialize and deserialize the GSet instances. This works with any type that has a registered Akka serializer. This is how such a serializer would look like for the `TwoPhaseSet`:

```

import docs.ddata.japi.TwoPhaseSet;
import docs.ddata.protobuf.msg.TwoPhaseSetMessages;
import docs.ddata.protobuf.msg.TwoPhaseSetMessages.TwoPhaseSet2.Builder;
import java.util.ArrayList;
import java.util.Collections;

import akka.actor.ExtendedActorSystem;
import akka.cluster.ddata.GSet;
import akka.cluster.ddata.protobuf.AbstractSerializationSupport;
import akka.cluster.ddata.protobuf.ReplicatedDataSerializer;

public class TwoPhaseSetSerializer2 extends AbstractSerializationSupport {

```

```

private final ExtendedActorSystem system;
private final ReplicatedDataSerializer replicatedDataSerializer;

public TwoPhaseSetSerializer2(ExtendedActorSystem system) {
    this.system = system;
    this.replicatedDataSerializer = new ReplicatedDataSerializer(system);
}

@Override
public ExtendedActorSystem system() {
    return this.system;
}

@Override
public boolean includeManifest() {
    return false;
}

@Override
public int identifier() {
    return 99998;
}

@Override
public byte[] toBinary(Object obj) {
    if (obj instanceof TwoPhaseSet) {
        return twoPhaseSetToProto((TwoPhaseSet) obj).toByteArray();
    } else {
        throw new IllegalArgumentException(
            "Can't serialize object of type " + obj.getClass());
    }
}

@Override
public Object fromBinaryJava(byte[] bytes, Class<?> manifest) {
    return twoPhaseSetFromBinary(bytes);
}

protected TwoPhaseSetMessages.TwoPhaseSet2 twoPhaseSetToProto(TwoPhaseSet twoPhaseSet) {
    Builder b = TwoPhaseSetMessages.TwoPhaseSet2.newBuilder();
    if (!twoPhaseSet.adds.isEmpty())
        b.setAdds(otherMessageToProto(twoPhaseSet.adds).toByteString());
    if (!twoPhaseSet.removals.isEmpty())
        b.setRemovals(otherMessageToProto(twoPhaseSet.removals).toByteString());
    return b.build();
}

@SuppressWarnings("unchecked")
protected TwoPhaseSet twoPhaseSetFromBinary(byte[] bytes) {
    try {
        TwoPhaseSetMessages.TwoPhaseSet2 msg =
            TwoPhaseSetMessages.TwoPhaseSet2.parseFrom(bytes);

        GSet<String> adds = GSet.create();
        if (msg.hasAdds())
            adds = (GSet<String>) otherMessageFromBinary(msg.getAdds().toByteArray());

        GSet<String> removals = GSet.create();
        if (msg.hasRemovals())
            adds = (GSet<String>) otherMessageFromBinary(msg.getRemovals().toByteArray());

        return new TwoPhaseSet(adds, removals);
    } catch (Exception e) {

```

```

    throw new RuntimeException(e.getMessage(), e);
  }
}
}

```

CRDT Garbage

One thing that can be problematic with CRDTs is that some data types accumulate history (garbage). For example a `GCounter` keeps track of one counter per node. If a `GCounter` has been updated from one node it will associate the identifier of that node forever. That can become a problem for long running systems with many cluster nodes being added and removed. To solve this problem the `Replicator` performs pruning of data associated with nodes that have been removed from the cluster. Data types that need pruning have to implement the `RemovedNodePruning` trait.

6.8.3 Samples

Several interesting samples are included and described in the [Typesafe Activator tutorial](#) named [Akka Distributed Data Samples with Java](#).

- [Low Latency Voting Service](#)
- [Highly Available Shopping Cart](#)
- [Distributed Service Registry](#)
- [Replicated Cache](#)
- [Replicated Metrics](#)

6.8.4 Limitations

There are some limitations that you should be aware of.

CRDTs cannot be used for all types of problems, and eventual consistency does not fit all domains. Sometimes you need strong consistency.

It is not intended for *Big Data*. The number of top level entries should not exceed 100000. When a new node is added to the cluster all these entries are transferred (gossiped) to the new node. The entries are split up in chunks and all existing nodes collaborate in the gossip, but it will take a while (tens of seconds) to transfer all entries and this means that you cannot have too many top level entries. The current recommended limit is 100000. We will be able to improve this if needed, but the design is still not intended for billions of entries.

All data is held in memory, which is another reason why it is not intended for *Big Data*.

When a data entry is changed the full state of that entry is replicated to other nodes. For example, if you add one element to a `Set` with 100 existing elements, all 101 elements are transferred to other nodes. This means that you cannot have too large data entries, because then the remote message size will be too large. We might be able to make this more efficient by implementing [Efficient State-based CRDTs by Delta-Mutation](#).

The data is only kept in memory. It is redundant since it is replicated to other nodes in the cluster, but if you stop all nodes the data is lost, unless you have saved it elsewhere. Making the data durable is a possible future feature, but even if we implement that it is not intended to be a full featured database.

6.8.5 Learn More about CRDTs

- [The Final Causal Frontier](#) talk by Sean Cribbs
- [Eventually Consistent Data Structures](#) talk by Sean Cribbs
- [Strong Eventual Consistency and Conflict-free Replicated Data Types](#) talk by Mark Shapiro

- A comprehensive study of Convergent and Commutative Replicated Data Types paper by Mark Shapiro et al.

Dependencies

To use Distributed Data you must add the following dependency in your project.

sbt:

```
"com.typesafe.akka" %% "akka-distributed-data-experimental" % "2.4.0"
```

maven:

```
<dependency>
  <groupId>com.typesafe.akka</groupId>
  <artifactId>akka-distributed-data-experimental_2.11</artifactId>
  <version>2.4.0</version>
</dependency>
```

6.8.6 Configuration

The DistributedData extension can be configured with the following properties:

```
# Settings for the DistributedData extension
akka.cluster.distributed-data {
  # Actor name of the Replicator actor, /system/ddataReplicator
  name = ddataReplicator

  # Replicas are running on members tagged with this role.
  # All members are used if undefined or empty.
  role = ""

  # How often the Replicator should send out gossip information
  gossip-interval = 2 s

  # How often the subscribers will be notified of changes, if any
  notify-subscribers-interval = 500 ms

  # Maximum number of entries to transfer in one gossip message when synchronizing
  # the replicas. Next chunk will be transferred in next round of gossip.
  max-delta-elements = 1000

  # The id of the dispatcher to use for Replicator actors. If not specified
  # default dispatcher is used.
  # If specified you need to define the settings of the actual dispatcher.
  use-dispatcher = ""

  # How often the Replicator checks for pruning of data associated with
  # removed cluster nodes.
  pruning-interval = 30 s

  # How long time it takes (worst case) to spread the data to all other replica nodes.
  # This is used when initiating and completing the pruning process of data associated
  # with removed cluster nodes. The time measurement is stopped when any replica is
  # unreachable, so it should be configured to worst case in a healthy cluster.
  max-pruning-dissemination = 60 s
}
```

6.9 Remoting

For an introduction of remoting capabilities of Akka please see *Location Transparency*.

Note: As explained in that chapter Akka remoting is designed for communication in a peer-to-peer fashion and it has limitations for client-server setups. In particular Akka Remoting does not work transparently with Network Address Translation, Load Balancers, or in Docker containers. For symmetric communication in these situations network and/or Akka configuration will have to be changed as described in *Peer-to-Peer vs. Client-Server*.

6.9.1 Preparing your ActorSystem for Remoting

The Akka remoting is a separate jar file. Make sure that you have the following dependency in your project:

```
<dependency>
  <groupId>com.typesafe.akka</groupId>
  <artifactId>akka-remote_2.11</artifactId>
  <version>2.4.0</version>
</dependency>
```

To enable remote capabilities in your Akka project you should, at a minimum, add the following changes to your `application.conf` file:

```
akka {
  actor {
    provider = "akka.remote.RemoteActorRefProvider"
  }
  remote {
    enabled-transport = ["akka.remote.netty.tcp"]
    netty.tcp {
      hostname = "127.0.0.1"
      port = 2552
    }
  }
}
```

As you can see in the example above there are four things you need to add to get started:

- Change `provider` from `akka.actor.LocalActorRefProvider` to `akka.remote.RemoteActorRefProvider`
- Add host name - the machine you want to run the actor system on; this host name is exactly what is passed to remote systems in order to identify this system and consequently used for connecting back to this system if need be, hence set it to a reachable IP address or resolvable name in case you want to communicate across the network.
- Add port number - the port the actor system should listen on, set to 0 to have it chosen automatically

Note: The port number needs to be unique for each actor system on the same machine even if the actor systems have different names. This is because each actor system has its own network subsystem listening for connections and handling messages as not to interfere with other actor systems.

The example above only illustrates the bare minimum of properties you have to add to enable remoting. All settings are described in *Remote Configuration*.

6.9.2 Looking up Remote Actors

`actorSelection(path)` will obtain an `ActorSelection` to an Actor on a remote node:

```
ActorSelection selection =
    context.actorSelection("akka.tcp://app@10.0.0.1:2552/user/serviceA/worker");
```

As you can see from the example above the following pattern is used to find an actor on a remote node:

```
akka.<protocol>://<actorsystemname>@<hostname>:<port>/<actor path>
```

Once you obtained a selection to the actor you can interact with it the same way you would with a local actor, e.g.:

```
selection.tell("Pretty awesome feature", getSelf());
```

To acquire an `ActorRef` for an `ActorSelection` you need to send a message to the selection and use the `getSender` reference of the reply from the actor. There is a built-in `Identify` message that all Actors will understand and automatically reply to with a `ActorIdentity` message containing the `ActorRef`. This can also be done with the `resolveOne` method of the `ActorSelection`, which returns a `Future` of the matching `ActorRef`.

Note: For more details on how actor addresses and paths are formed and used, please refer to *Actor References, Paths and Addresses*.

6.9.3 Creating Actors Remotely

If you want to use the creation functionality in Akka remoting you have to further amend the `application.conf` file in the following way (only showing deployment section):

```
akka {
  actor {
    deployment {
      /sampleActor {
        remote = "akka.tcp://sampleActorSystem@127.0.0.1:2553"
      }
    }
  }
}
```

The configuration above instructs Akka to react when an actor with path `/sampleActor` is created, i.e. using `system.actorOf(new Props(...), "sampleActor")`. This specific actor will not be directly instantiated, but instead the remote daemon of the remote system will be asked to create the actor, which in this sample corresponds to `sampleActorSystem@127.0.0.1:2553`.

Once you have configured the properties above you would do the following in code:

```
ActorRef actor = system.actorOf(Props.create(SampleActor.class), "sampleActor");
actor.tell("Pretty slick", ActorRef.noSender());
```

The actor class `SampleActor` has to be available to the runtimes using it, i.e. the classloader of the actor systems has to have a JAR containing the class.

Note: In order to ensure serializability of `Props` when passing constructor arguments to the actor being created, do not make the factory a non-static inner class: this will inherently capture a reference to its enclosing object, which in most cases is not serializable. It is best to make a static inner class which implements `Creator`.

Serializability of all `Props` can be tested by setting the configuration item `akka.actor.serialize-creators=on`. Only `Props` whose `deploy` has `LocalScope` are exempt from this check.

Note: You can use asterisks as wildcard matches for the actor path sections, so you could specify: `/*/sampleActor` and that would match all `sampleActor` on that level in the hierarchy. You can also use

wildcard in the last position to match all actors at a certain level: `/someParent/*`. Non-wildcard matches always have higher priority to match than wildcards, so: `/foo/bar` is considered **more specific** than `/foo/*` and only the highest priority match is used. Please note that it **cannot** be used to partially match section, like this: `/foo*/bar`, `/f*/bar` etc.

Programmatic Remote Deployment

To allow dynamically deployed systems, it is also possible to include deployment configuration in the `Props` which are used to create an actor: this information is the equivalent of a deployment section from the configuration file, and if both are given, the external configuration takes precedence.

With these imports:

```
import akka.actor.ActorRef;
import akka.actor.Address;
import akka.actor.AddressFromURIStrng;
import akka.actor.Deploy;
import akka.actor.Props;
import akka.actor.ActorSystem;
import akka.remote.RemoteScope;
```


and a remote address like this:

```
Address addr = new Address("akka.tcp", "sys", "host", 1234);
addr = AddressFromURIStrng.parse("akka.tcp://sys@host:1234"); // the same
```

you can advise the system to create a child on that remote node like so:

```
ActorRef ref = system.actorOf(Props.create(SampleActor.class).withDeploy(
    new Deploy(new RemoteScope(addr))));
```

6.9.4 Lifecycle and Failure Recovery Model



Each link with a remote system can be in one of the four states as illustrated above. Before any communication happens with a remote system at a given `Address` the state of the association is `Idle`. The first time a message is attempted to be sent to the remote system or an inbound connection is accepted the state of the link transitions to `Active` denoting that the two systems has messages to send or receive and no failures were encountered so far. When a communication failure happens and the connection is lost between the two systems the link becomes `Gated`.

In this state the system will not attempt to connect to the remote host and all outbound messages will be dropped. The time while the link is in the `Gated` state is controlled by the setting `akka.remote.retry-gate-closed-for`: after this time elapses the link state transitions to `Idle` again. Gate is one-sided in the sense that whenever a successful *inbound* connection is accepted from a remote system during Gate it automatically transitions to `Active` and communication resumes immediately.

In the face of communication failures that are unrecoverable because the state of the participating systems are inconsistent, the remote system becomes `Quarantined`. Unlike Gate, quarantining is permanent and lasts until one of the systems is restarted. After a restart communication can be resumed again and the link can become `Active` again.

6.9.5 Watching Remote Actors

Watching a remote actor is not different than watching a local actor, as described in *Lifecycle Monitoring aka DeathWatch*.

Failure Detector

Under the hood remote death watch uses heartbeat messages and a failure detector to generate `Terminated` message from network failures and JVM crashes, in addition to graceful termination of watched actor.

The heartbeat arrival times is interpreted by an implementation of [The Phi Accrual Failure Detector](#).

The suspicion level of failure is given by a value called *phi*. The basic idea of the phi failure detector is to express the value of *phi* on a scale that is dynamically adjusted to reflect current network conditions.

The value of *phi* is calculated as:


```
phi = -log10(1 - F(timeSinceLastHeartbeat))
```

where *F* is the cumulative distribution function of a normal distribution with mean and standard deviation estimated from historical heartbeat inter-arrival times.


In the *Remote Configuration* you can adjust the `akka.remote.watch-failure-detector.threshold` to define when a *phi* value is considered to be a failure.

A low `threshold` is prone to generate many false positives but ensures a quick detection in the event of a real crash. Conversely, a high `threshold` generates fewer mistakes but needs more time to detect actual crashes. The default `threshold` is 10 and is appropriate for most situations. However in cloud environments, such as Amazon EC2, the value could be increased to 12 in order to account for network issues that sometimes occur on such platforms.


The following chart illustrates how *phi* increase with increasing time since the previous heartbeat.



Phi is calculated from the mean and standard deviation of historical inter arrival times. The previous chart is an example for standard deviation of 200 ms. If the heartbeats arrive with less deviation the curve becomes steeper, i.e. it is possible to determine failure more quickly. The curve looks like this for a standard deviation of 100 ms.



To be able to survive sudden abnormalities, such as garbage collection pauses and transient network failures the failure detector is configured with a margin, `akka.remote.watch-failure-detector.acceptable-heartbeat-pause`. You may want to adjust the *Remote Configuration* of this depending on you environment. This is how the curve looks like for `acceptable-heartbeat-pause` configured to 3 seconds.



6.9.6 Serialization

When using remoting for actors you must ensure that the `props` and `messages` used for those actors are serializable. Failing to do so will cause the system to behave in an unintended way.

For more information please see *Serialization*

6.9.7 Routers with Remote Destinations

It is absolutely feasible to combine remoting with *Routing*.

A pool of remote deployed routees can be configured as:

```
akka.actor.deployment {
  /parent/remotePool {
    router = round-robin-pool
    nr-of-instances = 10
    target.nodes = ["akka.tcp://app@10.0.0.2:2552", "akka://app@10.0.0.3:2552"]
  }
}
```

This configuration setting will clone the actor defined in the `Props` of the `remotePool` 10 times and deploy it evenly distributed across the two given target nodes.

A group of remote actors can be configured as:

```
akka.actor.deployment {
  /parent/remoteGroup {
    router = round-robin-group
    routees.paths = [
      "akka.tcp://app@10.0.0.1:2552/user/workers/w1",
      "akka.tcp://app@10.0.0.2:2552/user/workers/w1",
      "akka.tcp://app@10.0.0.3:2552/user/workers/w1"]
  }
}
```

This configuration setting will send messages to the defined remote actor paths. It requires that you create the destination actors on the remote nodes with matching paths. That is not done by the router.

6.9.8 Remoting Sample

There is a more extensive remote example that comes with *Typesafe Activator*. The tutorial named *Akka Remote Samples with Java* demonstrates both remote deployment and look-up of remote actors.

Pluggable transport support

Akka can be configured to use various transports to communicate with remote systems. The core component of this feature is the `akka.remote.transport.Transport` SPI. Transport implementations must extend this trait. Transports can be loaded by setting the `akka.remote.enabled-transports` configuration key to point to one or more configuration sections containing driver descriptions.

An example of setting up the default Netty based SSL driver as default:

```
akka {
  remote {
    enabled-transports = [akka.remote.netty.ssl]

    netty.ssl.security {
      key-store = "mykeystore"
      trust-store = "mytruststore"
      key-store-password = "changeme"
    }
  }
}
```

```

key-password = "changeme"
trust-store-password = "changeme"
protocol = "TLSv1"
random-number-generator = "AES128CounterSecureRNG"
enabled-algorithms = [TLS_RSA_WITH_AES_128_CBC_SHA]
}
}
}

```

An example of setting up a custom transport implementation:

```

akka {
  remote {
    applied-transport = ["akka.remote.mytransport"]

    mytransport {
      # The transport-class configuration entry is required, and
      # it must contain the fully qualified name of the transport
      # implementation
      transport-class = "my.package.MyTransport"

      # It is possible to decorate Transports with additional services.
      # Adapters should be registered in the "adapters" sections to
      # be able to apply them to transports
      applied-adapters = []

      # Driver specific configuration options has to be in the same
      # section:
      some-config = foo
      another-config = bar
    }
  }
}

```

Remote Events

It is possible to listen to events that occur in Akka Remote, and to subscribe/unsubscribe to these events you simply register as listener to the below described types in on the `ActorSystem.eventStream`.

Note: To subscribe to any remote event, subscribe to `RemotingLifecycleEvent`. To subscribe to events related only to the lifecycle of associations, subscribe to `akka.remote.AssociationEvent`.

Note: The use of term “Association” instead of “Connection” reflects that the remoting subsystem may use connectionless transports, but an association similar to transport layer connections is maintained between endpoints by the Akka protocol.

By default an event listener is registered which logs all of the events described below. This default was chosen to help setting up a system, but it is quite common to switch this logging off once that phase of the project is finished.

Note: In order to switch off the logging, set `akka.remote.log-remote-lifecycle-events = off` in your `application.conf`.

To be notified when an association is over (“disconnected”) listen to `DisassociatedEvent` which holds the direction of the association (inbound or outbound) and the addresses of the involved parties.

To be notified when an association is successfully established (“connected”) listen to `AssociatedEvent` which holds the direction of the association (inbound or outbound) and the addresses of the involved parties.

To intercept errors directly related to associations, listen to `AssociationErrorEvent` which holds the direction of the association (inbound or outbound), the addresses of the involved parties and the `Throwable` cause.

To be notified when the remoting subsystem is ready to accept associations, listen to `RemotingListenEvent` which contains the addresses the remoting listens on.

To be notified when the remoting subsystem has been shut down, listen to `RemotingShutdownEvent`.

To intercept generic remoting related errors, listen to `RemotingErrorEvent` which holds the `Throwable` cause.

6.9.9 Remote Security

Akka provides a couple of ways to enhance security between remote nodes (client/server):

- Untrusted Mode
- Security Cookie Handshake

Untrusted Mode

As soon as an actor system can connect to another remotely, it may in principle send any possible message to any actor contained within that remote system. One example may be sending a `PoisonPill` to the system guardian, shutting that system down. This is not always desired, and it can be disabled with the following setting:

```
akka.remote.untrusted-mode = on
```

This disallows sending of system messages (actor life-cycle commands, `DeathWatch`, etc.) and any message extending `PossiblyHarmful` to the system on which this flag is set. Should a client send them nonetheless they are dropped and logged (at `DEBUG` level in order to reduce the possibilities for a denial of service attack). `PossiblyHarmful` covers the predefined messages like `PoisonPill` and `Kill`, but it can also be added as a marker trait to user-defined messages.

Messages sent with actor selection are by default discarded in untrusted mode, but permission to receive actor selection messages can be granted to specific actors defined in configuration:

```
akka.remote.trusted-selection-paths = ["/user/receptionist", "/user/namingService"]
```

The actual message must still not be of type `PossiblyHarmful`.

In summary, the following operations are ignored by a system configured in untrusted mode when incoming via the remoting layer:

- remote deployment (which also means no remote supervision)
- remote `DeathWatch`
- `system.stop()`, `PoisonPill`, `Kill`
- sending any message which extends from the `PossiblyHarmful` marker interface, which includes `Terminated`
- messages sent with actor selection, unless destination defined in `trusted-selection-paths`.

Note: Enabling the untrusted mode does not remove the capability of the client to freely choose the target of its message sends, which means that messages not prohibited by the above rules can be sent to any actor in the remote system. It is good practice for a client-facing system to only contain a well-defined set of entry point actors, which then forward requests (possibly after performing validation) to another actor system containing the actual worker actors. If messaging between these two server-side systems is done using local `ActorRef` (they can be exchanged safely between actor systems within the same JVM), you can restrict the messages on this interface by marking them `PossiblyHarmful` so that a client cannot forge them.

SSL

SSL can be used as the remote transport by adding `akka.remote.netty.ssl` to the `enabled-transport` configuration section. See a description of the settings in the [Remote Configuration](#) section.

The SSL support is implemented with Java Secure Socket Extension, please consult the official [Java Secure Socket Extension documentation](#) and related resources for troubleshooting.

Note: When using SHA1PRNG on Linux it's recommended specify `-Djava.security.egd=file:/dev/./urandom` as argument to the JVM to prevent blocking. It is NOT as secure because it reuses the seed. Use `'/dev/urandom'`, not `'/dev/urandom'` as that doesn't work according to [Bug ID: 6202721](#).

6.9.10 Remote Configuration

There are lots of configuration properties that are related to remoting in Akka. We refer to the [reference configuration](#) for more information.

Note: Setting properties like the listening IP and port number programmatically is best done by using something like the following:

```
ConfigFactory.parseString("akka.remote.netty.tcp.hostname=\"1.2.3.4\"")
    .withFallback(ConfigFactory.load());
```

Remote configuration for NAT and Docker

In setups involving Network Address Translation (NAT), Load Balancers or Docker containers the hostname and port pair that akka binds to will be different than the “logical” host name and port pair that is used to connect to the system from the outside. This requires special configuration that sets both the logical and the bind pairs for remoting.

```
akka {
  remote {
    netty.tcp {
      hostname = my.domain.com      # external (logical) hostname
      port = 8000                  # external (logical) port

      bind-hostname = local.address # internal (bind) hostname
      bind-port = 2552             # internal (bind) port
    }
  }
}
```

6.10 Serialization

Akka has a built-in Extension for serialization, and it is both possible to use the built-in serializers and to write your own.

The serialization mechanism is both used by Akka internally to serialize messages, and available for ad-hoc serialization of whatever you might need it for.

6.10.1 Usage

Configuration

For Akka to know which Serializer to use for what, you need edit your *Configuration*, in the “akka.actor.serializers”-section you bind names to implementations of the akka.serialization.Serializer you wish to use, like this:

```
akka {
  actor {
    serializers {
      java = "akka.serialization.JavaSerializer"
      proto = "akka.remote.serialization.ProtobufSerializer"
      myown = "docs.serialization.MyOwnSerializer"
    }
  }
}
```

After you’ve bound names to different implementations of Serializer you need to wire which classes should be serialized using which Serializer, this is done in the “akka.actor.serialization-bindings”-section:

```
akka {
  actor {
    serializers {
      java = "akka.serialization.JavaSerializer"
      proto = "akka.remote.serialization.ProtobufSerializer"
      myown = "docs.serialization.MyOwnSerializer"
    }

    serialization-bindings {
      "java.lang.String" = java
      "docs.serialization.Customer" = java
      "com.google.protobuf.Message" = proto
      "docs.serialization.MyOwnSerializable" = myown
      "java.lang.Boolean" = myown
    }
  }
}
```

You only need to specify the name of an interface or abstract base class of the messages. In case of ambiguity, i.e. the message implements several of the configured classes, the most specific configured class will be used, i.e. the one of which all other candidates are superclasses. If this condition cannot be met, because e.g. java.io.Serializable and MyOwnSerializable both apply and neither is a subtype of the other, a warning will be issued.

Akka provides serializers for java.io.Serializable and protobuf com.google.protobuf.GeneratedMessage by default (the latter only if depending on the akka-remote module), so normally you don’t need to add configuration for that; since com.google.protobuf.GeneratedMessage implements java.io.Serializable, protobuf messages will always be serialized using the protobuf protocol unless specifically overridden. In order to disable a default serializer, map its marker type to “none”:

```
akka.actor.serialization-bindings {
  "java.io.Serializable" = none
}
```

Verification

If you want to verify that your messages are serializable you can enable the following config option:

```
akka {
  actor {
    serialize-messages = on
  }
}
```

Warning: We only recommend using the config option turned on when you're running tests. It is completely pointless to have it turned on in other scenarios.

If you want to verify that your Props are serializable you can enable the following config option:

```
akka {
  actor {
    serialize-creators = on
  }
}
```

Warning: We only recommend using the config option turned on when you're running tests. It is completely pointless to have it turned on in other scenarios.

Programmatic

If you want to programmatically serialize/deserialize using Akka Serialization, here's some examples:

```
import akka.actor.*;
import akka.serialization.*;
```

```
ActorSystem system = ActorSystem.create("example");

// Get the Serialization Extension
Serialization serialization = SerializationExtension.get(system);

// Have something to serialize
String original = "woohoo";

// Find the Serializer for it
Serializer serializer = serialization.findSerializerFor(original);

// Turn it into bytes
byte[] bytes = serializer.toBinary(original);

// Turn it back into an object,
// the nulls are for the class manifest and for the classloader
String back = (String) serializer.fromBinary(bytes);

// Voilà!
assertEquals(original, back);
```

For more information, have a look at the [ScalaDoc](#) for `akka.serialization._`

6.10.2 Customization

So, lets say that you want to create your own Serializer, you saw the `docs.serialization.MyOwnSerializer` in the config example above?

Creating new Serializers

First you need to create a class definition of your `Serializer`, which is done by extending `akka.serialization.JSerializer`, like this:

```
import akka.actor.*;
import akka.serialization.*;

public class MyOwnSerializer extends JSerializer {

    // This is whether "fromBinary" requires a "clazz" or not
    @Override public boolean includeManifest() {
        return false;
    }

    // Pick a unique identifier for your Serializer,
    // you've got a couple of billions to choose from,
    // 0 - 16 is reserved by Akka itself
    @Override public int identifier() {
        return 1234567;
    }

    // "toBinary" serializes the given object to an Array of Bytes
    @Override public byte[] toBinary(Object obj) {
        // Put the code that serializes the object here
        // ... ..
    }

    // "fromBinary" deserializes the given array,
    // using the type hint (if any, see "includeManifest" above)
    @Override public Object fromBinaryJava(byte[] bytes,
                                           Class<?> clazz) {
        // Put your code that deserializes here
        // ... ..
    }
}
```

The manifest is a type hint so that the same serializer can be used for different classes. The manifest parameter in `fromBinaryJava` is the class of the object that was serialized. In `fromBinary` you can match on the class and deserialize the bytes to different objects.

Then you only need to fill in the blanks, bind it to a name in your *Configuration* and then list which classes that should be serialized using it.

Serializer with String Manifest

The `Serializer` illustrated above supports a class based manifest (type hint). For serialization of data that need to evolve over time the `SerializerWithStringManifest` is recommended instead of `Serializer` because the manifest (type hint) is a `String` instead of a `Class`. That means that the class can be moved/removed and the serializer can still deserialize old data by matching on the `String`. This is especially useful for *Persistence*.

The manifest string can also encode a version number that can be used in `fromBinary` to deserialize in different ways to migrate old data to new domain objects.

If the data was originally serialized with `Serializer` and in a later version of the system you change to `SerializerWithStringManifest` the manifest string will be the full class name if you used `includeManifest=true`, otherwise it will be the empty string.

This is how a `SerializerWithStringManifest` looks like:

```
public class MyOwnSerializer2 extends SerializerWithStringManifest {
```

```

private static final String CUSTOMER_MANIFEST = "customer";
private static final String USER_MANIFEST = "user";
private static final String UTF_8 = StandardCharsets.UTF_8.name();

// Pick a unique identifier for your Serializer,
// you've got a couple of billions to choose from,
// 0 - 16 is reserved by Akka itself
@Override public int identifier() {
    return 1234567;
}

@Override public String manifest(Object obj) {
    if (obj instanceof Customer)
        return CUSTOMER_MANIFEST;
    else if (obj instanceof User)
        return USER_MANIFEST;
    else
        throw new IllegalArgumentException("Unknow type: " + obj);
}

// "toBinary" serializes the given object to an Array of Bytes
@Override public byte[] toBinary(Object obj) {
    // Put the real code that serializes the object here
    try {
        if (obj instanceof Customer)
            return ((Customer) obj).name.getBytes(UTF_8);
        else if (obj instanceof User)
            return ((User) obj).name.getBytes(UTF_8);
        else
            throw new IllegalArgumentException("Unknow type: " + obj);
    } catch (UnsupportedEncodingException e) {
        throw new RuntimeException(e.getMessage(), e);
    }
}

// "fromBinary" deserializes the given array,
// using the type hint
@Override public Object fromBinary(byte[] bytes, String manifest) {
    // Put the real code that deserializes here
    try {
        if (manifest.equals(CUSTOMER_MANIFEST))
            return new Customer(new String(bytes, UTF_8));
        else if (manifest.equals(USER_MANIFEST))
            return new User(new String(bytes, UTF_8));
        else
            throw new IllegalArgumentException("Unknow manifest: " + manifest);
    } catch (UnsupportedEncodingException e) {
        throw new RuntimeException(e.getMessage(), e);
    }
}
}

```

You must also bind it to a name in your *Configuration* and then list which classes that should be serialized using it.

Serializing ActorRefs

All ActorRefs are serializable using `JavaSerializer`, but in case you are writing your own serializer, you might want to know how to serialize and deserialize them properly. In the general case, the local address to be used depends on the type of remote address which shall be the recipient of the serialized information. Use `Serialization.serializedActorPath(actorRef)` like this:

```

import akka.actor.*;
import akka.serialization.*;

// Serialize
// (beneath toBinary)
String identifier = Serialization.serializedActorPath(theActorRef);

// Then just serialize the identifier however you like

// Deserialize
// (beneath fromBinary)
final ActorRef deserializedActorRef = extendedSystem.provider().resolveActorRef(
    identifier);
// Then just use the ActorRef

```

This assumes that serialization happens in the context of sending a message through the remote transport. There are other uses of serialization, though, e.g. storing actor references outside of an actor application (database, etc.). In this case, it is important to keep in mind that the address part of an actor's path determines how that actor is communicated with. Storing a local actor path might be the right choice if the retrieval happens in the same logical context, but it is not enough when deserializing it on a different network host: for that it would need to include the system's remote transport address. An actor system is not limited to having just one remote transport per se, which makes this question a bit more interesting. To find out the appropriate address to use when sending to remoteAddr you can use `ActorRefProvider.getExternalAddressFor(remoteAddr)` like this:

```

public class ExternalAddressExt implements Extension {
    private final ExtendedActorSystem system;

    public ExternalAddressExt(ExtendedActorSystem system) {
        this.system = system;
    }

    public Address getAddressFor(Address remoteAddress) {
        final scala.Option<Address> optAddr = system.provider()
            .getExternalAddressFor(remoteAddress);
        if (optAddr.isDefined()) {
            return optAddr.get();
        } else {
            throw new UnsupportedOperationException(
                "cannot send to remote address " + remoteAddress);
        }
    }
}

public class ExternalAddress extends
    AbstractExtensionId<ExternalAddressExt> implements ExtensionIdProvider {
    public static final ExternalAddress ID = new ExternalAddress();

    public ExternalAddress lookup() {
        return ID;
    }

    public ExternalAddressExt createExtension(ExtendedActorSystem system) {
        return new ExternalAddressExt(system);
    }
}

public class ExternalAddressExample {
    public String serializeTo(ActorRef ref, Address remote) {
        return ref.path().toSerializationFormatWithAddress(
            ExternalAddress.ID.get(system).getAddressFor(remote));
    }
}

```

Note: `ActorPath.toSerializationFormatWithAddress` differs from `toString` if the address does not already have host and port components, i.e. it only inserts address information for local addresses.

`toSerializationFormatWithAddress` also adds the unique id of the actor, which will change when the actor is stopped and then created again with the same name. Sending messages to a reference pointing the old actor will not be delivered to the new actor. If you do not want this behavior, e.g. in case of long term storage of the reference, you can use `toStringWithAddress`, which does not include the unique id.

This requires that you know at least which type of address will be supported by the system which will deserialize the resulting actor reference; if you have no concrete address handy you can create a dummy one for the right protocol using `new Address(protocol, "", "", 0)` (assuming that the actual transport used is as lenient as Akka's `RemoteActorRefProvider`).

There is also a default remote address which is the one used by cluster support (and typical systems have just this one); you can get it like this:

```
public class DefaultAddressExt implements Extension {
    private final ExtendedActorSystem system;

    public DefaultAddressExt(ExtendedActorSystem system) {
        this.system = system;
    }

    public Address getAddress() {
        return system.provider().getDefaultAddress();
    }
}

public class DefaultAddress extends
    AbstractExtensionId<DefaultAddressExt> implements ExtensionIdProvider {
    public static final DefaultAddress ID = new DefaultAddress();

    public DefaultAddress lookup() {
        return ID;
    }

    public DefaultAddressExt createExtension(ExtendedActorSystem system) {
        return new DefaultAddressExt(system);
    }
}
```

Deep serialization of Actors

The recommended approach to do deep serialization of internal actor state is to use Akka *Persistence*.

6.10.3 A Word About Java Serialization

When using Java serialization without employing the `JavaSerializer` for the task, you must make sure to supply a valid `ExtendedActorSystem` in the dynamic variable `JavaSerializer.currentSystem`. This is used when reading in the representation of an `ActorRef` for turning the string representation into a real reference. `DynamicVariable` is a thread-local variable, so be sure to have it set while deserializing anything which might contain actor references.

6.10.4 External Akka Serializers

[Akka-protostuff](#) by Roman Levenstein

[Akka-quickser](#) by Roman Levenstein

Akka-kryo by Roman Levenstein

6.11 I/O

6.11.1 Introduction

The `akka.io` package has been developed in collaboration between the Akka and [spray.io](#) teams. Its design combines experiences from the `spray-io` module with improvements that were jointly developed for more general consumption as an actor-based service.

The guiding design goal for this I/O implementation was to reach extreme scalability, make no compromises in providing an API correctly matching the underlying transport mechanism and to be fully event-driven, non-blocking and asynchronous. The API is meant to be a solid foundation for the implementation of network protocols and building higher abstractions; it is not meant to be a full-service high-level NIO wrapper for end users.

6.11.2 Terminology, Concepts

The I/O API is completely actor based, meaning that all operations are implemented with message passing instead of direct method calls. Every I/O driver (TCP, UDP) has a special actor, called a *manager* that serves as an entry point for the API. I/O is broken into several drivers. The manager for a particular driver is accessible by querying an `ActorSystem`. For example the following code looks up the TCP manager and returns its `ActorRef`:

```
final ActorRef tcpManager = Tcp.get(getContext().system()).manager();
```

The manager receives I/O command messages and instantiates worker actors in response. The worker actors present themselves to the API user in the reply to the command that was sent. For example after a `Connect` command sent to the TCP manager the manager creates an actor representing the TCP connection. All operations related to the given TCP connections can be invoked by sending messages to the connection actor which announces itself by sending a `Connected` message.

DeathWatch and Resource Management

I/O worker actors receive commands and also send out events. They usually need a user-side counterpart actor listening for these events (such events could be inbound connections, incoming bytes or acknowledgements for writes). These worker actors *watch* their listener counterparts. If the listener stops then the worker will automatically release any resources that it holds. This design makes the API more robust against resource leaks.

Thanks to the completely actor based approach of the I/O API the opposite direction works as well: a user actor responsible for handling a connection can watch the connection actor to be notified if it unexpectedly terminates.

Write models (Ack, Nack)

I/O devices have a maximum throughput which limits the frequency and size of writes. When an application tries to push more data than a device can handle, the driver has to buffer bytes until the device is able to write them. With buffering it is possible to handle short bursts of intensive writes — but no buffer is infinite. “Flow control” is needed to avoid overwhelming device buffers.

Akka supports two types of flow control:

- *Ack-based*, where the driver notifies the writer when writes have succeeded.
- *Nack-based*, where the driver notifies the writer when writes have failed.

Each of these models is available in both the TCP and the UDP implementations of Akka I/O.

Individual writes can be acknowledged by providing an `ack` object in the write message (`Write` in the case of TCP and `Send` for UDP). When the write is complete the worker will send the `ack` object to the writing actor. This can be used to implement *ack-based* flow control; sending new data only when old data has been acknowledged.

If a write (or any other command) fails, the driver notifies the actor that sent the command with a special message (`CommandFailed` in the case of UDP and TCP). This message will also notify the writer of a failed write, serving as a nack for that write. Please note, that in a nack-based flow-control setting the writer has to be prepared for the fact that the failed write might not be the most recent write it sent. For example, the failure notification for a write `W1` might arrive after additional write commands `W2` and `W3` have been sent. If the writer wants to resend any nacked messages it may need to keep a buffer of pending messages.

Warning: An acknowledged write does not mean acknowledged delivery or storage; receiving an ack for a write simply signals that the I/O driver has successfully processed the write. The Ack/Nack protocol described here is a means of flow control not error handling. In other words, data may still be lost, even if every write is acknowledged.

ByteString

To maintain isolation, actors should communicate with immutable objects only. `ByteString` is an immutable container for bytes. It is used by Akka's I/O system as an efficient, immutable alternative the traditional byte containers used for I/O on the JVM, such as `byte[]` and `ByteBuffer`.

`ByteString` is a *rope-like* data structure that is immutable and provides fast concatenation and slicing operations (perfect for I/O). When two `ByteStrings` are concatenated together they are both stored within the resulting `ByteString` instead of copying both to a new array. Operations such as `drop` and `take` return `ByteStrings` that still reference the original array, but just change the offset and length that is visible. Great care has also been taken to make sure that the internal array cannot be modified. Whenever a potentially unsafe array is used to create a new `ByteString` a defensive copy is created. If you require a `ByteString` that only blocks a much memory as necessary for its content, use the `compact` method to get a `CompactByteString` instance. If the `ByteString` represented only a slice of the original array, this will result in copying all bytes in that slice.

`ByteString` inherits all methods from `IndexedSeq`, and it also has some new ones. For more information, look up the `akka.util.ByteString` class and its companion object in the `ScalaDoc`.

`ByteString` also comes with its own optimized builder and iterator classes `ByteStringBuilder` and `ByteIterator` which provide extra features in addition to those of normal builders and iterators.

Compatibility with java.io

A `ByteStringBuilder` can be wrapped in a `java.io.OutputStream` via the `asOutputStream` method. Likewise, `ByteIterator` can be wrapped in a `java.io.InputStream` via `asInputStream`. Using these, `akka.io` applications can integrate legacy code based on `java.io` streams.

6.11.3 Architecture in-depth

For further details on the design and internal architecture see *I/O Layer Design*.

6.12 Using TCP

The code snippets through-out this section assume the following imports:

```
import java.net.InetSocketAddress;
import akka.actor.ActorRef;
import akka.actor.ActorSystem;
import akka.actor.Props;
import akka.actor.UntypedActor;
import akka.io.Tcp;
import akka.io.Tcp.Bound;
```

```
import akka.io.Tcp.CommandFailed;
import akka.io.Tcp.Connected;
import akka.io.Tcp.ConnectionClosed;
import akka.io.Tcp.Received;
import akka.io.TcpMessage;
import akka.japi.Procedure;
import akka.util.ByteString;
```

All of the Akka I/O APIs are accessed through manager objects. When using an I/O API, the first step is to acquire a reference to the appropriate manager. The code below shows how to acquire a reference to the `Tcp` manager.

```
final ActorRef tcpManager = Tcp.get(getContext().system()).manager();
```

The manager is an actor that handles the underlying low level I/O resources (selectors, channels) and instantiates workers for specific tasks, such as listening to incoming connections.

6.12.1 Connecting

```
public class Client extends UntypedActor {

    final InetSocketAddress remote;
    final ActorRef listener;

    public static Props props(InetSocketAddress remote, ActorRef listener) {
        return Props.create(Client.class, remote, listener);
    }

    public Client(InetSocketAddress remote, ActorRef listener) {
        this.remote = remote;
        this.listener = listener;

        final ActorRef tcp = Tcp.get(getContext().system()).manager();
        tcp.tell(TcpMessage.connect(remote), getSelf());
    }

    @Override
    public void onReceive(Object msg) throws Exception {
        if (msg instanceof CommandFailed) {
            listener.tell("failed", getSelf());
            getContext().stop(getSelf());
        } else if (msg instanceof Connected) {
            listener.tell(msg, getSelf());
            getSender().tell(TcpMessage.register(getSelf()), getSelf());
            getContext().become(connected(getSender()));
        }
    }

    private Procedure<Object> connected(final ActorRef connection) {
        return new Procedure<Object>() {
            @Override
            public void apply(Object msg) throws Exception {

                if (msg instanceof ByteString) {
                    connection.tell(TcpMessage.write((ByteString) msg), getSelf());
                } else if (msg instanceof CommandFailed) {
                    // OS kernel socket buffer was full
                } else if (msg instanceof Received) {
                    listener.tell(((Received) msg).data(), getSelf());
                }
            }
        };
    }
}
```

```

    } else if (msg.equals("close")) {
        connection.tell(TcpMessage.close(), getSelf());

    } else if (msg instanceof ConnectionClosed) {
        getContext().stop(getSelf());
    }
}
};
}
}
}

```

The first step of connecting to a remote address is sending a `Connect` message to the TCP manager; in addition to the simplest form shown above there is also the possibility to specify a local `InetSocketAddress` to bind to and a list of socket options to apply.

Note: The `SO_NODELAY` (`TCP_NODELAY` on Windows) socket option defaults to true in Akka, independently of the OS default settings. This setting disables Nagle's algorithm, considerably improving latency for most applications. This setting could be overridden by passing `SO.TcpNoDelay(false)` in the list of socket options of the `Connect` message.

The TCP manager will then reply either with a `CommandFailed` or it will spawn an internal actor representing the new connection. This new actor will then send a `Connected` message to the original sender of the `Connect` message.

In order to activate the new connection a `Register` message must be sent to the connection actor, informing that one about who shall receive data from the socket. Before this step is done the connection cannot be used, and there is an internal timeout after which the connection actor will shut itself down if no `Register` message is received.

The connection actor watches the registered handler and closes the connection when that one terminates, thereby cleaning up all internal resources associated with that connection.

The actor in the example above uses `become` to switch from unconnected to connected operation, demonstrating the commands and events which are observed in that state. For a discussion on `CommandFailed` see [Throttling Reads and Writes](#) below. `ConnectionClosed` is a trait, which marks the different connection close events. The last line handles all connection close events in the same way. It is possible to listen for more fine-grained connection close events, see [Closing Connections](#) below.

