



Talend ESB Container

Administration Guide

6.0.1

Publication date: September 10, 2015
Copyright © 2011-2015 Talend Inc. All rights reserved.

Copyleft

This documentation is provided under the terms of the Creative Commons Public License (CCPL). For more information about what you can and cannot do with this documentation in accordance with the CCPL, please read: <http://creativecommons.org/licenses/by-nc-sa/2.0/>

This document may include documentation produced at The Apache Software Foundation which is licensed under The Apache License 2.0.

Notices

Talend and Talend ESB are trademarks of Talend, Inc.

Apache CXF, CXF, Apache Karaf, Karaf, Apache Camel, Camel, Apache Maven, Maven, Apache Archiva, Archiva, Apache Syncope, Syncope, Apache ActiveMQ, ActiveMQ, Apache Log4j, Log4j, Apache Felix, Felix, Apache ServiceMix, ServiceMix, Apache Ant, Ant, Apache Derby, Derby, Apache Tomcat, Tomcat, Apache ZooKeeper, ZooKeeper, Apache Jackrabbit, Jackrabbit, Apache Santuario, Santuario, Apache DS, DS, Apache Avro, Avro, Apache Abdera, Abdera, Apache Chemistry, Chemistry, Apache CouchDB, CouchDB, Apache Kafka, Kafka, Apache Lucene, Lucene, Apache MINA, MINA, Apache Velocity, Velocity, Apache FOP, FOP, Apache HBase, HBase, Apache Hadoop, Hadoop, Apache Shiro, Shiro, Apache Axiom, Axiom, Apache Neethi, Neethi, Apache WSS4J, WSS4J are trademarks of The Apache Foundation. Eclipse Equinox is a trademark of the Eclipse Foundation, Inc. SoapUI is a trademark of SmartBear Software. Hyperic is a trademark of VMware, Inc. Nagios is a trademark of Nagios Enterprises, LLC.

All other brands, product names, company names, trademarks and service marks are the properties of their respective owners.

Table of Contents

1. Introduction	1
1.1. Structure of this manual	1
2. Release directory structure	3
3. Starting/Stopping Talend Runtime and its ESB Infrastructure Components	5
3.1. Starting Talend Runtime	5
3.2. Stopping Talend Runtime	6
3.3. Starting and stopping Talend ESB infrastructure services	6
4. Using the console commands	9
4.1. Config Scope	10
4.2. Dev Scope	10
4.3. Feature Scope	10
4.4. JAAS Scope	11
4.5. Log Scope	12
4.6. Shell Scope	13
4.7. Miscellaneous Scopes	14
5. Deploying Multiple Karaf Containers	15
5.1. Deploying multiple containers using configuration adaption scripts	15
5.2. Troubleshooting	17
6. Remote Console	19
6.1. Configuring remote instances	19
6.2. Connecting and disconnecting remotely	20
6.3. Stopping a remote instance	21
7. Apache CXF and Camel, commands and tuning	23
7.1. Commands supplied by CXF	24
7.2. Commands supplied by Camel	24
7.3. Configuring CXF workqueues	25
8. Security	29
8.1. Managing users and passwords	29
8.2. Managing roles	30
8.3. Enabling password encryption	30
8.4. Managing realms	30
8.5. Deploying security providers	35
9. HTTP Configuration	37
9.1. Server HTTP Configuration	37
9.2. Client HTTP Configuration	41
10. Logging System	49
10.1. Configuration	49
10.2. Console Log Commands	50
10.3. Advanced Configuration	50
11. Deployer	53
11.1. Features deployer	53
11.2. Spring deployer	54
11.3. Wrap deployer	54
11.4. War deployer	56
12. Servlet Context	57
12.1. Configuration	57
12.2. Configuring the context	57
13. Provisioning	59
13.1. Example: Deploying a sample feature	59
13.2. Repositories	60
13.3. Commands	64
13.4. Service configuration	65
14. Web Applications	67
14.1. Installing WAR support	67
14.2. Deploying a WAR to the installed web feature	68
15. Monitoring and Administration using JMX	71

16. Installing the Talend Runtime container as a service	75
16.1. Supported platforms	75
16.2. Installing the wrapper	76
16.3. Installing the service	78
17. Setting up the Talend ESB Active/Passive configuration	81
17.1. HA/failover (active/passive)	81
18. Troubleshooting Talend ESB	87
18.1. Memory Allocation Parameters	87
18.2. On Windows	88
18.3. On Linux/Solaris	88



Chapter 1. Introduction

Talend ESB provides an Apache Karaf-based ESB container preconfigured to support Apache Camel routing and Apache CXF-based services (both REST and SOAP-based). This container administration manual is intended to provide information on Karaf tooling commands and general administration. It is NOT a tutorial on getting started with Karaf - please see Talend ESB Hands-on Guide for this.



As the Talend ESB container is based on Apache Karaf, the two terms are used interchangeably within this document, except when explicitly stressed otherwise.

1.1. Structure of this manual

The following chapters are in this manual:

- *Release directory structure* describes the directory layout of a Talend ESB installation.
- *Starting/Stopping Talend Runtime and its ESB Infrastructure Components* describes the various options that are available to start and stop Talend ESB.
- *Using the console commands* gives a list of the console commands available.
- *Deploying Multiple Karaf Containers* describes how to adapt the configuration to deploy multiple containers at the same machine.
- *Remote Console* describes how to access Talend ESB using a remote console.
- *Apache CXF and Camel, commands and tuning* describes commands provided by CXF and Camel.
- *Security* describes how to configure security with Talend ESB.
- *HTTP Configuration* describes how to configure HTTP in Talend Runtime.
- *Logging System* describes configuring and using the logging system.

- *Deployer* describes deploying bundles or re-defined groups of bundles (called features).
- *Servlet Context* describes the servlet context configuration.
- *Provisioning* describes how the provisioning system uses xml repositories that define sets of features (bundles). Thus when you install a feature and it will install all of its bundles defined under it.
- *Web Applications* describes deploying WAR-based web applications into Talend Runtime.
- *Monitoring and Administration using JMX* lists the MBeans available to monitor and administrate Karaf using any JMX client.
- *Installing the Talend Runtime container as a service* describes how to install Talend Runtime container as a service.
- *Troubleshooting Talend ESB* describes how to fix problems related to JVM memory allocation parameters running Talend ESB.



Chapter 2. Release directory structure

The Talend ESB software may be downloaded from <http://www.talend.com/download.php>.



The term `<TalendRuntimePath>` is used for the directory where Talend Runtime is installed. This is typically the full path of either `Runtime_ESBSE` or `Talend-ESB-V6.0.x`, depending on the version of the software that is being used. Please substitute appropriately.

The Talend Runtime container is in `<TalendRuntimePath