6.12.2 Accepting connections

```

public class Server extends UntypedActor {

    final ActorRef manager;

    public Server(ActorRef manager) {
        this.manager = manager;
    }

    public static Props props(ActorRef manager) {
        return Props.create(Server.class, manager);
    }

    @Override
    public void preStart() throws Exception {
        final ActorRef tcp = Tcp.get(getContext().system()).manager();
        tcp.tell(TcpMessage.bind(getSelf(),
            new InetSocketAddress("localhost", 0), 100), getSelf());
    }
}

```

```

@Override
public void onReceive(Object msg) throws Exception {
    if (msg instanceof Bound) {
        manager.tell(msg, getSelf());

    } else if (msg instanceof CommandFailed) {
        getContext().stop(getSelf());

    } else if (msg instanceof Connected) {
        final Connected conn = (Connected) msg;
        manager.tell(conn, getSelf());
        final ActorRef handler = getContext().actorOf(
            Props.create(SimplisticHandler.class));
        getSender().tell(TcpMessage.register(handler), getSelf());
    }
}
}
}

```

To create a TCP server and listen for inbound connections, a `Bind` command has to be sent to the TCP manager. This will instruct the TCP manager to listen for TCP connections on a particular `InetSocketAddress`; the port may be specified as 0 in order to bind to a random port.

The actor sending the `Bind` message will receive a `Bound` message signaling that the server is ready to accept incoming connections; this message also contains the `InetSocketAddress` to which the socket was actually bound (i.e. resolved IP address and correct port number).

From this point forward the process of handling connections is the same as for outgoing connections. The example demonstrates that handling the reads from a certain connection can be delegated to another actor by naming it as the handler when sending the `Register` message. Writes can be sent from any actor in the system to the connection actor (i.e. the actor which sent the `Connected` message). The simplistic handler is defined as:

```

public class SimplisticHandler extends UntypedActor {
    @Override
    public void onReceive(Object msg) throws Exception {
        if (msg instanceof Received) {
            final ByteString data = ((Received) msg).data();
            System.out.println(data);
            getSender().tell(TcpMessage.write(data), getSelf());
        } else if (msg instanceof ConnectionClosed) {
            getContext().stop(getSelf());
        }
    }
}
}

```

For a more complete sample which also takes into account the possibility of failures when sending please see [Throttling Reads and Writes](#) below.

The only difference to outgoing connections is that the internal actor managing the listen port—the sender of the `Bound` message—watches the actor which was named as the recipient for `Connected` messages in the `Bind` message. When that actor terminates the listen port will be closed and all resources associated with it will be released; existing connections will not be terminated at this point.

6.12.3 Closing connections

A connection can be closed by sending one of the commands `Close`, `ConfirmedClose` or `Abort` to the connection actor.

`Close` will close the connection by sending a FIN message, but without waiting for confirmation from the remote endpoint. Pending writes will be flushed. If the close is successful, the listener will be notified with `Closed`.

`ConfirmedClose` will close the sending direction of the connection by sending a FIN message, but data will continue to be received until the remote endpoint closes the connection, too. Pending writes will be flushed. If the

close is successful, the listener will be notified with `ConfirmedClosed`.

`Abort` will immediately terminate the connection by sending a RST message to the remote endpoint. Pending writes will be not flushed. If the close is successful, the listener will be notified with `Aborted`.

`PeerClosed` will be sent to the listener if the connection has been closed by the remote endpoint. Per default, the connection will then automatically be closed from this endpoint as well. To support half-closed connections set the `keepOpenOnPeerClosed` member of the `Register` message to `true` in which case the connection stays open until it receives one of the above close commands.

`ErrorClosed` will be sent to the listener whenever an error happened that forced the connection to be closed.

All close notifications are sub-types of `ConnectionClosed` so listeners who do not need fine-grained close events may handle all close events in the same way.

6.12.4 Writing to a connection

Once a connection has been established data can be sent to it from any actor in the form of a `Tcp.WriteCommand`. `Tcp.WriteCommand` is an abstract class with three concrete implementations:

Tcp.Write The simplest `WriteCommand` implementation which wraps a `ByteString` instance and an “ack” event. A `ByteString` (as explained in *this section*) models one or more chunks of immutable in-memory data with a maximum (total) size of 2 GB (2^{31} bytes).

Tcp.WriteFile If you want to send “raw” data from a file you can do so efficiently with the `Tcp.WriteFile` command. This allows you do designate a (contiguous) chunk of on-disk bytes for sending across the connection without the need to first load them into the JVM memory. As such `Tcp.WriteFile` can “hold” more than 2GB of data and an “ack” event if required.

Tcp.CompoundWrite Sometimes you might want to group (or interleave) several `Tcp.Write` and/or `Tcp.WriteFile` commands into one atomic write command which gets written to the connection in one go. The `Tcp.CompoundWrite` allows you to do just that and offers three benefits:

1. As explained in the following section the TCP connection actor can only handle one single write command at a time. By combining several writes into one `CompoundWrite` you can have them be sent across the connection with minimum overhead and without the need to spoon feed them to the connection actor via an *ACK-based* message protocol.
2. Because a `WriteCommand` is atomic you can be sure that no other actor can “inject” other writes into your series of writes if you combine them into one single `CompoundWrite`. In scenarios where several actors write to the same connection this can be an important feature which can be somewhat hard to achieve otherwise.
3. The “sub writes” of a `CompoundWrite` are regular `Write` or `WriteFile` commands that themselves can request “ack” events. These ACKs are sent out as soon as the respective “sub write” has been completed. This allows you to attach more than one ACK to a `Write` or `WriteFile` (by combining it with an empty write that itself requests an ACK) or to have the connection actor acknowledge the progress of transmitting the `CompoundWrite` by sending out intermediate ACKs at arbitrary points.

6.12.5 Throttling Reads and Writes

The basic model of the TCP connection actor is that it has no internal buffering (i.e. it can only process one write at a time, meaning it can buffer one write until it has been passed on to the O/S kernel in full). Congestion needs to be handled at the user level, for both writes and reads.

For back-pressuring writes there are three modes of operation

- *ACK-based*: every `Write` command carries an arbitrary object, and if this object is not `Tcp.NoAck` then it will be returned to the sender of the `Write` upon successfully writing all contained data to the socket. If no other write is initiated before having received this acknowledgement then no failures can happen due to buffer overrun.

- *NACK-based*: every write which arrives while a previous write is not yet completed will be replied to with a `CommandFailed` message containing the failed write. Just relying on this mechanism requires the implemented protocol to tolerate skipping writes (e.g. if each write is a valid message on its own and it is not required that all are delivered). This mode is enabled by setting the `useResumeWriting` flag to `false` within the `Register` message during connection activation.
- *NACK-based with write suspending*: this mode is very similar to the NACK-based one, but once a single write has failed no further writes will succeed until a `ResumeWriting` message is received. This message will be answered with a `WritingResumed` message once the last accepted write has completed. If the actor driving the connection implements buffering and resends the NACK'ed messages after having awaited the `WritingResumed` signal then every message is delivered exactly once to the network socket.

These write models (with the exception of the second which is rather specialised) are demonstrated in complete examples below. The full and contiguous source is available [on github](#).

For back-pressuring reads there are two modes of operation

- *Push-reading*: in this mode the connection actor sends the registered reader actor incoming data as soon as available as `Received` events. Whenever the reader actor wants to signal back-pressure to the remote TCP endpoint it can send a `SuspendReading` message to the connection actor to indicate that it wants to suspend the reception of new data. No `Received` events will arrive until a corresponding `ResumeReading` is sent indicating that the receiver actor is ready again.
- *Pull-reading*: after sending a `Received` event the connection actor automatically suspends accepting data from the socket until the reader actor signals with a `ResumeReading` message that it is ready to process more input data. Hence new data is “pulled” from the connection by sending `ResumeReading` messages.

Note: It should be obvious that all these flow control schemes only work between one writer/reader and one connection actor; as soon as multiple actors send write commands to a single connection no consistent result can be achieved.

6.12.6 ACK-Based Write Back-Pressure

For proper function of the following example it is important to configure the connection to remain half-open when the remote side closed its writing end: this allows the example `EchoHandler` to write all outstanding data back to the client before fully closing the connection. This is enabled using a flag upon connection activation (observe the `Register` message):

```
connection.tell(TcpMessage.register(handler,
    true, // <-- keepOpenOnPeerClosed flag
    true), getSelf());
```

With this preparation let us dive into the handler itself:

```
public class SimpleEchoHandler extends UntypedActor {

    final LoggingAdapter log = Logging
        .getLogger(getContext().system(), getSelf());

    final ActorRef connection;
    final InetSocketAddress remote;

    public static final long maxStored = 100000000;
    public static final long highWatermark = maxStored * 5 / 10;
    public static final long lowWatermark = maxStored * 2 / 10;

    public SimpleEchoHandler(ActorRef connection, InetSocketAddress remote) {
        this.connection = connection;
        this.remote = remote;

        // sign death pact: this actor stops when the connection is closed
```

```

    getContext().watch(connection);
}

@Override
public void onReceive(Object msg) throws Exception {
    if (msg instanceof Received) {
        final ByteString data = ((Received) msg).data();
        buffer(data);
        connection.tell(TcpMessage.write(data, ACK), getSelf());
        // now switch behavior to "waiting for acknowledgement"
        getContext().become(buffering, false);
    } else if (msg instanceof ConnectionClosed) {
        getContext().stop(getSelf());
    }
}

private final Procedure<Object> buffering = new Procedure<Object>() {
    @Override
    public void apply(Object msg) throws Exception {
        if (msg instanceof Received) {
            buffer(((Received) msg).data());
        } else if (msg == ACK) {
            acknowledge();
        } else if (msg instanceof ConnectionClosed) {
            if (((ConnectionClosed) msg).isPeerClosed()) {
                closing = true;
            } else {
                // could also be ErrorClosed, in which case we just give up
                getContext().stop(getSelf());
            }
        }
    }
};

// storage omitted ...
}

```

The principle is simple: when having written a chunk always wait for the Ack to come back before sending the next chunk. While waiting we switch behavior such that new incoming data are buffered. The helper functions used are a bit lengthy but not complicated:

```

protected void buffer(ByteString data) {
    storage.add(data);
    stored += data.size();

    if (stored > maxStored) {
        log.warning("drop connection to [{}] (buffer overrun)", remote);
        getContext().stop(getSelf());
    } else if (stored > highWatermark) {
        log.debug("suspending reading");
        connection.tell(TcpMessage.suspendReading(), getSelf());
        suspended = true;
    }
}

protected void acknowledge() {
    final ByteString acked = storage.remove();
    stored -= acked.size();
    transferred += acked.size();
}

```

```

if (suspended && stored < lowWatermark) {
    log.debug("resuming reading");
    connection.tell(TcpMessage.resumeReading(), getSelf());
    suspended = false;
}

if (storage.isEmpty()) {
    if (closing) {
        getContext().stop(getSelf());
    } else {
        getContext().unbecome();
    }
} else {
    connection.tell(TcpMessage.write(storage.peek(), ACK), getSelf());
}
}
}

```

The most interesting part is probably the last: an Ack removes the oldest data chunk from the buffer, and if that was the last chunk then we either close the connection (if the peer closed its half already) or return to the idle behavior; otherwise we just send the next buffered chunk and stay waiting for the next Ack.

Back-pressure can be propagated also across the reading side back to the writer on the other end of the connection by sending the `SuspendReading` command to the connection actor. This will lead to no data being read from the socket anymore (although this does happen after a delay because it takes some time until the connection actor processes this command, hence appropriate head-room in the buffer should be present), which in turn will lead to the O/S kernel buffer filling up on our end, then the TCP window mechanism will stop the remote side from writing, filling up its write buffer, until finally the writer on the other side cannot push any data into the socket anymore. This is how end-to-end back-pressure is realized across a TCP connection.

6.12.7 NACK-Based Write Back-Pressure with Suspending

```

public class EchoHandler extends UntypedActor {

    final LoggingAdapter log = Logging
        .getLogger(getContext().system(), getSelf());

    final ActorRef connection;
    final InetSocketAddress remote;

    public static final long MAX_STORED = 100000000;
    public static final long HIGH_WATERMARK = MAX_STORED * 5 / 10;
    public static final long LOW_WATERMARK = MAX_STORED * 2 / 10;

    private static class Ack implements Event {
        public final int ack;
        public Ack(int ack) {
            this.ack = ack;
        }
    }

    public EchoHandler(ActorRef connection, InetSocketAddress remote) {
        this.connection = connection;
        this.remote = remote;

        // sign death pact: this actor stops when the connection is closed
        getContext().watch(connection);

        // start out in optimistic write-through mode
        getContext().become(writing);
    }
}

```

```

private final Procedure<Object> writing = new Procedure<Object>() {
    @Override
    public void apply(Object msg) throws Exception {
        if (msg instanceof Received) {
            final ByteString data = ((Received) msg).data();
            connection.tell(TcpMessage.write(data, new Ack(currentOffset())), getSelf());
            buffer(data);

        } else if (msg instanceof Integer) {
            acknowledge((Integer) msg);

        } else if (msg instanceof CommandFailed) {
            final Write w = (Write) ((CommandFailed) msg).cmd();
            connection.tell(TcpMessage.resumeWriting(), getSelf());
            getContext().become(buffering((Ack) w.ack()));

        } else if (msg instanceof ConnectionClosed) {
            final ConnectionClosed cl = (ConnectionClosed) msg;
            if (cl.isPeerClosed()) {
                if (storage.isEmpty()) {
                    getContext().stop(getSelf());
                } else {
                    getContext().become(closing);
                }
            }
        }
    }
};

// buffering ...

// closing ...

// storage omitted ...
}

```

The principle here is to keep writing until a `CommandFailed` is received, using acknowledgements only to prune the resend buffer. When a such a failure was received, transition into a different state for handling and handle resending of all queued data:

```

protected Procedure<Object> buffering(final Ack nack) {
    return new Procedure<Object>() {

        private int toAck = 10;
        private boolean peerClosed = false;

        @Override
        public void apply(Object msg) throws Exception {
            if (msg instanceof Received) {
                buffer(((Received) msg).data());

            } else if (msg instanceof WritingResumed) {
                writeFirst();

            } else if (msg instanceof ConnectionClosed) {
                if (((ConnectionClosed) msg).isPeerClosed())
                    peerClosed = true;
                else
                    getContext().stop(getSelf());

            } else if (msg instanceof Integer) {
                final int ack = (Integer) msg;
                acknowledge(ack);
            }
        }
    };
}

```



```
};
```

Closing the connection while still sending all data is a bit more involved than in the ACK-based approach: the idea is to always send all outstanding messages and acknowledge all successful writes, and if a failure happens then switch behavior to await the `WritingResumed` event and start over.

The helper functions are very similar to the ACK-based case:

```
protected void buffer(ByteString data) {
    storage.add(data);
    stored += data.size();

    if (stored > MAX_STORED) {
        log.warning("drop connection to [{}] (buffer overrun)", remote);
        getContext().stop(getSelf());
    } else if (stored > HIGH_WATERMARK) {
        log.debug("suspending reading at {}", currentOffset());
        connection.tell(TcpMessage.suspendReading(), getSelf());
        suspended = true;
    }
}

protected void acknowledge(int ack) {
    assert ack == storageOffset();
    assert !storage.isEmpty();

    final ByteString acked = storage.remove();
    stored -= acked.size();
    transferred += acked.size();
    storageOffset += 1;

    if (suspended && stored < LOW_WATERMARK) {
        log.debug("resuming reading");
        connection.tell(TcpMessage.resumeReading(), getSelf());
        suspended = false;
    }
}
```

6.12.8 Read Back-Pressure with Pull Mode

When using push based reading, data coming from the socket is sent to the actor as soon as it is available. In the case of the previous Echo server example this meant that we needed to maintain a buffer of incoming data to keep it around since the rate of writing might be slower than the rate of the arrival of new data.

With the Pull mode this buffer can be completely eliminated as the following snippet demonstrates:

```
@Override
public void preStart() throws Exception {
    connection.tell(TcpMessage.resumeReading(), getSelf());
}

@Override
public void onReceive(Object message) throws Exception {
    if (message instanceof Tcp.Received) {
        ByteString data = ((Tcp.Received) message).data();
        connection.tell(TcpMessage.write(data, new Ack()), getSelf());
    } else if (message instanceof Ack) {
        connection.tell(TcpMessage.resumeReading(), getSelf());
    }
}
```

The idea here is that reading is not resumed until the previous write has been completely acknowledged by the connection actor. Every pull mode connection actor starts from suspended state. To start the flow of data we send a `ResumeReading` in the `preStart` method to tell the connection actor that we are ready to receive the first chunk of data. Since we only resume reading when the previous data chunk has been completely written there is no need for maintaining a buffer.

To enable pull reading on an outbound connection the `pullMode` parameter of the `Connect` should be set to `true`:

```
final List<Inet.SocketOption> options = new ArrayList<Inet.SocketOption>();
tcp.tell(
    TcpMessage.connect(new InetSocketAddress("localhost", 3000), null, options, null, true),
    getSelf()
);
```

Pull Mode Reading for Inbound Connections

The previous section demonstrated how to enable pull reading mode for outbound connections but it is possible to create a listener actor with this mode of reading by setting the `pullMode` parameter of the `Bind` command to `true`:

```
tcp = Tcp.get(getContext().system()).manager();
final List<Inet.SocketOption> options = new ArrayList<Inet.SocketOption>();
tcp.tell(
    TcpMessage.bind(getSelf(), new InetSocketAddress("localhost", 0), 100, options, true),
    getSelf()
);
```

One of the effects of this setting is that all connections accepted by this listener actor will use pull mode reading.

Another effect of this setting is that in addition of setting all inbound connections to pull mode, accepting connections becomes pull based, too. This means that after handling one (or more) `Connected` events the listener actor has to be resumed by sending it a `ResumeAccepting` message.

Listener actors with pull mode start suspended so to start accepting connections a `ResumeAccepting` command has to be sent to the listener actor after binding was successful:

```
public void onReceive(Object message) throws Exception {
    if (message instanceof Tcp.Bound) {
        listener = getSender();
        // Accept connections one by one
        listener.tell(TcpMessage.resumeAccepting(1), getSelf());
    } else if (message instanceof Tcp.Connected) {
        ActorRef handler = getContext().actorOf(Props.create(PullEcho.class, getSender()));
        getSender().tell(TcpMessage.register(handler), getSelf());
        // Resume accepting connections
        listener.tell(TcpMessage.resumeAccepting(1), getSelf());
    }
}
```

As shown in the example after handling an incoming connection we need to resume accepting again. The `ResumeAccepting` message accepts a `batchSize` parameter that specifies how many new connections are accepted before a next `ResumeAccepting` message is needed to resume handling of new connections.

6.13 Using UDP

UDP is a connectionless datagram protocol which offers two different ways of communication on the JDK level:

- sockets which are free to send datagrams to any destination and receive datagrams from any origin
- sockets which are restricted to communication with one specific remote socket address

In the low-level API the distinction is made—confusingly—by whether or not `connect` has been called on the socket (even when `connect` has been called the protocol is still connectionless). These two forms of UDP usage are offered using distinct IO extensions described below.

6.13.1 Unconnected UDP

Simple Send

```
public static class SimpleSender extends UntypedActor {
    final InetSocketAddress remote;

    public SimpleSender(InetSocketAddress remote) {
        this.remote = remote;

        // request creation of a SimpleSender
        final ActorRef mgr = Udp.get(getContext().system()).getManager();
        mgr.tell(UdpMessage.simpleSender(), getSelf());
    }

    @Override
    public void onReceive(Object msg) {
        if (msg instanceof Udp.SimpleSenderReady) {
            getContext().become(ready(getSender()));
        } else unhandled(msg);
    }

    private Procedure<Object> ready(final ActorRef send) {
        return new Procedure<Object>() {
            @Override
            public void apply(Object msg) throws Exception {
                if (msg instanceof String) {
                    final String str = (String) msg;
                    send.tell(UdpMessage.send(ByteString.fromString(str), remote), getSelf());
                } else unhandled(msg);
            }
        };
    }
}
```

The simplest form of UDP usage is to just send datagrams without the need of getting a reply. To this end a “simple sender” facility is provided as demonstrated above. The UDP extension is queried using the `simpleSender` message, which is answered by a `SimpleSenderReady` notification. The sender of this message is the newly created sender actor which from this point onward can be used to send datagrams to arbitrary destinations; in this example it will just send any UTF-8 encoded `String` it receives to a predefined remote address.

Note: The simple sender will not shut itself down because it cannot know when you are done with it. You will need to send it a `PoisonPill` when you want to close the ephemeral port the sender is bound to.

Bind (and Send)

```
public static class Listener extends UntypedActor {
    final ActorRef nextActor;

    public Listener(ActorRef nextActor) {
        this.nextActor = nextActor;

        // request creation of a bound listen socket
    }
}
```

```

final ActorRef mgr = Udp.get(getContext().system()).getManager();
mgr.tell(
    UdpMessage.bind(getSelf(), new InetSocketAddress("localhost", 0)),
    getSelf());
}

@Override
public void onReceive(Object msg) {
    if (msg instanceof Udp.Bound) {
        final Udp.Bound b = (Udp.Bound) msg;
        getContext().become(ready(getSender()));
    } else unhandled(msg);
}

private Procedure<Object> ready(final ActorRef socket) {
    return new Procedure<Object>() {
        @Override
        public void apply(Object msg) throws Exception {
            if (msg instanceof Udp.Received) {
                final Udp.Received r = (Udp.Received) msg;
                // echo server example: send back the data
                socket.tell(UdpMessage.send(r.data(), r.sender()), getSelf());
                // or do some processing and forward it on
                final Object processed = // parse data etc., e.g. using PipelineStage
                nextActor.tell(processed, getSelf());

            } else if (msg.equals(UdpMessage.unbind())) {
                socket.tell(msg, getSelf());
            } else if (msg instanceof Udp.Unbound) {
                getContext().stop(getSelf());
            } else unhandled(msg);
        }
    };
}
}

```

If you want to implement a UDP server which listens on a socket for incoming datagrams then you need to use the `bind` command as shown above. The local address specified may have a zero port in which case the operating system will automatically choose a free port and assign it to the new socket. Which port was actually bound can be found out by inspecting the `Bound` message.

The sender of the `Bound` message is the actor which manages the new socket. Sending datagrams is achieved by using the `send` message type and the socket can be closed by sending a `unbind` command, in which case the socket actor will reply with a `Unbound` notification.

Received datagrams are sent to the actor designated in the `bind` message, whereas the `Bound` message will be sent to the sender of the `bind`.

6.13.2 Connected UDP

The service provided by the connection based UDP API is similar to the `bind-and-send` service we saw earlier, but the main difference is that a connection is only able to send to the `remoteAddress` it was connected to, and will receive datagrams only from that address.

```

public static class Connected extends UntypedActor {
    final InetSocketAddress remote;

    public Connected(InetSocketAddress remote) {
        this.remote = remote;
    }
}

```

```

// create a restricted a.k.a. "connected" socket
final ActorRef mgr = UdpConnected.get(getContext().system()).getManager();
mgr.tell(UdpConnectedMessage.connect(getSelf(), remote), getSelf());
}

@Override
public void onReceive(Object msg) {
    if (msg instanceof UdpConnected.Connected) {
        getContext().become(ready(getSender()));
    } else unhandled(msg);
}

private Procedure<Object> ready(final ActorRef connection) {
    return new Procedure<Object>() {
        @Override
        public void apply(Object msg) throws Exception {
            if (msg instanceof UdpConnected.Received) {
                final UdpConnected.Received r = (UdpConnected.Received) msg;
                // process data, send it on, etc.
                // #connected
                if (r.data().utf8String().equals("hello")) {
                    connection.tell(
                        UdpConnectedMessage.send(ByteString.fromString("world"),
                            getSelf());
                }
                // #connected
            } else if (msg instanceof String) {
                final String str = (String) msg;
                connection
                    .tell(UdpConnectedMessage.send(ByteString.fromString(str)),
                        getSelf());
            } else if (msg.equals(UdpConnectedMessage.disconnect())) {
                connection.tell(msg, getSelf());
            } else if (msg instanceof UdpConnected.Disconnected) {
                getContext().stop(getSelf());
            } else unhandled(msg);
        }
    };
}
}

```

Consequently the example shown here looks quite similar to the previous one, the biggest difference is the absence of remote address information in `send` and `Received` messages.

Note: There is a small performance benefit in using connection based UDP API over the connectionless one. If there is a `SecurityManager` enabled on the system, every connectionless message send has to go through a security check, while in the case of connection-based UDP the security check is cached after connect, thus writes do not suffer an additional performance penalty.

6.13.3 UDP Multicast

If you want to use UDP multicast you will need to use Java 7. Akka provides a way to control various options of `DatagramChannel` through the `akka.io.Inet.SocketOption` interface. The example below shows how to setup a receiver of multicast messages using IPv6 protocol.

To select a Protocol Family you must extend `akka.io.Inet.DatagramChannelCreator` class which

implements `akka.io.Inet.SocketOption`. Provide custom logic for opening a datagram channel by overriding `create` method.

```
public static class Inet6ProtocolFamily extends Inet.DatagramChannelCreator {
    @Override
    public DatagramChannel create() throws Exception {
        return DatagramChannel.open(StandardProtocolFamily.INET6);
    }
}
```

Another socket option will be needed to join a multicast group.

```
public static class MulticastGroup extends Inet.AbstractSocketOptionV2 {
    private String address;
    private String interf;

    public MulticastGroup(String address, String interf) {
        this.address = address;
        this.interf = interf;
    }

    @Override
    public void afterBind(DatagramSocket s) {
        try {
            InetAddress group = InetAddress.getByName(address);
            NetworkInterface networkInterface = NetworkInterface.getByName(interf);
            s.getChannel().join(group, networkInterface);
        } catch (Exception ex) {
            System.out.println("Unable to join multicast group.");
        }
    }
}
```

Socket options must be provided to `UdpMessage.bind` command.

```
List<Inet.SocketOption> options = new ArrayList<>();
options.add(new Inet6ProtocolFamily());
options.add(new MulticastGroup(group, iface));

final ActorRef mgr = Udp.get(getContext().system()).getManager();
// listen for datagrams on this address
InetSocketAddress endpoint = new InetSocketAddress(port);
mgr.tell(UdpMessage.bind(getSelf(), endpoint, options), getSelf());
```

6.14 Camel

6.14.1 Introduction

The akka-camel module allows Untyped Actors to receive and send messages over a great variety of protocols and APIs. In addition to the native Scala and Java actor API, actors can now exchange messages with other systems over large number of protocols and APIs such as HTTP, SOAP, TCP, FTP, SMTP or JMS, to mention a few. At the moment, approximately 80 protocols and APIs are supported.

Apache Camel

The akka-camel module is based on [Apache Camel](#), a powerful and light-weight integration framework for the JVM. For an introduction to Apache Camel you may want to read this [Apache Camel article](#). Camel comes with a large number of [components](#) that provide bindings to different protocols and APIs. The [camel-extra](#) project provides further components.

Consumer

Here's an example of using Camel's integration components in Akka.

```
import akka.camel.CamelMessage;
import akka.camel.javaapi.UntypedConsumerActor;

public class MyEndpoint extends UntypedConsumerActor{
    private String uri;

    public String getEndpointUri() {
        return uri;
    }

    public void onReceive(Object message) throws Exception {
        if (message instanceof CamelMessage) {
            /* ... */
        } else
            unhandled(message);
    }

    // Extra constructor to change the default uri,
    // for instance to "jetty:http://localhost:8877/example"
    public MyEndpoint(String uri) {
        this.uri = uri;
    }

    public MyEndpoint() {
        this.uri = "mina2:tcp://localhost:6200?textline=true";
    }
}
```

The above example exposes an actor over a TCP endpoint via Apache Camel's [Mina component](#). The actor implements the `getEndpointUri` method to define an endpoint from which it can receive messages. After starting the actor, TCP clients can immediately send messages to and receive responses from that actor. If the message exchange should go over HTTP (via Camel's [Jetty component](#)), the actor's `getEndpointUri` method should return a different URI, for instance "jetty:http://localhost:8877/example". In the above case an extra constructor is added that can set the endpoint URI, which would result in the `getEndpointUri` returning the URI that was set using this constructor.

Producer

Actors can also trigger message exchanges with external systems i.e. produce to Camel endpoints.

```
import akka.camel.javaapi.UntypedProducerActor;

public class Orders extends UntypedProducerActor {
    public String getEndpointUri() {
        return "jms:queue:Orders";
    }
}
```

In the above example, any message sent to this actor will be sent to the JMS queue `Orders`. Producer actors may choose from the same set of Camel components as Consumer actors do. Below an example of how to send a message to the `Orders` producer.

```
ActorSystem system = ActorSystem.create("some-system");
Props props = Props.create(Orders.class);
ActorRef producer = system.actorOf(props, "jmsproducer");
producer.tell("<order amount=\"100\" currency=\"PLN\" itemId=\"12345\"/>",
    ActorRef.noSender());
```

CamelMessage

The number of Camel components is constantly increasing. The akka-camel module can support these in a plug-and-play manner. Just add them to your application's classpath, define a component-specific endpoint URI and use it to exchange messages over the component-specific protocols or APIs. This is possible because Camel components bind protocol-specific message formats to a Camel-specific [normalized message format](#). The normalized message format hides protocol-specific details from Akka and makes it therefore very easy to support a large number of protocols through a uniform Camel component interface. The akka-camel module further converts mutable Camel messages into immutable representations which are used by Consumer and Producer actors for pattern matching, transformation, serialization or storage. In the above example of the Orders Producer, the XML message is put in the body of a newly created Camel Message with an empty set of headers. You can also create a CamelMessage yourself with the appropriate body and headers as you see fit.

CamelExtension

The akka-camel module is implemented as an Akka Extension, the CamelExtension object. Extensions will only be loaded once per ActorSystem, which will be managed by Akka. The CamelExtension object provides access to the Camel interface. The Camel interface in turn provides access to two important Apache Camel objects, the CamelContext and the ProducerTemplate. Below you can see how you can get access to these Apache Camel objects.

```
ActorSystem system = ActorSystem.create("some-system");
Camel camel = CamelExtension.get(system);
CamelContext camelContext = camel.context();
ProducerTemplate producerTemplate = camel.template();
```

One CamelExtension is only loaded once for every one ActorSystem, which makes it safe to call the CamelExtension at any point in your code to get to the Apache Camel objects associated with it. There is one CamelContext and one ProducerTemplate for every one ActorSystem that uses a CamelExtension. By Default, a new CamelContext is created when the CamelExtension starts. If you want to inject your own context instead, you can implement the ContextProvider interface and add the FQCN of your implementation in the config, as the value of the "akka.camel.context-provider". This interface define a single method getContext() used to load the CamelContext.

Below an example on how to add the ActiveMQ component to the CamelContext, which is required when you would like to use the ActiveMQ component.

```
ActorSystem system = ActorSystem.create("some-system");
Camel camel = CamelExtension.get(system);
CamelContext camelContext = camel.context();
// camelContext.addComponent("activemq", ActiveMQComponent.activemqComponent(
//     "vm://localhost?broker.persistent=false"));
```

The CamelContext joins the lifecycle of the ActorSystem and CamelExtension it is associated with; the CamelContext is started when the CamelExtension is created, and it is shut down when the associated ActorSystem is shut down. The same is true for the ProducerTemplate.

The CamelExtension is used by both *Producer* and *Consumer* actors to interact with Apache Camel internally. You can access the CamelExtension inside a *Producer* or a *Consumer* using the camel method, or get straight at the CamelContext using the getCamelContext method or to the ProducerTemplate using the getProducerTemplate method. Actors are created and started asynchronously. When a *Consumer* actor is created, the *Consumer* is published at its Camel endpoint (more precisely, the route is added to the CamelContext from the Endpoint to the actor). When a *Producer* actor is created, a SendProcessor and Endpoint are created so that the Producer can send messages to it. Publication is done asynchronously; setting up an endpoint may still be in progress after you have requested the actor to be created. Some Camel components can take a while to startup, and in some cases you might want to know when the endpoints are activated and ready to be used. The Camel interface allows you to find out when the endpoint is activated or deactivated.

```
import akka.actor.ActorRef;
import akka.actor.ActorSystem;
import akka.actor.Props;
```

```

import akka.camel.Camel;
import akka.camel.CamelExtension;
import akka.camel.javaapi.UntypedConsumerActor;
import akka.testkit.JavaTestKit;
import akka.testkit.TestKit;
import akka.util.Timeout;
import scala.concurrent.Future;
import scala.concurrent.duration.Duration;
import static java.util.concurrent.TimeUnit.SECONDS;

// ..
ActorSystem system = ActorSystem.create("some-system");
Props props = Props.create(MyConsumer.class);
ActorRef producer = system.actorOf(props, "myproducer");
Camel camel = CamelExtension.get(system);
// get a future reference to the activation of the endpoint of the Consumer Actor
Timeout timeout = new Timeout(Duration.create(10, SECONDS));
Future<ActorRef> activationFuture = camel.activationFutureFor(producer,
    timeout, system.dispatcher());

```

The above code shows that you can get a `Future` to the activation of the route from the endpoint to the actor, or you can wait in a blocking fashion on the activation of the route. An `ActivationTimeoutException` is thrown if the endpoint could not be activated within the specified timeout. Deactivation works in a similar fashion:

```

// ..
system.stop(producer);
// get a future reference to the deactivation of the endpoint of the Consumer Actor
Future<ActorRef> deactivationFuture = camel.deactivationFutureFor(producer,
    timeout, system.dispatcher());

```

Deactivation of a Consumer or a Producer actor happens when the actor is terminated. For a Consumer, the route to the actor is stopped. For a Producer, the `SendProcessor` is stopped. A `DeActivationTimeoutException` is thrown if the associated camel objects could not be deactivated within the specified timeout.

6.14.2 Consumer Actors

For objects to receive messages, they must inherit from the `UntypedConsumerActor` class. For example, the following actor class (`Consumer1`) implements the `getEndpointUri` method, which is declared in the `UntypedConsumerActor` class, in order to receive messages from the `file:data/input/actor` Camel endpoint.

```

import akka.camel.CamelMessage;
import akka.camel.javaapi.UntypedConsumerActor;
import akka.event.Logging;
import akka.event.LoggingAdapter;

public class Consumer1 extends UntypedConsumerActor {
    LoggingAdapter log = Logging.getLogger(getContext().system(), this);

    public String getEndpointUri() {
        return "file:data/input/actor";
    }

    public void onReceive(Object message) {
        if (message instanceof CamelMessage) {
            CamelMessage camelMessage = (CamelMessage) message;
            String body = camelMessage.getBodyAs(String.class, getCamelContext());
            log.info("Received message: {}", body);
        } else {
            unhandled(message);
        }
    }
}

```

Whenever a file is put into the `data/input/actor` directory, its content is picked up by the Camel [file component](#) and sent as message to the actor. Messages consumed by actors from Camel endpoints are of type `CamelMessage`. These are immutable representations of Camel messages.

Here's another example that sets the `endpointUri` to `jetty:http://localhost:8877/camel/default`. It causes Camel's [Jetty component](#) to start an embedded [Jetty server](#), accepting HTTP connections from localhost on port 8877.

```
import akka.camel.CamelMessage;
import akka.camel.javaapi.UntypedConsumerActor;

public class Consumer2 extends UntypedConsumerActor {
    public String getEndpointUri() {
        return "jetty:http://localhost:8877/camel/default";
    }

    public void onReceive(Object message) {
        if (message instanceof CamelMessage) {
            CamelMessage camelMessage = (CamelMessage) message;
            String body = camelMessage.getBodyAs(String.class, getCamelContext());
            getSender().tell(String.format("Received message: %s", body), getSelf());
        } else
            unhandled(message);
    }
}
```

After starting the actor, clients can send messages to that actor by POSTing to `http://localhost:8877/camel/default`. The actor sends a response by using the `getSender().tell` method. For returning a message body and headers to the HTTP client the response type should be `CamelMessage`. For any other response type, a new `CamelMessage` object is created by akka-camel with the actor response as message body.

Delivery acknowledgements

With in-out message exchanges, clients usually know that a message exchange is done when they receive a reply from a consumer actor. The reply message can be a `CamelMessage` (or any object which is then internally converted to a `CamelMessage`) on success, and a `Failure` message on failure.

With in-only message exchanges, by default, an exchange is done when a message is added to the consumer actor's mailbox. Any failure or exception that occurs during processing of that message by the consumer actor cannot be reported back to the endpoint in this case. To allow consumer actors to positively or negatively acknowledge the receipt of a message from an in-only message exchange, they need to override the `autoAck` method to return `false`. In this case, consumer actors must reply either with a special `akka.camel.Ack` message (positive acknowledgement) or a `akka.actor.Status.Failure` (negative acknowledgement).

```
import akka.actor.Status;
import akka.camel.Ack;
import akka.camel.CamelMessage;
import akka.camel.javaapi.UntypedConsumerActor;

public class Consumer3 extends UntypedConsumerActor{

    @Override
    public boolean autoAck() {
        return false;
    }

    public String getEndpointUri() {
        return "jms:queue:test";
    }

    public void onReceive(Object message) {
```

```

if (message instanceof CamelMessage) {
    getSender().tell(Ack.getInstance(), getSelf());
    // on success
    // ..
    Exception someException = new Exception("e1");
    // on failure
    getSender().tell(new Status.Failure(someException), getSelf());
} else
    unhandled(message);
}
}

```

Consumer timeout

Camel Exchanges (and their corresponding endpoints) that support two-way communications need to wait for a response from an actor before returning it to the initiating client. For some endpoint types, timeout values can be defined in an endpoint-specific way which is described in the documentation of the individual [Camel components](#). Another option is to configure timeouts on the level of consumer actors.

Two-way communications between a Camel endpoint and an actor are initiated by sending the request message to the actor with the `ask` pattern and the actor replies to the endpoint when the response is ready. The ask request to the actor can timeout, which will result in the [Exchange](#) failing with a `TimeoutException` set on the failure of the [Exchange](#). The timeout on the consumer actor can be overridden with the `replyTimeout`, as shown below.

```

import akka.camel.CamelMessage;
import akka.camel.javaapi.UntypedConsumerActor;
import scala.concurrent.duration.Duration;
import scala.concurrent.duration.FiniteDuration;

import java.util.concurrent.TimeUnit;

public class Consumer4 extends UntypedConsumerActor {
    private final static FiniteDuration timeout =
        Duration.create(500, TimeUnit.MILLISECONDS);

    @Override
    public FiniteDuration replyTimeout() {
        return timeout;
    }

    public String getEndpointUri() {
        return "jetty:http://localhost:8877/camel/default";
    }

    public void onReceive(Object message) {
        if (message instanceof CamelMessage) {
            CamelMessage camelMessage = (CamelMessage) message;
            String body = camelMessage.getBodyAs(String.class, getCamelContext());
            getSender().tell(String.format("Hello %s", body), getSelf());
        } else
            unhandled(message);
    }
}

```

6.14.3 Producer Actors

For sending messages to Camel endpoints, actors need to inherit from the [UntypedProducerActor](#) class and implement the `getEndpointUri` method.

```
import akka.camel.javaapi.UntypedProducerActor;

public class Producer1 extends UntypedProducerActor {
    public String getEndpointUri() {
        return "http://localhost:8080/news";
    }
}
```

Producer1 inherits a default implementation of the `onReceive` method from the `UntypedProducerActor` class. To customize a producer actor's default behavior you must override the `UntypedProducerActor.onTransformResponse` and `UntypedProducerActor.onTransformOutgoingMessage` methods. This is explained later in more detail. Producer Actors cannot override the `UntypedProducerActor.onReceive` method.

Any message sent to a Producer actor will be sent to the associated Camel endpoint, in the above example to `http://localhost:8080/news`. The `UntypedProducerActor` always sends messages asynchronously. Response messages (if supported by the configured endpoint) will, by default, be returned to the original sender. The following example uses the `ask` pattern to send a message to a Producer actor and waits for a response.

```
ActorSystem system = ActorSystem.create("some-system");
Props props = Props.create(FirstProducer.class);
ActorRef producer = system.actorOf(props, "myproducer");
Future<Object> future = Patterns.ask(producer, "some request", 1000);
```

The future contains the response `CamelMessage`, or an `AkkaCamelException` when an error occurred, which contains the headers of the response.

Custom Processing

Instead of replying to the initial sender, producer actors can implement custom response processing by overriding the `onRouteResponse` method. In the following example, the response message is forwarded to a target actor instead of being replied to the original sender.

```
import akka.actor.UntypedActor;
import akka.camel.CamelMessage;

public class ResponseReceiver extends UntypedActor{
    public void onReceive(Object message) {
        if(message instanceof CamelMessage) {
            // do something with the forwarded response
        }
    }
}
```

```
import akka.actor.ActorRef;
import akka.camel.javaapi.UntypedProducerActor;

public class Forwarder extends UntypedProducerActor {
    private String uri;
    private ActorRef target;

    public Forwarder(String uri, ActorRef target) {
        this.uri = uri;
        this.target = target;
    }

    public String getEndpointUri() {
        return uri;
    }

    @Override
    public void onRouteResponse(Object message) {
        target.forward(message, getContext());
    }
}
```

```

}
}

ActorSystem system = ActorSystem.create("some-system");
Props receiverProps = Props.create(ResponseReceiver.class);
final ActorRef receiver = system.actorOf(receiverProps, "responseReceiver");
ActorRef forwardResponse = system.actorOf(Props.create(
    Forwarder.class, "http://localhost:8080/news/akka", receiver));
// the Forwarder sends out a request to the web page and forwards the response to
// the ResponseReceiver
forwardResponse.tell("some request", ActorRef.noSender());

```

Before producing messages to endpoints, producer actors can pre-process them by overriding the `UntypedProducerActor.onTransformOutgoingMessage` method.

```

import akka.camel.CamelMessage;
import akka.camel.javaapi.UntypedProducerActor;
import akka.dispatch.Mapper;
import akka.japi.Function;

public class Transformer extends UntypedProducerActor{
    private String uri;

    public Transformer(String uri) {
        this.uri = uri;
    }

    public String getEndpointUri() {
        return uri;
    }

    private CamelMessage upperCase(CamelMessage msg) {
        return msg.mapBody(new Mapper<String, String>() {
            @Override
            public String apply(String body) {
                return body.toUpperCase();
            }
        });
    }

    @Override
    public Object onTransformOutgoingMessage(Object message) {
        if(message instanceof CamelMessage) {
            CamelMessage camelMessage = (CamelMessage) message;
            return upperCase(camelMessage);
        } else {
            return message;
        }
    }
}

```

Producer configuration options

The interaction of producer actors with Camel endpoints can be configured to be one-way or two-way (by initiating in-only or in-out message exchanges, respectively). By default, the producer initiates an in-out message exchange with the endpoint. For initiating an in-only exchange, producer actors have to override the `isOneway` method to return `true`.

```

import akka.camel.javaapi.UntypedProducerActor;

public class OnewaySender extends UntypedProducerActor{
    private String uri;

```

```

public OnewaySender(String uri) {
    this.uri = uri;
}
public String getEndpointUri() {
    return uri;
}

@Override
public boolean isOneway() {
    return true;
}
}

```

Message correlation

To correlate request with response messages, applications can set the *Message.MessageExchangeId* message header.

```

ActorSystem system = ActorSystem.create("some-system");
Props props = Props.create(Orders.class);
ActorRef producer = system.actorOf(props, "jmsproducer");
Map<String, Object> headers = new HashMap<String, Object>();
headers.put(CamelMessage.MessageExchangeId(), "123");
producer.tell(new CamelMessage("<order amount=\"100\" currency=\"PLN\" \" +
    "itemId=\"12345\"/>", headers), ActorRef.noSender());

```

ProducerTemplate

The *UntypedProducerActor* class is a very convenient way for actors to produce messages to Camel endpoints. Actors may also use a Camel *ProducerTemplate* for producing messages to endpoints.

```

import akka.actor.UntypedActor;
import akka.camel.Camel;
import akka.camel.CamelExtension;
import org.apache.camel.ProducerTemplate;

public class MyActor extends UntypedActor {
    public void onReceive(Object message) {
        Camel camel = CamelExtension.get(getContext().system());
        ProducerTemplate template = camel.template();
        template.sendBody("direct:news", message);
    }
}

```

For initiating a a two-way message exchange, one of the *ProducerTemplate.request** methods must be used.

```

import akka.actor.UntypedActor;
import akka.camel.Camel;
import akka.camel.CamelExtension;
import org.apache.camel.ProducerTemplate;

public class RequestBodyActor extends UntypedActor {
    public void onReceive(Object message) {
        Camel camel = CamelExtension.get(getContext().system());
        ProducerTemplate template = camel.template();
        getSender().tell(template.requestBody("direct:news", message), getSelf());
    }
}

```

6.14.4 Asynchronous routing

In-out message exchanges between endpoints and actors are designed to be asynchronous. This is the case for both, consumer and producer actors.

- A consumer endpoint sends request messages to its consumer actor using the `tell` method and the actor returns responses with `getSender().tell` once they are ready.
- A producer actor sends request messages to its endpoint using Camel's asynchronous routing engine. Asynchronous responses are wrapped and added to the producer actor's mailbox for later processing. By default, response messages are returned to the initial sender but this can be overridden by Producer implementations (see also description of the `onRouteResponse` method in *Custom Processing*).

However, asynchronous two-way message exchanges, without allocating a thread for the full duration of exchange, cannot be generically supported by Camel's asynchronous routing engine alone. This must be supported by the individual [Camel components](#) (from which endpoints are created) as well. They must be able to suspend any work started for request processing (thereby freeing threads to do other work) and resume processing when the response is ready. This is currently the case for a [subset of components](#) such as the [Jetty component](#). All other Camel components can still be used, of course, but they will cause allocation of a thread for the duration of an in-out message exchange. There's also [Examples](#) that implements both, an asynchronous consumer and an asynchronous producer, with the [jetty component](#).

If the used Camel component is blocking it might be necessary to use a separate *dispatcher* for the producer. The Camel processor is invoked by a child actor of the producer and the dispatcher can be defined in the deployment section of the configuration. For example, if your producer actor has path `/user/integration/output` the dispatcher of the child actor can be defined with:

```
akka.actor.deployment {
  /integration/output/* {
    dispatcher = my-dispatcher
  }
}
```

6.14.5 Custom Camel routes

In all the examples so far, routes to consumer actors have been automatically constructed by akka-camel, when the actor was started. Although the default route construction templates, used by akka-camel internally, are sufficient for most use cases, some applications may require more specialized routes to actors. The akka-camel module provides two mechanisms for customizing routes to actors, which will be explained in this section. These are:

- Usage of *Akka Camel components* to access actors. Any Camel route can use these components to access Akka actors.
- *Intercepting route construction* to actors. This option gives you the ability to change routes that have already been added to Camel. Consumer actors have a hook into the route definition process which can be used to change the route.

Akka Camel components

Akka actors can be accessed from Camel routes using the `actor` Camel component. This component can be used to access any Akka actor (not only consumer actors) from Camel routes, as described in the following sections.

Access to actors

To access actors from custom Camel routes, the `actor` Camel component should be used. It fully supports Camel's [asynchronous routing engine](#).

This component accepts the following endpoint URI format:

- `[<actor-path>]?<options>`

where `<actor-path>` is the `ActorPath` to the actor. The `<options>` are name-value pairs separated by `&` (i.e. `name1=value1&name2=value2&...`).

URI options

The following URI options are supported:

Name	Type	Default	Description
<code>replyTimeout</code>	Duration	false	The reply timeout, specified in the same way that you use the duration in akka, for instance <code>10 seconds</code> except that in the url it is handy to use a <code>+</code> between the amount and the unit, like for example <code>200+millis</code> . See also Consumer timeout .
<code>autoAck</code>	Boolean	true	If set to true, in-only message exchanges are auto-acknowledged when the message is added to the actor's mailbox. If set to false, actors must acknowledge the receipt of the message. See also Delivery acknowledgements .

Here's an actor endpoint URI example containing an actor path:

```
akka://some-system/user/myconsumer?autoAck=false&replyTimeout=100+millis
```

In the following example, a custom route to an actor is created, using the actor's path.

```
import akka.actor.UntypedActor;
import akka.camel.CamelMessage;
import akka.dispatch.Mapper;
import akka.japi.Function;

public class Responder extends UntypedActor{

    public void onReceive(Object message) {
        if (message instanceof CamelMessage) {
            CamelMessage camelMessage = (CamelMessage) message;
            getSender().tell(createResponse(camelMessage), getSelf());
        } else
            unhandled(message);
    }

    private CamelMessage createResponse(CamelMessage msg) {
        return msg.mapBody(new Mapper<String,String>() {
            @Override
            public String apply(String body) {
                return String.format("received %s", body);
            }
        });
    }
}
```

```
import akka.actor.ActorRef;
import akka.camel.internal.component.CamelPath;
import org.apache.camel.builder.RouteBuilder;

public class CustomRouteBuilder extends RouteBuilder{
    private String uri;

    public CustomRouteBuilder(ActorRef responder) {
        uri = CamelPath.toUri(responder);
    }

    public void configure() throws Exception {
        from("jetty:http://localhost:8877/camel/custom").to(uri);
    }
}
```

```

    }
}

ActorSystem system = ActorSystem.create("some-system");
Camel camel = CamelExtension.get(system);
ActorRef responder = system.actorOf(Props.create(Responder.class), "TestResponder");
camel.context().addRoutes(new CustomRouteBuilder(responder));

```

The `CamelPath.toCamelUri` converts the `ActorRef` to the Camel actor component URI format which points to the actor endpoint as described above. When a message is received on the jetty endpoint, it is routed to the Responder actor, which in return replies back to the client of the HTTP request.

Intercepting route construction

The previous section, *Akka Camel components*, explained how to setup a route to an actor manually. It was the application's responsibility to define the route and add it to the current `CamelContext`. This section explains a more convenient way to define custom routes: akka-camel is still setting up the routes to consumer actors (and adds these routes to the current `CamelContext`) but applications can define extensions to these routes. Extensions can be defined with Camel's [Java DSL](#) or [Scala DSL](#). For example, an extension could be a custom error handler that redelivers messages from an endpoint to an actor's bounded mailbox when the mailbox was full.

The following examples demonstrate how to extend a route to a consumer actor for handling exceptions thrown by that actor.

```

import akka.actor.Status;
import akka.camel.CamelMessage;
import akka.camel.javaapi.UntypedConsumerActor;
import akka.dispatch.Mapper;
import org.apache.camel.builder.Builder;
import org.apache.camel.model.ProcessorDefinition;
import org.apache.camel.model.RouteDefinition;
import scala.Option;

public class ErrorThrowingConsumer extends UntypedConsumerActor{
    private String uri;

    private static Mapper<RouteDefinition, ProcessorDefinition<?>> mapper =
        new Mapper<RouteDefinition, ProcessorDefinition<?>>() {
            public ProcessorDefinition<?> apply(RouteDefinition rd) {
                // Catch any exception and handle it by returning the exception message
                // as response
                return rd.onException(Exception.class).handled(true).
                    transform(Builder.exceptionMessage()).end();
            }
        };

    public ErrorThrowingConsumer(String uri){
        this.uri = uri;
    }

    public String getEndpointUri() {
        return uri;
    }

    public void onReceive(Object message) throws Exception{
        if (message instanceof CamelMessage) {
            CamelMessage camelMessage = (CamelMessage) message;
            String body = camelMessage.getBodyAs(String.class, getCamelContext());
            throw new Exception(String.format("error: %s",body));
        } else
            unhandled(message);
    }
}

```

```

@Override
public Mapper<RouteDefinition,
    ProcessorDefinition<?>> getRouteDefinitionHandler() {
    return mapper;
}

@Override
public void preRestart(Throwable reason, Option<Object> message) {
    getSender().tell(new Status.Failure(reason), getSelf());
}
}

```

The above `ErrorThrowingConsumer` sends the `Failure` back to the sender in `preRestart` because the `Exception` that is thrown in the actor would otherwise just crash the actor, by default the actor would be restarted, and the response would never reach the client of the `Consumer`.

The akka-camel module creates a `RouteDefinition` instance by calling `from(endpointUri)` on a Camel `RouteBuilder` (where `endpointUri` is the endpoint URI of the consumer actor) and passes that instance as argument to the route definition handler `*`). The route definition handler then extends the route and returns a `ProcessorDefinition` (in the above example, the `ProcessorDefinition` returned by the `end` method. See the `org.apache.camel.model` package for details). After executing the route definition handler, akka-camel finally calls `a.to(targetActorUri)` on the returned `ProcessorDefinition` to complete the route to the consumer actor (where `targetActorUri` is the actor component URI as described in [Access to actors](#)). If the actor cannot be found, a `ActorNotRegisteredException` is thrown.

`*`) Before passing the `RouteDefinition` instance to the route definition handler, akka-camel may make some further modifications to it.

6.14.6 Examples

The [Typesafe Activator](#) tutorial named [Akka Camel Samples with Java](#) contains 3 samples:

- Asynchronous routing and transformation - This example demonstrates how to implement consumer and producer actors that support *Asynchronous routing* with their Camel endpoints.
- Custom Camel route - Demonstrates the combined usage of a `Producer` and a `Consumer` actor as well as the inclusion of a custom Camel route.
- Quartz Scheduler Example - Showing how simple is to implement a cron-style scheduler by using the Camel Quartz component

6.14.7 Configuration

There are several configuration properties for the Camel module, please refer to the [reference configuration](#).

6.14.8 Additional Resources

For an introduction to akka-camel 2, see also the Peter Gabryanczyk's talk [Migrating akka-camel module to Akka 2.x](#).

For an introduction to akka-camel 1, see also the [Appendix E - Akka and Camel](#) (pdf) of the book [Camel in Action](#).

Other, more advanced external articles (for version 1) are:

- [Akka Consumer Actors: New Features and Best Practices](#)
- [Akka Producer Actors: New Features and Best Practices](#)

UTILITIES

7.1 Event Bus

Originally conceived as a way to send messages to groups of actors, the `EventBus` has been generalized into a set of abstract base classes implementing a simple interface:

```
/**
 * Attempts to register the subscriber to the specified Classifier
 * @return true if successful and false if not (because it was already
 *     subscribed to that Classifier, or otherwise)
 */
public boolean subscribe(Subscriber subscriber, Classifier to);

/**
 * Attempts to deregister the subscriber from the specified Classifier
 * @return true if successful and false if not (because it wasn't subscribed
 *     to that Classifier, or otherwise)
 */
public boolean unsubscribe(Subscriber subscriber, Classifier from);

/**
 * Attempts to deregister the subscriber from all Classifiers it may be subscribed to
 */
public void unsubscribe(Subscriber subscriber);

/**
 * Publishes the specified Event to this bus
 */
public void publish(Event event);
```

Note: Please note that the `EventBus` does not preserve the sender of the published messages. If you need a reference to the original sender you have to provide it inside the message.

This mechanism is used in different places within Akka, e.g. the [Event Stream](#). Implementations can make use of the specific building blocks presented below.

An event bus must define the following three type parameters:

- `Event (E)` is the type of all events published on that bus
- `Subscriber (S)` is the type of subscribers allowed to register on that event bus
- `Classifier (C)` defines the classifier to be used in selecting subscribers for dispatching events

The traits below are still generic in these types, but they need to be defined for any concrete implementation.

7.1.1 Classifiers

The classifiers presented here are part of the Akka distribution, but rolling your own in case you do not find a perfect match is not difficult, check the implementation of the existing ones on [github](#)

Lookup Classification

The simplest classification is just to extract an arbitrary classifier from each event and maintaining a set of subscribers for each possible classifier. This can be compared to tuning in on a radio station. The trait `LookupClassification` is still generic in that it abstracts over how to compare subscribers and how exactly to classify.

The necessary methods to be implemented are illustrated with the following example:

```
import akka.event.japi.LookupEventBus;
import java.util.concurrent.TimeUnit;

public class MsgEnvelope {
    public final String topic;
    public final Object payload;

    public MsgEnvelope(String topic, Object payload) {
        this.topic = topic;
        this.payload = payload;
    }
}

/**
 * Publishes the payload of the MsgEnvelope when the topic of the
 * MsgEnvelope equals the String specified when subscribing.
 */
public class LookupBusImpl extends LookupEventBus<MsgEnvelope, ActorRef, String> {

    // is used for extracting the classifier from the incoming events
    @Override public String classify(MsgEnvelope event) {
        return event.topic;
    }

    // will be invoked for each event for all subscribers which registered themselves
    // for the event's classifier
    @Override public void publish(MsgEnvelope event, ActorRef subscriber) {
        subscriber.tell(event.payload, ActorRef.noSender());
    }

    // must define a full order over the subscribers, expressed as expected from
    // `java.lang.Comparable.compare`
    @Override public int compareSubscribers(ActorRef a, ActorRef b) {
        return a.compareTo(b);
    }

    // determines the initial size of the index data structure
    // used internally (i.e. the expected number of different classifiers)
    @Override public int mapSize() {
        return 128;
    }
}
```

A test for this implementation may look like this:

```
LookupBusImpl lookupBus = new LookupBusImpl();
lookupBus.subscribe(getTestActor(), "greetings");
```

```
lookupBus.publish(new MsgEnvelope("time", System.currentTimeMillis()));
lookupBus.publish(new MsgEnvelope("greetings", "hello"));
expectMsgEquals("hello");
```

This classifier is efficient in case no subscribers exist for a particular event.

Subchannel Classification

If classifiers form a hierarchy and it is desired that subscription be possible not only at the leaf nodes, this classification may be just the right one. It can be compared to tuning in on (possibly multiple) radio channels by genre. This classification has been developed for the case where the classifier is just the JVM class of the event and subscribers may be interested in subscribing to all subclasses of a certain class, but it may be used with any classifier hierarchy.

The necessary methods to be implemented are illustrated with the following example:

```
import akka.event.japi.SubchannelEventBus;
import akka.util.Subclassification;

public class StartsWithSubclassification implements Subclassification<String> {
    @Override public boolean isEqual(String x, String y) {
        return x.equals(y);
    }

    @Override public boolean isSubclass(String x, String y) {
        return x.startsWith(y);
    }
}

/**
 * Publishes the payload of the MsgEnvelope when the topic of the
 * MsgEnvelope starts with the String specified when subscribing.
 */
public class SubchannelBusImpl extends SubchannelEventBus<MsgEnvelope, ActorRef, String> {

    // Subclassification is an object providing 'isEqual' and 'isSubclass'
    // to be consumed by the other methods of this classifier
    @Override public Subclassification<String> subclassification() {
        return new StartsWithSubclassification();
    }

    // is used for extracting the classifier from the incoming events
    @Override public String classify(MsgEnvelope event) {
        return event.topic;
    }

    // will be invoked for each event for all subscribers which registered themselves
    // for the event's classifier
    @Override public void publish(MsgEnvelope event, ActorRef subscriber) {
        subscriber.tell(event.payload, ActorRef.noSender());
    }
}
```

A test for this implementation may look like this:

```
SubchannelBusImpl subchannelBus = new SubchannelBusImpl();
subchannelBus.subscribe(getTestActor(), "abc");
subchannelBus.publish(new MsgEnvelope("xyzabc", "x"));
subchannelBus.publish(new MsgEnvelope("bcdef", "b"));
subchannelBus.publish(new MsgEnvelope("abc", "c"));
expectMsgEquals("c");
```

```
subchannelBus.publish(new MsgEnvelope("abcdef", "d"));
expectMsgEquals("d");
```

This classifier is also efficient in case no subscribers are found for an event, but it uses conventional locking to synchronize an internal classifier cache, hence it is not well-suited to use cases in which subscriptions change with very high frequency (keep in mind that “opening” a classifier by sending the first message will also have to re-check all previous subscriptions).

Scanning Classification

The previous classifier was built for multi-classifier subscriptions which are strictly hierarchical, this classifier is useful if there are overlapping classifiers which cover various parts of the event space without forming a hierarchy. It can be compared to tuning in on (possibly multiple) radio stations by geographical reachability (for old-school radio-wave transmission).

The necessary methods to be implemented are illustrated with the following example:

```
import akka.event.japi.ScanningEventBus;

/**
 * Publishes String messages with length less than or equal to the length
 * specified when subscribing.
 */
public class ScanningBusImpl extends ScanningEventBus<String, ActorRef, Integer> {

    // is needed for determining matching classifiers and storing them in an
    // ordered collection
    @Override public int compareClassifiers(Integer a, Integer b) {
        return a.compareTo(b);
    }

    // is needed for storing subscribers in an ordered collection
    @Override public int compareSubscribers(ActorRef a, ActorRef b) {
        return a.compareTo(b);
    }

    // determines whether a given classifier shall match a given event; it is invoked
    // for each subscription for all received events, hence the name of the classifier
    @Override public boolean matches(Integer classifier, String event) {
        return event.length() <= classifier;
    }

    // will be invoked for each event for all subscribers which registered themselves
    // for the event's classifier
    @Override public void publish(String event, ActorRef subscriber) {
        subscriber.tell(event, ActorRef.noSender());
    }
}
```

A test for this implementation may look like this:

```
ScanningBusImpl scanningBus = new ScanningBusImpl();
scanningBus.subscribe(getTestActor(), 3);
scanningBus.publish("xyzabc");
scanningBus.publish("ab");
expectMsgEquals("ab");
scanningBus.publish("abc");
expectMsgEquals("abc");
```

This classifier takes always a time which is proportional to the number of subscriptions, independent of how many actually match.

Actor Classification

This classification was originally developed specifically for implementing *DeathWatch*: subscribers as well as classifiers are of type `ActorRef`.

This classification requires an `ActorSystem` in order to perform book-keeping operations related to the subscribers being `Actors`, which can terminate without first unsubscribing from the `EventBus`. `ManagedActorClassification` maintains a system `Actor` which takes care of unsubscribing terminated actors automatically.

The necessary methods to be implemented are illustrated with the following example:

```
import akka.event.japi.ManagedActorEventBus;

public class Notification {
    public final ActorRef ref;
    public final int id;

    public Notification(ActorRef ref, int id) {
        this.ref = ref;
        this.id = id;
    }
}

public class ActorBusImpl extends ManagedActorEventBus<Notification> {

    // the ActorSystem will be used for book-keeping operations, such as subscribers terminating
    public ActorBusImpl(ActorSystem system) {
        super(system);
    }

    // is used for extracting the classifier from the incoming events
    @Override public ActorRef classify(Notification event) {
        return event.ref;
    }

    // determines the initial size of the index data structure
    // used internally (i.e. the expected number of different classifiers)
    @Override public int mapSize() {
        return 128;
    }
}
```

A test for this implementation may look like this:

```
ActorRef observer1 = new JavaTestKit(system).getRef();
ActorRef observer2 = new JavaTestKit(system).getRef();
JavaTestKit probe1 = new JavaTestKit(system);
JavaTestKit probe2 = new JavaTestKit(system);
ActorRef subscriber1 = probe1.getRef();
ActorRef subscriber2 = probe2.getRef();
ActorBusImpl actorBus = new ActorBusImpl(system);
actorBus.subscribe(subscriber1, observer1);
actorBus.subscribe(subscriber2, observer1);
actorBus.subscribe(subscriber2, observer2);
Notification n1 = new Notification(observer1, 100);
actorBus.publish(n1);
probe1.expectMsgEquals(n1);
probe2.expectMsgEquals(n1);
Notification n2 = new Notification(observer2, 101);
actorBus.publish(n2);
probe2.expectMsgEquals(n2);
probe1.expectNoMsg(FiniteDuration.create(500, TimeUnit.MILLISECONDS));
```

This classifier is still is generic in the event type, and it is efficient for all use cases.

7.1.2 Event Stream

The event stream is the main event bus of each actor system: it is used for carrying *log messages* and *Dead Letters* and may be used by the user code for other purposes as well. It uses *Subchannel Classification* which enables registering to related sets of channels (as is used for `RemotingLifecycleEvent`). The following example demonstrates how a simple subscription works. Given a simple actor:

```
import akka.actor.ActorRef;
import akka.actor.ActorSystem;

public class DeadLetterActor extends UntypedActor {
    public void onReceive(Object message) {
        if (message instanceof DeadLetter) {
            System.out.println(message);
        }
    }
}
```

it can be subscribed like this:

```
final ActorSystem system = ActorSystem.create("DeadLetters");
final ActorRef actor = system.actorOf(Props.create(DeadLetterActor.class));
system.eventStream().subscribe(actor, DeadLetter.class);
```

It is also worth pointing out that thanks to the way the subchannel classification is implemented in the event stream, it is possible to subscribe to a group of events, by subscribing to their common superclass as demonstrated in the following example:

```
interface AllKindsOfMusic { }

class Jazz implements AllKindsOfMusic {
    final public String artist;
    public Jazz(String artist) {
        this.artist = artist;
    }
}

class Electronic implements AllKindsOfMusic {
    final public String artist;
    public Electronic(String artist) {
        this.artist = artist;
    }
}

class Listener extends UntypedActor {
    @Override
    public void onReceive(Object message) throws Exception {
        if (message instanceof Jazz) {
            System.out.printf("%s is listening to: %s%n", self().path().name(), message);
        } else if (message instanceof Electronic) {
            System.out.printf("%s is listening to: %s%n", self().path().name(), message);
        }
    }
}

final ActorRef actor = system.actorOf(Props.create(DeadLetterActor.class));
system.eventStream().subscribe(actor, DeadLetter.class);

final ActorRef jazzListener = system.actorOf(Props.create(Listener.class));
final ActorRef musicListener = system.actorOf(Props.create(Listener.class));
system.eventStream().subscribe(jazzListener, Jazz.class);
system.eventStream().subscribe(musicListener, AllKindsOfMusic.class);
```

```
// only musicListener gets this message, since it listens to *all* kinds of music:
system.eventStream().publish(new Electronic("Parov Stelar"));

// jazzListener and musicListener will be notified about Jazz:
system.eventStream().publish(new Jazz("Sonny Rollins"));
```

Similarly to [Actor Classification](#), `EventStream` will automatically remove subscribers when they terminate.

Note: The event stream is a *local facility*, meaning that it will *not* distribute events to other nodes in a clustered environment (unless you subscribe a Remote Actor to the stream explicitly). If you need to broadcast events in an Akka cluster, *without* knowing your recipients explicitly (i.e. obtaining their ActorRefs), you may want to look into: [Distributed Publish Subscribe in Cluster](#).

Default Handlers

Upon start-up the actor system creates and subscribes actors to the event stream for logging: these are the handlers which are configured for example in `application.conf`:

```
akka {
  loggers = ["akka.event.Logging$DefaultLogger"]
}
```

The handlers listed here by fully-qualified class name will be subscribed to all log event classes with priority higher than or equal to the configured log-level and their subscriptions are kept in sync when changing the log-level at runtime:

```
system.eventStream.setLogLevel(Logging.DebugLevel());
```

This means that log events for a level which will not be logged are typically not dispatched at all (unless manual subscriptions to the respective event class have been done)

Dead Letters

As described at [Stopping actors](#), messages queued when an actor terminates or sent after its death are re-routed to the dead letter mailbox, which by default will publish the messages wrapped in `DeadLetter`. This wrapper holds the original sender, receiver and message of the envelope which was redirected.

Some internal messages (marked with the `DeadLetterSuppression` trait) will not end up as dead letters like normal messages. These are by design safe and expected to sometimes arrive at a terminated actor and since they are nothing to worry about, they are suppressed from the default dead letters logging mechanism.

However, in case you find yourself in need of debugging these kinds of low level suppressed dead letters, it's still possible to subscribe to them explicitly:

```
system.eventStream().subscribe(actor, SuppressedDeadLetter.class);
```

or all dead letters (including the suppressed ones):

```
system.eventStream().subscribe(actor, AllDeadLetters.class);
```

Other Uses

The event stream is always there and ready to be used, just publish your own events (it accepts `Object`) and subscribe listeners to the corresponding JVM classes.

7.2 Logging

Logging in Akka is not tied to a specific logging backend. By default log messages are printed to STDOUT, but you can plug-in a SLF4J logger or your own logger. Logging is performed asynchronously to ensure that logging has minimal performance impact. Logging generally means IO and locks, which can slow down the operations of your code if it was performed synchronously.

7.2.1 How to Log

Create a `LoggingAdapter` and use the `error`, `warning`, `info`, or `debug` methods, as illustrated in this example:

```
import akka.actor.*;
import akka.event.Logging;
import akka.event.LoggingAdapter;

class MyActor extends UntypedActor {
  LoggingAdapter log = Logging.getLogger(getContext().system(), this);

  @Override
  public void preStart() {
    log.debug("Starting");
  }

  @Override
  public void preRestart(Throwable reason, Option<Object> message) {
    log.error(reason, "Restarting due to [{}] when processing [{}]",
      reason.getMessage(), message.isDefined() ? message.get() : "");
  }

  public void onReceive(Object message) {
    if (message.equals("test")) {
      log.info("Received test");
    } else {
      log.warning("Received unknown message: {}", message);
    }
  }
}
```

The first parameter to `Logging.getLogger` could also be any `LoggingBus`, specifically `system.eventStream()`; in the demonstrated case, the actor system's address is included in the `akkaSource` representation of the log source (see [Logging Thread](#), [Akka Source](#) and [Actor System in MDC](#)) while in the second case this is not automatically done. The second parameter to `Logging.getLogger` is the source of this logging channel. The source object is translated to a `String` according to the following rules:

- if it is an `Actor` or `ActorRef`, its path is used
- in case of a `String` it is used as is
- in case of a class an approximation of its `simpleName`
- and in all other cases the `simpleName` of its class

The log message may contain argument placeholders `{}`, which will be substituted if the log level is enabled. Giving more arguments as there are placeholders results in a warning being appended to the log statement (i.e. on the same line with the same severity). You may pass a Java array as the only substitution argument to have its elements be treated individually:

```
final Object[] args = new Object[] { "The", "brown", "fox", "jumps", 42 };
system.log().debug("five parameters: {}, {}, {}, {}, {}", args);
```

The Java Class of the log source is also included in the generated `LogEvent`. In case of a simple string this is replaced with a “marker” class `akka.event.DummyClassForStringSources` in order to allow special

treatment of this case, e.g. in the SLF4J event listener which will then use the string instead of the class' name for looking up the logger instance to use.

Logging of Dead Letters

By default messages sent to dead letters are logged at info level. Existence of dead letters does not necessarily indicate a problem, but it might be, and therefore they are logged by default. After a few messages this logging is turned off, to avoid flooding the logs. You can disable this logging completely or adjust how many dead letters that are logged. During system shutdown it is likely that you see dead letters, since pending messages in the actor mailboxes are sent to dead letters. You can also disable logging of dead letters during shutdown.

```
akka {
  log-dead-letters = 10
  log-dead-letters-during-shutdown = on
}
```

To customize the logging further or take other actions for dead letters you can subscribe to the *Event Stream*.

Auxiliary logging options

Akka has a couple of configuration options for very low level debugging, that makes most sense in for developers and not for operations.

You almost definitely need to have logging set to DEBUG to use any of the options below:

```
akka {
  loglevel = "DEBUG"
}
```

This config option is very good if you want to know what config settings are loaded by Akka:

```
akka {
  # Log the complete configuration at INFO level when the actor system is started.
  # This is useful when you are uncertain of what configuration is used.
  log-config-on-start = on
}
```

If you want very detailed logging of all automatically received messages that are processed by Actors:

```
akka {
  actor {
    debug {
      # enable DEBUG logging of all AutoReceiveMessages (Kill, PoisonPill et.c.)
      autoreceive = on
    }
  }
}
```

If you want very detailed logging of all lifecycle changes of Actors (restarts, deaths etc):

```
akka {
  actor {
    debug {
      # enable DEBUG logging of actor lifecycle changes
      lifecycle = on
    }
  }
}
```

If you want unhandled messages logged at DEBUG:

```
akka {
  actor {
    debug {
      # enable DEBUG logging of unhandled messages
      unhandled = on
    }
  }
}
```

If you want very detailed logging of all events, transitions and timers of FSM Actors that extend LoggingFSM:

```
akka {
  actor {
    debug {
      # enable DEBUG logging of all LoggingFSMs for events, transitions and timers
      fsm = on
    }
  }
}
```

If you want to monitor subscriptions (subscribe/unsubscribe) on the ActorSystem.eventStream:

```
akka {
  actor {
    debug {
      # enable DEBUG logging of subscription changes on the eventStream
      event-stream = on
    }
  }
}
```

Auxiliary remote logging options

If you want to see all messages that are sent through remoting at DEBUG log level: (This is logged as they are sent by the transport layer, not by the Actor)

```
akka {
  remote {
    # If this is "on", Akka will log all outbound messages at DEBUG level,
    # if off then they are not logged
    log-sent-messages = on
  }
}
```

If you want to see all messages that are received through remoting at DEBUG log level: (This is logged as they are received by the transport layer, not by any Actor)

```
akka {
  remote {
    # If this is "on", Akka will log all inbound messages at DEBUG level,
    # if off then they are not logged
    log-received-messages = on
  }
}
```

If you want to see message types with payload size in bytes larger than a specified limit at INFO log level:

```
akka {
  remote {
    # Logging of message types with payload size in bytes larger than
    # this value. Maximum detected size per message type is logged once,
    # with an increase threshold of 10%.
    # By default this feature is turned off. Activate it by setting the property to
```

```
# a value in bytes, such as 1000b. Note that for all messages larger than this
# limit there will be extra performance and scalability cost.
log-frame-size-exceeding = 1000b
}
}
```

Also see the logging options for TestKit: *Tracing Actor Invocations*.

Turn Off Logging

To turn off logging you can configure the log levels to be OFF like this.

```
akka {
  stdout-loglevel = "OFF"
  loglevel = "OFF"
}
```

The `stdout-loglevel` is only in effect during system startup and shutdown, and setting it to OFF as well, ensures that nothing gets logged during system startup or shutdown.

7.2.2 Loggers

Logging is performed asynchronously through an event bus. Log events are processed by an event handler actor and it will receive the log events in the same order as they were emitted.

Note: The event handler actor does not have a bounded inbox and is run on the default dispatcher. This means that logging extreme amounts of data may affect your application badly. It can be somewhat mitigated by making sure to use an async logging backend though. (See *Using the SLF4J API directly*)

You can configure which event handlers are created at system start-up and listen to logging events. That is done using the `loggers` element in the *Configuration*. Here you can also define the log level. More fine grained filtering based on the log source can be implemented in a custom `LoggingFilter`, which can be defined in the `logging-filter` configuration property.

```
akka {
  # Loggers to register at boot time (akka.event.Logging$DefaultLogger logs
  # to STDOUT)
  loggers = ["akka.event.Logging$DefaultLogger"]
  # Options: OFF, ERROR, WARNING, INFO, DEBUG
  loglevel = "DEBUG"
}
```

The default one logs to STDOUT and is registered by default. It is not intended to be used for production. There is also an *SLF4J* logger available in the 'akka-slf4j' module.

Example of creating a listener:

```
import akka.actor.*;
import akka.event.Logging;
import akka.event.LoggingAdapter;

import akka.event.Logging.InitializeLogger;
import akka.event.Logging.Error;
import akka.event.Logging.Warning;
import akka.event.Logging.Info;
import akka.event.Logging.Debug;

class MyEventListener extends UntypedActor {
  public void onReceive(Object message) {
    if (message instanceof InitializeLogger) {
```

```

    getSender().tell(Logging.loggerInitialized(), getSelf());
  } else if (message instanceof Error) {
    // ...
  } else if (message instanceof Warning) {
    // ...
  } else if (message instanceof Info) {
    // ...
  } else if (message instanceof Debug) {
    // ...
  }
}
}
}

```

7.2.3 Logging to stdout during startup and shutdown

While the actor system is starting up and shutting down the configured `loggers` are not used. Instead log messages are printed to stdout (`System.out`). The default log level for this stdout logger is `WARNING` and it can be silenced completely by setting `akka.stdout-loglevel=OFF`.

7.2.4 SLF4J

Akka provides a logger for [SL4FJ](#). This module is available in the ‘`akka-slf4j.jar`’. It has one single dependency; the `slf4j-api.jar`. In runtime you also need a SLF4J backend, we recommend [Logback](#):

```

<dependency>
  <groupId>ch.qos.logback</groupId>
  <artifactId>logback-classic</artifactId>
  <version>1.1.3</version>
</dependency>

```

You need to enable the `Slf4jLogger` in the `loggers` element in the *Configuration*. Here you can also define the log level of the event bus. More fine grained log levels can be defined in the configuration of the SLF4J backend (e.g. `logback.xml`). You should also define `akka.event.slf4j.Slf4jLoggingFilter` in the `logging-filter` configuration property. It will filter the log events using the backend configuration (e.g. `logback.xml`) before they are published to the event bus.

Warning: If you set the `loglevel` to a higher level than “DEBUG”, any DEBUG events will be filtered out already at the source and will never reach the logging backend, regardless of how the backend is configured.

```

akka {
  loggers = ["akka.event.slf4j.Slf4jLogger"]
  loglevel = "DEBUG"
  logging-filter = "akka.event.slf4j.Slf4jLoggingFilter"
}

```

One gotcha is that the timestamp is attributed in the event handler, not when actually doing the logging.

The SLF4J logger selected for each log event is chosen based on the `Class` of the log source specified when creating the `LoggingAdapter`, unless that was given directly as a string in which case that string is used (i.e. `LoggerFactory.getLogger(Class c)` is used in the first case and `LoggerFactory.getLogger(String s)` in the second).

Note: Beware that the actor system’s name is appended to a `String` log source if the `LoggingAdapter` was created giving an `ActorSystem` to the factory. If this is not intended, give a `LoggingBus` instead as shown below:

```
final LoggingAdapter log = Logging.getLogger(system.eventStream(), "my.string");
```

Using the SLF4J API directly

If you use the SLF4J API directly in your application, remember that the logging operations will block while the underlying infrastructure writes the log statements.

This can be avoided by configuring the logging implementation to use a non-blocking appender. Logback provides [AsyncAppender](#) that does this. It also contains a feature which will drop INFO and DEBUG messages if the logging load is high.

Logging Thread, Akka Source and Actor System in MDC

Since the logging is done asynchronously the thread in which the logging was performed is captured in Mapped Diagnostic Context (MDC) with attribute name `sourceThread`. With Logback the thread name is available with `%X{sourceThread}` specifier within the pattern layout configuration:

```
<appender name="STDOUT" class="ch.qos.logback.core.ConsoleAppender">
  <encoder>
    <pattern>%date{ISO8601} %-5level %logger{36} %X{sourceThread} - %msg%n</pattern>
  </encoder>
</appender>
```

Note: It will probably be a good idea to use the `sourceThread` MDC value also in non-Akka parts of the application in order to have this property consistently available in the logs.

Another helpful facility is that Akka captures the actor's address when instantiating a logger within it, meaning that the full instance identification is available for associating log messages e.g. with members of a router. This information is available in the MDC with attribute name `akkaSource`:

```
<appender name="STDOUT" class="ch.qos.logback.core.ConsoleAppender">
  <encoder>
    <pattern>%date{ISO8601} %-5level %logger{36} %X{akkaSource} - %msg%n</pattern>
  </encoder>
</appender>
```

Finally, the actor system in which the logging was performed is available in the MDC with attribute name `sourceActorSystem`:

```
<appender name="STDOUT" class="ch.qos.logback.core.ConsoleAppender">
  <encoder>
    <pattern>%date{ISO8601} %-5level %logger{36} %X{sourceActorSystem} - %msg%n</pattern>
  </encoder>
</appender>
```

For more details on what this attribute contains—also for non-actors—please see [How to Log](#).

More accurate timestamps for log output in MDC

Akka's logging is asynchronous which means that the timestamp of a log entry is taken from when the underlying logger implementation is called, which can be surprising at first. If you want to more accurately output the timestamp, use the MDC attribute `akkaTimestamp`:

```
<appender name="STDOUT" class="ch.qos.logback.core.ConsoleAppender">
  <encoder>
    <pattern>%X{akkaTimestamp} %-5level %logger{36} %X{akkaSource} - %msg%n</pattern>
  </encoder>
</appender>
```

MDC values defined by the application

One useful feature available in Slf4j is MDC, Akka has a way for let the application specify custom values, you just need to get a specialized `LoggingAdapter`, the `DiagnosticLoggingAdapter`. In order to get it you will use the factory receiving an `UntypedActor` as `logSource`:

```
// Within your UntypedActor
final DiagnosticLoggingAdapter log = Logging.getLogger(this);
```

Once you have the logger, you just need to add the custom values before you log something. This way, the values will be put in the SLF4J MDC right before appending the log and removed after.

Note: The cleanup (removal) should be done in the actor at the end, otherwise, next message will log with same mdc values, if it is not set to a new map. Use `log.clearMDC()`.

```
import akka.event.Logging;
import akka.event.DiagnosticLoggingAdapter;
import java.util.HashMap;
import java.util.Map;

class MdcActor extends UntypedActor {

    final DiagnosticLoggingAdapter log = Logging.getLogger(this);

    public void onReceive(Object message) {

        Map<String, Object> mdc;
        mdc = new HashMap<String, Object>();
        mdc.put("requestId", 1234);
        mdc.put("visitorId", 5678);
        log.setMDC(mdc);

        log.info("Starting new request");

        log.clearMDC();
    }
}
```

Now, the values will be available in the MDC, so you can use them in the layout pattern:

```
<appender name="STDOUT" class="ch.qos.logback.core.ConsoleAppender">
  <encoder>
    <pattern>
      %-5level %logger{36} [req: %X{requestId}, visitor: %X{visitorId}] - %msg%n
    </pattern>
  </encoder>
</appender>
```

7.3 Scheduler

Sometimes the need for making things happen in the future arises, and where do you go look then? Look no further than `ActorSystem`! There you find the `scheduler` method that returns an instance of `akka.actor.Scheduler`, this instance is unique per `ActorSystem` and is used internally for scheduling things to happen at specific points in time.

You can schedule sending of messages to actors and execution of tasks (functions or `Runnable`). You will get a `Cancellable` back that you can call `cancel` on to cancel the execution of the scheduled operation.

Warning: The default implementation of `Scheduler` used by Akka is based on job buckets which are emptied according to a fixed schedule. It does not execute tasks at the exact time, but on every tick, it will run everything that is (over)due. The accuracy of the default `Scheduler` can be modified by the `akka.scheduler.tick-duration` configuration property.

7.3.1 Some examples

Schedule to send the “foo”-message to the `testActor` after 50ms:

```
import akka.actor.Props;
import scala.concurrent.duration.Duration;
import java.util.concurrent.TimeUnit;

system.scheduler().scheduleOnce(Duration.create(50, TimeUnit.MILLISECONDS),
    testActor, "foo", system.dispatcher(), null);
```

Schedule a `Runnable`, that sends the current time to the `testActor`, to be executed after 50ms:

```
system.scheduler().scheduleOnce(Duration.create(50, TimeUnit.MILLISECONDS),
    new Runnable() {
        @Override
        public void run() {
            testActor.tell(System.currentTimeMillis(), ActorRef.noSender());
        }
    }, system.dispatcher());
```

Warning: If you schedule `Runnable` instances you should be extra careful to not pass or close over unstable references. In practice this means that you should not call methods on the enclosing `Actor` from within the `Runnable`. If you need to schedule an invocation it is better to use the `schedule()` variant accepting a message and an `ActorRef` to schedule a message to self (containing the necessary parameters) and then call the method when the message is received.

Schedule to send the “Tick”-message to the `tickActor` after 0ms repeating every 50ms:

```
import akka.actor.Props;
import scala.concurrent.duration.Duration;
import java.util.concurrent.TimeUnit;
import akka.actor.UntypedActor;
import akka.actor.Cancellable;

class Ticker extends UntypedActor {
    @Override
    public void onReceive(Object message) {
        if (message.equals("Tick")) {
            // Do something
        } else {
            unhandled(message);
        }
    }
}

ActorRef tickActor = system.actorOf(Props.create(Ticker.class, this));

//This will schedule to send the Tick-message
//to the tickActor after 0ms repeating every 50ms
Cancellable cancellable = system.scheduler().schedule(Duration.Zero(),
    Duration.create(50, TimeUnit.MILLISECONDS), tickActor, "Tick",
    system.dispatcher(), null);
```

```
//This cancels further Ticks to be sent
cancellable.cancel();
```

7.3.2 From akka.actor.ActorSystem

```
/**
 * Light-weight scheduler for running asynchronous tasks after some deadline
 * in the future. Not terribly precise but cheap.
 */
def scheduler: Scheduler
```

Warning: All scheduled task will be executed when the ActorSystem is terminated, i.e. the task may execute before its timeout.

7.3.3 The Scheduler Interface for Implementors

The actual scheduler implementation is loaded reflectively upon ActorSystem start-up, which means that it is possible to provide a different one using the akka.scheduler.implementation configuration property. The referenced class must implement the following interface:

```
/**
 * An Akka scheduler service. This one needs one special behavior: if
 * Closeable, it MUST execute all outstanding tasks upon .close() in order
 * to properly shutdown all dispatchers.
 *
 * Furthermore, this timer service MUST throw IllegalStateException if it
 * cannot schedule a task. Once scheduled, the task MUST be executed. If
 * executed upon close(), the task may execute before its timeout.
 *
 * Scheduler implementation are loaded reflectively at ActorSystem start-up
 * with the following constructor arguments:
 * 1) the system's com.typesafe.config.Config (from system.settings.config)
 * 2) a akka.event.LoggingAdapter
 * 3) a java.util.concurrent.ThreadFactory
 */
public abstract class AbstractScheduler extends AbstractSchedulerBase {

    /**
     * Schedules a function to be run repeatedly with an initial delay and
     * a frequency. E.g. if you would like the function to be run after 2
     * seconds and thereafter every 100ms you would set delay = Duration(2,
     * TimeUnit.SECONDS) and interval = Duration(100, TimeUnit.MILLISECONDS)
     */
    @Override
    public abstract Cancellable schedule(FiniteDuration initialDelay,
        FiniteDuration interval, Runnable runnable, ExecutionContext executor);

    /**
     * Schedules a Runnable to be run once with a delay, i.e. a time period that
     * has to pass before the runnable is executed.
     */
    @Override
    public abstract Cancellable scheduleOnce(FiniteDuration delay, Runnable runnable,
        ExecutionContext executor);

    /**
     * The maximum supported task frequency of this scheduler, i.e. the inverse
     * of the minimum time interval between executions of a recurring task, in Hz.
     */
}
```

```
@Override
public abstract double maxFrequency();
}
```

7.3.4 The Cancellable interface

Scheduling a task will result in a `Cancellable` (or throw an `IllegalStateException` if attempted after the scheduler's shutdown). This allows you to cancel something that has been scheduled for execution.

Warning: This does not abort the execution of the task, if it had already been started. Check the return value of `cancel` to detect whether the scheduled task was canceled or will (eventually) have run.

```
/**
 * Signifies something that can be cancelled
 * There is no strict guarantee that the implementation is thread-safe,
 * but it should be good practice to make it so.
 */
trait Cancellable {
  /**
   * Cancels this Cancellable and returns true if that was successful.
   * If this cancellable was (concurrently) cancelled already, then this method
   * will return false although isCancelled will return true.
   *
   * Java & Scala API
   */
  def cancel(): Boolean

  /**
   * Returns true if and only if this Cancellable has been successfully cancelled
   *
   * Java & Scala API
   */
  def isCancelled: Boolean
}
```

7.4 Duration

Durations are used throughout the Akka library, wherefore this concept is represented by a special data type, `scala.concurrent.duration.Duration`. Values of this type may represent infinite (`Duration.Inf`, `Duration.MinusInf`) or finite durations, or be `Duration.Undefined`.

7.4.1 Finite vs. Infinite

Since trying to convert an infinite duration into a concrete time unit like seconds will throw an exception, there are different types available for distinguishing the two kinds at compile time:

- `FiniteDuration` is guaranteed to be finite, calling `toNanos` and friends is safe
- `Duration` can be finite or infinite, so this type should only be used when finite-ness does not matter; this is a supertype of `FiniteDuration`

7.4.2 Scala

In Scala durations are constructable using a mini-DSL and support all expected arithmetic operations:

```
import scala.concurrent.duration._

val fivesec = 5.seconds
val threemillis = 3.millis
val diff = fivesec - threemillis
assert(diff < fivesec)
val fourmillis = threemillis * 4 / 3 // you cannot write it the other way around
val n = threemillis / (1 millisecond)
```

Note: You may leave out the dot if the expression is clearly delimited (e.g. within parentheses or in an argument list), but it is recommended to use it if the time unit is the last token on a line, otherwise semi-colon inference might go wrong, depending on what starts the next line.

7.4.3 Java

Java provides less syntactic sugar, so you have to spell out the operations as method calls instead:

```
import scala.concurrent.duration.Duration;
import scala.concurrent.duration.Deadline;

final Duration fivesec = Duration.create(5, "seconds");
final Duration threemillis = Duration.create("3 millis");
final Duration diff = fivesec.minus(threemillis);
assert diff.lt(fivesec);
assert Duration.Zero().lt(Duration.Inf());
```

7.4.4 Deadline

Durations have a brother named `Deadline`, which is a class holding a representation of an absolute point in time, and support deriving a duration from this by calculating the difference between now and the deadline. This is useful when you want to keep one overall deadline without having to take care of the book-keeping wrt. the passing of time yourself:

```
val deadline = 10.seconds.fromNow
// do something
val rest = deadline.timeLeft
```

In Java you create these from durations:

```
final Deadline deadline = Duration.create(10, "seconds").fromNow();
final Duration rest = deadline.timeLeft();
```

7.5 Circuit Breaker

7.5.1 Why are they used?

A circuit breaker is used to provide stability and prevent cascading failures in distributed systems. These should be used in conjunction with judicious timeouts at the interfaces between remote systems to prevent the failure of a single component from bringing down all components.

As an example, we have a web application interacting with a remote third party web service. Let's say the third party has oversold their capacity and their database melts down under load. Assume that the database fails in such a way that it takes a very long time to hand back an error to the third party web service. This in turn makes calls fail after a long period of time. Back to our web application, the users have noticed that their form submissions take much longer seeming to hang. Well the users do what they know to do which is use the refresh button, adding more requests to their already running requests. This eventually causes the failure of the web application due to


resource exhaustion. This will affect all users, even those who are not using functionality dependent on this third party web service.

Introducing circuit breakers on the web service call would cause the requests to begin to fail-fast, letting the user know that something is wrong and that they need not refresh their request. This also confines the failure behavior to only those users that are using functionality dependent on the third party, other users are no longer affected as there is no resource exhaustion. Circuit breakers can also allow savvy developers to mark portions of the site that use the functionality unavailable, or perhaps show some cached content as appropriate while the breaker is open.

The Akka library provides an implementation of a circuit breaker called `akka.pattern.CircuitBreaker` which has the behavior described below.

7.5.2 What do they do?

- **During normal operation, a circuit breaker is in the *Closed* state:**
 - Exceptions or calls exceeding the configured `callTimeout` increment a failure counter
 - Successes reset the failure count to zero
 - When the failure counter reaches a `maxFailures` count, the breaker is tripped into *Open* state
- **While in *Open* state:**
 - All calls fail-fast with a `CircuitBreakerOpenException`
 - After the configured `resetTimeout`, the circuit breaker enters a *Half-Open* state
- **In *Half-Open* state:**
 - The first call attempted is allowed through without failing fast
 - All other calls fail-fast with an exception just as in *Open* state
 - If the first call succeeds, the breaker is reset back to *Closed* state
 - If the first call fails, the breaker is tripped again into the *Open* state for another full `resetTimeout`
- **State transition listeners:**
 - Callbacks can be provided for every state entry via `onOpen`, `onClose`, and `onHalfOpen`
 - These are executed in the `ExecutionContext` provided.



7.5.3 Examples

Initialization

Here's how a `CircuitBreaker` would be configured for:

- 5 maximum failures
- a call timeout of 10 seconds

- a reset timeout of 1 minute

Scala

```
import scala.concurrent.duration._
import akka.pattern.CircuitBreaker
import akka.pattern.pipe
import akka.actor.Actor
import akka.actor.ActorLogging
import scala.concurrent.Future
import akka.event.Logging

class DangerousActor extends Actor with ActorLogging {
  import context.dispatcher

  val breaker =
    new CircuitBreaker(context.system.scheduler,
      maxFailures = 5,
      callTimeout = 10.seconds,
      resetTimeout = 1.minute).onOpen(notifyMeOnOpen())

  def notifyMeOnOpen(): Unit =
    log.warning("My CircuitBreaker is now open, and will not close for one minute")
}
```

Java

```
import akka.actor.UntypedActor;
import scala.concurrent.Future;
import akka.event.LoggingAdapter;
import scala.concurrent.duration.Duration;
import akka.pattern.CircuitBreaker;
import akka.event.Logging;

import static akka.pattern.Patterns.pipe;
import static akka.dispatch.Futures.future;

import java.util.concurrent.Callable;

public class DangerousJavaActor extends UntypedActor {

  private final CircuitBreaker breaker;
  private final LoggingAdapter log = Logging.getLogger(getContext().system(), this);

  public DangerousJavaActor() {
    this.breaker = new CircuitBreaker(
      getContext().dispatcher(), getContext().system().scheduler(),
      5, Duration.create(10, "s"), Duration.create(1, "m"))
      .onOpen(new Runnable() {
        public void run() {
          notifyMeOnOpen();
        }
      });
  }

  public void notifyMeOnOpen() {
    log.warning("My CircuitBreaker is now open, and will not close for one minute");
  }
}
```

Call Protection

Here's how the `CircuitBreaker` would be used to protect an asynchronous call as well as a synchronous one:

Scala

```
def dangerousCall: String = "This really isn't that dangerous of a call after all"

def receive = {
  case "is my middle name" =>
    breaker.withCircuitBreaker(Future(dangerousCall)) pipeTo sender()
  case "block for me" =>
    sender() ! breaker.withSyncCircuitBreaker(dangerousCall)
}
```

Java

```
public String dangerousCall() {
    return "This really isn't that dangerous of a call after all";
}

@Override
public void onReceive(Object message) {
    if (message instanceof String) {
        String m = (String) message;
        if ("is my middle name".equals(m)) {
            pipe(breaker.callWithCircuitBreaker(
                new Callable<Future<String>>() {
                    public Future<String> call() throws Exception {
                        return future(
                            new Callable<String>() {
                                public String call() {
                                    return dangerousCall();
                                }
                            }, getContext().dispatcher());
                    }
                }
            ), getContext().dispatcher().to(getSender()));
        }
        if ("block for me".equals(m)) {
            getSender().tell(breaker
                .callWithSyncCircuitBreaker(
                    new Callable<String>() {
                        @Override
                        public String call() throws Exception {
                            return dangerousCall();
                        }
                    }
                ), getSelf());
        }
    }
}
```

Note: Using the `CircuitBreaker` companion object's `apply` or `create` methods will return a `CircuitBreaker` where callbacks are executed in the caller's thread. This can be useful if the asynchronous `Future` behavior is unnecessary, for example invoking a synchronous-only API.

7.6 Akka Extensions

If you want to add features to Akka, there is a very elegant, but powerful mechanism for doing so. It's called Akka Extensions and is comprised of 2 basic components: an `Extension` and an `ExtensionId`.

Extensions will only be loaded once per `ActorSystem`, which will be managed by Akka. You can choose to have your `Extension` loaded on-demand or at `ActorSystem` creation time through the Akka configuration. Details on how to make that happens are below, in the "Loading from Configuration" section.

Warning: Since an extension is a way to hook into Akka itself, the implementor of the extension needs to ensure the thread safety of his/her extension.

7.6.1 Building an Extension

So let's create a sample extension that just lets us count the number of times something has happened.

First, we define what our `Extension` should do:

```
import akka.actor.*;
import java.util.concurrent.atomic.AtomicLong;

public class CountExtensionImpl implements Extension {
    //Since this Extension is a shared instance
    // per ActorSystem we need to be threadsafe
    private final AtomicLong counter = new AtomicLong(0);

    //This is the operation this Extension provides
    public long increment() {
        return counter.incrementAndGet();
    }
}
```

Then we need to create an `ExtensionId` for our extension so we can grab a hold of it.

```
import akka.actor.*;
import java.util.concurrent.atomic.AtomicLong;

public class CountExtension extends AbstractExtensionId<CountExtensionImpl>
    implements ExtensionIdProvider {
    //This will be the identifier of our CountExtension
    public final static CountExtension CountExtensionProvider = new CountExtension();

    private CountExtension() {}

    //The lookup method is required by ExtensionIdProvider,
    // so we return ourselves here, this allows us
    // to configure our extension to be loaded when
    // the ActorSystem starts up
    public CountExtension lookup() {
        return CountExtension.CountExtensionProvider; //The public static final
    }

    //This method will be called by Akka
    // to instantiate our Extension
    public CountExtensionImpl createExtension(ExtendedActorSystem system) {
        return new CountExtensionImpl();
    }
}
```

Wicked! Now all we need to do is to actually use it:

```
// typically you would use static import of the
// CountExtension.CountExtensionProvider field
CountExtension.CountExtensionProvider.get(system).increment();
```

Or from inside of an Akka Actor:

```
public class MyActor extends UntypedActor {
    public void onReceive(Object msg) {
        // typically you would use static import of the
        // CountExtension.CountExtensionProvider field
        CountExtension.CountExtensionProvider.get(getContext().system()).increment();
    }
}
```

That's all there is to it!

7.6.2 Loading from Configuration

To be able to load extensions from your Akka configuration you must add FQCNs of implementations of either `ExtensionId` or `ExtensionIdProvider` in the “akka.extensions” section of the config you provide to your `ActorSystem`.

```
akka {
  extensions = ["docs.extension.ExtensionDocTest.CountExtension"]
}
```

7.6.3 Applicability

The sky is the limit! By the way, did you know that Akka's Typed Actors, Serialization and other features are implemented as Akka Extensions?

Application specific settings

The *Configuration* can be used for application specific settings. A good practice is to place those settings in an Extension.

Sample configuration:

```
myapp {
  db {
    uri = "mongodb://example1.com:27017,example2.com:27017"
  }
  circuit-breaker {
    timeout = 30 seconds
  }
}
```

The Extension:

```
import akka.actor.Extension;
import akka.actor.AbstractExtensionId;
import akka.actor.ExtensionIdProvider;
import akka.actor.ActorSystem;
import akka.actor.ExtendedActorSystem;
import scala.concurrent.duration.Duration;
import com.typesafe.config.Config;
import java.util.concurrent.TimeUnit;
```

```

public class SettingsImpl implements Extension {

    public final String DB_URI;
    public final Duration CIRCUIT_BREAKER_TIMEOUT;

    public SettingsImpl(Config config) {
        DB_URI = config.getString("myapp.db.uri");
        CIRCUIT_BREAKER_TIMEOUT =
            Duration.create(config.getDuration("myapp.circuit-breaker.timeout",
                TimeUnit.MILLISECONDS), TimeUnit.MILLISECONDS);
    }

}

public class Settings extends AbstractExtensionId<SettingsImpl>
    implements ExtensionIdProvider {
    public final static Settings SettingsProvider = new Settings();

    private Settings() {}

    public Settings lookup() {
        return Settings.SettingsProvider;
    }

    public SettingsImpl createExtension(ExtendedActorSystem system) {
        return new SettingsImpl(system.settings().config());
    }
}

```

Use it:

```

public class MyActor extends UntypedActor {
    // typically you would use static import of the Settings.SettingsProvider field
    final SettingsImpl settings =
        Settings.SettingsProvider.get(getContext().system());
    Connection connection =
        connect(settings.DB_URI, settings.CIRCUIT_BREAKER_TIMEOUT);
}

```

7.7 Use-case and Deployment Scenarios

7.7.1 How can I use and deploy Akka?

Akka can be used in different ways:

- As a library: used as a regular JAR on the classpath and/or in a web app, to be put into `WEB-INF/lib`
- Package with [sbt-native-packager](#)
- Package and deploy using [Typesafe ConductR](#).

7.7.2 Native Packager

[sbt-native-packager](#) is a tool for creating distributions of any type of application, including an Akka applications.

Define sbt version in `project/build.properties` file:

```
sbt.version=0.13.7
```

Add [sbt-native-packager](#) in `project/plugins.sbt` file:

```
addSbtPlugin("com.typesafe.sbt" % "sbt-native-packager" % "1.0.0-RC1")
```

Use the package settings and optionally specify the mainClass in build.sbt file:

```
import NativePackagerHelper._

name := "akka-sample-main-scala"

version := "2.4.0"

scalaVersion := "2.11.7"

libraryDependencies += Seq(
  "com.typesafe.akka" %% "akka-actor" % "2.4.0"
)

enablePlugins(JavaServerAppPackaging)

mainClass in Compile := Some("sample.hello.Main")

mappings in Universal += {
  // optional example illustrating how to copy additional directory
  directory("scripts") ++
  // copy configuration files to config directory
  contentOf("src/main/resources").toMap.mapValues("config/" + _)
}

// add 'config' directory first in the classpath of the start script,
// an alternative is to set the config file locations via CLI parameters
// when starting the application
scriptClasspath := Seq("../config/") ++ scriptClasspath.value
```

Note: Use the `JavaServerAppPackaging`. Don't use `AkkaAppPackaging` (previously named `packageArchetype.akka_application`, since it doesn't have the same flexibility and quality as the `JavaServerAppPackaging`.

Use `sbt task dist` package the application.

To start the application (on a unix-based system):

```
cd target/universal/
unzip akka-sample-main-scala-2.4.0.zip
chmod u+x akka-sample-main-scala-2.4.0/bin/akka-sample-main-scala
akka-sample-main-scala-2.4.0/bin/akka-sample-main-scala sample.hello.Main
```

Use `Ctrl-C` to interrupt and exit the application.

On a Windows machine you can also use the `bin\akka-sample-main-scala.bat` script.

HOWTO: COMMON PATTERNS

This section lists common actor patterns which have been found to be useful, elegant or instructive. Anything is welcome, example topics being message routing strategies, supervision patterns, restart handling, etc. As a special bonus, additions to this section are marked with the contributor's name, and it would be nice if every Akka user who finds a recurring pattern in his or her code could share it for the profit of all. Where applicable it might also make sense to add to the `akka.pattern` package for creating an [OTP-like library](#).

You might find some of the patterns described in the Scala chapter of *howto-scala* useful even though the example code is written in Scala.

8.1 Scheduling Periodic Messages

This pattern describes how to schedule periodic messages to yourself in two different ways.

The first way is to set up periodic message scheduling in the constructor of the actor, and cancel that scheduled sending in `postStop` or else we might have multiple registered message sends to the same actor.

Note: With this approach the scheduled periodic message send will be restarted with the actor on restarts. This also means that the time period that elapses between two tick messages during a restart may drift off based on when you restart the scheduled message sends relative to the time that the last message was sent, and how long the initial delay is. Worst case scenario is `interval` plus `initialDelay`.

```
public class ScheduleInConstructor extends UntypedActor {

    private final Cancellable tick = getContext().system().scheduler().schedule(
        Duration.create(500, TimeUnit.MILLISECONDS),
        Duration.create(1000, TimeUnit.MILLISECONDS),
        getSelf(), "tick", getContext().dispatcher(), null);

    @Override
    public void postStop() {
        tick.cancel();
    }

    @Override
    public void onReceive(Object message) throws Exception {
        if (message.equals("tick")) {
            // do something useful here
        }
        else {
            unhandled(message);
        }
    }
}
```

The second variant sets up an initial one shot message send in the `preStart` method of the actor, and then the actor when it receives this message sets up a new one shot message send. You also have to override `postRestart` so we don't call `preStart` and schedule the initial message send again.

Note: With this approach we won't fill up the mailbox with tick messages if the actor is under pressure, but only schedule a new tick message when we have seen the previous one.

```
public class ScheduleInReceive extends UntypedActor {

    @Override
    public void preStart() {
        getContext().system().scheduler().scheduleOnce(
            Duration.create(500, TimeUnit.MILLISECONDS),
            getSelf(), "tick", getContext().dispatcher(), null);
    }

    // override postRestart so we don't call preStart and schedule a new message
    @Override
    public void postRestart(Throwable reason) {
    }

    @Override
    public void onReceive(Object message) throws Exception {
        if (message.equals("tick")) {
            // send another periodic tick after the specified delay
            getContext().system().scheduler().scheduleOnce(
                Duration.create(1000, TimeUnit.MILLISECONDS),
                getSelf(), "tick", getContext().dispatcher(), null);
            // do something useful here
        }
        else {
            unhandled(message);
        }
    }
}
```

8.2 Single-Use Actor Trees with High-Level Error Reporting

Contributed by: Rick Latrine

A nice way to enter the actor world from java is the use of `Patterns.ask()`. This method starts a temporary actor to forward the message and collect the result from the actor to be "asked". In case of errors within the asked actor the default supervision handling will take over. The caller of `Patterns.ask()` will *not* be notified.

If that caller is interested in such an exception, he must make sure that the asked actor replies with `Status.Failure(Throwable)`. Behind the asked actor a complex actor hierarchy might be spawned to accomplish asynchronous work. Then supervision is the established way to control error handling.

Unfortunately the asked actor must know about supervision and must catch the exceptions. Such an actor is unlikely to be reused in a different actor hierarchy and contains crippled try/catch blocks.

This pattern provides a way to encapsulate supervision and error propagation to the temporary actor. Finally the promise returned by `Patterns.ask()` is fulfilled as a failure, including the exception.

Let's have a look at the example code:

```
package docs.pattern;

import java.util.concurrent.TimeoutException;

import scala.concurrent.Future;
```

```

import scala.concurrent.duration.Duration;

import akka.actor.Actor;
import akka.actor.ActorKilledException;
import akka.actor.ActorRef;
import akka.actor.ActorRefFactory;
import akka.actor.Cancellable;
import akka.actor.OneForOneStrategy;
import akka.actor.Props;
import akka.actor.Scheduler;
import akka.actor.Status;
import akka.actor.SupervisorStrategy;
import akka.actor.SupervisorStrategy.Directive;
import akka.actor.Terminated;
import akka.actor.UntypedActor;
import akka.japi.Function;
import akka.pattern.Patterns;
import akka.util.Timeout;

public class SupervisedAsk {

    private static class AskParam {
        Props props;
        Object message;
        Timeout timeout;

        AskParam(Props props, Object message, Timeout timeout) {
            this.props = props;
            this.message = message;
            this.timeout = timeout;
        }
    }

    private static class AskTimeout {
    }

    public static class AskSupervisorCreator extends UntypedActor {

        @Override
        public void onReceive(Object message) throws Exception {
            if (message instanceof AskParam) {
                ActorRef supervisor = getContext().actorOf(
                    Props.create(AskSupervisor.class));
                supervisor.forward(message, getContext());
            } else {
                unhandled(message);
            }
        }
    }

    public static class AskSupervisor extends UntypedActor {
        private ActorRef targetActor;
        private ActorRef caller;
        private AskParam askParam;
        private Cancellable timeoutMessage;

        @Override
        public SupervisorStrategy supervisorStrategy() {
            return new OneForOneStrategy(0, Duration.Zero(),
                new Function<Throwable, Directive>() {
                    public Directive apply(Throwable cause) {
                        caller.tell(new Status.Failure(cause), self());
                    }
                },
                return SupervisorStrategy.stop();
        }
    }

```

```

    }
    });
}

@Override
public void onReceive(Object message) throws Exception {
    if (message instanceof AskParam) {
        askParam = (AskParam) message;
        caller = getSender();
        targetActor = getContext().actorOf(askParam.props);
        getContext().watch(targetActor);
        targetActor.forward(askParam.message, getContext());
        Scheduler scheduler = getContext().system().scheduler();
        timeoutMessage = scheduler.scheduleOnce(askParam.timeout.duration(),
            self(), new AskTimeout(), context().dispatcher(), null);
    } else if (message instanceof Terminated) {
        Throwable ex = new ActorKilledException("Target actor terminated.");
        caller.tell(new Status.Failure(ex), self());
        timeoutMessage.cancel();
        getContext().stop(self());
    } else if (message instanceof AskTimeout) {
        Throwable ex = new TimeoutException("Target actor timed out after "
            + askParam.timeout.toString());
        caller.tell(new Status.Failure(ex), self());
        getContext().stop(self());
    } else
        unhandled(message);
}
}

public static Future<Object> askOf(ActorRef supervisorCreator, Props props,
    Object message, Timeout timeout) {
    AskParam param = new AskParam(props, message, timeout);
    return Patterns.ask(supervisorCreator, param, timeout);
}

synchronized public static ActorRef createSupervisorCreator(
    ActorRefFactory factory) {
    return factory.actorOf(Props.create(AskSupervisorCreator.class));
}
}
}

```

In the `askOf` method the `SupervisorCreator` is sent the user message. The `SupervisorCreator` creates a `SupervisorActor` and forwards the message. This prevents the actor system from overloading due to actor creations. The `SupervisorActor` is responsible to create the user actor, forwards the message, handles actor termination and supervision. Additionally the `SupervisorActor` stops the user actor if execution time expired.

In case of an exception the supervisor tells the temporary actor which exception was thrown. Afterwards the actor hierarchy is stopped.

Finally we are able to execute an actor and receive the results or exceptions.

```

package docs.pattern;

import scala.concurrent.Await;
import scala.concurrent.Future;
import akka.actor.ActorRef;
import akka.actor.ActorRefFactory;
import akka.actor.Props;
import akka.actor.UntypedActor;
import akka.util.Timeout;

public class SupervisedAskSpec {

```

```
public Object execute(Class<? extends UntypedActor> someActor,
    Object message, Timeout timeout, ActorRefFactory actorSystem)
    throws Exception {
    // example usage
    try {
        ActorRef supervisorCreator = SupervisedAsk
            .createSupervisorCreator(actorSystem);
        Future<Object> finished = SupervisedAsk.askOf(supervisorCreator,
            Props.create(someActor), message, timeout);
        return Await.result(finished, timeout.duration());
    } catch (Exception e) {
        // exception propagated by supervision
        throw e;
    }
}
```

8.3 Template Pattern

Contributed by: N. N.

This is an especially nice pattern, since it does even come with some empty example code:

```
class JavaTemplate {
    public JavaTemplate() {
        System.out.println("Hello, Template!");
    }
    // uninteresting stuff ...
}
```

Note: Spread the word: this is the easiest way to get famous!

Please keep this pattern at the end of this file.

EXPERIMENTAL MODULES

The following modules of Akka are marked as experimental, which means that they are in early access mode, which also means that they are not covered by commercial support. The purpose of releasing them early, as experimental, is to make them easily available and improve based on feedback, or even discover that the module wasn't useful.

An experimental module doesn't have to obey the rule of staying binary compatible between micro releases. Breaking API changes may be introduced in minor releases without notice as we refine and simplify based on your feedback. An experimental module may be dropped in minor releases without prior deprecation.

9.1 Multi Node Testing

9.1.1 Multi Node Testing Concepts

When we talk about multi node testing in Akka we mean the process of running coordinated tests on multiple actor systems in different JVMs. The multi node testing kit consist of three main parts.


- [The Test Conductor](#). that coordinates and controls the nodes under test.
- [The Multi Node Spec](#). that is a convenience wrapper for starting the `TestConductor` and letting all nodes connect to it.
- [The SbtMultiJvm Plugin](#). that starts tests in multiple JVMs possibly on multiple machines.

9.1.2 The Test Conductor

The basis for the multi node testing is the `TestConductor`. It is an Akka Extension that plugs in to the network stack and it is used to coordinate the nodes participating in the test and provides several features including:

- **Node Address Lookup:** Finding out the full path to another test node (No need to share configuration between test nodes)
- **Node Barrier Coordination:** Waiting for other nodes at named barriers.
- **Network Failure Injection:** Throttling traffic, dropping packets, unplugging and plugging nodes back in.

This is a schematic overview of the test conductor.



The test conductor server is responsible for coordinating barriers and sending commands to the test conductor clients that act upon them, e.g. throttling network traffic to/from another client. More information on the possible operations is available in the `akka.remote.testconductor.Conductor` API documentation.

9.1.3 The Multi Node Spec

The Multi Node Spec consists of two parts. The `MultiNodeConfig` that is responsible for common configuration and enumerating and naming the nodes under test. The `MultiNodeSpec` that contains a number of convenience functions for making the test nodes interact with each other. More information on the possible operations is available in the `akka.remote.testkit.MultiNodeSpec` API documentation.

The setup of the `MultiNodeSpec` is configured through java system properties that you set on all JVMs that's going to run a node under test. These can easily be set on the JVM command line with `-Dproperty=value`.

These are the available properties:

- `multinode.max-nodes`
The maximum number of nodes that a test can have.
- `multinode.host`
The host name or IP for this node. Must be resolvable using `InetAddress.getByName`.
- `multinode.port`
The port number for this node. Defaults to 0 which will use a random port.
- `multinode.server-host`
The host name or IP for the server node. Must be resolvable using `InetAddress.getByName`.
- `multinode.server-port`
The port number for the server node. Defaults to 4711.
- `multinode.index`
The index of this node in the sequence of roles defined for the test. The index 0 is special and that machine will be the server. All failure injection and throttling must be done from this node.

9.1.4 The SbtMultiJvm Plugin

The *SbtMultiJvm Plugin* has been updated to be able to run multi node tests, by automatically generating the relevant `multinode.*` properties. This means that you can easily run multi node tests on a single machine without any special configuration by just running them as normal multi-jvm tests. These tests can then be run distributed over multiple machines without any changes simply by using the multi-node additions to the plugin.

Multi Node Specific Additions

The plugin also has a number of new `multi-node-*` sbt tasks and settings to support running tests on multiple machines. The necessary test classes and dependencies are packaged for distribution to other machines with [SbtAssembly](#) into a jar file with a name on the format `<projectName>_<scalaVersion>-<projectVersion>-multi-jvm-assembly.jar`

Note: To be able to distribute and kick off the tests on multiple machines, it is assumed that both host and target systems are POSIX like systems with `ssh` and `rsync` available.

These are the available sbt multi-node settings:

- `multiNodeHosts`
A sequence of hosts to use for running the test, on the form `user@host:java` where `host` is the only required part. Will override settings from file.
- `multiNodeHostsFileName`
A file to use for reading in the hosts to use for running the test. One per line on the same format as above. Defaults to `multi-node-test.hosts` in the base project directory.
- `multiNodeTargetDirName`
A name for the directory on the target machine, where to copy the jar file. Defaults to `multi-node-test` in the base directory of the ssh user used to rsync the jar file.
- `multiNodeJavaName`
The name of the default Java executable on the target machines. Defaults to `java`.

Here are some examples of how you define hosts:

- `localhost`
The current user on localhost using the default java.
- `user1@host1`
User `user1` on host `host1` with the default java.
- `user2@host2:/usr/lib/jvm/java-7-openjdk-amd64/bin/java`
User `user2` on host `host2` using java 7.
- `host3:/usr/lib/jvm/java-6-openjdk-amd64/bin/java`
The current user on host `host3` using java 6.

Running the Multi Node Tests

To run all the multi node test in multi-node mode (i.e. distributing the jar files and kicking off the tests remotely) from inside sbt, use the `multi-node-test` task:

```
multi-node-test
```

To run all of them in multi-jvm mode (i.e. all JVMs on the local machine) do:

```
multi-jvm:test
```

To run individual tests use the `multi-node-test-only` task:

```
multi-node-test-only your.MultiNodeTest
```

To run individual tests in the `multi-jvm` mode do:

```
multi-jvm:test-only your.MultiNodeTest
```

More than one test name can be listed to run multiple specific tests. Tab completion in `sbt` makes it easy to complete the test names.

9.1.5 Preparing Your Project for Multi Node Testing

The multi node testing kit is a separate jar file. Make sure that you have the following dependency in your project:

```
"com.typesafe.akka" %% "akka-multi-node-testkit" % "2.4.0"
```

If you are using the latest nightly build you should pick a timestamped Akka version from http://repo.akka.io/snapshots/com/typesafe/akka/akka-multi-node-testkit_2.11/. We recommend against using SNAPSHOT in order to obtain stable builds.

9.1.6 A Multi Node Testing Example

First we need some scaffolding to hook up the `MultiNodeSpec` with your favorite test framework. Lets define a trait `STMultiNodeSpec` that uses `ScalaTest` to start and stop `MultiNodeSpec`.

```
package sample.multinode

import org.scalatest.{ BeforeAndAfterAll, WordSpecLike }
import org.scalatest.Matchers
import akka.remote.testkit.MultiNodeSpecCallbacks

/**
 * Hooks up MultiNodeSpec with ScalaTest
 */
trait STMultiNodeSpec extends MultiNodeSpecCallbacks
  with WordSpecLike with Matchers with BeforeAndAfterAll {

  override def beforeAll() = multiNodeSpecBeforeAll()

  override def afterAll() = multiNodeSpecAfterAll()
}
```

Then we need to define a configuration. Lets use two nodes "node1" and "node2" and call it `MultiNodeSampleConfig`.

```
package sample.multinode
import akka.remote.testkit.MultiNodeConfig

object MultiNodeSampleConfig extends MultiNodeConfig {
  val node1 = role("node1")
  val node2 = role("node2")
}
```

And then finally to the node test code. That starts the two nodes, and demonstrates a barrier, and a remote actor message send/receive.

```
package sample.multinode
import akka.remote.testkit.MultiNodeSpec
```

```

import akka.testkit.ImplicitSender
import akka.actor.{ Props, Actor }

class MultiNodeSampleSpecMultiJvmNode1 extends MultiNodeSample
class MultiNodeSampleSpecMultiJvmNode2 extends MultiNodeSample

object MultiNodeSample {
  class Ponger extends Actor {
    def receive = {
      case "ping" => sender() ! "pong"
    }
  }
}

class MultiNodeSample extends MultiNodeSpec(MultiNodeSampleConfig)
  with STMultiNodeSpec with ImplicitSender {

  import MultiNodeSampleConfig._
  import MultiNodeSample._

  def initialParticipants = roles.size

  "A MultiNodeSample" must {

    "wait for all nodes to enter a barrier" in {
      enterBarrier("startup")
    }

    "send to and receive from a remote node" in {
      runOn(node1) {
        enterBarrier("deployed")
        val ponger = system.actorSelection(node(node2) / "user" / "ponger")
        ponger ! "ping"
        import scala.concurrent.duration._
        expectMsg(10.seconds, "pong")
      }

      runOn(node2) {
        system.actorOf(Props[Ponger], "ponger")
        enterBarrier("deployed")
      }

      enterBarrier("finished")
    }
  }
}

```

The easiest way to run this example yourself is to download [Typesafe Activator](#) and open the tutorial named [Akka Multi-Node Testing Sample with Scala](#).

9.1.7 Things to Keep in Mind

There are a couple of things to keep in mind when writing multi node tests or else your tests might behave in surprising ways.

- Don't issue a shutdown of the first node. The first node is the controller and if it shuts down your test will break.
- To be able to use `blackhole`, `passThrough`, and `throttle` you must activate the failure injector and throttler transport adapters by specifying `testTransport(on = true)` in your `MultiNodeConfig`.
- Throttling, shutdown and other failure injections can only be done from the first node, which again is the

controller.

- Don't ask for the address of a node using `node(address)` after the node has been shut down. Grab the address before shutting down the node.
- Don't use `MultiNodeSpec` methods like `address lookup`, `barrier entry` et.c. from other threads than the main test thread. This also means that you shouldn't use them from inside an actor, a future, or a scheduled task.

9.1.8 Configuration

There are several configuration properties for the Multi-Node Testing module, please refer to the *reference configuration*.

9.2 Actors (Java with Lambda Support)

The *Actor Model* provides a higher level of abstraction for writing concurrent and distributed systems. It alleviates the developer from having to deal with explicit locking and thread management, making it easier to write correct concurrent and parallel systems. Actors were defined in the 1973 paper by Carl Hewitt but have been popularized by the Erlang language, and used for example at Ericsson with great success to build highly concurrent and reliable telecom systems.

The API of Akka's Actors is similar to Scala Actors which has borrowed some of its syntax from Erlang.

Warning: The Java with lambda support part of Akka is marked as “**experimental**” as of its introduction in Akka 2.3.0. We will continue to improve this API based on our users' feedback, which implies that while we try to keep incompatible changes to a minimum, but the binary compatibility guarantee for maintenance releases does not apply to the `akka.actor.AbstractActor`, related classes and the `akka.japi.pf` package.

9.2.1 Creating Actors

Note: Since Akka enforces parental supervision every actor is supervised and (potentially) the supervisor of its children, it is advisable that you familiarize yourself with *Actor Systems* and *Supervision and Monitoring* and it may also help to read *Actor References, Paths and Addresses*.

Defining an Actor class

Actor classes are implemented by extending the `AbstractActor` class and setting the “initial behavior” in the constructor by calling the `receive` method in the `AbstractActor`.

The argument to the `receive` method is a `PartialFunction<Object, BoxedUnit>` that defines which messages your Actor can handle, along with the implementation of how the messages should be processed.

Don't let the type signature scare you. To allow you to easily build up a partial function there is a builder named `ReceiveBuilder` that you can use.

Here is an example:

```
import akka.actor.AbstractActor;
import akka.event.Logging;
import akka.event.LoggingAdapter;
import akka.japi.pf.ReceiveBuilder;

public class MyActor extends AbstractActor {
    private final LoggingAdapter log = Logging.getLogger(context().system(), this);
```

```

public MyActor() {
    receive(ReceiveBuilder.
        match(String.class, s -> {
            log.info("Received String message: {}", s);
        }).
        matchAny(o -> log.info("received unknown message")).build()
    );
}
}

```

Please note that the Akka Actor `receive` message loop is exhaustive, which is different compared to Erlang and the late Scala Actors. This means that you need to provide a pattern match for all messages that it can accept and if you want to be able to handle unknown messages then you need to have a default case as in the example above. Otherwise an `akka.actor.UnhandledMessage(message, sender, recipient)` will be published to the ActorSystem's `EventStream`.

Note further that the return type of the behavior defined above is `Unit`; if the actor shall reply to the received message then this must be done explicitly as explained below.

The argument to the `receive` method is a partial function object, which is stored within the actor as its “initial behavior”, see [Become/Unbecome](#) for further information on changing the behavior of an actor after its construction.

Props

`Props` is a configuration class to specify options for the creation of actors, think of it as an immutable and thus freely shareable recipe for creating an actor including associated deployment information (e.g. which dispatcher to use, see more below). Here are some examples of how to create a `Props` instance.

```

import akka.actor.Props;

Props props1 = Props.create(MyActor.class);
Props props2 = Props.create(ActorWithArgs.class,
    () -> new ActorWithArgs("arg")); // careful, see below
Props props3 = Props.create(ActorWithArgs.class, "arg");

```

The second variant shows how to pass constructor arguments to the `Actor` being created, but it should only be used outside of actors as explained below.

The last line shows a possibility to pass constructor arguments regardless of the context it is being used in. The presence of a matching constructor is verified during construction of the `Props` object, resulting in an `IllegalArgumentException` if no or multiple matching constructors are found.

Dangerous Variants

```

// NOT RECOMMENDED within another actor:
// encourages to close over enclosing class
Props props7 = Props.create(ActorWithArgs.class,
    () -> new ActorWithArgs("arg"));

```

This method is not recommended to be used within another actor because it encourages to close over the enclosing scope, resulting in non-serializable `Props` and possibly race conditions (breaking the actor encapsulation). On the other hand using this variant in a `Props` factory in the actor's companion object as documented under “Recommended Practices” below is completely fine.

There were two use-cases for these methods: passing constructor arguments to the actor—which is solved by the newly introduced `Props.create(clazz, args)` method above or the recommended practice below—and creating actors “on the spot” as anonymous classes. The latter should be solved by making these actors named classes instead (if they are not declared within a top-level object then the enclosing instance's `this` reference needs to be passed as the first argument).

Warning: Declaring one actor within another is very dangerous and breaks actor encapsulation. Never pass an actor's `this` reference into `Props`!

Recommended Practices

It is a good idea to provide factory methods on the companion object of each `Actor` which help keeping the creation of suitable `Props` as close to the actor definition as possible. This also avoids the pitfalls associated with using the `Props.create(...)` method which takes a by-name argument, since within a companion object the given code block will not retain a reference to its enclosing scope:

```
public class DemoActor extends AbstractActor {
    /**
     * Create Props for an actor of this type.
     * @param magicNumber The magic number to be passed to this actor's constructor.
     * @return a Props for creating this actor, which can then be further configured
     *         (e.g. calling .withDispatcher() on it)
     */
    static Props props(Integer magicNumber) {
        // You need to specify the actual type of the returned actor
        // since Java 8 lambdas have some runtime type information erased
        return Props.create(DemoActor.class, () -> new DemoActor(magicNumber));
    }

    private final Integer magicNumber;

    DemoActor(Integer magicNumber) {
        this.magicNumber = magicNumber;
        receive(ReceiveBuilder.
            match(Integer.class, i -> {
                sender().tell(i + magicNumber, self());
            }).build()
        );
    }
}

public class SomeOtherActor extends AbstractActor {
    // Props(new DemoActor(42)) would not be safe
    ActorRef demoActor = context().actorOf(DemoActor.props(42), "demo");
    // ...
}
```

Another good practice is to declare what messages an `Actor` can receive as close to the actor definition as possible (e.g. as static classes inside the `Actor` or using other suitable class), which makes it easier to know what it can receive.

```
public class DemoMessagesActor extends AbstractLoggingActor {

    static public class Greeting {
        private final String from;

        public Greeting(String from) {
            this.from = from;
        }

        public String getGreeter() {
            return from;
        }
    }

    DemoMessagesActor() {
        receive(ReceiveBuilder.

```

```

    match(Greeting.class, g -> {
        log().info("I was greeted by {}", g.getGreeter());
    }).build()
    );
};
}

```

Creating Actors with Props

Actors are created by passing a `Props` instance into the `actorOf` factory method which is available on `ActorSystem` and `ActorContext`.

```

import akka.actor.ActorRef;
import akka.actor.ActorSystem;

// ActorSystem is a heavy object: create only one per application
final ActorSystem system = ActorSystem.create("MySystem", config);
final ActorRef myActor = system.actorOf(Props.create(MyActor.class), "myactor");

```

Using the `ActorSystem` will create top-level actors, supervised by the actor system's provided guardian actor, while using an actor's context will create a child actor.

```

public class FirstActor extends AbstractActor {
    final ActorRef child = context().actorOf(Props.create(MyActor.class), "myChild");
    // plus some behavior ...
}

```

It is recommended to create a hierarchy of children, grand-children and so on such that it fits the logical failure-handling structure of the application, see [Actor Systems](#).

The call to `actorOf` returns an instance of `ActorRef`. This is a handle to the actor instance and the only way to interact with it. The `ActorRef` is immutable and has a one to one relationship with the Actor it represents. The `ActorRef` is also serializable and network-aware. This means that you can serialize it, send it over the wire and use it on a remote host and it will still be representing the same Actor on the original node, across the network.

The name parameter is optional, but you should preferably name your actors, since that is used in log messages and for identifying actors. The name must not be empty or start with `$`, but it may contain URL encoded characters (eg. `%20` for a blank space). If the given name is already in use by another child to the same parent an `InvalidActorNameException` is thrown.

Actors are automatically started asynchronously when created.

Dependency Injection

If your `UntypedActor` has a constructor that takes parameters then those need to be part of the `Props` as well, as described [above](#). But there are cases when a factory method must be used, for example when the actual constructor arguments are determined by a dependency injection framework.

```

import akka.actor.Actor;
import akka.actor.IndirectActorProducer;

class DependencyInjector implements IndirectActorProducer {
    final Object applicationContext;
    final String beanName;

    public DependencyInjector(Object applicationContext, String beanName) {
        this.applicationContext = applicationContext;
        this.beanName = beanName;
    }

    @Override

```

```

public Class<? extends Actor> actorClass() {
    return MyActor.class;
}

@Override
public MyActor produce() {
    MyActor result;
    // obtain fresh Actor instance from DI framework ...
    return result;
}
}

final ActorRef myActor = getContext().actorOf(
    Props.create(DependencyInjector.class, applicationContext, "MyActor"),
    "myactor3");

```

Warning: You might be tempted at times to offer an `IndirectActorProducer` which always returns the same instance, e.g. by using a static field. This is not supported, as it goes against the meaning of an actor restart, which is described here: [What Restarting Means](#).
When using a dependency injection framework, actor beans *MUST NOT* have singleton scope.

Techniques for dependency injection and integration with dependency injection frameworks are described in more depth in the [Using Akka with Dependency Injection](#) guideline and the [Akka Java Spring](#) tutorial in Typesafe Activator.

The Inbox

When writing code outside of actors which shall communicate with actors, the `ask` pattern can be a solution (see below), but there are two things it cannot do: receiving multiple replies (e.g. by subscribing an `ActorRef` to a notification service) and watching other actors' lifecycle. For these purposes there is the `Inbox` class:

```

final Inbox inbox = Inbox.create(system);
inbox.send(target, "hello");
try {
    assert inbox.receive(Duration.create(1, TimeUnit.SECONDS)).equals("world");
} catch (java.util.concurrent.TimeoutException e) {
    // timeout
}

```

The `send` method wraps a normal `tell` and supplies the internal actor's reference as the sender. This allows the reply to be received on the last line. Watching an actor is quite simple as well:

```

final Inbox inbox = Inbox.create(system);
inbox.watch(target);
target.tell(PoisonPill.getInstance(), ActorRef.noSender());
try {
    assert inbox.receive(Duration.create(1, TimeUnit.SECONDS)) instanceof Terminated;
} catch (java.util.concurrent.TimeoutException e) {
    // timeout
}

```

9.2.2 Actor API

The `AbstractActor` class defines a method called `receive`, that is used to set the "initial behavior" of the actor.

If the current actor behavior does not match a received message, `unhandled` is called, which by default publishes an `akka.actor.UnhandledMessage(message, sender, recipient)` on the actor system's event

stream (set configuration item `akka.actor.debug.unhandled` to `on` to have them converted into actual Debug messages).

In addition, it offers:

- `self` reference to the `ActorRef` of the actor
- `sender` reference `sender Actor` of the last received message, typically used as described in [Reply to messages](#)
- `supervisorStrategy` user overridable definition the strategy to use for supervising child actors

This strategy is typically declared inside the actor in order to have access to the actor's internal state within the decider function: since failure is communicated as a message sent to the supervisor and processed like other messages (albeit outside of the normal behavior), all values and variables within the actor are available, as is the `sender` reference (which will be the immediate child reporting the failure; if the original failure occurred within a distant descendant it is still reported one level up at a time).

- `context` exposes contextual information for the actor and the current message, such as:
 - factory methods to create child actors (`actorOf`)
 - system that the actor belongs to
 - parent supervisor
 - supervised children
 - lifecycle monitoring
 - hotswap behavior stack as described in [Become/Unbecome](#)

The remaining visible methods are user-overridable life-cycle hooks which are described in the following:

```
public void preStart() {
}


public void preRestart(Throwable reason, scala.Option<Object> message) {
    for (ActorRef each : getContext().getChildren()) {
        getContext().unwatch(each);
        getContext().stop(each);
    }
    postStop();
}

public void postRestart(Throwable reason) {
    preStart();
}

public void postStop() {
}
```

The implementations shown above are the defaults provided by the `AbstractActor` class.

Actor Lifecycle



A path in an actor system represents a “place” which might be occupied by a living actor. Initially (apart from system initialized actors) a path is empty. When `actorOf()` is called it assigns an *incarnation* of the actor described by the passed `Props` to the given path. An actor incarnation is identified by the path *and a UID*. A restart only swaps the `Actor` instance defined by the `Props` but the incarnation and hence the UID remains the same.

The lifecycle of an incarnation ends when the actor is stopped. At that point the appropriate lifecycle events are called and watching actors are notified of the termination. After the incarnation is stopped, the path can be reused again by creating an actor with `actorOf()`. In this case the name of the new incarnation will be the same as the previous one but the UIDs will differ.

An `ActorRef` always represents an incarnation (path and UID) not just a given path. Therefore if an actor is stopped and a new one with the same name is created an `ActorRef` of the old incarnation will not point to the new one.

`ActorSelection` on the other hand points to the path (or multiple paths if wildcards are used) and is completely oblivious to which incarnation is currently occupying it. `ActorSelection` cannot be watched for this reason. It is possible to resolve the current incarnation’s `ActorRef` living under the path by sending an `Identify` message to the `ActorSelection` which will be replied to with an `ActorIdentity` containing the correct reference (see *Identifying Actors via Actor Selection*). This can also be done with the `resolveOne` method of the `ActorSelection`, which returns a `Future` of the matching `ActorRef`.

Lifecycle Monitoring aka DeathWatch

In order to be notified when another actor terminates (i.e. stops permanently, not temporary failure and restart), an actor may register itself for reception of the `Terminated` message dispatched by the other actor upon termination (see *Stopping Actors*). This service is provided by the `DeathWatch` component of the actor system.

Registering a monitor is easy:

```
public class WatchActor extends AbstractActor {
    private final ActorRef child = context().actorOf(Props.empty(), "target");
    private ActorRef lastSender = system.deadLetters();

    public WatchActor() {
        context().watch(child); // <-- this is the only call needed for registration

        receive(ReceiveBuilder.
            matchEquals("kill", s -> {
                context().stop(child);
                lastSender = sender();
            }).
            match(Terminated.class, t -> t.actor().equals(child), t -> {
                lastSender.tell("finished", self());
            }).build()
        );
    }
}
```

It should be noted that the `Terminated` message is generated independent of the order in which registration and termination occur. In particular, the watching actor will receive a `Terminated` message even if the watched actor has already been terminated at the time of registration.

Registering multiple times does not necessarily lead to multiple messages being generated, but there is no guarantee that only exactly one such message is received: if termination of the watched actor has generated and queued the message, and another registration is done before this message has been processed, then a second message will be queued, because registering for monitoring of an already terminated actor leads to the immediate generation of the `Terminated` message.

It is also possible to deregister from watching another actor's liveliness using `context.unwatch(target)`. This works even if the `Terminated` message has already been enqueued in the mailbox; after calling `unwatch` no `Terminated` message for that actor will be processed anymore.

Start Hook

Right after starting the actor, its `preStart` method is invoked.

```
@Override
public void preStart() {
    target = context().actorOf(Props.create(MyActor.class, "target"));
}
```

This method is called when the actor is first created. During restarts it is called by the default implementation of `postRestart`, which means that by overriding that method you can choose whether the initialization code in this method is called only exactly once for this actor or for every restart. Initialization code which is part of the actor's constructor will always be called when an instance of the actor class is created, which happens at every restart.

Restart Hooks

All actors are supervised, i.e. linked to another actor with a fault handling strategy. Actors may be restarted in case an exception is thrown while processing a message (see *Supervision and Monitoring*). This restart involves the hooks mentioned above:

1. The old actor is informed by calling `preRestart` with the exception which caused the restart and the message which triggered that exception; the latter may be `None` if the restart was not caused by processing a message, e.g. when a supervisor does not trap the exception and is restarted in turn by its supervisor, or if an actor is restarted due to a sibling's failure. If the message is available, then that message's sender is also accessible in the usual way (i.e. by calling `sender`).

This method is the best place for cleaning up, preparing hand-over to the fresh actor instance, etc. By default it stops all children and calls `postStop`.

2. The initial factory from the `actorOf` call is used to produce the fresh instance.
3. The new actor's `postRestart` method is invoked with the exception which caused the restart. By default the `preStart` is called, just as in the normal start-up case.

An actor restart replaces only the actual actor object; the contents of the mailbox is unaffected by the restart, so processing of messages will resume after the `postRestart` hook returns. The message that triggered the exception will not be received again. Any message sent to an actor while it is being restarted will be queued to its mailbox as usual.

Warning: Be aware that the ordering of failure notifications relative to user messages is not deterministic. In particular, a parent might restart its child before it has processed the last messages sent by the child before the failure. See *Discussion: Message Ordering* for details.

Stop Hook

After stopping an actor, its `postStop` hook is called, which may be used e.g. for deregistering this actor from other services. This hook is guaranteed to run after message queuing has been disabled for this actor, i.e. messages sent to a stopped actor will be redirected to the `deadLetters` of the `ActorSystem`.

9.2.3 Identifying Actors via Actor Selection

As described in *Actor References, Paths and Addresses*, each actor has a unique logical path, which is obtained by following the chain of actors from child to parent until reaching the root of the actor system, and it has a physical path, which may differ if the supervision chain includes any remote supervisors. These paths are used by the system to look up actors, e.g. when a remote message is received and the recipient is searched, but they are also useful more directly: actors may look up other actors by specifying absolute or relative paths—logical or physical—and receive back an `ActorSelection` with the result:

```
// will look up this absolute path
context().actorSelection("/user/serviceA/actor");
// will look up sibling beneath same supervisor
context().actorSelection("../joe");
```

Note: It is always preferable to communicate with other Actors using their `ActorRef` instead of relying upon `ActorSelection`. Exceptions are

- sending messages using the *At-Least-Once Delivery* facility
- initiating first contact with a remote system

In all other cases `ActorRefs` can be provided during Actor creation or initialization, passing them from parent to child or introducing Actors by sending their `ActorRefs` to other Actors within messages.

The supplied path is parsed as a `java.net.URI`, which basically means that it is split on `/` into path elements. If the path starts with `/`, it is absolute and the look-up starts at the root guardian (which is the parent of `/user`); otherwise it starts at the current actor. If a path element equals `..`, the look-up will take a step “up” towards the supervisor of the currently traversed actor, otherwise it will step “down” to the named child. It should be noted that the `..` in actor paths here always means the logical structure, i.e. the supervisor.

The path elements of an actor selection may contain wildcard patterns allowing for broadcasting of messages to that section:

```
// will look all children to serviceB with names starting with worker
context().actorSelection("/user/serviceB/worker*");
// will look up all siblings beneath same supervisor
context().actorSelection("../*");
```

Messages can be sent via the `ActorSelection` and the path of the `ActorSelection` is looked up when delivering each message. If the selection does not match any actors the message will be dropped.

To acquire an `ActorRef` for an `ActorSelection` you need to send a message to the selection and use the `sender()` reference of the reply from the actor. There is a built-in `Identify` message that all Actors will understand and automatically reply to with a `ActorIdentity` message containing the `ActorRef`. This message is handled specially by the actors which are traversed in the sense that if a concrete name lookup fails (i.e. a non-wildcard path element does not correspond to a live actor) then a negative result is generated. Please note that this does not mean that delivery of that reply is guaranteed, it still is a normal message.

```
import akka.actor.ActorIdentity;
import akka.actor.ActorSelection;
import akka.actor.Identify;

public class Follower extends AbstractActor {
    final Integer identifyId = 1;

    public Follower() {
        ActorSelection selection = context().actorSelection("/user/another");
        selection.tell(new Identify(identifyId), self());

        receive(ReceiveBuilder.
            match(ActorIdentity.class, id -> id.getRef() != null, id -> {
                ActorRef ref = id.getRef();
                context().watch(ref);
                context().become(active(ref));
            }).
            match(ActorIdentity.class, id -> id.getRef() == null, id -> {
                context().stop(self());
            }).build()
        );
    }

    final PartialFunction<Object, BoxedUnit> active(final ActorRef another) {
        return ReceiveBuilder.
            match(Terminated.class, t -> t.actor().equals(another), t -> {
                context().stop(self());
            }).build();
    }
}
```

You can also acquire an `ActorRef` for an `ActorSelection` with the `resolveOne` method of the `ActorSelection`. It returns a `Future` of the matching `ActorRef` if such an actor exists. It is completed with failure `[[akka.actor.ActorNotFound]]` if no such actor exists or the identification didn't complete within the supplied *timeout*.

Remote actor addresses may also be looked up, if *remoting* is enabled:

```
context().actorSelection("akka.tcp://app@otherhost:1234/user/serviceB");
```

An example demonstrating actor look-up is given in *Remoting Sample*.

9.2.4 Messages and immutability

IMPORTANT: Messages can be any kind of object but have to be immutable. Akka can't enforce immutability (yet) so this has to be by convention.

Here is an example of an immutable message:

```
public class ImmutableMessage {
    private final int sequenceNumber;
    private final List<String> values;
```

```

public ImmutableMessage(int sequenceNumber, List<String> values) {
    this.sequenceNumber = sequenceNumber;
    this.values = Collections.unmodifiableList(new ArrayList<String>(values));
}

public int getSequenceNumber() {
    return sequenceNumber;
}

public List<String> getValues() {
    return values;
}
}

```

9.2.5 Send messages

Messages are sent to an Actor through one of the following methods.

- `tell` means “fire-and-forget”, e.g. send a message asynchronously and return immediately.
- `ask` sends a message asynchronously and returns a `Future` representing a possible reply.

Message ordering is guaranteed on a per-sender basis.

Note: There are performance implications of using `ask` since something needs to keep track of when it times out, there needs to be something that bridges a `Promise` into an `ActorRef` and it also needs to be reachable through remoting. So always prefer `tell` for performance, and only `ask` if you must.

In all these methods you have the option of passing along your own `ActorRef`. Make it a practice of doing so because it will allow the receiver actors to be able to respond to your message, since the sender reference is sent along with the message.

Tell: Fire-forget

This is the preferred way of sending messages. No blocking waiting for a message. This gives the best concurrency and scalability characteristics.

```

// don't forget to think about who is the sender (2nd argument)
target.tell(message, self());

```

The sender reference is passed along with the message and available within the receiving actor via its `sender` method while processing this message. Inside of an actor it is usually `self` who shall be the sender, but there can be cases where replies shall be routed to some other actor—e.g. the parent—in which the second argument to `tell` would be a different one. Outside of an actor and if no reply is needed the second argument can be `null`; if a reply is needed outside of an actor you can use the `ask`-pattern described next..

Ask: Send-And-Receive-Future

The `ask` pattern involves actors as well as futures, hence it is offered as a use pattern rather than a method on `ActorRef`:

```

import static akka.pattern.Patterns.ask;
import static akka.pattern.Patterns.pipe;
import scala.concurrent.Future;
import scala.concurrent.duration.Duration;
import akka.dispatch.Futures;
import akka.dispatch.Mapper;
import akka.util.Timeout;

```

```

final Timeout t = new Timeout(Duration.create(5, TimeUnit.SECONDS));

final ArrayList<Future<Object>> futures = new ArrayList<Future<Object>>();
futures.add(ask(actorA, "request", 1000)); // using 1000ms timeout
futures.add(ask(actorB, "another request", t)); // using timeout from
// above

final Future<Iterable<Object>> aggregate = Futures.sequence(futures,
    system.dispatcher());

final Future<Result> transformed = aggregate.map(
    new Mapper<Iterable<Object>, Result>() {
        public Result apply(Iterable<Object> coll) {
            final Iterator<Object> it = coll.iterator();
            final String x = (String) it.next();
            final String s = (String) it.next();
            return new Result(x, s);
        }
    }, system.dispatcher());

pipe(transformed, system.dispatcher()).to(actorC);

```

This example demonstrates `ask` together with the `pipe` pattern on futures, because this is likely to be a common combination. Please note that all of the above is completely non-blocking and asynchronous: `ask` produces a `Future`, two of which are composed into a new future using the `Futures.sequence` and `map` methods and then `pipe` installs an `onComplete`-handler on the future to effect the submission of the aggregated `Result` to another actor.

Using `ask` will send a message to the receiving Actor as with `tell`, and the receiving actor must reply with `sender().tell(reply, self())` in order to complete the returned `Future` with a value. The `ask` operation involves creating an internal actor for handling this reply, which needs to have a timeout after which it is destroyed in order not to leak resources; see more below.

Warning: To complete the future with an exception you need send a `Failure` message to the sender. This is *not done automatically* when an actor throws an exception while processing a message.

```

try {
    String result = operation();
    sender().tell(result, self());
} catch (Exception e) {
    sender().tell(new akka.actor.Status.Failure(e), self());
    throw e;
}

```

If the actor does not complete the future, it will expire after the timeout period, specified as parameter to the `ask` method; this will complete the `Future` with an `AskTimeoutException`.

See [Futures](#) for more information on how to await or query a future.

The `onComplete`, `onSuccess`, or `onFailure` methods of the `Future` can be used to register a callback to get a notification when the `Future` completes. Gives you a way to avoid blocking.

Warning: When using future callbacks, inside actors you need to carefully avoid closing over the containing actor's reference, i.e. do not call methods or access mutable state on the enclosing actor from within the callback. This would break the actor encapsulation and may introduce synchronization bugs and race conditions because the callback will be scheduled concurrently to the enclosing actor. Unfortunately there is not yet a way to detect these illegal accesses at compile time. See also: [Actors and shared mutable state](#)

Forward message

You can forward a message from one actor to another. This means that the original sender address/reference is maintained even though the message is going through a ‘mediator’. This can be useful when writing actors that work as routers, load-balancers, replicators etc.

```
target.forward(result, context());
```

9.2.6 Receive messages

An Actor either has to set its initial receive behavior in the constructor by calling the `receive` method in the `AbstractActor`:

```
public SomeActor() {
    receive(ReceiveBuilder.
        // and some behavior ...
        build());
}
```

or by implementing the `receive` method in the `Actor` interface:

```
public abstract PartialFunction<Object, BoxedUnit> receive();
```

Both the argument to the `AbstractActor` `receive` method and the return type of the `Actor` `receive` method is a `PartialFunction<Object, BoxedUnit>` that defines which messages your Actor can handle, along with the implementation of how the messages should be processed.

Don't let the type signature scare you. To allow you to easily build up a partial function there is a builder named `ReceiveBuilder` that you can use.

Here is an example:

```
import akka.actor.AbstractActor;
import akka.event.Logging;
import akka.event.LoggingAdapter;
import akka.japi.pf.ReceiveBuilder;

public class MyActor extends AbstractActor {
    private final LoggingAdapter log = Logging.getLogger(context().system(), this);

    public MyActor() {
        receive(ReceiveBuilder.
            match(String.class, s -> {
                log.info("Received String message: {}", s);
            }).
            matchAny(o -> log.info("received unknown message")).build()
        );
    }
}
```

9.2.7 Reply to messages

If you want to have a handle for replying to a message, you can use `sender()`, which gives you an `ActorRef`. You can reply by sending to that `ActorRef` with `sender().tell(replyMsg, self())`. You can also store the `ActorRef` for replying later, or passing on to other actors. If there is no sender (a message was sent without an actor or future context) then the sender defaults to a ‘dead-letter’ actor ref.

```
sender().tell(s, self());
```

9.2.8 Receive timeout

The `ActorContext` `setReceiveTimeout` defines the inactivity timeout after which the sending of a `ReceiveTimeout` message is triggered. When specified, the receive function should be able to handle an `akka.actor.ReceiveTimeout` message. 1 millisecond is the minimum supported timeout.

Please note that the receive timeout might fire and enqueue the `ReceiveTimeout` message right after another message was enqueued; hence it is **not guaranteed** that upon reception of the receive timeout there must have been an idle period beforehand as configured via this method.

Once set, the receive timeout stays in effect (i.e. continues firing repeatedly after inactivity periods). Pass in `Duration.Undefined` to switch off this feature.

```
public class ReceiveTimeoutActor extends AbstractActor {
  public ReceiveTimeoutActor() {
    // To set an initial delay
    context().setReceiveTimeout(Duration.create("10 seconds"));

    receive(ReceiveBuilder.
      matchEquals("Hello", s -> {
        // To set in a response to a message
        context().setReceiveTimeout(Duration.create("1 second"));
      }).
      match(ReceiveTimeout.class, r -> {
        // To turn it off
        context().setReceiveTimeout(Duration.Undefined());
      }).build()
    );
  }
}
```

Messages marked with `NotInfluenceReceiveTimeout` will not reset the timer. This can be useful when `ReceiveTimeout` should be fired by external inactivity but not influenced by internal activity, e.g. scheduled tick messages.

9.2.9 Stopping actors

Actors are stopped by invoking the `stop` method of a `ActorRefFactory`, i.e. `ActorContext` or `ActorSystem`. Typically the context is used for stopping child actors and the system for stopping top level actors. The actual termination of the actor is performed asynchronously, i.e. `stop` may return before the actor is stopped.

Processing of the current message, if any, will continue before the actor is stopped, but additional messages in the mailbox will not be processed. By default these messages are sent to the `deadLetters` of the `ActorSystem`, but that depends on the mailbox implementation.

Termination of an actor proceeds in two steps: first the actor suspends its mailbox processing and sends a stop command to all its children, then it keeps processing the internal termination notifications from its children until the last one is gone, finally terminating itself (invoking `postStop`, dumping mailbox, publishing `Terminated` on the `DeathWatch`, telling its supervisor). This procedure ensures that actor system sub-trees terminate in an orderly fashion, propagating the stop command to the leaves and collecting their confirmation back to the stopped supervisor. If one of the actors does not respond (i.e. processing a message for extended periods of time and therefore not receiving the stop command), this whole process will be stuck.

Upon `ActorSystem.terminate`, the system guardian actors will be stopped, and the aforementioned process will ensure proper termination of the whole system.

The `postStop` hook is invoked after an actor is fully stopped. This enables cleaning up of resources:

```
@Override
public void postStop() {
  // clean up some resources ...
}
```

Note: Since stopping an actor is asynchronous, you cannot immediately reuse the name of the child you just stopped; this will result in an `InvalidActorNameException`. Instead, watch the terminating actor and create its replacement in response to the `Terminated` message which will eventually arrive.

PoisonPill

You can also send an actor the `akka.actor.PoisonPill` message, which will stop the actor when the message is processed. `PoisonPill` is enqueued as ordinary messages and will be handled after messages that were already queued in the mailbox.

Graceful Stop

`gracefulStop` is useful if you need to wait for termination or compose ordered termination of several actors:

```
import static akka.pattern.Patterns.gracefulStop;
import scala.concurrent.Await;
import scala.concurrent.Future;
import scala.concurrent.duration.Duration;
import akka.pattern.AskTimeoutException;
```

```
try {
    Future<Boolean> stopped =
        gracefulStop(actorRef, Duration.create(5, TimeUnit.SECONDS), Manager.SHUTDOWN);
    Await.result(stopped, Duration.create(6, TimeUnit.SECONDS));
    // the actor has been stopped
} catch (AskTimeoutException e) {
    // the actor wasn't stopped within 5 seconds
}
```

```
public class Manager extends AbstractActor {
    private static enum Shutdown {
        Shutdown
    }
    public static final Shutdown SHUTDOWN = Shutdown.Shutdown;

    private ActorRef worker =
        context().watch(context().actorOf(Props.create(Cruncher.class), "worker"));

    public Manager() {
        receive(ReceiveBuilder.
            matchEquals("job", s -> {
                worker.tell("crunch", self());
            }).
            matchEquals(SHUTDOWN, x -> {
                worker.tell(PoisonPill.getInstance(), self());
                context().become(shuttingDown);
            }).build()
        );
    }

    public PartialFunction<Object, BoxedUnit> shuttingDown =
        ReceiveBuilder.
            matchEquals("job", s -> {
                sender().tell("service unavailable, shutting down", self());
            }).
            match(Terminated.class, t -> t.actor().equals(worker), t -> {
                context().stop(self());
            }).build();
}
```

When `gracefulStop()` returns successfully, the actor's `postStop()` hook will have been executed: there exists a happens-before edge between the end of `postStop()` and the return of `gracefulStop()`.

In the above example a custom `Manager.Shutdown` message is sent to the target actor to initiate the process of stopping the actor. You can use `PoisonPill` for this, but then you have limited possibilities to perform interactions with other actors before stopping the target actor. Simple cleanup tasks can be handled in `postStop`.

Warning: Keep in mind that an actor stopping and its name being deregistered are separate events which happen asynchronously from each other. Therefore it may be that you will find the name still in use after `gracefulStop()` returned. In order to guarantee proper deregistration, only reuse names from within a supervisor you control and only in response to a `Terminated` message, i.e. not for top-level actors.

9.2.10 Become/Unbecome

Upgrade

Akka supports hotswapping the Actor's message loop (e.g. its implementation) at runtime: invoke the `context.become` method from within the Actor. `become` takes a `PartialFunction<Object, BoxedUnit>` that implements the new message handler. The hotswapped code is kept in a `Stack` which can be pushed and popped.

Warning: Please note that the actor will revert to its original behavior when restarted by its Supervisor.

To hotswap the Actor behavior using `become`:

```
public class HotSwapActor extends AbstractActor {
    private PartialFunction<Object, BoxedUnit> angry;
    private PartialFunction<Object, BoxedUnit> happy;

    public HotSwapActor() {
        angry =
            ReceiveBuilder.
                matchEquals("foo", s -> {
                    sender().tell("I am already angry?", self());
                }).
                matchEquals("bar", s -> {
                    context().become(happy);
                }).build();

        happy = ReceiveBuilder.
            matchEquals("bar", s -> {
                sender().tell("I am already happy :-)", self());
            }).
            matchEquals("foo", s -> {
                context().become(angry);
            }).build();

        receive(ReceiveBuilder.
            matchEquals("foo", s -> {
                context().become(angry);
            }).
            matchEquals("bar", s -> {
                context().become(happy);
            }).build()
        );
    }
}
```

This variant of the `become` method is useful for many different things, such as to implement a Finite State Machine (FSM, for an example see [Dining Hakkers](#)). It will replace the current behavior (i.e. the top of the behavior stack), which means that you do not use `unbecome`, instead always the next behavior is explicitly installed.

The other way of using `become` does not replace but add to the top of the behavior stack. In this case care must be taken to ensure that the number of “pop” operations (i.e. `unbecome`) matches the number of “push” ones in the long run, otherwise this amounts to a memory leak (which is why this behavior is not the default).

```
public class Swapper extends AbstractLoggingActor {
  public Swapper() {
    receive(ReceiveBuilder.
      matchEquals(Swap, s -> {
        log().info("Hi");
        context().become(ReceiveBuilder.
          matchEquals(Swap, x -> {
            log().info("Ho");
            context().unbecome(); // resets the latest 'become' (just for fun)
          }).build(), false); // push on top instead of replace
        }).build()
    );
  }
}

public class SwapperApp {
  public static void main(String[] args) {
    ActorSystem system = ActorSystem.create("SwapperSystem");
    ActorRef swapper = system.actorOf(Props.create(Swapper.class), "swapper");
    swapper.tell(Swap, ActorRef.noSender()); // logs Hi
    swapper.tell(Swap, ActorRef.noSender()); // logs Ho
    swapper.tell(Swap, ActorRef.noSender()); // logs Hi
    swapper.tell(Swap, ActorRef.noSender()); // logs Ho
    swapper.tell(Swap, ActorRef.noSender()); // logs Hi
    swapper.tell(Swap, ActorRef.noSender()); // logs Ho
    system.terminate();
  }
}
```

9.2.11 Stash

The `AbstractActorWithStash` class enables an actor to temporarily stash away messages that can not or should not be handled using the actor’s current behavior. Upon changing the actor’s message handler, i.e., right before invoking `context().become()` or `context().unbecome()`, all stashed messages can be “unstashed”, thereby prepending them to the actor’s mailbox. This way, the stashed messages can be processed in the same order as they have been received originally. An actor that extends `AbstractActorWithStash` will automatically get a deque-based mailbox.

Note: The abstract class `AbstractActorWithStash` implements the marker interface `RequiresMessageQueue<DequeBasedMessageQueueSemantics>` which requests the system to automatically choose a deque based mailbox implementation for the actor. If you want more control over the mailbox, see the documentation on mailboxes: [Mailboxes](#).

Here is an example of the `AbstractActorWithStash` class in action:

```
public class ActorWithProtocol extends AbstractActorWithStash {
  public ActorWithProtocol() {
    receive(ReceiveBuilder.
      matchEquals("open", s -> {
        context().become(ReceiveBuilder.
          matchEquals("write", ws -> { /* do writing */ }).

```

```

    matchEquals("close", cs -> {
        unstashAll();
        context().unbecome();
    }).
    matchAny(msg -> stash().build(), false);
  }).
  matchAny(msg -> stash().build()
  );
}
}

```

Invoking `stash()` adds the current message (the message that the actor received last) to the actor's stash. It is typically invoked when handling the default case in the actor's message handler to stash messages that aren't handled by the other cases. It is illegal to stash the same message twice; to do so results in an `IllegalStateException` being thrown. The stash may also be bounded in which case invoking `stash()` may lead to a capacity violation, which results in a `StashOverflowException`. The capacity of the stash can be configured using the `stash-capacity` setting (an `Int`) of the mailbox's configuration.

Invoking `unstashAll()` enqueues messages from the stash to the actor's mailbox until the capacity of the mailbox (if any) has been reached (note that messages from the stash are prepended to the mailbox). In case a bounded mailbox overflows, a `MessageQueueAppendFailedException` is thrown. The stash is guaranteed to be empty after calling `unstashAll()`.

The stash is backed by a `scala.collection.immutable.Vector`. As a result, even a very large number of messages may be stashed without a major impact on performance.

Note that the stash is part of the ephemeral actor state, unlike the mailbox. Therefore, it should be managed like other parts of the actor's state which have the same property. The `AbstractActorWithStash` implementation of `preRestart` will call `unstashAll()`, which is usually the desired behavior.

Note: If you want to enforce that your actor can only work with an unbounded stash, then you should use the `AbstractActorWithUnboundedStash` class instead.

9.2.12 Killing an Actor

You can kill an actor by sending a `Kill` message. This will cause the actor to throw a `ActorKilledException`, triggering a failure. The actor will suspend operation and its supervisor will be asked how to handle the failure, which may mean resuming the actor, restarting it or terminating it completely. See *What Supervision Means* for more information.

Use `Kill` like this:

```
victim.tell(akka.actor.Kill.getInstance(), ActorRef.noSender());
```

9.2.13 Actors and exceptions

It can happen that while a message is being processed by an actor, that some kind of exception is thrown, e.g. a database exception.

What happens to the Message

If an exception is thrown while a message is being processed (i.e. taken out of its mailbox and handed over to the current behavior), then this message will be lost. It is important to understand that it is not put back on the mailbox. So if you want to retry processing of a message, you need to deal with it yourself by catching the exception and retry your flow. Make sure that you put a bound on the number of retries since you don't want a system to livelock (so consuming a lot of cpu cycles without making progress). Another possibility would be to have a look at the *PeekMailbox pattern*.

What happens to the mailbox

If an exception is thrown while a message is being processed, nothing happens to the mailbox. If the actor is restarted, the same mailbox will be there. So all messages on that mailbox will be there as well.

What happens to the actor

If code within an actor throws an exception, that actor is suspended and the supervision process is started (see *Supervision and Monitoring*). Depending on the supervisor's decision the actor is resumed (as if nothing happened), restarted (wiping out its internal state and starting from scratch) or terminated.

9.2.14 Initialization patterns

The rich lifecycle hooks of Actors provide a useful toolkit to implement various initialization patterns. During the lifetime of an `ActorRef`, an actor can potentially go through several restarts, where the old instance is replaced by a fresh one, invisibly to the outside observer who only sees the `ActorRef`.

One may think about the new instances as “incarnations”. Initialization might be necessary for every incarnation of an actor, but sometimes one needs initialization to happen only at the birth of the first instance when the `ActorRef` is created. The following sections provide patterns for different initialization needs.

Initialization via constructor

Using the constructor for initialization has various benefits. First of all, it makes it possible to use `val` fields to store any state that does not change during the life of the actor instance, making the implementation of the actor more robust. The constructor is invoked for every incarnation of the actor, therefore the internals of the actor can always assume that proper initialization happened. This is also the drawback of this approach, as there are cases when one would like to avoid reinitializing internals on restart. For example, it is often useful to preserve child actors across restarts. The following section provides a pattern for this case.

Initialization via `preStart`

The method `preStart()` of an actor is only called once directly during the initialization of the first instance, that is, at creation of its `ActorRef`. In the case of restarts, `preStart()` is called from `postRestart()`, therefore if not overridden, `preStart()` is called on every incarnation. However, overriding `postRestart()` one can disable this behavior, and ensure that there is only one call to `preStart()`.

One useful usage of this pattern is to disable creation of new `ActorRefs` for children during restarts. This can be achieved by overriding `preRestart()`:

```
@Override
public void preStart() {
    // Initialize children here
}

// Overriding postRestart to disable the call to preStart()
// after restarts
@Override
public void postRestart(Throwable reason) {
}

// The default implementation of preRestart() stops all the children
// of the actor. To opt-out from stopping the children, we
// have to override preRestart()
@Override
public void preRestart(Throwable reason, Option<Object> message)
    throws Exception {
    // Keep the call to postStop(), but no stopping of children
}
```

```
postStop();
}
```

Please note, that the child actors are *still restarted*, but no new `ActorRef` is created. One can recursively apply the same principles for the children, ensuring that their `preStart()` method is called only at the creation of their refs.

For more information see [What Restarting Means](#).

Initialization via message passing

There are cases when it is impossible to pass all the information needed for actor initialization in the constructor, for example in the presence of circular dependencies. In this case the actor should listen for an initialization message, and use `become()` or a finite state-machine state transition to encode the initialized and uninitialized states of the actor.

```
receive(ReceiveBuilder.
  matchEquals("init", m1 -> {
    initializeMe = "Up and running";
    context().become(ReceiveBuilder.
      matchEquals("U OK?", m2 -> {
        sender().tell(initializeMe, self());
      }).build());
  }).build())
```

If the actor may receive messages before it has been initialized, a useful tool can be the `Stash` to save messages until the initialization finishes, and replaying them after the actor became initialized.

Warning: This pattern should be used with care, and applied only when none of the patterns above are applicable. One of the potential issues is that messages might be lost when sent to remote actors. Also, publishing an `ActorRef` in an uninitialized state might lead to the condition that it receives a user message before the initialization has been done.

9.2.15 Lambdas and Performance

There is one big difference between the optimized partial functions created by the Scala compiler and the ones created by the `ReceiveBuilder`. The partial functions created by the `ReceiveBuilder` consist of multiple lambda expressions for every match statement, where each lambda is an object referencing the code to be run. This is something that the JVM can have problems optimizing and the resulting code might not be as performant as the Scala equivalent or the corresponding *untyped actor* version.

9.3 FSM (Java with Lambda Support)

9.3.1 Overview

The FSM (Finite State Machine) is available as an abstract base class that implements an akka `Actor` and is best described in the [Erlang design principles](#)

A FSM can be described as a set of relations of the form:

State(S) x Event(E) -> Actions (A), State(S')

These relations are interpreted as meaning:

If we are in state S and the event E occurs, we should perform the actions A and make a transition to the state S'.

Warning: The Java with lambda support part of Akka is marked as “**experimental**” as of its introduction in Akka 2.3.0. We will continue to improve this API based on our users’ feedback, which implies that while we try to keep incompatible changes to a minimum, but the binary compatibility guarantee for maintenance releases does not apply to the `akka.actor.AbstractFSM`, related classes and the `akka.japi.pf` package.

9.3.2 A Simple Example

To demonstrate most of the features of the `AbstractFSM` class, consider an actor which shall receive and queue messages while they arrive in a burst and send them on after the burst ended or a flush request is received.

First, consider all of the below to use these import statements:

```
import akka.actor.AbstractFSM;
import akka.actor.ActorRef;
import akka.japi.pf.UnitMatch;
import java.util.Arrays;
import java.util.LinkedList;
import java.util.List;
import scala.concurrent.duration.Duration;
```

The contract of our “Buncher” actor is that it accepts or produces the following messages:

```
public final class SetTarget {
    private final ActorRef ref;

    public SetTarget(ActorRef ref) {
        this.ref = ref;
    }

    public ActorRef getRef() {
        return ref;
    }
    // boilerplate ...
}

public final class Queue {
    private final Object obj;

    public Queue(Object obj) {
        this.obj = obj;
    }

    public Object getObj() {
        return obj;
    }
    // boilerplate ...
}

public final class Batch {
    private final List<Object> list;

    public Batch(List<Object> list) {
        this.list = list;
    }

    public List<Object> getList() {
        return list;
    }
    // boilerplate ...
}
```

```
public enum Flush {
    Flush
}
```

SetTarget is needed for starting it up, setting the destination for the Batches to be passed on; Queue will add to the internal queue while Flush will mark the end of a burst.

The actor can be in two states: no message queued (aka Idle) or some message queued (aka Active). The states and the state data is defined like this:

```
// states
enum State {
    Idle, Active
}

// state data
interface Data {
}

enum Uninitialized implements Data {
    Uninitialized
}

final class Todo implements Data {
    private final ActorRef target;
    private final List<Object> queue;

    public Todo(ActorRef target, List<Object> queue) {
        this.target = target;
        this.queue = queue;
    }

    public ActorRef getTarget() {
        return target;
    }

    public List<Object> getQueue() {
        return queue;
    }
    // boilerplate ...
}
```

The actor starts out in the idle state. Once a message arrives it will go to the active state and stay there as long as messages keep arriving and no flush is requested. The internal state data of the actor is made up of the target actor reference to send the batches to and the actual queue of messages.

Now let's take a look at the skeleton for our FSM actor:

```
public class Buncher extends AbstractFSM<State, Data> {
    {
        startWith(Idle, Uninitialized);

        when(Idle,
            matchEvent(SetTarget.class, Uninitialized.class,
                (setTarget, uninitialized) ->
                    stay().using(new Todo(setTarget.getRef(), new LinkedList<>()))));

        // transition elided ...

        when(Active, Duration.create(1, "second"),
            matchEvent(Arrays.asList(Flush.class, StateTimeout()), Todo.class,
                (event, todo) -> goTo(Idle).using(todo.copy(new LinkedList<>()))));

        // unhandled elided ...
    }
}
```

```

    initialize();
  }
}

```

The basic strategy is to declare the actor, by inheriting the `AbstractFSM` class and specifying the possible states and data values as type parameters. Within the body of the actor a DSL is used for declaring the state machine:

- `startWith` defines the initial state and initial data
- then there is one `when(<state>) { ... }` declaration per state to be handled (could potentially be multiple ones, the passed `PartialFunction` will be concatenated using `orElse`)
- finally starting it up using `initialize`, which performs the transition into the initial state and sets up timers (if required).

In this case, we start out in the `Idle` and `Uninitialized` state, where only the `SetTarget()` message is handled; `stay` prepares to end this event's processing for not leaving the current state, while the `using` modifier makes the FSM replace the internal state (which is `Uninitialized` at this point) with a fresh `Todo()` object containing the target actor reference. The `Active` state has a state timeout declared, which means that if no message is received for 1 second, a `FSM.StateTimeout` message will be generated. This has the same effect as receiving the `Flush` command in this case, namely to transition back into the `Idle` state and resetting the internal queue to the empty vector. But how do messages get queued? Since this shall work identically in both states, we make use of the fact that any event which is not handled by the `when()` block is passed to the `whenUnhandled()` block:

```

whenUnhandled(
  matchEvent(Queue.class, Todo.class,
    (queue, todo) -> goTo(Active).using(todo.addElement(queue.getObj()))).
  anyEvent((event, state) -> {
    log().warning("received unhandled request {} in state {}/{}",
      event, stateName(), state);
    return stay();
  }));

```

The first case handled here is adding `Queue()` requests to the internal queue and going to the `Active` state (this does the obvious thing of staying in the `Active` state if already there), but only if the FSM data are not `Uninitialized` when the `Queue()` event is received. Otherwise—and in all other non-handled cases—the second case just logs a warning and does not change the internal state.

The only missing piece is where the `Batches` are actually sent to the target, for which we use the `onTransition` mechanism: you can declare multiple such blocks and all of them will be tried for matching behavior in case a state transition occurs (i.e. only when the state actually changes).

```

onTransition(
  matchState(Active, Idle, () -> {
    // reuse this matcher
    final UnitMatch<Data> m = UnitMatch.create(
      matchData(Todo.class,
        todo -> todo.getTarget().tell(new Batch(todo.getQueue()), self()));
    m.match(stateData());
  }).
  state(Idle, Active, () -> { /* Do something here */ });

```

The transition callback is a partial function which takes as input a pair of states—the current and the next state. During the state change, the old state data is available via `stateData` as shown, and the new state data would be available as `nextStateData`.

To verify that this buncher actually works, it is quite easy to write a test using the *akka-testkit*, here using JUnit as an example:

```

public class BuncherTest {

    static ActorSystem system;

```

```

@BeforeClass
public static void setup() {
    system = ActorSystem.create("BuncherTest");
}

@AfterClass
public static void tearDown() {
    JavaTestKit.shutdownActorSystem(system);
    system = null;
}

@Test
public void testBuncherActorBatchesCorrectly() {
    new JavaTestKit(system) {{
        final ActorRef buncher =
            system.actorOf(Props.create(Buncher.class));
        final ActorRef probe = getRef();

        buncher.tell(new SetTarget(probe), probe);
        buncher.tell(new Queue(42), probe);
        buncher.tell(new Queue(43), probe);
        LinkedList<Object> list1 = new LinkedList<>();
        list1.add(42);
        list1.add(43);
        expectMsgEquals(new Batch(list1));
        buncher.tell(new Queue(44), probe);
        buncher.tell(Flush, probe);
        buncher.tell(new Queue(45), probe);
        LinkedList<Object> list2 = new LinkedList<>();
        list2.add(44);
        expectMsgEquals(new Batch(list2));
        LinkedList<Object> list3 = new LinkedList<>();
        list3.add(45);
        expectMsgEquals(new Batch(list3));
        system.stop(buncher);
    }};
}

@Test
public void testBuncherActorDoesntBatchUninitialized() {
    new JavaTestKit(system) {{
        final ActorRef buncher =
            system.actorOf(Props.create(Buncher.class));
        final ActorRef probe = getRef();

        buncher.tell(new Queue(42), probe);
        expectNoMsg();
        system.stop(buncher);
    }};
}
}

```

9.3.3 Reference

The AbstractFSM Class

The `AbstractFSM` abstract class is the base class used to implement an FSM. It implements `Actor` since an `Actor` is created to drive the FSM.

```

public class Buncher extends AbstractFSM<State, Data> {
    {

```

```
// fsm body ...
}
}
```

Note: The `AbstractFSM` class defines a `receive` method which handles internal messages and passes everything else through to the FSM logic (according to the current state). When overriding the `receive` method, keep in mind that e.g. state timeout handling depends on actually passing the messages through the FSM logic.

The `AbstractFSM` class takes two type parameters:

1. the supertype of all state names, usually an enum,
2. the type of the state data which are tracked by the `AbstractFSM` module itself.

Note: The state data together with the state name describe the internal state of the state machine; if you stick to this scheme and do not add mutable fields to the FSM class you have the advantage of making all changes of the internal state explicit in a few well-known places.

Defining States

A state is defined by one or more invocations of the method

```
when(<name>[, stateTimeout = <timeout>])(stateFunction).
```

The given name must be an object which is type-compatible with the first type parameter given to the `AbstractFSM` class. This object is used as a hash key, so you must ensure that it properly implements `equals` and `hashCode`; in particular it must not be mutable. The easiest fit for these requirements are case objects.

If the `stateTimeout` parameter is given, then all transitions into this state, including staying, receive this timeout by default. Initiating the transition with an explicit timeout may be used to override this default, see [Initiating Transitions](#) for more information. The state timeout of any state may be changed during action processing with `setStateTimeout(state, duration)`. This enables runtime configuration e.g. via external message.

The `stateFunction` argument is a `PartialFunction[Event, State]`, which is conveniently given using the state function builder syntax as demonstrated below:

```
when(Idle,
  matchEvent(SetTarget.class, Uninitialized.class,
    (setTarget, uninitialized) ->
      stay().using(new Todo(setTarget.getRef(), new LinkedList<>())));
```

Warning: It is required that you define handlers for each of the possible FSM states, otherwise there will be failures when trying to switch to undeclared states.

It is recommended practice to declare the states as an enum and then verify that there is a `when` clause for each of the states. If you want to leave the handling of a state “unhandled” (more below), it still needs to be declared like this:

```
when(SomeState, AbstractFSM.NullFunction());
```

Defining the Initial State

Each FSM needs a starting point, which is declared using

```
startWith(state, data[, timeout])
```

The optionally given timeout argument overrides any specification given for the desired initial state. If you want to cancel a default timeout, use `Duration.Inf`.

Unhandled Events

If a state doesn't handle a received event a warning is logged. If you want to do something else in this case you can specify that with `whenUnhandled(stateFunction)`:

```
whenUnhandled(
  matchEvent(X.class, (x, data) -> {
    log().info("Received unhandled event: " + x);
    return stay();
  }).
  anyEvent((event, data) -> {
    log().warning("Received unknown event: " + event);
    return goto(Error);
  }));
}
```

Within this handler the state of the FSM may be queried using the `stateName` method.

IMPORTANT: This handler is not stacked, meaning that each invocation of `whenUnhandled` replaces the previously installed handler.

Initiating Transitions

The result of any `stateFunction` must be a definition of the next state unless terminating the FSM, which is described in [Termination from Inside](#). The state definition can either be the current state, as described by the `stay` directive, or it is a different state as given by `goto(state)`. The resulting object allows further qualification by way of the modifiers described in the following:

- `forMax(duration)`

This modifier sets a state timeout on the next state. This means that a timer is started which upon expiry sends a `StateTimeout` message to the FSM. This timer is canceled upon reception of any other message in the meantime; you can rely on the fact that the `StateTimeout` message will not be processed after an intervening message.

This modifier can also be used to override any default timeout which is specified for the target state. If you want to cancel the default timeout, use `Duration.Inf`.

- `using(data)`

This modifier replaces the old state data with the new data given. If you follow the advice *above*, this is the only place where internal state data are ever modified.

- `replying(msg)`

This modifier sends a reply to the currently processed message and otherwise does not modify the state transition.

All modifiers can be chained to achieve a nice and concise description:

```
when(SomeState, matchAnyEvent((msg, data) -> {
  return goto(Processing).using(newData).
    forMax(Duration.create(5, SECONDS)).replying(WillDo);
}));
```

The parentheses are not actually needed in all cases, but they visually distinguish between modifiers and their arguments and therefore make the code even more pleasant to read for foreigners.

Note: Please note that the `return` statement may not be used in `when` blocks or similar; this is a Scala restriction. Either refactor your code using `if () ... else ...` or move it into a method definition.

Monitoring Transitions

Transitions occur “between states” conceptually, which means after any actions you have put into the event handling block; this is obvious since the next state is only defined by the value returned by the event handling logic. You do not need to worry about the exact order with respect to setting the internal state variable, as everything within the FSM actor is running single-threaded anyway.

Internal Monitoring

Up to this point, the FSM DSL has been centered on states and events. The dual view is to describe it as a series of transitions. This is enabled by the method

```
onTransition(handler)
```

which associates actions with a transition instead of with a state and event. The handler is a partial function which takes a pair of states as input; no resulting state is needed as it is not possible to modify the transition in progress.

```
onTransition(
  matchState(Active, Idle, () -> setTimer("timeout",
    Tick, Duration.create(1, SECONDS), true)).
  state(Active, null, () -> cancelTimer("timeout")).
  state(null, Idle, (f, t) -> log().info("entering Idle from " + f));
```

It is also possible to pass a function object accepting two states to `onTransition`, in case your transition handling logic is implemented as a method:

```
public void handler(StateType from, StateType to) {
  // handle transition here
}

onTransition(this::handler);
```

The handlers registered with this method are stacked, so you can intersperse `onTransition` blocks with `when` blocks as suits your design. It should be noted, however, that *all handlers will be invoked for each transition*, not only the first matching one. This is designed specifically so you can put all transition handling for a certain aspect into one place without having to worry about earlier declarations shadowing later ones; the actions are still executed in declaration order, though.

Note: This kind of internal monitoring may be used to structure your FSM according to transitions, so that for example the cancellation of a timer upon leaving a certain state cannot be forgot when adding new target states.

External Monitoring

External actors may be registered to be notified of state transitions by sending a message `SubscribeTransitionCallback(actorRef)`. The named actor will be sent a `CurrentState(self, stateName)` message immediately and will receive `Transition(actorRef, oldState, newState)` messages whenever a new state is reached. External monitors may be unregistered by sending `UnsubscribeTransitionCallback(actorRef)` to the FSM actor.

Stopping a listener without unregistering will not remove the listener from the subscription list; use `UnsubscribeTransitionCallback` before stopping the listener.

Timers

Besides state timeouts, FSM manages timers identified by `String` names. You may set a timer using

```
setTimer(name, msg, interval, repeat)
```

where `msg` is the message object which will be sent after the duration `interval` has elapsed. If `repeat` is `true`, then the timer is scheduled at fixed rate given by the `interval` parameter. Any existing timer with the same name will automatically be canceled before adding the new timer.

Timers may be canceled using

```
cancelTimer(name)
```

which is guaranteed to work immediately, meaning that the scheduled message will not be processed after this call even if the timer already fired and queued it. The status of any timer may be inquired with

```
isTimerActive(name)
```

These named timers complement state timeouts because they are not affected by intervening reception of other messages.

Termination from Inside

The FSM is stopped by specifying the result state as

```
stop([reason[, data]])
```

The reason must be one of `Normal` (which is the default), `Shutdown` or `Failure` (`reason`), and the second argument may be given to change the state data which is available during termination handling.

Note: It should be noted that `stop` does not abort the actions and stop the FSM immediately. The stop action must be returned from the event handler in the same way as a state transition (but note that the `return` statement may not be used within a `when` block).

```
when(Error, matchEventEquals("stop", (event, data) -> {
  // do cleanup ...
  return stop();
}));
```

You can use `onTermination(handler)` to specify custom code that is executed when the FSM is stopped. The handler is a partial function which takes a `StopEvent(reason, stateName, stateData)` as argument:

```
onTermination(
  matchStop(Normal(),
    (state, data) -> { /* Do something here */}).
  stop(Shutdown(),
    (state, data) -> { /* Do something here */}).
  stop(Failure.class,
    (reason, state, data) -> { /* Do something here */}));
```

As for the `whenUnhandled` case, this handler is not stacked, so each invocation of `onTermination` replaces the previously installed handler.

Termination from Outside

When an `ActorRef` associated to a FSM is stopped using the `stop` method, its `postStop` hook will be executed. The default implementation by the `AbstractFSM` class is to execute the `onTermination` handler if that is prepared to handle a `StopEvent(Shutdown, ...)`.

Warning: In case you override `postStop` and want to have your `onTermination` handler called, do not forget to call `super.postStop`.

9.3.4 Testing and Debugging Finite State Machines

During development and for trouble shooting FSMs need care just as any other actor. There are specialized tools available as described in *TestFSMRef* and in the following.

Event Tracing

The setting `akka.actor.debug.fsm` in *Configuration* enables logging of an event trace by `LoggingFSM` instances:

```
public class MyFSM extends AbstractLoggingFSM<StateType, Data> {
  // body elided ...
}
```

This FSM will log at DEBUG level:

- all processed events, including `StateTimeout` and scheduled timer messages
- every setting and cancellation of named timers
- all state transitions

Life cycle changes and special messages can be logged as described for *Actors*.

Rolling Event Log

The `AbstractLoggingFSM` class adds one more feature to the FSM: a rolling event log which may be used during debugging (for tracing how the FSM entered a certain failure state) or for other creative uses:

```
public class MyFSM extends AbstractLoggingFSM<StateType, Data> {
  @Override
  public int logDepth() { return 12; }
  {
    onTermination(
      matchStop(Failure.class, (reason, state, data) -> {
        String lastEvents = getLog().mkString("\n\t");
        log().warning("Failure in state " + state + " with data " + data + "\n" +
          "Events leading up to this point:\n\t" + lastEvents);
      })
    );
    //...
  }
}
```

The `logDepth` defaults to zero, which turns off the event log.

Warning: The log buffer is allocated during actor creation, which is why the configuration is done using a virtual method call. If you want to override with a `val`, make sure that its initialization happens before the initializer of `LoggingFSM` runs, and do not change the value returned by `logDepth` after the buffer has been allocated.

The contents of the event log are available using method `getLog`, which returns an `IndexedSeq[LogEntry]` where the oldest entry is at index zero.

9.3.5 Examples

A bigger FSM example contrasted with Actor's `become/unbecome` can be found in the *Typesafe Activator* template named *Akka FSM in Scala*

9.4 Persistence Query

Akka persistence query complements *Persistence* by providing a universal asynchronous stream based query interface that various journal plugins can implement in order to expose their query capabilities.

The most typical use case of persistence query is implementing the so-called query side (also known as “read side”) in the popular CQRS architecture pattern - in which the writing side of the application (e.g. implemented using akka persistence) is completely separated from the “query side”. Akka Persistence Query itself is *not* directly the query side of an application, however it can help to migrate data from the write side to the query side database. In very simple scenarios Persistence Query may be powerful enough to fulfill the query needs of your app, however we highly recommend (in the spirit of CQRS) of splitting up the write/read sides into separate datastores as the need arises.

Warning: This module is marked as “**experimental**” as of its introduction in Akka 2.4.0. We will continue to improve this API based on our users’ feedback, which implies that while we try to keep incompatible changes to a minimum the binary compatibility guarantee for maintenance releases does not apply to the contents of the `akka.persistence.query` package.

9.4.1 Dependencies

Akka persistence query is a separate jar file. Make sure that you have the following dependency in your project:

```
<dependency>
  <groupId>com.typesafe.akka</groupId>
  <artifactId>akka-persistence-query-experimental_2.11</artifactId>
  <version>2.4.0</version>
</dependency>
```

9.4.2 Design overview

Akka persistence query is purposely designed to be a very loosely specified API. This is in order to keep the provided APIs general enough for each journal implementation to be able to expose its best features, e.g. a SQL journal can use complex SQL queries or if a journal is able to subscribe to a live event stream this should also be possible to expose the same API - a typed stream of events.

Each read journal must explicitly document which types of queries it supports. Refer to the your journal’s plugins documentation for details on which queries and semantics it supports.

While Akka Persistence Query does not provide actual implementations of ReadJournals, it defines a number of pre-defined query types for the most common query scenarios, that most journals are likely to implement (however they are not required to).

9.4.3 Read Journals

In order to issue queries one has to first obtain an instance of a ReadJournal. Read journals are implemented as [Community plugins](#), each targeting a specific datastore (for example Cassandra or JDBC databases). For example, given a library that provides a `akka.persistence.query.my-read-journal` obtaining the related journal is as simple as:

```
// obtain read journal by plugin id
final MyJavadslReadJournal readJournal =
  PersistenceQuery.get(system).getReadJournalFor(MyJavadslReadJournal.class,
    "akka.persistence.query.my-read-journal");

// issue query to journal
Source<EventEnvelope, BoxedUnit> source =
  readJournal.eventsByPersistenceId("user-1337", 0, Long.MAX_VALUE);
```

```
// materialize stream, consuming events
ActorMaterializer mat = ActorMaterializer.create(system);
source.runForeach(event -> System.out.println("Event: " + event), mat);
```

Journal implementers are encouraged to put this identifier in a variable known to the user, such that one can access it via `getJournalFor(NoopJournal.class, NoopJournal.identifier)`, however this is not enforced.

Read journal implementations are available as [Community plugins](#).

Predefined queries

Akka persistence query comes with a number of query interfaces built in and suggests Journal implementors to implement them according to the semantics described below. It is important to notice that while these query types are very common a journal is not obliged to implement all of them - for example because in a given journal such query would be significantly inefficient.

Note: Refer to the documentation of the `ReadJournal` plugin you are using for a specific list of supported query types. For example, Journal plugins should document their stream completion strategies.

The predefined queries are:

AllPersistenceIdsQuery and CurrentPersistenceIdsQuery

`allPersistenceIds` which is designed to allow users to subscribe to a stream of all persistent ids in the system. By default this stream should be assumed to be a “live” stream, which means that the journal should keep emitting new persistence ids as they come into the system:

```
readJournal.allPersistenceIds();
```

If your usage does not require a live stream, you can use the `currentPersistenceIds` query:

```
readJournal.currentPersistenceIds();
```

EventsByPersistenceIdQuery and CurrentEventsByPersistenceIdQuery

`eventsByPersistenceId` is a query equivalent to replaying a *PersistentActor*, however, since it is a stream it is possible to keep it alive and watch for additional incoming events persisted by the persistent actor identified by the given `persistenceId`.

```
readJournal.eventsByPersistenceId("user-us-1337", 0L, Long.MAX_VALUE);
```

Most journals will have to revert to polling in order to achieve this, which can typically be configured with a `refresh-interval` configuration property.

If your usage does not require a live stream, you can use the `currentEventsByPersistenceId` query.

EventsByTag and CurrentEventsByTag

`eventsByTag` allows querying events regardless of which `persistenceId` they are associated with. This query is hard to implement in some journals or may need some additional preparation of the used data store to be executed efficiently. The goal of this query is to allow querying for all events which are “tagged” with a specific tag. That includes the use case to query all domain events of an Aggregate Root type. Please refer to your read journal plugin’s documentation to find out if and how it is supported.

Some journals may support tagging of events via an *Event Adapters* that wraps the events in a `akka.persistence.journal.Tagged` with the given tags. The journal may support other ways of doing tagging - again, how exactly this is implemented depends on the used journal. Here is an example of such a tagging event adapter:

```
public class MyTaggingEventAdapter implements WriteEventAdapter {

    @Override
    public Object toJournal(Object event) {
        if (event instanceof String) {
            String s = (String) event;
            Set<String> tags = new HashSet<String>();
            if (s.contains("green")) tags.add("green");
            if (s.contains("black")) tags.add("black");
            if (s.contains("blue")) tags.add("blue");
            if (tags.isEmpty())
                return event;
            else
                return new Tagged(event, tags);
        } else {
            return event;
        }
    }

    @Override
    public String manifest(Object event) {
        return "";
    }
}
```

Note: A very important thing to keep in mind when using queries spanning multiple persistenceIds, such as `EventsByTag` is that the order of events at which the events appear in the stream rarely is guaranteed (or stable between materializations).

Journals *may* choose to opt for strict ordering of the events, and should then document explicitly what kind of ordering guarantee they provide - for example “*ordered by timestamp ascending, independently of persistenceId*” is easy to achieve on relational databases, yet may be hard to implement efficiently on plain key-value datastores.

In the example below we query all events which have been tagged (we assume this was performed by the write-side using an *EventAdapter*, or that the journal is smart enough that it can figure out what we mean by this tag - for example if the journal stored the events as json it may try to find those with the field `tag` set to this value etc.).

```
// assuming journal is able to work with numeric offsets we can:
final Source<EventEnvelope, BoxedUnit> blueThings =
    readJournal.eventsByTag("blue", 0L);

// find top 10 blue things:
final Future<List<Object>> top10BlueThings =
    (Future<List<Object>>) blueThings
        .map(t -> t.event())
        .take(10) // cancels the query stream after pulling 10 elements
        .<List<Object>>runFold(new ArrayList<>(10), (acc, e) -> {
            acc.add(e);
            return acc;
        }, mat);

// start another query, from the known offset
Source<EventEnvelope, BoxedUnit> blue = readJournal.eventsByTag("blue", 10);
```

As you can see, we can use all the usual stream combinators available from [Akka Streams](#) on the resulting query stream, including for example taking the first 10 and cancelling the stream. It is worth pointing out that the built-in `EventsByTag` query has an optionally supported offset parameter (of type `Long`) which the journals can use

to implement resumable-streams. For example a journal may be able to use a WHERE clause to begin the read starting from a specific row, or in a datastore that is able to order events by insertion time it could treat the Long as a timestamp and select only older events.

If your usage does not require a live stream, you can use the `currentEventsByTag` query.

Materialized values of queries

Journals are able to provide additional information related to a query by exposing [materialized values](#), which are a feature of [Akka Streams](#) that allows to expose additional values at stream materialization time.

More advanced query journals may use this technique to expose information about the character of the materialized stream, for example if it's finite or infinite, strictly ordered or not ordered at all. The materialized value type is defined as the second type parameter of the returned `Source`, which allows journals to provide users with their specialised query object, as demonstrated in the sample below:

```
public class RichEvent {
    public final Set<String> tags;
    public final Object payload;

    public RichEvent(Set<String> tags, Object payload) {
        this.tags = tags;
        this.payload = payload;
    }
}

// a plugin can provide:
public final class QueryMetadata{
    public final boolean deterministicOrder;
    public final boolean infinite;

    public QueryMetadata(boolean deterministicOrder, boolean infinite) {
        this.deterministicOrder = deterministicOrder;
        this.infinite = infinite;
    }
}

public Source<RichEvent, QueryMetadata> byTagsWithMeta(Set<String> tags) {

    Set<String> tags = new HashSet<String>();
    tags.add("red");
    tags.add("blue");
    final Source<RichEvent, QueryMetadata> events = readJournal.byTagsWithMeta(tags)
        .mapMaterializedValue(meta -> {
            System.out.println("The query is: " +
                "ordered deterministically: " + meta.deterministicOrder + " " +
                "infinite: " + meta.infinite);
            return meta;
        });

    events.map(event -> {
        System.out.println("Event payload: " + event.payload);
        return event.payload;
    }).runWith(Sink.ignore(), mat);
}
```

9.4.4 Performance and denormalization

When building systems using [Event sourcing](#) and CQRS (Command & Query Responsibility Segregation) techniques it is tremendously important to realise that the write-side has completely different needs from the read-side, and separating those concerns into datastores that are optimised for either side makes it possible to offer the best experience for the write and read sides independently.

For example, in a bidding system it is important to “take the write” and respond to the bidder that we have accepted the bid as soon as possible, which means that write-throughput is of highest importance for the write-side – often this means that data stores which are able to scale to accommodate these requirements have a less expressive query side.

On the other hand the same application may have some complex statistics view or we may have analysts working with the data to figure out best bidding strategies and trends – this often requires some kind of expressive query capabilities like for example SQL or writing Spark jobs to analyse the data. Therefore the data stored in the write-side needs to be projected into the other read-optimised datastore.

Note: When referring to **Materialized Views** in Akka Persistence think of it as “some persistent storage of the result of a Query”. In other words, it means that the view is created once, in order to be afterwards queried multiple times, as in this format it may be more efficient or interesting to query it (instead of the source events directly).

Materialize view to Reactive Streams compatible datastore

If the read datastore exposes an [Reactive Streams](#) interface then implementing a simple projection is as simple as, using the read-journal and feeding it into the databases driver interface, for example like so:

```
final ReactiveStreamsCompatibleDBDriver driver = new ReactiveStreamsCompatibleDBDriver();
final Subscriber<List<Object>> dbBatchWriter = driver.batchWriter();

// Using an example (Reactive Streams) Database driver
readJournal
  .eventsByPersistenceId("user-1337", 0L, Long.MAX_VALUE)
  .map(envelope -> envelope.event())
  .grouped(20) // batch inserts into groups of 20
  .runWith(Sink.create(dbBatchWriter), mat); // write batches to read-side database
```

Materialize view using mapAsync

If the target database does not provide a reactive streams `Subscriber` that can perform writes, you may have to implement the write logic using plain functions or Actors instead.

In case your write logic is state-less and you just need to convert the events from one data type to another before writing into the alternative datastore, then the projection is as simple as:

```
class ExampleStore {
  Future<Void> save(Object any) {
    // ...
  }
}

final ExampleStore store = new ExampleStore();

readJournal
  .eventsByTag("bid", 0L)
  .mapAsync(1, store::save)
  .runWith(Sink.ignore(), mat);
```

Resumable projections

Sometimes you may need to implement “resumable” projections, that will not start from the beginning of time each time when run. In this case you will need to store the sequence number (or `offset`) of the processed event and use it the next time this projection is started. This pattern is not built-in, however is rather simple to implement yourself.

The example below additionally highlights how you would use Actors to implement the write side, in case you need to do some complex logic that would be best handled inside an Actor before persisting the event into the other datastore:

```
final Timeout timeout = Timeout.apply(3, TimeUnit.SECONDS);

final MyResumableProjection bidProjection = new MyResumableProjection("bid");

final Props writerProps = Props.create(TheOneWhoWritesToQueryJournal.class, "bid");
final ActorRef writer = system.actorOf(writerProps, "bid-projection-writer");

long startFromOffset = Await.result(bidProjection.latestOffset(), timeout.duration());

readJournal
  .eventsByTag("bid", startFromOffset)
  .<Long>mapAsync(8, envelope -> {
    final Future<Object> f = ask(writer, envelope.event(), timeout);
    return f.<Long>map(new Mapper<Object, Long>() {
      @Override public Long apply(Object in) {
        return envelope.offset();
      }
    }, system.dispatcher());
  })
  .mapAsync(1, offset -> bidProjection.saveProgress(offset))
  .runWith(Sink.ignore(), mat);
}
```

```
final class TheOneWhoWritesToQueryJournal extends AbstractActor {
  private final ExampleStore store;

  private ComplexState state = new ComplexState();

  public TheOneWhoWritesToQueryJournal() {
    store = new ExampleStore();

    receive(ReceiveBuilder.matchAny(message -> {
      state = updateState(state, message);

      // example saving logic that requires state to become ready:
      if (state.readyToSave())
        store.save(Record.of(state));

    }).build());
  }

  ComplexState updateState(ComplexState state, Object msg) {
    // some complicated aggregation logic here ...
    return state;
  }
}
```

9.4.5 Query plugins

Query plugins are various (mostly community driven) `ReadJournal` implementations for all kinds of available datastores. The complete list of available plugins is maintained on the Akka Persistence Query [Community Plugins](#) page.

The plugin for LevelDB is described in *Persistence Query for LevelDB*.

This section aims to provide tips and guide plugin developers through implementing a custom query plugin. Most users will not need to implement journals themselves, except if targeting a not yet supported datastore.

Note: Since different data stores provide different query capabilities journal plugins **must extensively document** their exposed semantics as well as handled query scenarios.

ReadJournal plugin API

A read journal plugin must implement `akka.persistence.query.ReadJournalProvider` which creates instances of `akka.persistence.query.scaladsl.ReadJournal` and `akka.persistence.query.javaadsl.ReadJournal`. The plugin must implement both the `scaladsl` and the `javaadsl` interfaces because the `akka.stream.scaladsl.Source` and `akka.stream.javaadsl.Source` are different types and even though those types can easily be converted to each other it is most convenient for the end user to get access to the Java or Scala `Source` directly. As illustrated below one of the implementations can delegate to the other.

Below is a simple journal implementation:

```
public class MyReadJournalProvider implements ReadJournalProvider {
    private final MyJavaadslReadJournal javaadslReadJournal;

    public MyReadJournalProvider(ExtendedActorSystem system, Config config) {
        this.javaadslReadJournal = new MyJavaadslReadJournal(system, config);
    }

    @Override
    public MyScaladslReadJournal scaladslReadJournal() {
        return new MyScaladslReadJournal(javaadslReadJournal);
    }

    @Override
    public MyJavaadslReadJournal javaadslReadJournal() {
        return this.javaadslReadJournal;
    }
}

public class MyJavaadslReadJournal implements
akka.persistence.query.javaadsl.ReadJournal,
akka.persistence.query.javaadsl.EventsByTagQuery,
akka.persistence.query.javaadsl.EventsByPersistenceIdQuery,
akka.persistence.query.javaadsl.AllPersistenceIdsQuery,
akka.persistence.query.javaadsl.CurrentPersistenceIdsQuery {

    private final FiniteDuration refreshInterval;

    public MyJavaadslReadJournal(ExtendedActorSystem system, Config config) {
        refreshInterval =
            FiniteDuration.create(config.getDuration("refresh-interval",
                TimeUnit.MILLISECONDS), TimeUnit.MILLISECONDS);
    }

    @Override
    public Source<EventEnvelope, BoxedUnit> eventsByTag(String tag, long offset) {
        final Props props = MyEventsByTagPublisher.props(tag, offset, refreshInterval);
        return Source.<EventEnvelope>actorPublisher(props).
            mapMaterializedValue(m -> BoxedUnit.UNIT);
    }

    @Override
    public Source<EventEnvelope, BoxedUnit> eventsByPersistenceId(String persistenceId,
        long fromSequenceNr, long toSequenceNr) {
        // implement in a similar way as eventsByTag
        throw new UnsupportedOperationException("Not implemented yet");
    }
}
```

```

@Override
public Source<String, BoxedUnit> allPersistenceIds() {
    // implement in a similar way as eventsByTag
    throw new UnsupportedOperationException("Not implemented yet");
}

@Override
public Source<String, BoxedUnit> currentPersistenceIds() {
    // implement in a similar way as eventsByTag
    throw new UnsupportedOperationException("Not implemented yet");
}

// possibility to add more plugin specific queries

public Source<RichEvent, QueryMetadata> byTagsWithMeta(Set<String> tags) {
    // implement in a similar way as eventsByTag
    throw new UnsupportedOperationException("Not implemented yet");
}
}

public class MyScaladslReadJournal implements
akka.persistence.query.scaladsl.ReadJournal,
akka.persistence.query.scaladsl.EventsByTagQuery,
akka.persistence.query.scaladsl.EventsByPersistenceIdQuery,
akka.persistence.query.scaladsl.AllPersistenceIdsQuery,
akka.persistence.query.scaladsl.CurrentPersistenceIdsQuery {

    private final MyJavadslReadJournal javadslReadJournal;

    public MyScaladslReadJournal(MyJavadslReadJournal javadslReadJournal) {
        this.javadslReadJournal = javadslReadJournal;
    }

    @Override
    public akka.stream.scaladsl.Source<EventEnvelope, BoxedUnit> eventsByTag(
        String tag, long offset) {
        return javadslReadJournal.eventsByTag(tag, offset).asScala();
    }

    @Override
    public akka.stream.scaladsl.Source<EventEnvelope, BoxedUnit> eventsByPersistenceId(
        String persistenceId, long fromSequenceNr, long toSequenceNr) {
        return javadslReadJournal.eventsByPersistenceId(persistenceId, fromSequenceNr,
            toSequenceNr).asScala();
    }

    @Override
    public akka.stream.scaladsl.Source<String, BoxedUnit> allPersistenceIds() {
        return javadslReadJournal.allPersistenceIds().asScala();
    }

    @Override
    public akka.stream.scaladsl.Source<String, BoxedUnit> currentPersistenceIds() {
        return javadslReadJournal.currentPersistenceIds().asScala();
    }

    // possibility to add more plugin specific queries

    public akka.stream.scaladsl.Source<RichEvent, QueryMetadata> byTagsWithMeta(
        scala.collection.Set<String> tags) {
        Set<String> jTags = scala.collection.JavaConversions.setAsJavaSet(tags);
        return javadslReadJournal.byTagsWithMeta(jTags).asScala();
    }
}

```

```

}
}

```

And the `EventsByTag` could be backed by such an Actor for example:

```

class MyEventsByTagJavaPublisher extends AbstractActorPublisher<EventEnvelope> {
  private final Serialization serialization =
    SerializationExtension.get(context().system());

  private final Connection connection;

  private final String tag;

  private final String CONTINUE = "CONTINUE";
  private final int LIMIT = 1000;
  private long currentOffset;
  private List<EventEnvelope> buf = new LinkedList<>();

  private Cancellable continueTask;

  public MyEventsByTagJavaPublisher(Connection connection,
                                    String tag,
                                    Long offset,
                                    FiniteDuration refreshInterval) {

    this.connection = connection;
    this.tag = tag;
    this.currentOffset = offset;

    final Scheduler scheduler = context().system().scheduler();
    this.continueTask = scheduler
      .schedule(refreshInterval, refreshInterval, self(), CONTINUE,
               context().dispatcher(), self());

    receive(ReceiveBuilder
      .matchEquals(CONTINUE, (in) -> {
        query();
        deliverBuf();
      })
      .match(Cancel.class, (in) -> {
        context().stop(self());
      })
      .build());
  }

  public static Props props(Connection conn, String tag, Long offset,
                            FiniteDuration refreshInterval) {
    return Props.create(() ->
      new MyEventsByTagJavaPublisher(conn, tag, offset, refreshInterval));
  }

  @Override
  public void postStop() {
    continueTask.cancel();
  }

  private void query() {
    if (buf.isEmpty()) {
      final String query = "SELECT id, persistent_repr " +
        "FROM journal WHERE tag = ? AND id >= ? " +
        "ORDER BY id LIMIT ?";

      try (PreparedStatement s = connection.prepareStatement(query)) {

```

```

s.setString(1, tag);
s.setLong(2, currentOffset);
s.setLong(3, LIMIT);
try (ResultSet rs = s.executeQuery()) {

    final List<Pair<Long, byte[]>> res = new ArrayList<>(LIMIT);
    while (rs.next())
        res.add(Pair.create(rs.getLong(1), rs.getBytes(2)));

    if (!res.isEmpty()) {
        currentOffset = res.get(res.size() - 1).first();
    }

    buf = res.stream().map(in -> {
        final Long id = in.first();
        final byte[] bytes = in.second();

        final PersistentRepr p =
            serialization.deserialize(bytes, PersistentRepr.class).get();

        return new EventEnvelope(id, p.persistenceId(), p.sequenceNr(), p.payload());
    }).collect(toList());
} catch (Exception e) {
    onErrorThenStop(e);
}
}

private void deliverBuf() {
    while (totalDemand() > 0 && !buf.isEmpty())
        onNext(buf.remove(0));
}
}

```

If the underlying datastore only supports queries that are completed when they reach the end of the “result set”, the journal has to submit new queries after a while in order to support “infinite” event streams that include events stored after the initial query has completed. It is recommended that the plugin use a configuration property named `refresh-interval` for defining such a refresh interval.

Plugin TCK

TODO, not available yet.

Another reason for marking a module as experimental is that it’s too early to tell if the module has a maintainer that can take the responsibility of the module over time. These modules live in the `akka-contrib` subproject:

9.5 External Contributions

This subproject provides a home to modules contributed by external developers which may or may not move into the officially supported code base over time. The conditions under which this transition can occur include:

- there must be enough interest in the module to warrant inclusion in the standard distribution,
- the module must be actively maintained and
- code quality must be good enough to allow efficient maintenance by the Akka core development team

If a contributions turns out to not “take off” it may be removed again at a later time.

9.5.1 Caveat Emptor

A module in this subproject doesn't have to obey the rule of staying binary compatible between minor releases. Breaking API changes may be introduced in minor releases without notice as we refine and simplify based on your feedback. A module may be dropped in any release without prior deprecation. The Typesafe subscription does not cover support for these modules.

9.5.2 The Current List of Modules

Reliable Proxy Pattern

Looking at *Message Delivery Reliability* one might come to the conclusion that Akka actors are made for blue-sky scenarios: sending messages is the only way for actors to communicate, and then that is not even guaranteed to work. Is the whole paradigm built on sand? Of course the answer is an emphatic “No!”.

A local message send—within the same JVM instance—is not likely to fail, and if it does the reason was one of


- it was meant to fail (due to consciously choosing a bounded mailbox, which upon overflow will have to drop messages)
- or it failed due to a catastrophic VM error, e.g. an `OutOfMemoryError`, a memory access violation (“segmentation fault”, GPF, etc.), JVM bug—or someone calling `System.exit()`.

In all of these cases, the actor was very likely not in a position to process the message anyway, so this part of the non-guarantee is not problematic.

It is a lot more likely for an unintended message delivery failure to occur when a message send crosses JVM boundaries, i.e. an intermediate unreliable network is involved. If someone unplugs an ethernet cable, or a power failure shuts down a router, messages will be lost while the actors would be able to process them just fine.

Note: This does not mean that message send semantics are different between local and remote operations, it just means that in practice there is a difference between how good the “best effort” is.

Introducing the Reliable Proxy



To bridge the disparity between “local” and “remote” sends is the goal of this pattern. When sending from A to B must be as reliable as in-JVM, regardless of the deployment, then you can interject a reliable tunnel and send through that instead. The tunnel consists of two end-points, where the ingress point P (the “proxy”) is a child of A and the egress point E is a child of P, deployed onto the same network node where B lives. Messages sent to P will be wrapped in an envelope, tagged with a sequence number and sent to E, who verifies that the received envelope

has the right sequence number (the next expected one) and forwards the contained message to B. When B receives this message, the `sender()` will be a reference to the sender() of the original message to P. Reliability is added by E replying to orderly received messages with an ACK, so that P can tick those messages off its resend list. If ACKs do not come in a timely fashion, P will try to resend until successful.

Exactly what does it guarantee?

Sending via a `ReliableProxy` makes the message send exactly as reliable as if the represented target were to live within the same JVM, provided that the remote actor system does not terminate. In effect, both ends (i.e. JVM and actor system) must be considered as one when evaluating the reliability of this communication channel. The benefit is that the network in-between is taken out of that equation.

Connecting to the target The `proxy` tries to connect to the `target` using the mechanism outlined in *actorSelection-scala*. Once connected, if the `tunnel` terminates the `proxy` will optionally try to reconnect to the target using using the same process.

Note that during the reconnection process there is a possibility that a message could be delivered to the `target` more than once. Consider the case where a message is delivered to the `target` and the target system crashes before the ACK is sent to the `sender`. After the `proxy` reconnects to the `target` it will start resending all of the messages that it has not received an ACK for, and the message that it never got an ACK for will be redelivered. Either this possibility should be considered in the design of the `target` or reconnection should be disabled.

How to use it

Since this implementation does not offer much in the way of configuration, simply instantiate a proxy wrapping a target `ActorPath`. From Java it looks like this:

```
import akka.contrib.pattern.ReliableProxy;

public class ProxyParent extends UntypedActor {
    private final ActorRef proxy;

    public ProxyParent(ActorPath targetPath) {
        proxy = getContext().actorOf(
            ReliableProxy.props(targetPath,
                Duration.create(100, TimeUnit.MILLISECONDS));
        }

    public void onReceive(Object msg) {
        if ("hello".equals(msg)) {
            proxy.tell("world!", getSelf());
        }
    }
}
```

And from Scala like this:

```
import akka.contrib.pattern.ReliableProxy

class ProxyParent(targetPath: ActorPath) extends Actor {
    val proxy = context.actorOf(ReliableProxy.props(targetPath, 100.millis))

    def receive = {
        case "hello" => proxy ! "world!"
    }
}
```

Since the `ReliableProxy` actor is an *fsm-scala*, it also offers the capability to subscribe to state transitions. If you need to know when all enqueued messages have been received by the remote end-point (and consequently

been forwarded to the target), you can subscribe to the FSM notifications and observe a transition from state `ReliableProxy.Active` to state `ReliableProxy.Idle`.

```
public class ProxyTransitionParent extends UntypedActor {
    private final ActorRef proxy;
    private ActorRef client = null;

    public ProxyTransitionParent(ActorPath targetPath) {
        proxy = getContext().actorOf(
            ReliableProxy.props(targetPath,
                Duration.create(100, TimeUnit.MILLISECONDS)));
        proxy.tell(new FSM.SubscribeTransitionCallBack(getSelf()), getSelf());
    }

    public void onReceive(Object msg) {
        if ("hello".equals(msg)) {
            proxy.tell("world!", getSelf());
            client = getSender();
        } else if (msg instanceof FSM.CurrentState<?>) {
            // get initial state
        } else if (msg instanceof FSM.Transition<?>) {
            @SuppressWarnings("unchecked")
            final FSM.Transition<ReliableProxy.State> transition =
                (FSM.Transition<ReliableProxy.State>) msg;
            assert transition.fsmRef().equals(proxy);
            if (transition.from().equals(ReliableProxy.active()) &&
                transition.to().equals(ReliableProxy.idle())) {
                client.tell("done", getSelf());
            }
        }
    }
}
```

From Scala it would look like so:

```
class ProxyTransitionParent(targetPath: ActorPath) extends Actor {
    val proxy = context.actorOf(ReliableProxy.props(targetPath, 100.millis))
    proxy ! FSM.SubscribeTransitionCallBack(self)

    var client: ActorRef = _

    def receive = {
        case "go" =>
            proxy ! 42
            client = sender()
        case FSM.CurrentState(`proxy`, initial) =>
        case FSM.Transition(`proxy`, from, to) =>
            if (to == ReliableProxy.Idle)
                client ! "done"
    }
}
```

Configuration

- Set `akka.reliable-proxy.debug` to `on` to turn on extra debug logging for your `ReliableProxy` actors.
- `akka.reliable-proxy.default-connect-interval` is used only if you create a `ReliableProxy` with no reconnections (that is, `reconnectAfter == None`). The default value is the value of the configuration property `akka.remote.retry-gate-closed-for`. For example, if `akka.remote.retry-gate-closed-for` is `5 s` case the `ReliableProxy` will send an `Identify` message to the *target* every 5 seconds to try to resolve the `ActorPath` to an `ActorRef` so that messages can be sent to the *target*.

The Actor Contract

Message it Processes

- `FSM.SubscribeTransitionCallBack` and `FSM.UnsubscribeTransitionCallBack`, see *fsm-scala*
- `ReliableProxy.Unsent`, see the API documentation for details.
- any other message is transferred through the reliable tunnel and forwarded to the designated target actor

Messages it Sends

- `FSM.CurrentState` and `FSM.Transition`, see *fsm-scala*
- `ReliableProxy.TargetChanged` is sent to the FSM transition subscribers if the proxy reconnects to a new target.
- `ReliableProxy.ProxyTerminated` is sent to the FSM transition subscribers if the proxy is stopped.

Exceptions it Escalates

- no specific exception types
- any exception encountered by either the local or remote end-point are escalated (only fatal VM errors)

Arguments it Takes

- *target* is the `ActorPath` to the actor to which the tunnel shall reliably deliver messages, B in the above illustration.
- *retryAfter* is the timeout for receiving ACK messages from the remote end-point; once it fires, all outstanding message sends will be retried.
- *reconnectAfter* is an optional interval between connection attempts. It is also used as the interval between receiving a `Terminated` for the tunnel and attempting to reconnect to the target actor.
- *maxConnectAttempts* is an optional maximum number of attempts to connect to the target while in the `Connecting` state.

Throttling Actor Messages

Introduction

Suppose you are writing an application that makes HTTP requests to an external web service and that this web service has a restriction in place: you may not make more than 10 requests in 1 minute. You will get blocked or need to pay if you don't stay under this limit. In such a scenario you will want to employ a *message throttler*.

This extension module provides a simple implementation of a throttling actor, the `TimerBasedThrottler`.

How to use it

You can use a `TimerBasedThrottler` as follows. From Java it looks like this:

```
// A simple actor that prints whatever it receives
ActorRef printer = system.actorOf(Props.create(Printer.class));
// The throttler for this example, setting the rate
ActorRef throttler = system.actorOf(Props.create(TimerBasedThrottler.class,
    new Throttler.Rate(3, Duration.create(1, TimeUnit.SECONDS))));
// Set the target
throttler.tell(new Throttler.SetTarget(printer), null);
```

```
// These three messages will be sent to the target immediately
throttler.tell("1", null);
throttler.tell("2", null);
throttler.tell("3", null);
// These two will wait until a second has passed
throttler.tell("4", null);
throttler.tell("5", null);

//A simple actor that prints whatever it receives
public class Printer extends UntypedActor {
  @Override
  public void onReceive(Object msg) {
    System.out.println(msg);
  }
}
```

And from Scala like this:

```
// A simple actor that prints whatever it receives
class PrintActor extends Actor {
  def receive = {
    case x => println(x)
  }
}

val printer = system.actorOf(Props[PrintActor])
// The throttler for this example, setting the rate
val throttler = system.actorOf(Props(classOf[TimerBasedThrottler],
  3 msgsPer 1.second))
// Set the target
throttler ! SetTarget(Some(printer))
// These three messages will be sent to the target immediately
throttler ! "1"
throttler ! "2"
throttler ! "3"
// These two will wait until a second has passed
throttler ! "4"
throttler ! "5"
```

Please refer to the JavaDoc/ScalaDoc documentation for the details.

The guarantees

`TimerBasedThrottler` uses a timer internally. When the throttler's rate is 3 msg/s, for example, the throttler will start a timer that triggers every second and each time will give the throttler exactly three "vouchers"; each voucher gives the throttler a right to deliver a message. In this way, at most 3 messages will be sent out by the throttler in each interval.

It should be noted that such timer-based throttlers provide relatively **weak guarantees**:

- Only *start times* are taken into account. This may be a problem if, for example, the throttler is used to throttle requests to an external web service. If a web request takes very long on the server then the rate *observed on the server* may be higher.
- A timer-based throttler only makes guarantees for the intervals of its own timer. In our example, no more than 3 messages are delivered within such intervals. Other intervals on the timeline, however, may contain more calls.

The two cases are illustrated in the two figures below, each showing a timeline and three intervals of the timer. The message delivery times chosen by the throttler are indicated by dots, and as you can see, each interval contains at most 3 point, so the throttler works correctly. Still, there is in each example an interval (the red one) that is problematic. In the first scenario, this is because the delivery times are merely the start times of longer requests (indicated by the four bars above the timeline that start at the dots), so that the server observes four requests during

the red interval. In the second scenario, the messages are centered around one of the points in time where the timer triggers, causing the red interval to contain too many messages.



For some application scenarios, the guarantees provided by a timer-based throttler might be too weak. Charles Cordingley's [blog post](#) discusses a throttler with stronger guarantees (it solves problem 2 from above). Future versions of this module may feature throttlers with better guarantees.

Java Logging (JUL)

This extension module provides a logging backend which uses the *java.util.logging* (j.u.l) API to do the endpoint logging for *akka.event.Logging*.

Provided with this module is an implementation of *akka.event.LoggingAdapter* which is independent of any *ActorSystem* being in place. This means that j.u.l can be used as the backend, via the Akka Logging API, for both Actor and non-Actor codebases.

To enable j.u.l as the *akka.event.Logging* backend, use the following Akka config:

```
loggers = ["akka.contrib.jul.JavaLogger"]
```

To access the *akka.event.Logging* API from non-Actor code, mix in *akka.contrib.jul.JavaLogging*.

This module is preferred over SLF4J with its JDK14 backend, due to integration issues resulting in the incorrect handling of *threadId*, *className* and *methodName*.

This extension module was contributed by Sam Halliday.

Mailbox with Explicit Acknowledgement

When an Akka actor is processing a message and an exception occurs, the normal behavior is for the actor to drop that message, and then continue with the next message after it has been restarted. This is in some cases not the desired solution, e.g. when using failure and supervision to manage a connection to an unreliable resource; the actor could after the restart go into a buffering mode (i.e. change its behavior) and retry the real processing later, when the unreliable resource is back online.

One way to do this is by sending all messages through the supervisor and buffering them there, acknowledging successful processing in the child; another way is to build an explicit acknowledgement mechanism into the mailbox. The idea with the latter is that a message is reprocessed in case of failure until the mailbox is told that processing was successful.

The pattern is implemented [here](#). A demonstration of how to use it (although for brevity not a perfect example) is the following:

```
class MyActor extends Actor {
  def receive = {
    case msg =>
      println(msg)
      doStuff(msg) // may fail
      PeekMailboxExtension.ack()
  }
}
```

```

}

// business logic elided ...
}

object MyApp extends App {
  val system = ActorSystem("MySystem", ConfigFactory.parseString("""
    peek-dispatcher {
      mailbox-type = "akka.contrib.mailbox.PeekMailboxType"
      max-retries = 2
    }
    """))

  val myActor = system.actorOf(Props[MyActor].withDispatcher("peek-dispatcher"),
    name = "myActor")

  myActor ! "Hello"
  myActor ! "World"
  myActor ! PoisonPill
}

```

Running this application (try it in the Akka sources by saying `sbt akka-contrib/test:run`) may produce the following output (note the processing of “World” on lines 2 and 16):

```

Hello
World
[ERROR] [12/17/2012 16:28:36.581] [MySystem-peek-dispatcher-5] [akka://MySystem/user/myActor] DONTW
java.lang.Exception: DONTWANNA
    at akka.contrib.mailbox.MyActor.doStuff( PeekMailbox.scala:105)
    at akka.contrib.mailbox.MyActor$$anonfun$receive$1.applyOrElse( PeekMailbox.scala:98)
    at akka.actor.ActorCell.receiveMessage( ActorCell.scala:425)
    at akka.actor.ActorCell.invoke( ActorCell.scala:386)
    at akka.dispatch.Mailbox.processMailbox( Mailbox.scala:230)
    at akka.dispatch.Mailbox.run( Mailbox.scala:212)
    at akka.dispatch.ForkJoinExecutorConfigurator$MailboxExecutionTask.exec( AbstractDispatcher.s
    at scala.concurrent.forkjoin.ForkJoinTask.doExec( ForkJoinTask.java:262)
    at scala.concurrent.forkjoin.ForkJoinPool$WorkQueue.runTask( ForkJoinPool.java:975)
    at scala.concurrent.forkjoin.ForkJoinPool.runWorker( ForkJoinPool.java:1478)
    at scala.concurrent.forkjoin.ForkJoinWorkerThread.run( ForkJoinWorkerThread.java:104)
World

```

Normally one would want to make processing idempotent (i.e. it does not matter if a message is processed twice) or `context.become` a different behavior upon restart; the above example included the `println(msg)` call just to demonstrate the re-processing.

Aggregator Pattern

The aggregator pattern supports writing actors that aggregate data from multiple other actors and updates its state based on those responses. It is even harder to optionally aggregate more data based on the runtime state of the actor or take certain actions (sending another message and get another response) based on two or more previous responses.

A common thought is to use the `ask` pattern to request information from other actors. However, `ask` creates another actor specifically for the ask. We cannot use a callback from the future to update the state as the thread executing the callback is not defined. This will likely close-over the current actor.

The aggregator pattern solves such scenarios. It makes sure we’re acting from the same actor in the scope of the actor receive.

Introduction

The aggregator pattern allows match patterns to be dynamically added to and removed from an actor from inside the message handling logic. All match patterns are called from the receive loop and run in the thread handling the incoming message. These dynamically added patterns and logic can safely read and/or modify this actor's mutable state without risking integrity or concurrency issues.

Usage

To use the aggregator pattern, you need to extend the `Aggregator` trait. The trait takes care of `receive` and actors extending this trait should not override `receive`. The trait provides the `expect`, `expectOnce`, and `unexpect` calls. The `expect` and `expectOnce` calls return a handle that can be used for later de-registration by passing the handle to `unexpect`.

`expect` is often used for standing matches such as catching error messages or timeouts.

```
expect {
  case TimedOut => collectBalances(force = true)
}
```

`expectOnce` is used for matching the initial message as well as other requested messages

```
expectOnce {
  case GetCustomerAccountBalances(id, types) =>
    new AccountAggregator(sender(), id, types)
  case _ =>
    sender() ! CantUnderstand
    context.stop(self)
}
```

```
def fetchCheckingAccountsBalance() {
  context.actorOf(Props[CheckingAccountProxy]) ! GetAccountBalances(id)
  expectOnce {
    case CheckingAccountBalances(balances) =>
      results += (Checking -> balances)
      collectBalances()
  }
}
```

`unexpect` can be used for expecting multiple responses until a timeout or when the logic dictates such an expect no longer applies.

```
val handle = expect {
  case Response(name, value) =>
    values += value
    if (values.size > 3) processList()
  case TimedOut => processList()
}

def processList() {
  unexpect(handle)

  if (values.size > 0) {
    context.actorSelection("/user/evaluator") ! values.toList
    expectOnce {
      case EvaluationResults(name, eval) => processFinal(eval)
    }
  } else processFinal(List.empty[Int])
}
```

As the name eludes, `expect` keeps the partial function matching any received messages until `unexpect` is called or the actor terminates, whichever comes first. On the other hand, `expectOnce` removes the partial

function once a match has been established.

It is a common pattern to register the initial `expectOnce` from the construction of the actor to accept the initial message. Once that message is received, the actor starts doing all aggregations and sends the response back to the original requester. The aggregator should terminate after the response is sent (or timed out). A different original request should use a different actor instance.

As you can see, aggregator actors are generally stateful, short lived actors.

Sample Use Case - AccountBalanceRetriever

This example below shows a typical and intended use of the aggregator pattern.

```
import scala.collection._
import scala.concurrent.duration._
import scala.math.BigDecimal.int2bigDecimal

import akka.actor._
/**
 * Sample and test code for the aggregator patter.
 * This is based on Jamie Allen's tutorial at
 * http://jaxenter.com/tutorial-asynchronous-programming-with-akka-actors-46220.html
 */

sealed trait AccountType
case object Checking extends AccountType
case object Savings extends AccountType
case object MoneyMarket extends AccountType

final case class GetCustomerAccountBalances(id: Long, accountTypes: Set[AccountType])
final case class GetAccountBalances(id: Long)

final case class AccountBalances(accountType: AccountType,
                                  balance: Option[List[(Long, BigDecimal)]])

final case class CheckingAccountBalances(balances: Option[List[(Long, BigDecimal)]])
final case class SavingsAccountBalances(balances: Option[List[(Long, BigDecimal)]])
final case class MoneyMarketAccountBalances(balances: Option[List[(Long, BigDecimal)]])

case object TimedOut
case object CantUnderstand

class SavingsAccountProxy extends Actor {
  def receive = {
    case GetAccountBalances(id: Long) =>
      sender() ! SavingsAccountBalances(Some(List((1, 150000), (2, 29000))))
  }
}

class CheckingAccountProxy extends Actor {
  def receive = {
    case GetAccountBalances(id: Long) =>
      sender() ! CheckingAccountBalances(Some(List((3, 15000))))
  }
}

class MoneyMarketAccountProxy extends Actor {
  def receive = {
    case GetAccountBalances(id: Long) =>
      sender() ! MoneyMarketAccountBalances(None)
  }
}

class AccountBalanceRetriever extends Actor with Aggregator {
```

```

import context._

expectOnce {
  case GetCustomerAccountBalances(id, types) =>
    new AccountAggregator(sender(), id, types)
  case _ =>
    sender() ! CantUnderstand
    context.stop(self)
}

class AccountAggregator(originalSender: ActorRef,
                        id: Long, types: Set[AccountType]) {

  val results =
    mutable.ArrayBuffer.empty[(AccountType, Option[List[(Long, BigDecimal)])])

  if (types.size > 0)
    types foreach {
      case Checking    => fetchCheckingAccountsBalance()
      case Savings     => fetchSavingsAccountsBalance()
      case MoneyMarket => fetchMoneyMarketAccountsBalance()
    }
  else collectBalances() // Empty type list yields empty response

  context.system.scheduler.scheduleOnce(1.second, self, TimedOut)
  expect {
    case TimedOut => collectBalances(force = true)
  }

  def fetchCheckingAccountsBalance() {
    context.actorOf(Props[CheckingAccountProxy]) ! GetAccountBalances(id)
    expectOnce {
      case CheckingAccountBalances(balances) =>
        results += (Checking -> balances)
        collectBalances()
    }
  }

  def fetchSavingsAccountsBalance() {
    context.actorOf(Props[SavingsAccountProxy]) ! GetAccountBalances(id)
    expectOnce {
      case SavingsAccountBalances(balances) =>
        results += (Savings -> balances)
        collectBalances()
    }
  }

  def fetchMoneyMarketAccountsBalance() {
    context.actorOf(Props[MoneyMarketAccountProxy]) ! GetAccountBalances(id)
    expectOnce {
      case MoneyMarketAccountBalances(balances) =>
        results += (MoneyMarket -> balances)
        collectBalances()
    }
  }

  def collectBalances(force: Boolean = false) {
    if (results.size == types.size || force) {
      originalSender ! results.toList // Make sure it becomes immutable
      context.stop(self)
    }
  }
}

```

```
}

```

Sample Use Case - Multiple Response Aggregation and Chaining

A shorter example showing aggregating responses and chaining further requests.

```
final case class InitialRequest(name: String)
final case class Request(name: String)
final case class Response(name: String, value: String)
final case class EvaluationResults(name: String, eval: List[Int])
final case class FinalResponse(qualifiedValues: List[String])

/**
 * An actor sample demonstrating use of unexpect and chaining.
 * This is just an example and not a complete test case.
 */
class ChainingSample extends Actor with Aggregator {

  expectOnce {
    case InitialRequest(name) => new MultipleResponseHandler(sender(), name)
  }

  class MultipleResponseHandler(originalSender: ActorRef, propName: String) {

    import context.dispatcher
    import collection.mutable.ArrayBuffer

    val values = ArrayBuffer.empty[String]

    context.actorSelection("/user/request_proxies") ! Request(propName)
    context.system.scheduler.scheduleOnce(50.milliseconds, self, TimedOut)

    val handle = expect {
      case Response(name, value) =>
        values += value
        if (values.size > 3) processList()
      case TimedOut => processList()
    }

    def processList() {
      unexpect(handle)

      if (values.size > 0) {
        context.actorSelection("/user/evaluator") ! values.toList
        expectOnce {
          case EvaluationResults(name, eval) => processFinal(eval)
        }
      } else processFinal(List.empty[Int])
    }

    def processFinal(eval: List[Int]) {
      // Select only the entries coming back from eval
      originalSender ! FinalResponse(eval map values)
      context.stop(self)
    }
  }
}

```

Pitfalls

- The current implementation does not match the sender of the message. This is designed to work with `ActorSelection` as well as `ActorRef`. Without the `sender()`, there is a chance a received message can be matched by more than one partial function. The partial function that was registered via `expect` or `expectOnce` first (chronologically) and is not yet de-registered by `unexpected` takes precedence in this case. Developers should make sure the messages can be uniquely matched or the wrong logic can be executed for a certain message.
- The `sender` referenced in any `expect` or `expectOnce` logic refers to the `sender()` of that particular message and not the `sender()` of the original message. The original `sender()` still needs to be saved so a final response can be sent back.
- `context.become` is not supported when extending the `Aggregator` trait.
- We strongly recommend against overriding `receive`. If your use case really dictates, you may do so with extreme caution. Always provide a pattern match handling aggregator messages among your `receive` pattern matches, as follows:

```
case msg if handleMessage(msg) => // noop
// side effects of handleMessage does the actual match
```

Sorry, there is not yet a Java implementation of the aggregator pattern available.

Receive Pipeline Pattern

The Receive Pipeline Pattern lets you define general interceptors for your messages and plug an arbitrary amount of them into your Actors. It's useful for defining cross aspects to be applied to all or many of your Actors.

Some Possible Use Cases

- Measure the time spent for processing the messages
- Audit messages with an associated author
- Log all important messages
- Secure restricted messages
- Text internationalization

Interceptors

Multiple interceptors can be added to actors that mixin the `ReceivePipeline` trait. These interceptors internally define layers of decorators around the actor's behavior. The first interceptor defines an outer decorator which delegates to a decorator corresponding to the second interceptor and so on, until the last interceptor which defines a decorator for the actor's `Receive`.

The first or outermost interceptor receives messages sent to the actor.

For each message received by an interceptor, the interceptor will typically perform some processing based on the message and decide whether or not to pass the received message (or some other message) to the next interceptor.

An `Interceptor` is a type alias for `PartialFunction[Any, Delegation]`. The `Any` input is the message it receives from the previous interceptor (or, in the case of the first interceptor, the message that was sent to the actor). The `Delegation` return type is used to control what (if any) message is passed on to the next interceptor.

A simple example

To pass a transformed message to the actor (or next inner interceptor) an interceptor can return `Inner(newMsg)` where `newMsg` is the transformed message.

The following interceptor shows this. It intercepts `Int` messages, adds one to them and passes on the incremented value to the next interceptor.

```
val incrementInterceptor: Interceptor = {
  case i: Int => Inner(i + 1)
}
```

Building the Pipeline

To give your Actor the ability to pipeline received messages in this way, you'll need to mixin with the `ReceivePipeline` trait. It has two methods for controlling the pipeline, `pipelineOuter` and `pipelineInner`, both receiving an `Interceptor`. The first one adds the interceptor at the beginning of the pipeline and the second one adds it at the end, just before the current Actor's behavior.

In this example we mixin our Actor with the `ReceivePipeline` trait and we add `Increment` and `Double` interceptors with `pipelineInner`. They will be applied in this very order.

```
class PipelinedActor extends Actor with ReceivePipeline {

  // Increment
  pipelineInner { case i: Int => Inner(i + 1) }
  // Double
  pipelineInner { case i: Int => Inner(i * 2) }

  def receive: Receive = { case any => println(any) }
}

actor ! 5 // prints 12 = (5 + 1) * 2
```

If we add `Double` with `pipelineOuter` it will be applied before `Increment` so the output is 11

```
// Increment
pipelineInner { case i: Int => Inner(i + 1) }
// Double
pipelineOuter { case i: Int => Inner(i * 2) }

// prints 11 = (5 * 2) + 1
```

Interceptors Mixin

Defining all the pipeline inside the Actor implementation is good for showing up the pattern, but it isn't very practical. The real flexibility of this pattern comes when you define every interceptor in its own trait and then you mixin any of them into your Actors.

Let's see it in an example. We have the following model:

```
val texts = Map(
  "that.rug_EN" -> "That rug really tied the room together.",
  "your.opinion_EN" -> "Yeah, well, you know, that's just, like, your opinion, man.",
  "that.rug_ES" -> "Esa alfombra realmente completaba la sala.",
  "your.opinion_ES" -> "Sí, bueno, ya sabes, eso es solo, como, tu opinion, amigo.")

case class I18nText(locale: String, key: String)
case class Message(author: Option[String], text: Any)
```

and these two interceptors defined, each one in its own trait:

```

trait I18nInterceptor {
  this: ReceivePipeline =>

  pipelineInner {
    case m @ Message(_, I18nText(loc, key)) =>
      Inner(m.copy(text = texts(s"${key}_$loc")))
  }
}

trait AuditInterceptor {
  this: ReceivePipeline =>

  pipelineOuter {
    case m @ Message(Some(author), text) =>
      println(s"$author is about to say: $text")
      Inner(m)
  }
}

```

The first one intercepts any messages having an internationalized text and replaces it with the resolved text before resuming with the chain. The second one intercepts any message with an author defined and prints it before resuming the chain with the message unchanged. But since `I18n` adds the interceptor with `pipelineInner` and `Audit` adds it with `pipelineOuter`, the audit will happen before the internationalization.

So if we mixin both interceptors in our Actor, we will see the following output for these example messages:

```

class PrinterActor extends Actor with ReceivePipeline
  with I18nInterceptor with AuditInterceptor {

  override def receive: Receive = {
    case Message(author, text) =>
      println(s"${author.getOrElse("Unknown")} says '$text'")
  }
}

printerActor ! Message(Some("The Dude"), I18nText("EN", "that.rug"))
// The Dude is about to say: I18nText(EN,that.rug)
// The Dude says 'That rug really tied the room together.'

printerActor ! Message(Some("The Dude"), I18nText("EN", "your.opinion"))
// The Dude is about to say: I18nText(EN,your.opinion)
// The Dude says 'Yeah, well, you know, that's just, like, your opinion, man.'

```

Unhandled Messages

With all that behaviors chaining occurring, what happens to unhandled messages? Let me explain it with a simple rule.

Note: Every message not handled by an interceptor will be passed to the next one in the chain. If none of the interceptors handles a message, the current Actor's behavior will receive it, and if the behavior doesn't handle it either, it will be treated as usual with the unhandled method.

But sometimes it is desired for interceptors to break the chain. You can do it by explicitly indicating that the message has been completely handled by the interceptor by returning `HandledCompletely`.

```

case class PrivateMessage(userId: Option[Long], msg: Any)

trait PrivateInterceptor {
  this: ReceivePipeline =>

```

```

pipelineInner {
  case PrivateMessage(Some(userId), msg) =>
    if (isGranted(userId))
      Inner(msg)
    else
      HandledCompletely
}
}

```

Processing after delegation

But what if you want to perform some action after the actor has processed the message (for example to measure the message processing time)?

In order to support such use cases, the `Inner` return type has a method `andAfter` which accepts a code block that can perform some action after the message has been processed by subsequent inner interceptors.

The following is a simple interceptor that times message processing:

```

trait TimerInterceptor extends ActorLogging {
  this: ReceivePipeline =>

  def logTimeTaken(time: Long) = log.debug(s"Time taken: $time ns")

  pipelineOuter {
    case e =>
      val start = System.nanoTime
      Inner(e).andAfter {
        val end = System.nanoTime
        logTimeTaken(end - start)
      }
  }
}

```

Note: The `andAfter` code blocks are run on return from handling the message with the next inner handler and on the same thread. It is therefore safe for the `andAfter` logic to close over the interceptor's state.

Using Receive Pipelines with Persistence

When using `ReceivePipeline` together with `PersistentActor` make sure to mix-in the traits in the following order for them to properly co-operate:

```

class ExampleActor extends PersistentActor with ReceivePipeline {
  /* ... */
}

```

The order is important here because of how both traits use internal “around” methods to implement their features, and if mixed-in the other way around it would not work as expected. If you want to learn more about how exactly this works, you can read up on Scala's [type linearization mechanism](#);

9.5.3 Suggested Way of Using these Contributions

Since the Akka team does not restrict updates to this subproject even during otherwise binary compatible releases, and modules may be removed without deprecation, it is suggested to copy the source files into your own code base, changing the package name. This way you can choose when to update or which fixes to include (to keep binary compatibility if needed) and later releases of Akka do not potentially break your application.

9.5.4 Suggested Format of Contributions

Each contribution should be a self-contained unit, consisting of one source file or one exclusively used package, without dependencies to other modules in this subproject; it may depend on everything else in the Akka distribution, though. This ensures that contributions may be moved into the standard distribution individually. The module shall be within a subpackage of `akka.contrib`.

Each module must be accompanied by a test suite which verifies that the provided features work, possibly complemented by integration and unit tests. The tests should follow the *Developer Guidelines* and go into the `src/test/scala` or `src/test/java` directories (with package name matching the module which is being tested). As an example, if the module were called `akka.contrib.pattern.ReliableProxy`, then the test suite should be called `akka.contrib.pattern.ReliableProxySpec`.

Each module must also have proper documentation in *reStructured Text* format. The documentation should be a single `<module>.rst` file in the `akka-contrib/docs` directory, including a link from `index.rst` (this file).

INFORMATION FOR AKKA DEVELOPERS

10.1 Building Akka

This page describes how to build and run Akka from the latest source code.

10.1.1 Get the Source Code

Akka uses [Git](#) and is hosted at [Github](#).

You first need [Git](#) installed on your machine. You can then clone the source repository from <http://github.com/akka/akka>.

For example:

```
git clone git://github.com/akka/akka.git
```

If you have already cloned the repository previously then you can update the code with `git pull`:

```
git pull origin master
```

10.1.2 sbt - Simple Build Tool

Akka is using the excellent [sbt](#) build system. So the first thing you have to do is to download and install [sbt](#). You can read more about how to do that in the [sbt setup](#) documentation.

The [sbt](#) commands that you'll need to build Akka are all included below. If you want to find out more about [sbt](#) and using it for your own projects do read the [sbt documentation](#).

The Akka [sbt](#) build file is `project/AkkaBuild.scala`.

10.1.3 Building Akka

First make sure that you are in the akka code directory:

```
cd akka
```

Building

To compile all the Akka core modules use the `compile` command:

```
sbt compile
```

You can run all tests with the `test` command:

```
sbt test
```

If compiling and testing are successful then you have everything working for the latest Akka development version.

Parallel Execution

By default the tests are executed sequentially. They can be executed in parallel to reduce build times, if hardware can handle the increased memory and cpu usage. Add the following system property to sbt launch script to activate parallel execution:

```
-Dakka.parallelExecution=true
```

Long Running and Time Sensitive Tests

By default are the long running tests (mainly cluster tests) and time sensitive tests (dependent on the performance of the machine it is running on) disabled. You can enable them by adding one of the flags:

```
-Dakka.test.tags.include=long-running
-Dakka.test.tags.include=timing
```

Or if you need to enable them both:

```
-Dakka.test.tags.include=long-running,timing
```

Publish to Local Ivy Repository

If you want to deploy the artifacts to your local Ivy repository (for example, to use from an sbt project) use the `publish-local` command:

```
sbt publish-local
```

sbt Interactive Mode

Note that in the examples above we are calling `sbt compile` and `sbt test` and so on, but sbt also has an interactive mode. If you just run `sbt` you enter the interactive sbt prompt and can enter the commands directly. This saves starting up a new JVM instance for each command and can be much faster and more convenient.

For example, building Akka as above is more commonly done like this:

```
% sbt
[info] Set current project to default (in build file:/.../akka/project/plugins/)
[info] Set current project to akka (in build file:/.../akka/)
> compile
...
> test
...
```

sbt Batch Mode

It's also possible to combine commands in a single call. For example, testing, and publishing Akka to the local Ivy repository can be done with:

```
sbt test publish-local
```

10.1.4 Dependencies

You can look at the Ivy dependency resolution information that is created on `sbt update` and found in `~/.ivy2/cache`. For example, the `~/.ivy2/cache/com.typesafe.akka-akka-remote-compile.xml` file contains the resolution information for the akka-remote module compile dependencies. If you open this file in a web browser you will get an easy to navigate view of dependencies.

10.1.5 Scaladoc Dependencies

Akka generates class diagrams for the API documentation using ScalaDoc. This needs the `dot` command from the Graphviz software package to be installed to avoid errors. You can disable the diagram generation by adding the flag `-Dakka.scaladoc.diagrams=false`. After installing Graphviz, make sure you add the toolset to the PATH (definitely on Windows).

10.2 Multi JVM Testing

Supports running applications (objects with main methods) and `ScalaTest` tests in multiple JVMs at the same time. Useful for integration testing where multiple systems communicate with each other.

10.2.1 Setup

The multi-JVM testing is an sbt plugin that you can find at <http://github.com/typesafehub/sbt-multi-jvm>.

You can add it as a plugin by adding the following to your `project/plugins.sbt`:

```
addSbtPlugin("com.typesafe.sbt" % "sbt-multi-jvm" % "0.3.8")
```

You can then add multi-JVM testing to `build.sbt` or `project/Build.scala` by including the `MultiJvm` settings and config. Please note that `MultiJvm` test sources are located in `src/multi-jvm/...`, and not in `src/test/...`

Here is an example `build.sbt` file for sbt 0.13 that uses the `MultiJvm` plugin:

```
import com.typesafe.sbt.SbtMultiJvm
import com.typesafe.sbt.SbtMultiJvm.MultiJvmKeys.MultiJvm

val akkaVersion = "2.4.0"

val project = Project(
  id = "akka-sample-multi-node-scala",
  base = file("."),
  settings = Project.defaultSettings ++ SbtMultiJvm.multiJvmSettings ++ Seq(
    name := "akka-sample-multi-node-scala",
    version := "2.4.0",
    scalaVersion := "2.11.7",
    libraryDependencies ++= Seq(
      "com.typesafe.akka" %% "akka-actor" % akkaVersion,
      "com.typesafe.akka" %% "akka-remote" % akkaVersion,
      "com.typesafe.akka" %% "akka-multi-node-testkit" % akkaVersion,
      "org.scalatest" %% "scalatest" % "2.2.1" % "test"),
    // make sure that MultiJvm test are compiled by the default test compilation
    compile in MultiJvm <=<= (compile in MultiJvm) triggeredBy (compile in Test),
    // disable parallel tests
    parallelExecution in Test := false,
    // make sure that MultiJvm tests are executed by the default test target,
    // and combine the results from ordinary test and multi-jvm tests
    executeTests in Test <=<= (executeTests in Test, executeTests in MultiJvm) map {
```

```

    case (testResults, multiNodeResults) =>
      val overall =
        if (testResults.overall.id < multiNodeResults.overall.id)
          multiNodeResults.overall
        else
          testResults.overall
      Tests.Output(overall,
        testResults.events ++ multiNodeResults.events,
        testResults.summaries ++ multiNodeResults.summaries)
  }
)
) configs (MultiJvm)

```

You can specify JVM options for the forked JVMs:

```
jvmOptions in MultiJvm := Seq("-Xmx256M")
```

10.2.2 Running tests

The multi-JVM tasks are similar to the normal tasks: `test`, `test-only`, and `run`, but are under the `multi-jvm` configuration.

So in Akka, to run all the multi-JVM tests in the akka-remote project use (at the sbt prompt):

```
akka-remote-tests/multi-jvm:test
```

Or one can change to the akka-remote-tests project first, and then run the tests:

```
project akka-remote-tests
multi-jvm:test
```

To run individual tests use `test-only`:

```
multi-jvm:test-only akka.remote.RandomRoutedRemoteActor
```

More than one test name can be listed to run multiple specific tests. Tab-completion in sbt makes it easy to complete the test names.

It's also possible to specify JVM options with `test-only` by including those options after the test names and `--`. For example:

```
multi-jvm:test-only akka.remote.RandomRoutedRemoteActor -- -Dsome.option=something
```

10.2.3 Creating application tests

The tests are discovered, and combined, through a naming convention. MultiJvm test sources are located in `src/multi-jvm/...` A test is named with the following pattern:

```
{TestName}MultiJvm{NodeName}
```

That is, each test has `MultiJvm` in the middle of its name. The part before it groups together tests/applications under a single `TestName` that will run together. The part after, the `NodeName`, is a distinguishing name for each forked JVM.

So to create a 3-node test called `Sample`, you can create three applications like the following:

```

package sample

object SampleMultiJvmNode1 {
  def main(args: Array[String]) {
    println("Hello from node 1")
  }
}

```

```

}

object SampleMultiJvmNode2 {
  def main(args: Array[String]) {
    println("Hello from node 2")
  }
}

object SampleMultiJvmNode3 {
  def main(args: Array[String]) {
    println("Hello from node 3")
  }
}

```

When you call `multi-jvm:run sample.Sample` at the sbt prompt, three JVMs will be spawned, one for each node. It will look like this:

```

> multi-jvm:run sample.Sample
...
[info] * sample.Sample
[JVM-1] Hello from node 1
[JVM-2] Hello from node 2
[JVM-3] Hello from node 3
[success] Total time: ...

```

10.2.4 Changing Defaults

You can change the name of the multi-JVM test source directory by adding the following configuration to your project:

```

unmanagedSourceDirectories in MultiJvm <<=
  Seq(baseDirectory(_ / "src/some_directory_here")).join

```

You can change what the `MultiJvm` identifier is. For example, to change it to `ClusterTest` use the `multiJvmMarker` setting:

```

multiJvmMarker in MultiJvm := "ClusterTest"

```

Your tests should now be named `{TestName}ClusterTest{NodeName}`.

10.2.5 Configuration of the JVM instances

You can define specific JVM options for each of the spawned JVMs. You do that by creating a file named after the node in the test with suffix `.opts` and put them in the same directory as the test.

For example, to feed the JVM options `-Dakka.remote.port=9991` and `-Xmx256m` to the `SampleMultiJvmNode1` let's create three `*.opts` files and add the options to them. Separate multiple options with space.

`SampleMultiJvmNode1.opts:`

```
-Dakka.remote.port=9991 -Xmx256m
```

`SampleMultiJvmNode2.opts:`

```
-Dakka.remote.port=9992 -Xmx256m
```

`SampleMultiJvmNode3.opts:`

```
-Dakka.remote.port=9993 -Xmx256m
```

10.2.6 ScalaTest

There is also support for creating ScalaTest tests rather than applications. To do this use the same naming convention as above, but create ScalaTest suites rather than objects with main methods. You need to have ScalaTest on the classpath. Here is a similar example to the one above but using ScalaTest:

```
package sample

import org.scalatest.WordSpec
import org.scalatest.matchers.MustMatchers

class SpecMultiJvmNode1 extends WordSpec with MustMatchers {
  "A node" should {
    "be able to say hello" in {
      val message = "Hello from node 1"
      message must be("Hello from node 1")
    }
  }
}

class SpecMultiJvmNode2 extends WordSpec with MustMatchers {
  "A node" should {
    "be able to say hello" in {
      val message = "Hello from node 2"
      message must be("Hello from node 2")
    }
  }
}
```

To run just these tests you would call `multi-jvm:test-only sample.Spec` at the sbt prompt.

10.2.7 Multi Node Additions

There has also been some additions made to the `SbtMultiJvm` plugin to accommodate the *experimental* module *multi node testing*, described in that section.

10.3 I/O Layer Design

The `akka.io` package has been developed in collaboration between the Akka and `spray.io` teams. Its design incorporates the experiences with the `spray-io` module along with improvements that were jointly developed for more general consumption as an actor-based service.

10.3.1 Requirements

In order to form a general and extensible IO layer basis for a wide range of applications, with Akka remoting and spray HTTP being the initial ones, the following requirements were established as key drivers for the design:

- scalability to millions of concurrent connections
- lowest possible latency in getting data from an input channel into the target actor's mailbox
- maximal throughput
- optional back-pressure in both directions (i.e. throttling local senders as well as allowing local readers to throttle remote senders, where allowed by the protocol)
- a purely actor-based API with immutable data representation
- extensibility for integrating new transports by way of a very lean SPI; the goal is to not force I/O mechanisms into a lowest common denominator but instead allow completely protocol-specific user-level APIs.

10.3.2 Basic Architecture

Each transport implementation will be made available as a separate Akka extension, offering an `ActorRef` representing the initial point of contact for client code. This “manager” accepts requests for establishing a communications channel (e.g. connect or listen on a TCP socket). Each communications channel is represented by one dedicated actor, which is exposed to client code for all interaction with this channel over its entire lifetime.

The central element of the implementation is the transport-specific “selector” actor; in the case of TCP this would wrap a `java.nio.channels.Selector`. The channel actors register their interest in readability or writability of their channel by sending corresponding messages to their assigned selector actor. However, the actual channel reading and writing is performed by the channel actors themselves, which frees the selector actors from time-consuming tasks and thereby ensures low latency. The selector actor’s only responsibility is the management of the underlying selector’s key set and the actual select operation, which is the only operation to typically block.

The assignment of channels to selectors is performed by the manager actor and remains unchanged for the entire lifetime of a channel. Thereby the management actor “stripes” new channels across one or more selector actors based on some implementation-specific distribution logic. This logic may be delegated (in part) to the selectors actors, which could, for example, choose to reject the assignment of a new channel when they consider themselves to be at capacity.

The manager actor creates (and therefore supervises) the selector actors, which in turn create and supervise their channel actors. The actor hierarchy of one single transport implementation therefore consists of three distinct actor levels, with the management actor at the top-, the channel actors at the leaf- and the selector actors at the mid-level.

Back-pressure for output is enabled by allowing the user to specify within its `Write` messages whether it wants to receive an acknowledgement for enqueueing that write to the O/S kernel. Back-pressure for input is enabled by sending the channel actor a message which temporarily disables read interest for the channel until reading is re-enabled with a corresponding resume command. In the case of transports with flow control—like TCP—the act of not consuming data at the receiving end (thereby causing them to remain in the kernels read buffers) is propagated back to the sender, linking these two mechanisms across the network.

10.3.3 Design Benefits

Staying within the actor model for the whole implementation allows us to remove the need for explicit thread handling logic, and it also means that there are no locks involved (besides those which are part of the underlying transport library). Writing only actor code results in a cleaner implementation, while Akka’s efficient actor messaging does not impose a high tax for this benefit. In fact the event-based nature of I/O maps so well to the actor model that we expect clear performance and especially scalability benefits over traditional solutions with explicit thread management and synchronization.

Another benefit of supervision hierarchies is that clean-up of resources comes naturally: shutting down a selector actor will automatically clean up all channel actors, allowing proper closing of the channels and sending the appropriate messages to user-level client actors. `DeathWatch` allows the channel actors to notice the demise of their user-level handler actors and terminate in an orderly fashion in that case as well; this naturally reduces the chances of leaking open channels.

The choice of using `ActorRef` for exposing all functionality entails that these references can be distributed or delegated freely and in general handled as the user sees fit, including the use of remoting and life-cycle monitoring (just to name two).

10.3.4 How to go about Adding a New Transport

The best start is to study the TCP reference implementation to get a good grip on the basic working principle and then design an implementation, which is similar in spirit, but adapted to the new protocol in question. There are vast differences between I/O mechanisms (e.g. compare file I/O to a message broker) and the goal of this I/O layer is explicitly **not** to shoehorn all of them into a uniform API, which is why only the basic architecture ideas are documented here.

10.4 Developer Guidelines

Note: First read [The Akka Contributor Guidelines](#).

10.4.1 Code Style

The Akka code style follows the [Scala Style Guide](#). The only exception is the style of block comments:

```
/**
 * Style mandated by "Scala Style Guide"
 */

/**
 * Style adopted in the Akka codebase
 */
```

Akka is using `Scalariform` to format the source code as part of the build. So just hack away and then run `sbt compile` and it will reformat the code according to Akka standards.

10.4.2 Process

The full process is described in [The Akka Contributor Guidelines](#). In summary:

- Make sure you have signed the Akka CLA, if not, [sign it online](#).
- Pick a ticket, if there is no ticket for your work then create one first.
- Fork `akka/akka`. Start working in a feature branch.
- When you are done, create a GitHub Pull-Request towards the targeted branch.
- When there's consensus on the review, someone from the Akka Core Team will merge it.

10.4.3 Commit messages

Please follow the conventions described in [The Akka Contributor Guidelines](#) when creating public commits and writing commit messages.

10.4.4 Testing

All code that is checked in **should** have tests. All testing is done with `ScalaTest` and `ScalaCheck`.

- Name tests as **Test.scala** if they do not depend on any external stuff. That keeps surefire happy.
- Name tests as **Spec.scala** if they have external dependencies.

Actor TestKit

There is a useful test kit for testing actors: `akka.util.TestKit`. It enables assertions concerning replies received and their timing, there is more documentation in the `akka-testkit` module.

Multi-JVM Testing

Included in the example is an sbt trait for multi-JVM testing which will fork JVMs for multi-node testing. There is support for running applications (objects with main methods) and running `ScalaTest` tests.

NetworkFailureTest

You can use the ‘NetworkFailureTest’ trait to test network failure.

10.5 Documentation Guidelines

The Akka documentation uses `reStructuredText` as its markup language and is built using `Sphinx`.

10.5.1 Sphinx

For more details see [The Sphinx Documentation](#)

10.5.2 reStructuredText

For more details see [The reST Quickref](#)

Sections

Section headings are very flexible in reST. We use the following convention in the Akka documentation:

- # (over and under) for module headings
- = for sections
- – for subsections
- ^ for subsubsections
- ~ for subsubsubsections

Cross-referencing

Sections that may be cross-referenced across the documentation should be marked with a reference. To mark a section use `.. _ref-name:` before the section heading. The section can then be linked with `:ref: `ref-name``. These are unique references across the entire documentation.

For example:

```
.. _akka-module:

#####
Akka Module
#####

This is the module documentation.

.. _akka-section:

Akka Section
=====

Akka Subsection
-----

Here is a reference to "akka section": :ref:`akka-section` which will have the
name "Akka Section".
```

10.5.3 Build the documentation

First install [Sphinx](#). See below.

Building

For the html version of the docs:

```
sbt sphinx:generateHtml
open <project-dir>/akka-docs/target/sphinx/html/index.html
```

For the pdf version of the docs:

```
sbt sphinx:generatePdf
open <project-dir>/akka-docs/target/sphinx/latex/AkkaJava.pdf
or
open <project-dir>/akka-docs/target/sphinx/latex/AkkaScala.pdf
```

Installing Sphinx on OS X

Install [Homebrew](#)

Install Python with Homebrew:

```
brew install python
```

Homebrew will automatically add Python executable to your \$PATH and pip is a part of the default Python installation with Homebrew.

More information in case of trouble: <https://github.com/mxcl/homebrew/wiki/Homebrew-and-Python>

Install sphinx:

```
pip install sphinx
```

Install BasicTeX package from: <http://www.tug.org/mactex/morepackages.html>

Add texlive bin to \$PATH:

```
export TEXLIVE_PATH=/usr/local/texlive/2015basic/bin/universal-darwin
export PATH=$TEXLIVE_PATH:$PATH
```

Add missing tex packages:

```
sudo tlmgr update --self
sudo tlmgr install titlesec
sudo tlmgr install framed
sudo tlmgr install threeparttable
sudo tlmgr install wrapfig
sudo tlmgr install helvetic
sudo tlmgr install courier
sudo tlmgr install multirow
```

If you get the error “unknown locale: UTF-8” when generating the documentation the solution is to define the following environment variables:

```
export LANG=en_US.UTF-8
export LC_ALL=en_US.UTF-8
```

10.6 Team

Name	Role
Jonas Bonér	Founder, Despot, Committer
Viktor Klang	Honorary Member
Roland Kuhn	Project Lead
Patrik Nordwall	Core Team
Björn Antonsson	Core Team
Endre Varga	Core Team
Mathias Doenitz	Committer
Johannes Rudolph	Committer
Raymond Roestenburg	Committer
Piotr Gabryanczyk	Committer
Helena Edelson	Committer
Martin Krasser	Committer
Henrik Engström	Alumnus
Peter Vlugter	Alumnus
Derek Williams	Alumnus
Debasish Ghosh	Alumnus
Ross McDonald	Alumnus
Eckhart Hertzler	Alumnus
Mikael Höggqvist	Alumnus
Tim Perrett	Alumnus
Jeanfrancois Arcand	Alumnus
Jan Van Besien	Alumnus
Michael Kober	Alumnus
Peter Veentjer	Alumnus
Irmo Manie	Alumnus
Heiko Seeberger	Alumnus
Hiram Chirino	Alumnus
Scott Clasen	Alumnus

PROJECT INFORMATION

11.1 Migration Guides

11.1.1 Migration Guide 1.3.x to 2.0.x

Migration from 1.3.x to 2.0.x is described in the [documentation of 2.0](#).

11.1.2 Migration Guide 2.0.x to 2.1.x

Migration from 2.0.x to 2.1.x is described in the [documentation of 2.1](#).

11.1.3 Migration Guide 2.1.x to 2.2.x

Migration from 2.1.x to 2.2.x is described in the [documentation of 2.2](#).

11.1.4 Migration Guide 2.2.x to 2.3.x

Migration from 2.2.x to 2.3.x is described in the [documentation of 2.3](#).

11.1.5 Migration Guide Akka Persistence (experimental) 2.3.3 to 2.3.4 (and 2.4.x)

Akka Persistence is an **experimental module**, which means that neither Binary Compatibility nor API stability is provided for Persistence while under the *experimental* flag. The goal of this phase is to gather user feedback before we freeze the APIs in a major release.

defer renamed to deferAsync

The `defer` method in `PersistentActor` was renamed to `deferAsync` as it matches the semantics of `persistAsync` more closely than `persist`, which was causing confusion for users.

Its semantics remain unchanged.

Renamed EventsourcedProcessor to PersistentActor

`EventsourcedProcessor` is now deprecated and replaced by `PersistentActor` which provides the same (and more) API. Migrating to 2.4.x is as simple as changing all your classes to extending `PersistentActor`.

Replace all classes like:

```
class DeprecatedProcessor extends EventsourcedProcessor {
  def processorId = "id"
  /*...*/
}
```

To extend `PersistentActor`:

```
class NewPersistentProcessor extends PersistentActor {
  def persistenceId = "id"
  /*...*/
}
```

Read more about the persistent actor in the *documentation for Scala* and *documentation for Java*.

Changed `processorId` to (abstract) `persistenceId`

In Akka Persistence 2.3.3 and previously, the main building block of applications were `Processors`. `Persistent` messages, as well as processors implemented the `processorId` method to identify which persistent entity a message belonged to.

This concept remains the same in Akka 2.3.4, yet we rename `processorId` to `persistenceId` because `Processors` will be removed, and persistent messages can be used from different classes not only `PersistentActor` (`Views`, directly from `Journals` etc).

Please note that `persistenceId` is **abstract** in the new API classes (`PersistentActor` and `PersistentView`), and we do **not** provide a default (actor-path derived) value for it like we did for `processorId`. The rationale behind this change being stricter de-coupling of your Actor hierarchy and the logical “which persistent entity this actor represents”. A longer discussion on this subject can be found on [issue #15436](#) on github.

In case you want to preserve the old behavior of providing the actor’s path as the default `persistenceId`, you can easily implement it yourself either as a helper trait or simply by overriding `persistenceId` as follows:

```
override def persistenceId = self.path.toStringWithoutAddress
```

We provided the renamed method also on already deprecated classes (`Channels`), so you can simply apply a global rename of `processorId` to `persistenceId`.

Removed `Processor` in favour of extending `PersistentActor` with `persistAsync`

The `Processor` is now deprecated since 2.3.4 and will be removed in 2.4.x. It’s semantics replicated in `PersistentActor` in the form of an additional `persist` method: `persistAsync`.

In essence, the difference between `persist` and `persistAsync` is that the former will stash all incoming commands until all `persist` callbacks have been processed, whereas the latter does not stash any commands. The new `persistAsync` should be used in cases of low consistency yet high responsiveness requirements, the Actor can keep processing incoming commands, even though not all previous events have been handled.

When these `persist` and `persistAsync` are used together in the same `PersistentActor`, the `persist` logic will win over the `async` version so that all guarantees concerning `persist` still hold. This will however lower the throughput

Now deprecated code using `Processor`:

```
class OldProcessor extends Processor {
  override def processorId = "user-wallet-1337"

  def receive = {
    case Persistent(cmd) => sender() ! cmd
  }
}
```

Replacement code, with the same semantics, using `PersistentActor`:

```
class NewProcessor extends PersistentActor {
  override def persistenceId = "user-wallet-1337"

  def receiveCommand = {
    case cmd =>
      persistAsync(cmd) { e => sender() ! e }
  }

  def receiveRecover = {
    case _ => // logic for handling replay
  }
}
```

It is worth pointing out that using `sender()` inside the `persistAsync` callback block is **valid**, and does *not* suffer any of the problems `Futures` have when closing over the sender reference.

Using the `PersistentActor` instead of `Processor` also shifts the responsibility of deciding if a message should be persisted to the receiver instead of the sender of the message. Previously, using `Processor`, clients would have to wrap messages as `Persistent(cmd)` manually, as well as have to be aware of the receiver being a `Processor`, which didn't play well with transparency of the `ActorRefs` in general.

How to migrate data from `Processor` to `PersistentActor`

The recommended approach for migrating persisted messages from a `Processor` to events that can be replayed by a `PersistentActor` is to write a custom migration tool with a `PersistentView` and a `PersistentActor`. Connect the `PersistentView` to the `persistenceId` of the old `Processor` to replay the stored persistent messages. Send the messages from the view to a `PersistentActor` with another `persistenceId`. There you can transform the old messages to domain events that the real `PersistentActor` will be able to understand. Store the events with `persistAsync`.

Note that you can implement back-pressure between the writing `PersistentActor` and the reading `PersistentView` by turning off auto-update in the view and send custom `Update` messages to the view with a limited `replayMax` value.

Removed `deleteMessage`

`deleteMessage` is deprecated and will be removed. When using command sourced `Processor` the command was stored before it was received and could be validated and then there was a reason to remove faulty commands to avoid repeating the error during replay. When using `PersistentActor` you can always validate the command before persisting and therefore the stored event (or command) should always be valid for replay.

`deleteMessages` can still be used for pruning of all messages up to a sequence number.

Renamed `View` to `PersistentView`, which receives plain messages (`Persistent()` wrapper is gone)

Views used to receive messages wrapped as `Persistent(payload, seqNr)`, this is no longer the case and views receive the `payload` as message from the `Journal` directly. The rationale here is that the wrapper approach was inconsistent with the other Akka Persistence APIs and also is not easily “discoverable” (you have to *know* you will be getting this `Persistent` wrapper).

Instead, since 2.3.4, views get plain messages, and can use additional methods provided by the `View` to identify if a message was sent from the `Journal` (had been played back to the view). So if you had code like this:

```
class AverageView extends View {
  override def processorId = "average-view"

  def receive = {
    case Persistent(msg, seqNr) =>
```

```
// from Journal

case msg =>
  // from user-land
}
```

You should update it to extend `PersistentView` instead:

```
class AverageView extends PersistentView {
  override def persistenceId = "persistence-sample"
  override def viewId = "persistence-sample-average"

  def receive = {
    case msg if isPersistent =>
      // from Journal
      val seqNr = lastSequenceNr // in case you require the sequence number

    case msg =>
      // from user-land
  }
}
```

In case you need to obtain the current sequence number the view is looking at, you can use the `lastSequenceNr` method. It is equivalent to “current sequence number”, when `isPersistent` returns true, otherwise it yields the sequence number of the last persistent message that this view was updated with.

Removed Channel and PersistentChannel in favour of AtLeastOnceDelivery trait

One of the primary tasks of a `Channel` was to de-duplicate messages that were sent from a `Processor` during recovery. Performing external side effects during recovery is not encouraged with event sourcing and therefore the `Channel` is not needed for this purpose.

The `Channel` and `PersistentChannel` also performed at-least-once delivery of messages, but it did not free a sending actor from implementing retransmission or timeouts, since the acknowledgement from the channel is needed to guarantee safe hand-off. Therefore at-least-once delivery is provided in a new `AtLeastOnceDelivery` trait that is mixed-in to the persistent actor on the sending side.

Read more about at-least-once delivery in the [documentation for Scala](#) and [documentation for Java](#).

Default persistence plugins

Previously default `akka.persistence.journal.plugin` was set to the `LevelDB journal akka.persistence.journal.leveldb` and default `akka.persistence.snapshot-store.plugin` was set to the local file-system snapshot `akka.persistence.snapshot-store.local`. Now default journal and default snapshot-store plugins are set to empty “” in the persistence extension `reference.conf`, and require explicit user configuration via override in the user `application.conf`. This change was needed to decouple persistence extension from the LevelDB dependency, and to support multiple plugin configurations. Please see `persistence extension reference.conf` for details.

Converted LevelDB to an optional dependency

Persistence extension uses LevelDB based plugins for own development and keeps related code in the published jar. However previously LevelDB was a `compile` scope dependency, and now it is an `optional;provided` dependency. To continue using LevelDB based persistence plugins it is now required for related user projects to include an additional explicit dependency declaration for the LevelDB artifacts. This change allows production akka deployments to avoid need for the LevelDB provisioning. Please see `persistence extension reference.conf` for details.

11.1.6 Migration Guide Eventsourced to Akka Persistence 2.3.x

General notes

Eventsourced and Akka *persistence-scala* share many high-level concepts but strongly differ on design, implementation and usage level. This migration guide is more a higher-level comparison of Eventsourced and Akka Persistence rather than a sequence of low-level instructions how to transform Eventsourced application code into Akka Persistence application code. This should provide a good starting point for a migration effort. Development teams should consult the user documentation of both projects for further details.

Scope of this migration guide is code migration, not journal migration. Journals written by Eventsourced can neither be used directly Akka Persistence nor migrated to Akka Persistence compatible journals. Journal migration tools may be provided in the future but do not exist at the moment.

Extensions

Eventsourced and Akka Persistence are both *extending-akka-scala*.

Eventsourced: `EventsourcingExtension`

- Must be explicitly created with an actor system and an application-defined journal actor as arguments. (see [example usage](#)).
- `Processors` and `Channels` must be created with the factory methods `processorOf` and `channelOf` defined on `EventsourcingExtension`.
- Is a central registry of created processors and channels.

Akka Persistence: `Persistence` extension

- Must **not** be explicitly created by an application. A `Persistence` extension is implicitly created upon first `PersistentActor`' creation. Journal actors are automatically created from a journal plugin configuration (see [journal-plugin-api](#)).
- `PersistentActor` can be created like any other actor with `actorOf` without using the `Persistence` extension.
- Is **not** a central registry of persistent actors.

Processors / PersistentActor

Eventsourced: `Eventsourced`

- Stackable trait that adds journaling (write-ahead-logging) to actors (see processor [definition](#) and [creation](#)). Name `Eventsourced` caused some confusion in the past as many examples used `Eventsourced` processors for *command sourcing*. See also [this FAQ entry](#) for clarification.
- Must be explicitly [recovered](#) using recovery methods on `EventsourcingExtension`. Applications are responsible for avoiding an interference of replayed messages and new messages i.e. applications have to explicitly wait for recovery to complete. Recovery on processor re-start is not supported out-of-the box.
- Journaling-preserving [behavior changes](#) are only possible with special-purpose methods `become` and `unbecome`, in addition to non-journaling-preserving behavior changes with default methods `context.become` and `context.unbecome`.
- Writes messages of type `Message` to the journal (see processor [usage](#)). [Sender references](#) are not stored in the journal i.e. the sender reference of a replayed message is always `system.deadLetters`.
- Supports [snapshots](#).
- Identifiers are of type `Int` and must be application-defined.
- Does not support batch-writes of messages to the journal.
- Does not support stashing of messages.

Akka Persistence: `PersistentActor`

- Trait that adds journaling to actors (see *event-sourcing-scala*) and used by applications for *event sourcing* or *command sourcing*. Corresponds to `Eventsourced` processors in `Eventsourced` but is not a stackable trait.
- Automatically recovers on start and re-start, by default. *recovery* can be customized or turned off by overriding actor life cycle hooks `preStart` and `preRestart`. `Processor` takes care that new messages never interfere with replayed messages. New messages are internally buffered until recovery completes.
- No special-purpose behavior change methods. Default behavior change methods `context.become` and `context.unbecome` can be used and are journaling-preserving.
- Sender references are written to the journal. Sender references of type `PromiseActorRef` are not journaled, they are `system.deadLetters` on replay.
- Supports *snapshots*.
- *persistence-identifiers* are of type `String`, have a default value and can be overridden by applications.
- Supports *batch-writes*.
- Supports stashing of messages.

Channels**Eventsourced:** `DefaultChannel`

- Prevents redundant delivery of messages to a destination (see *usage example* and *default channel*).
- Is associated with a single destination actor reference. A new incarnation of the destination is not automatically resolved by the channel. In this case a new channel must be created.
- Must be explicitly activated using methods `deliver` or `recover` defined on `EventsourcingExtension`.
- Must be activated **after** all processors have been activated. Depending on the *recovery* method, this is either done automatically or must be followed by the application (see *non-blocking recovery*). This is necessary for a network of processors and channels to recover consistently.
- Does not redeliver messages on missing or negative delivery confirmation.
- Cannot be used standalone.

Eventsourced: `ReliableChannel`

- Provides `DefaultChannel` functionality plus persistence and recovery from sender JVM crashes (see *ReliableChannel*). Also provides message redelivery in case of missing or negative delivery confirmations.
- Delivers next message to a destination only if previous message has been successfully delivered (flow control is done by destination).
- Stops itself when the maximum number of redelivery attempts has been reached.
- Cannot reply on persistence.
- Can be used standalone.

Akka Persistence: `AtLeastOnceDelivery`

- `AtLeastOnceDelivery` trait is mixed in to a `PersistentActor`
- Does not prevent redundant delivery of messages to a destination entirely, but won't re-send messages whose delivery is confirmed during recovery.
- Is not associated with a single destination. A destination can be specified with each `deliver` request and is referred to by an actor path. A destination path is resolved to the current destination incarnation during delivery (via `actorSelection`).

- Redelivers messages on missing delivery confirmation. In contrast to `Eventsourced`, Akka Persistence doesn't distinguish between missing and negative confirmations. It only has a notion of missing confirmations using timeouts (which are closely related to negative confirmations as both trigger message redelivery).

Views

Eventsourced:

- No direct support for views. Only way to maintain a view is to use a channel and a processor as destination.

Akka Persistence: `View`

- Receives the message stream written by a `PersistentActor` by reading it directly from the journal (see *persistent-views*). Alternative to using channels. Useful in situations where actors shall receive a persistent message stream in correct order without duplicates.
- Supports *snapshots*.

Serializers

Eventsourced:

- Uses a protobuf serializer for serializing `Message` objects.
- Delegates to a configurable Akka serializer for serializing `Message` payloads.
- Delegates to a configurable, proprietary (stream) serializer for serializing snapshots.
- See [Serialization](#).

Akka Persistence:

- Uses a protobuf serializer for serializing `Persistent` objects.
- Delegates to a configurable Akka serializer for serializing `Persistent` payloads.
- Delegates to a configurable Akka serializer for serializing snapshots.
- See *custom-serialization*.

Sequence numbers

Eventsourced:

- Generated on a per-journal basis.

Akka Persistence:

- Generated on a per persistent actor basis.

Storage plugins

Eventsourced:

- Plugin API: `SynchronousWriteReplaySupport` and `AsynchronousWriteReplaySupport`
- No separation between journal and snapshot storage plugins.
- All plugins pre-packaged with project (see [journals](#) and [snapshot configuration](#))

Akka Persistence:

- Plugin API: `SyncWriteJournal`, `AsyncWriteJournal`, `AsyncRecovery`, `SnapshotStore` (see *storage-plugins*).
- Clear separation between journal and snapshot storage plugins.

- Limited number of *pre-packaged-plugins* (LevelDB journal and local snapshot store).
- Impressive list of [community plugins](#).

11.1.7 Migration Guide 2.3.x to 2.4.x

The 2.4 release contains some structural changes that require some simple, mechanical source-level changes in client code.

When migrating from earlier versions you should first follow the instructions for migrating *1.3.x to 2.0.x* and then *2.0.x to 2.1.x* and then *2.1.x to 2.2.x* and then *2.2.x to 2.3.x*.

Binary Compatibility

Akka 2.4.x is backwards binary compatible with previous 2.3.x versions apart from the following exceptions. This means that the new JARs are a drop-in replacement for the old one (but not the other way around) as long as your build does not enable the inliner (Scala-only restriction).

The following parts are not binary compatible with 2.3.x:

- akka-testkit and akka-remote-testkit
- experimental modules, such as akka-persistence and akka-contrib
- features, classes, methods that were deprecated in 2.3.x and removed in 2.4.x

The dependency to **Netty** has been updated from version 3.8.0.Final to 3.10.3.Final. The changes in those versions might not be fully binary compatible, but we believe that it will not be a problem in practice. No changes were needed to the Akka source code for this update. Users of libraries that depend on 3.8.0.Final that break with 3.10.3.Final should be able to manually downgrade the dependency to 3.8.0.Final and Akka will still work with that version.

Advanced Notice: TypedActors will go away

While technically not yet deprecated, the current `akka.actor.TypedActor` support will be superseded by the *typed-scala* project that is currently being developed in open preview mode. If you are using TypedActors in your projects you are advised to look into this, as it is superior to the Active Object pattern expressed in TypedActors. The generic ActorRefs in Akka Typed allow the same type-safety that is afforded by TypedActors while retaining all the other benefits of an explicit actor model (including the ability to change behaviors etc.).

It is likely that TypedActors will be officially deprecated in the next major update of Akka and subsequently removed.

Removed Deprecated Features

The following, previously deprecated, features have been removed:

- akka-dataflow
- akka-transactor
- durable mailboxes (akka-mailboxes-common, akka-file-mailbox)
- `Cluster.publishCurrentClusterState`
- akka.cluster.auto-down, replaced by akka.cluster.auto-down-unreachable-after in Akka 2.3
- Old routers and configuration.

Note that in router configuration you must now specify if it is a `pool` or a `group` in the way that was introduced in Akka 2.3.

- Timeout constructor without unit

- `JavaLoggingEventHandler`, replaced by `JavaLogger`
- `UntypedActorFactory`
- Java API `TestKit.dilated`, moved to `JavaTestKit.dilated`

Protobuf Dependency

The transitive dependency to Protobuf has been removed to make it possible to use any version of Protobuf for the application messages. If you use Protobuf in your application you need to add the following dependency with desired version number:

```
"com.google.protobuf" % "protobuf-java" % "2.5.0"
```

Internally Akka is using an embedded version of protobuf that corresponds to `com.google.protobuf/protobuf-java` version 2.5.0. The package name of the embedded classes has been changed to `akka.protobuf`.

Added parameter validation to `RootActorPath`

Previously `akka.actor.RootActorPath` allowed passing in arbitrary strings into its name parameter, which is meant to be the *name* of the root Actor. Subsequently, if constructed with an invalid name such as a full path for example `(/user/Full/Path)` some features using this path may transparently fail - such as using `actorSelection` on such invalid path.

In Akka 2.4.x the `RootActorPath` validates the input and may throw an `IllegalArgumentException` if the passed in name string is illegal (contains `/` elsewhere than in the beginning of the string or contains `#`).

`TestKit.remaining` throws `AssertionError`

In earlier versions of Akka `TestKit.remaining` returned the default timeout configurable under “`akka.test.single-expect-default`”. This was a bit confusing and thus it has been changed to throw an `AssertionError` if called outside of within. The old behavior however can still be achieved by calling `TestKit.remainingOrDefault` instead.

`EventStream` and `ManagedActorClassification` `EventBus` now require an `ActorSystem`

Both the `EventStream` (*Scala*, *Java*) and the `ManagedActorClassification`, `ManagedActorEventBus` (*Scala*, *Java*) now require an `ActorSystem` to properly operate. The reason for that is moving away from stateful internal lifecycle checks to a fully reactive model for unsubscribing actors that have `Terminated`. Therefore the `ActorClassification` and `ActorEventBus` was deprecated and replaced by `ManagedActorClassification` and `ManagedActorEventBus`

If you have implemented a custom event bus, you will need to pass in the actor system through the constructor now:

```
import akka.event.ActorEventBus
import akka.event.ManagedActorClassification
import akka.event.ActorClassifier

final case class Notification(ref: ActorRef, id: Int)

class ActorBusImpl(val system: ActorSystem) extends ActorEventBus with ActorClassifier with ManagedActorClassification
  type Event = Notification

  // is used for extracting the classifier from the incoming events
  override protected def classify(event: Event): ActorRef = event.ref

  // determines the initial size of the index data structure
  // used internally (i.e. the expected number of different classifiers)
```

```

    override protected def mapSize: Int = 128
  }

```

If you have been creating `EventStreams` manually, you now have to provide an actor system and *start the unsubscrber*:

```

val bus = new EventStream(system, true)
bus.startUnsubscriber()

```

Please note that this change affects you only if you have implemented your own buses, Akka's own `context.eventStream` is still there and does not require any attention from you concerning this change.

FSM notifies on same state transitions

When changing states in an Finite-State-Machine Actor (FSM), state transition events are emitted and can be handled by the user either by registering `onTransition` handlers or by subscribing to these events by sending it an `SubscribeTransitionCallBack` message.

Previously in 2.3.x when an FSM was in state A and performed an `goto(A)` transition, no state transition notification would be sent. This is because it would effectively stay in the same state, and was deemed to be semantically equivalent to calling `stay()`.

In 2.4.x when an FSM performs a any `goto(X)` transition, it will always trigger state transition events. Which turns out to be useful in many systems where same-state transitions actually should have an effect.

In case you do *not* want to trigger a state transition event when effectively performing an X->X transition, use `stay()` instead.

Circuit Breaker Timeout Change

In 2.3.x calls protected by the `CircuitBreaker` were allowed to run indefinitely and the check to see if the timeout had been exceeded was done after the call had returned.

In 2.4.x the `failureCount` of the Breaker will be increased as soon as the timeout is reached and a `Failure[TimeoutException]` will be returned immediately for asynchronous calls. Synchronous calls will now throw a `TimeoutException` after the call is finished.

Slf4j logging filter

If you use `Slf4jLogger` you should add the following configuration:

```
akka.logging-filter = "akka.event.slf4j.Slf4jLoggingFilter"
```

It will filter the log events using the backend configuration (e.g. `logback.xml`) before they are published to the event bus.

Inbox.receive Java API

`Inbox.receive` now throws a checked `java.util.concurrent.TimeoutException` exception if the receive timeout is reached.

Pool routers nrOfInstances method now takes ActorSystem

In order to make cluster routers smarter about when they can start local routees, `nrOfInstances` defined on `Pool` now takes `ActorSystem` as an argument. In case you have implemented a custom `Pool` you will have to update the method's signature, however the implementation can remain the same if you don't need to rely on an `ActorSystem` in your logic.

Group routers paths method now takes ActorSystem

In order to make cluster routers smarter about when they can start local routees, `paths` defined on `Group` now takes `ActorSystem` as an argument. In case you have implemented a custom `Group` you will have to update the method's signature, however the implementation can remain the same if you don't need to rely on an `ActorSystem` in your logic.

Cluster aware router max-total-nr-of-instances

In 2.3.x the deployment configuration property `nr-of-instances` was used for cluster aware routers to specify total number of routees in the cluster. This was confusing, especially since the default value is 1.

In 2.4.x there is a new deployment property `cluster.max-total-nr-of-instances` that defines total number of routees in the cluster. By default `max-total-nr-of-instances` is set to a high value (10000) that will result in new routees added to the router when nodes join the cluster. Set it to a lower value if you want to limit total number of routees.

For backwards compatibility reasons `nr-of-instances` is still used if defined by user, i.e. if defined it takes precedence over `max-total-nr-of-instances`.

Logger names use full class name

Previously, few places in akka used "simple" logger names, such as `Cluster` or `Remoting`. Now they use full class names, such as `akka.cluster.Cluster` or `akka.remote.Remoting`, in order to allow package level log level definitions and ease source code lookup. In case you used specific "simple" logger name based rules in your `logback.xml` configurations, please change them to reflect appropriate package name, such as `<logger name='akka.cluster' level='warn' />` or `<logger name='akka.remote' level='error' />`

Default interval for TestKit.awaitAssert changed to 100 ms

Default check interval changed from 800 ms to 100 ms. You can define the interval explicitly if you need a longer interval.

Secure Cookies

Secure cookies feature was deprecated.

AES128CounterInetRNG and AES256CounterInetRNG are Deprecated

Use `AES128CounterSecureRNG` or `AES256CounterSecureRNG` as `akka.remote.netty.ssl.security.random-number-generator`.

Microkernel is Deprecated

Akka Microkernel is deprecated and will be removed. It is replaced by using an ordinary user defined main class and packaging with `sbt-native-packager` or `Typesafe ConductR`. Please see *Use-case and Deployment Scenarios* for more information.

New Cluster Metrics Extension

Previously, cluster metrics functionality was located in the `akka-cluster` jar. Now it is split out and moved into a separate akka module: `akka-cluster-metrics` jar. The module comes with few enhancements, such as use of Kamon sigar-loader for native library provisioning as well as use of statistical averaging of metrics data. Note that both old and new metrics configuration entries in the `reference.conf` are still in the same name space `akka.cluster.metrics` but are not compatible. Make sure to disable legacy metrics in akka-cluster: `akka.cluster.metrics.enabled=off`, since it is still enabled in akka-cluster by default (for compatibility with past releases). Router configuration entries have also changed for the module, they use prefix `cluster-metrics-:` `cluster-metrics-adaptive-pool` and `cluster-metrics-adaptive-group` Metrics extension classes and objects are located in the new package `akka.cluster.metrics`. Please see *Scala, Java* for more information.

Cluster tools moved to separate module

The Cluster Singleton, Distributed Pub-Sub, and Cluster Client previously located in the `akka-contrib` jar is now moved to a separate module named `akka-cluster-tools`. You need to replace this dependency if you use any of these tools.

The classes changed package name from `akka.contrib.pattern` to `akka.cluster.singleton`, `akka.cluster.pubsub` and `akka.cluster.client`.

The configuration properties changed name to `akka.cluster.pub-sub` and `akka.cluster.client`.

Cluster sharding moved to separate module

The Cluster Sharding previously located in the `akka-contrib` jar is now moved to a separate module named `akka-cluster-sharding`. You need to replace this dependency if you use Cluster Sharding.

The classes changed package name from `akka.contrib.pattern` to `akka.cluster.sharding`.

The configuration properties changed name to `akka.cluster.sharding`.

ClusterSharding construction

Several parameters of the `start` method of the `ClusterSharding` extension are now defined in a settings object `ClusterShardingSettings`. It can be created from system configuration properties and also amended with API. These settings can be defined differently per entry type if needed.

Starting the `ShardRegion` in proxy mode is now done with the `startProxy` method of the `ClusterSharding` extension instead of the optional `entryProps` parameter.

Entry was renamed to `Entity`, for example in the `MessagesExtractor` in the Java API and the `EntityId` type in the Scala API.

`idExtractor` function was renamed to `extractEntityId`. `shardResolver` function was renamed to `extractShardId`.

Cluster Sharding Entry Path Change

Previously in 2.3.x entries were direct children of the local `ShardRegion`. In examples the `persistenceId` of entries included `self.path.parent.name` to include the cluster type name.

In 2.4.x entries are now children of a `Shard`, which in turn is a child of the local `ShardRegion`. To include the shard type in the `persistenceId` it is now accessed by `self.path.parent.parent.name` from each entry.

Asynchronous ShardAllocationStrategy

The methods of the `ShardAllocationStrategy` and `AbstractShardAllocationStrategy` in `Cluster Sharding` have changed return type to a `Future` to support asynchronous decision. For example you can ask an actor external actor of how to allocate shards or rebalance shards.

For the synchronous case you can return the result via `scala.concurrent.Future.successful` in Scala or `akka.dispatch.Futures.successful` in Java.

Cluster Sharding internal data

The Cluster Sharding coordinator stores the locations of the shards using Akka Persistence. This data can safely be removed when restarting the whole Akka Cluster.

The serialization format of the internal persistent events stored by the Cluster Sharding coordinator has been changed and it cannot load old data from 2.3.x or some 2.4 milestone.

The `persistenceId` of the Cluster Sharding coordinator has been changed since 2.3.x so it should not load such old data, but it can be a problem if you have used a 2.4 milestone release. In that case you should remove the persistent data that the Cluster Sharding coordinator stored. Note that this is not application data.

You can use the `RemoveInternalClusterShardingData` utility program to remove this data.

The new `persistenceId` is `s"/sharding/${typeName}Coordinator"`. The old `persistenceId` is `s"/user/sharding/${typeName}Coordinator/singleton/coordinator"`.

ClusterSingletonManager and ClusterSingletonProxy construction

Parameters to the `Props` factory methods have been moved to settings object `ClusterSingletonManagerSettings` and `ClusterSingletonProxySettings`. These can be created from system configuration properties and also amended with API as needed.

The buffer size of the `ClusterSingletonProxy` can be defined in the `ClusterSingletonProxySettings` instead of defining `stash-capacity` of the mailbox. Buffering can be disabled by using a buffer size of 0.

The `singletonPath` parameter of `ClusterSingletonProxy.props` has changed. It is now named `singletonManagerPath` and is the logical path of the singleton manager, e.g. `/user/singletonManager`, which ends with the name you defined in `actorOf` when creating the `ClusterSingletonManager`. In 2.3.x it was the path to singleton instance, which was error-prone because one had to provide both the name of the singleton manager and the singleton actor.

DistributedPubSub construction

Normally, the `DistributedPubSubMediator` actor is started by the `DistributedPubSubExtension`. This extension has been renamed to `DistributedPubSub`. It is also possible to start it as an ordinary actor if you need multiple instances of it with different settings. The parameters of the `Props` factory methods in the `DistributedPubSubMediator` companion has been moved to settings object `DistributedPubSubSettings`. This can be created from system configuration properties and also amended with API as needed.

ClusterClient construction

The parameters of the `Props` factory methods in the `ClusterClient` companion has been moved to settings object `ClusterClientSettings`. This can be created from system configuration properties and also amended with API as needed.

The buffer size of the `ClusterClient` can be defined in the `ClusterClientSettings` instead of defining `stash-capacity` of the mailbox. Buffering can be disabled by using a buffer size of 0.

Normally, the `ClusterReceptionist` actor is started by the `ClusterReceptionistExtension`. This extension has been renamed to `ClusterClientReceptionist`. It is also possible to start it as an ordinary actor if you need multiple instances of it with different settings. The parameters of the Props factory methods in the `ClusterReceptionist` companion has been moved to settings object `ClusterReceptionistSettings`. This can be created from system configuration properties and also amended with API as needed.

The `ClusterReceptionist` actor that is started by the `ClusterReceptionistExtension` is now started as a system actor instead of a user actor, i.e. the default path for the `ClusterClient` initial contacts has changed to `"akka.tcp://system@hostname:port/system/receptionist"`.

ClusterClient sender

In 2.3 the sender () of the response messages, as seen by the client, was the actor in cluster.

In 2.4 the sender () of the response messages, as seen by the client, is `deadLetters` since the client should normally send subsequent messages via the `ClusterClient`. It is possible to pass the the original sender inside the reply messages if the client is supposed to communicate directly to the actor in the cluster.

Akka Persistence

Experimental removed

The artifact name has changed from `akka-persistence-experimental` to `akka-persistence`.

New sbt dependency:

```
"com.typesafe.akka" %% "akka-persistence" % "2.4.0"
```

New Maven dependency:

```
<dependency>
  <groupId>com.typesafe.akka</groupId>
  <artifactId>akka-persistence_2.11</artifactId>
  <version>2.4.0</version>
</dependency>
```

The artefact name of the Persistent TCK has changed from `akka-persistence-tck-experimental` (`akka-persistence-experimental-tck`) to `akka-persistence-tck`.

Mandatory persistenceId

It is now mandatory to define the `persistenceId` in subclasses of `PersistentActor`, `UntypedPersistentActor` and `AbstractPersistentId`.

The rationale behind this change being stricter de-coupling of your Actor hierarchy and the logical “which persistent entity this actor represents”.

In case you want to preserve the old behavior of providing the actor’s path as the default `persistenceId`, you can easily implement it yourself either as a helper trait or simply by overriding `persistenceId` as follows:

```
override def persistenceId = self.path.toStringWithoutAddress
```

Persist sequence of events

The `persist` method that takes a `Seq` (Scala) or `Iterable` (Java) of events parameter was deprecated and renamed to `persistAll` to avoid mistakes of persisting other collection types as one single event by calling the overloaded `persist(event)` method.

non-permanent deletion

The `permanent` flag in `deleteMessages` was removed. non-permanent deletes are not supported any more. Events that were deleted with `permanent=false` with older version will still not be replayed in this version.

Recover message is gone, replaced by Recovery config

Previously the way to cause recover in `PersistentActors` was sending them a `Recover()` message. Most of the time it was the actor itself sending such message to `self` in its `preStart` method, however it was possible to send this message from an external source to any `PersistentActor` or `PersistentView` to make it start recovering.

This style of starting recovery does not fit well with usual Actor best practices: an Actor should be independent and know about its internal state, and also about its recovery or lack thereof. In order to guide users towards more independent Actors, the `Recovery()` object is now not used as a message, but as configuration option used by the Actor when it starts. In order to migrate previous code which customised its recovery mode use this example as reference:

```
// previously
class OldCookieMonster extends PersistentActor {
  def preStart() = self ! Recover(toSequenceNr = 42L)
  // ...
}
// now:
class NewCookieMonster extends PersistentActor {
  override def recovery = Recovery(toSequenceNr = 42L)
  // ...
}
```

Sender reference of replayed events is deadLetters

While undocumented, previously the `sender()` of the replayed messages would be the same sender that originally had sent the message. Since `sender` is an `ActorRef` and those events are often replayed in different incarnations of actor systems and during the entire lifetime of the app, relying on the existence of this reference is most likely not going to succeed. In order to avoid bugs in the style of “it worked last week”, the `sender()` reference is now not stored, in order to avoid potential bugs which this could have provoked.

The previous behaviour was never documented explicitly (nor was it a design goal), so it is unlikely that applications have explicitly relied on this behaviour, however if you find yourself with an application that did exploit this you should rewrite it to explicitly store the `ActorPath` of where such replies during replay may have to be sent to, instead of relying on the sender reference during replay.

max-message-batch-size config

Configuration property `akka.persistence.journal.max-message-batch-size` has been moved into the plugin configuration section, to allow different values for different journal plugins. See `reference.conf`.

PersistentView is deprecated

`PersistentView` is deprecated. Use `persistence-query-scala` instead. The corresponding query type is `EventsByPersistenceId`. There are several alternatives for connecting the `Source` to an actor corresponding to a previous `PersistentView` actor:

- `Sink.actorRef` is simple, but has the disadvantage that there is no back-pressure signal from the destination actor, i.e. if the actor is not consuming the messages fast enough the mailbox of the actor will grow

- `mapAsync` combined with *Ask: Send-And-Receive-Future* is almost as simple with the advantage of back-pressure being propagated all the way
- `ActorSubscriber` in case you need more fine grained control

The consuming actor may be a plain `Actor` or a `PersistentActor` if it needs to store its own state (e.g. `fromSequenceNr` offset).

Persistence Plugin APIs

SyncWriteJournal removed

`SyncWriteJournal` removed in favor of using `AsyncWriteJournal`.

If the storage backend API only supports synchronous, blocking writes, the methods can still be implemented in terms of the asynchronous API. Example of how to do that is included in the See *Journal plugin API for Scala* or *Journal plugin API for Java*.

SnapshotStore: Snapshots can now be deleted asynchronously (and report failures)

Previously the `SnapshotStore` plugin SPI did not allow for asynchronous deletion of snapshots, and failures of deleting a snapshot may have been even silently ignored.

Now `SnapshotStore` must return a `Future` representing the deletion of the snapshot. If this future completes successfully the `PersistentActor` which initiated the snapshotting will be notified via an `DeleteSnapshotSuccess` message. If the deletion fails for some reason a `DeleteSnapshotFailure` will be sent to the actor instead.

For criteria based deletion of snapshots (def `deleteSnapshots(criteria: SnapshotSelectionCriteria)`) equivalent `DeleteSnapshotsSuccess` and `DeleteSnapshotsFailure` messages are sent, which contain the specified criteria, instead of `SnapshotMetadata` as is the case with the single snapshot deletion messages.

SnapshotStore: Removed 'saved' callback

Snapshot Stores previously were required to implement a `def saved(meta: SnapshotMetadata): Unit` method which would be called upon successful completion of a `saveAsync` (`doSaveAsync` in Java API) snapshot write.

Currently all journals and snapshot stores perform asynchronous writes and deletes, thus all could potentially benefit from such callback methods. The only gain these callback give over composing an `onComplete` over `Future` returned by the journal or snapshot store is that it is executed in the Actors context, thus it can safely (without additional synchronization modify its internal state - for example a “pending writes” counter).

However, this feature was not used by many plugins, and expanding the API to accommodate all callbacks would have grown the API a lot. Instead, Akka Persistence 2.4.x introduces an additional (optionally overrideable) `receivePluginInternal: Actor.Receive` method in the plugin API, which can be used for handling those as well as any custom messages that are sent to the plugin Actor (imagine use cases like “wake up and continue reading” or custom protocols which your specialised journal can implement).

Implementations using the previous feature should adjust their code as follows:

```
// previously
class MySnapshots extends SnapshotStore {
  // old API:
  // def saved(meta: SnapshotMetadata): Unit = doThings()

  // new API:
  def saveAsync(metadata: SnapshotMetadata, snapshot: Any): Future[Unit] = {
    // completion or failure of the returned future triggers internal messages in receivePluginInternal
  }
}
```

```

val f: Future[Unit] = ???

// custom messages can be piped to self in order to be received in receivePluginInternal
f.map(MyCustomMessage(_)) pipeTo self

f
}

def receivePluginInternal = {
  case SaveSnapshotSuccess(metadata) => doThings()
  case MyCustomMessage(data)         => doOtherThings()
}

// ...
}

```

SnapshotStore: Java 8 Optional used in Java plugin APIs

In places where previously `akka.japi.Option` was used in Java APIs, including the return type of `doLoadAsync`, the Java 8 provided `Optional` type is used now.

Please remember that when creating a `java.util.Optional` instance from a (possibly) null value you will want to use the non-throwing `Optional.fromNullable` method, which converts a null into a `None` value - which is slightly different than its Scala counterpart (where `Option.apply(null)` returns `None`).

Atomic writes

`asyncWriteMessages` takes a `immutable.Seq[AtomicWrite]` parameter instead of `immutable.Seq[PersistentRepr]`.

Each `AtomicWrite` message contains the single `PersistentRepr` that corresponds to the event that was passed to the `persist` method of the `PersistentActor`, or it contains several `PersistentRepr` that corresponds to the events that were passed to the `persistAll` method of the `PersistentActor`. All `PersistentRepr` of the `AtomicWrite` must be written to the data store atomically, i.e. all or none must be stored.

If the journal (data store) cannot support atomic writes of multiple events it should reject such writes with a `Try Failure` with an `UnsupportedOperationException` describing the issue. This limitation should also be documented by the journal plugin.

Rejecting writes

`asyncWriteMessages` returns a `Future[immutable.Seq[Try[Unit]]]`.

The journal can signal that it rejects individual messages (`AtomicWrite`) by the returned `immutable.Seq[Try[Unit]]`. The returned `Seq` must have as many elements as the input messages `Seq`. Each `Try` element signals if the corresponding `AtomicWrite` is rejected or not, with an exception describing the problem. Rejecting a message means it was not stored, i.e. it must not be included in a later replay. Rejecting a message is typically done before attempting to store it, e.g. because of serialization error.

Read the *API documentation* of this method for more information about the semantics of rejections and failures.

asyncReplayMessages Java API

The signature of `asyncReplayMessages` in the Java API changed from `akka.japi.Procedure` to `java.util.function.Consumer`.

asyncDeleteMessagesTo

The permanent deletion flag was removed. Support for non-permanent deletions was removed. Events that were deleted with `permanent=false` with older version will still not be replayed in this version.

References to “replay” in names

Previously a number of classes and methods used the word “replay” interchangeably with the word “recover”. This led to slight inconsistencies in APIs, where a method would be called `recovery`, yet the signal for a completed recovery was named `ReplayMessagesSuccess`.

This is now fixed, and all methods use the same “recovery” wording consistently across the entire API. The old `ReplayMessagesSuccess` is now called `RecoverySuccess`, and an additional method called `onRecoveryFailure` has been introduced.

AtLeastOnceDelivery deliver signature

The signature of `deliver` changed slightly in order to allow both `ActorSelection` and `ActorPath` to be used with it.

Previously:

```
def deliver(destination: ActorPath, deliveryIdToMessage: Long ⇒ Any): Unit
```

Now:

```
def deliver(destination: ActorSelection)(deliveryIdToMessage: Long ⇒ Any): Unit
def deliver(destination: ActorPath)(deliveryIdToMessage: Long ⇒ Any): Unit
```

The Java API remains unchanged and has simply gained the 2nd overload which allows `ActorSelection` to be passed in directly (without converting to `ActorPath`).

11.2 Issue Tracking

Akka is using GitHub Issues as its issue tracking system.

11.2.1 Browsing

Tickets

Before filing a ticket, please check the existing [Akka tickets](#) for earlier reports of the same problem. You are very welcome to comment on existing tickets, especially if you have reproducible test cases that you can share.

Roadmaps

Please refer to the [Akka roadmap](#) in order to find out the general theme of work to be done for upcoming versions of Akka.

11.2.2 Creating tickets

Please include the versions of Scala and Akka and relevant configuration files.

You can create a [new ticket](#) if you have registered a GitHub user account.

Thanks a lot for reporting bugs and suggesting features!

11.3 Licenses

11.3.1 Akka License

This software is licensed under the Apache 2 license, quoted below.

```
Copyright 2009–2015 Typesafe Inc. <http://www.typesafe.com>
```

```
Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at
```

```
http://www.apache.org/licenses/LICENSE-2.0
```

```
Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.
```

11.3.2 Akka Committer License Agreement

All committers have signed this [CLA](#). It can be [signed online](#).

11.3.3 Licenses for Dependency Libraries

Each dependency and its license can be seen in the project build file (the comment on the side of each dependency): [AkkaBuild.scala](#)

11.4 Sponsors

11.4.1 Typesafe

Typesafe is the company behind the Akka Project, Scala Programming Language, Play Web Framework, Scala IDE, Simple Build Tool and many other open source projects. It also provides the Typesafe Stack, a full-featured development stack consisting of AKka, Play and Scala. Learn more at typesafe.com.

11.4.2 YourKit

YourKit is kindly supporting open source projects with its full-featured Java Profiler.

YourKit, LLC is the creator of innovative and intelligent tools for profiling Java and .NET applications. Take a look at YourKit's leading software products: [YourKit Java Profiler](#) and [YourKit .NET Profiler](#)

11.5 Project

11.5.1 Commercial Support

Commercial support is provided by [Typesafe](#). Akka is part of the [Typesafe Reactive Platform](#).

11.5.2 Mailing List

Akka User Google Group

Akka Developer Google Group

11.5.3 Downloads

<http://akka.io/downloads>

11.5.4 Source Code

Akka uses Git and is hosted at [Github](#).

- Akka: clone the Akka repository from <http://github.com/akka/akka>

11.5.5 Releases Repository

All Akka releases are published via Sonatype to Maven Central, see search.maven.org

11.5.6 Snapshots Repository

Nightly builds are available in <http://repo.akka.io/snapshots/> as both SNAPSHOT and timestamped versions.

For timestamped versions, pick a timestamp from http://repo.akka.io/snapshots/com/typesafe/akka/akka-actor_2.11/. All Akka modules that belong to the same build have the same timestamp.

sbt definition of snapshot repository

Make sure that you add the repository to the sbt resolvers:

```
resolvers += "Typesafe Snapshots" at "http://repo.akka.io/snapshots/"
```

Define the library dependencies with the timestamp as version. For example:

```
libraryDependencies += "com.typesafe.akka" % "akka-remote_2.11" %  
  "2.1-20121016-001042"
```

maven definition of snapshot repository

Make sure that you add the repository to the maven repositories in pom.xml:

```
<repositories>  
  <repository>  
    <id>akka-snapshots</id>  
    <name>Akka Snapshots</name>  
    <url>http://repo.akka.io/snapshots/</url>  
    <layout>default</layout>  
  </repository>  
</repositories>
```

Define the library dependencies with the timestamp as version. For example:

```
<dependencies>
  <dependency>
    <groupId>com.typesafe.akka</groupId>
    <artifactId>akka-remote_2.11</artifactId>
    <version>2.1-20121016-001042</version>
  </dependency>
</dependencies>
```

ADDITIONAL INFORMATION

12.1 Frequently Asked Questions

12.1.1 Akka Project

Where does the name Akka come from?

It is the name of a beautiful Swedish [mountain](#) up in the northern part of Sweden called Laponia. The mountain is also sometimes called ‘The Queen of Laponia’.

Akka is also the name of a goddess in the Sámi (the native Swedish population) mythology. She is the goddess that stands for all the beauty and good in the world. The mountain can be seen as the symbol of this goddess.

Also, the name AKKA is the a palindrome of letters A and K as in Actor Kernel.

Akka is also:

- the name of the goose that Nils traveled across Sweden on in [The Wonderful Adventures of Nils](#) by the Swedish writer Selma Lagerlöf.
- the Finnish word for ‘nasty elderly woman’ and the word for ‘elder sister’ in the Indian languages Tamil, Telugu, Kannada and Marathi.
- a [font](#)
- a town in Morocco
- a near-earth asteroid

12.1.2 Actors in General

`sender()/getSender()` disappears when I use Future in my Actor, why?

When using future callbacks, inside actors you need to carefully avoid closing over the containing actor’s reference, i.e. do not call methods or access mutable state on the enclosing actor from within the callback. This breaks the actor encapsulation and may introduce synchronization bugs and race conditions because the callback will be scheduled concurrently to the enclosing actor. Unfortunately there is not yet a way to detect these illegal accesses at compile time.

Read more about it in the docs for [Actors and shared mutable state](#).

Why OutOfMemoryError?

It can be many reasons for OutOfMemoryError. For example, in a pure push based system with message consumers that are potentially slower than corresponding message producers you must add some kind of message flow control. Otherwise messages will be queued in the consumers’ mailboxes and thereby filling up the heap memory.

Some articles for inspiration:

- [Balancing Workload across Nodes with Akka 2.](#)
- [Work Pulling Pattern to prevent mailbox overflow, throttle and distribute work](#)

12.1.3 Actors Scala API

How can I get compile time errors for missing messages in *receive*?

One solution to help you get a compile time warning for not handling a message that you should be handling is to define your actors input and output messages implementing base traits, and then do a match that will be checked for exhaustiveness.

Here is an example where the compiler will warn you that the match in *receive* isn't exhaustive:

```
object MyActor {
  // these are the messages we accept
  sealed abstract trait Message
  final case class FooMessage(foo: String) extends Message
  final case class BarMessage(bar: Int) extends Message

  // these are the replies we send
  sealed abstract trait Reply
  final case class BazMessage(baz: String) extends Reply
}

class MyActor extends Actor {
  import MyActor._
  def receive = {
    case message: Message => message match {
      case BarMessage(bar) => sender() ! BazMessage("Got " + bar)
      // warning here:
      // "match may not be exhaustive. It would fail on the following input: FooMessage(_)"
    }
  }
}
```

12.1.4 Remoting

I want to send to a remote system but it does not do anything

Make sure that you have remoting enabled on both ends: client and server. Both do need hostname and port configured, and you will need to know the port of the server; the client can use an automatic port in most cases (i.e. configure port zero). If both systems are running on the same network host, their ports must be different

If you still do not see anything, look at what the logging of remote life-cycle events tells you (normally logged at INFO level) or switch on *Auxiliary remote logging options* to see all sent and received messages (logged at DEBUG level).

Which options shall I enable when debugging remoting issues?

Have a look at the *Remote Configuration*, the typical candidates are:

- `akka.remote.log-sent-messages`
- `akka.remote.log-received-messages`
- `akka.remote.log-remote-lifecycle-events` (this also includes deserialization errors)

What is the name of a remote actor?

When you want to send messages to an actor on a remote host, you need to know its *full path*, which is of the form:

```
akka.protocol://system@host:1234/user/my/actor/hierarchy/path
```

Observe all the parts you need here:

- **protocol is the protocol to be used to communicate with the remote system.** Most of the cases this is *tcp*.
- *system* is the remote system's name (must match exactly, case-sensitive!)
- *host* is the remote system's IP address or DNS name, and it must match that system's configuration (i.e. *akka.remote.netty.tcp.hostname*)
- *1234* is the port number on which the remote system is listening for connections and receiving messages
- */user/my/actor/hierarchy/path* is the absolute path of the remote actor in the remote system's supervision hierarchy, including the system's guardian (i.e. */user*, there are others e.g. */system* which hosts loggers, */temp* which keeps temporary actor refs used with *ask*, */remote* which enables remote deployment, etc.); this matches how the actor prints its own *self* reference on the remote host, e.g. in log output.

Why are replies not received from a remote actor?

The most common reason is that the local system's name (i.e. the *system@host:1234* part in the answer above) is not reachable from the remote system's network location, e.g. because *host* was configured to be *0.0.0.0*, *localhost* or a NAT'ed IP address.

If you are running an ActorSystem under a NAT or inside a docker container, make sure to set *akka.remote.netty.tcp.hostname* and *akka.remote.netty.tcp.port* to the address it is reachable at from other ActorSystems. If you need to bind your network interface to a different address - use *akka.remote.netty.tcp.bind-hostname* and *akka.remote.netty.tcp.bind-port* settings. Also make sure your network is configured to translate from the address your ActorSystem is reachable at to the address your ActorSystem network interface is bound to.

How reliable is the message delivery?

The general rule is **at-most-once delivery**, i.e. no guaranteed delivery. Stronger reliability can be built on top, and Akka provides tools to do so.

Read more in [Message Delivery Reliability](#).

12.1.5 Debugging

How do I turn on debug logging?

To turn on debug logging in your actor system add the following to your configuration:

```
akka.loglevel = DEBUG
```

To enable different types of debug logging add the following to your configuration:

- *akka.actor.debug.receive* will log all messages sent to an actor if that actors *receive* method is a *LoggingReceive*
- *akka.actor.debug.autoreceive* will log all *special* messages like *Kill*, *PoisonPill* e.t.c. sent to all actors
- *akka.actor.debug.lifecycle* will log all actor lifecycle events of all actors

Read more about it in the docs for [Logging](#) and [actor.logging-scala](#).

12.2 Books

- [Akka in Action](#), by Raymond Roestenburg and Rob Bakker, Manning Publications Co., ISBN: 9781617291012, est fall 2015
- [Reactive Messaging Patterns with the Actor Model](#), by Vaughn Vernon, Addison-Wesley Professional, ISBN: 0133846830, August 2015
- [Developing an Akka Edge](#), by Thomas Lockney and Raymond Tay, Bleeding Edge Press, ISBN: 9781939902054, April 2014
- [Effective Akka](#), by Jamie Allen, O'Reilly Media, ISBN: 1449360076, August 2013
- [Akka Concurrency](#), by Derek Wyatt, artima developer, ISBN: 0981531660, May 2013
- [Akka Essentials](#), by Munish K. Gupta, PACKT Publishing, ISBN: 1849518289, October 2012

12.3 Other Language Bindings

12.3.1 JRuby

Read more here: <https://github.com/iconara/mikka>.

12.3.2 Groovy/Groovy++

Read more here: <https://gist.github.com/620439>.

12.3.3 Clojure

Read more here: <http://blog.darevay.com/2011/06/clojure-and-akka-a-match-made-in/>.

12.4 Akka in OSGi

12.4.1 Configuring the OSGi Framework

To use Akka in an OSGi environment, the `org.osgi.framework.bootdelegation` property must be set to always delegate the `sun.misc` package to the boot classloader instead of resolving it through the normal OSGi class space.

12.4.2 Activator

To bootstrap Akka inside an OSGi environment, you can use the `akka.osgi.ActorSystemActivator` class to conveniently set up the ActorSystem.

```
import akka.actor.{ Props, ActorSystem }
import org.osgi.framework.BundleContext
import akka.osgi.ActorSystemActivator

class Activator extends ActorSystemActivator {

  def configure(context: BundleContext, system: ActorSystem) {
    // optionally register the ActorSystem in the OSGi Service Registry
    registerService(context, system)
  }
}
```

```
val someActor = system.actorOf(Props[SomeActor], name = "someName")
someActor ! SomeMessage
}
}
```

The `ActorSystemActivator` creates the actor system with a class loader that finds resources (reference.conf files) and classes from the application bundle and all transitive dependencies.

The `ActorSystemActivator` class is included in the akka-osgi artifact:

```
<dependency>
  <groupId>com.typesafe.akka</groupId>
  <artifactId>akka-osgi_2.11</artifactId>
  <version>2.4.0</version>
</dependency>
```

12.4.3 Sample

A complete sample project is provided in akka-sample-osgi-dining-hackers.

12.5 Incomplete List of HTTP Frameworks

12.5.1 Play

The `Play framework` is built using Akka, and is well suited for building both full web applications as well as REST services.

12.5.2 Spray

The `Spray toolkit` is built using Akka, and is a minimalistic HTTP/REST layer.

12.5.3 Akka Mist

If you are using Akka Mist (Akka's old HTTP/REST module) with Akka 1.x and wish to upgrade to 2.x there is now a port of Akka Mist to Akka 2.x. You can find it [here](#).

12.5.4 Other Alternatives

There are a bunch of other alternatives for using Akka with HTTP/REST. You can find some of them among the [Community Projects](#).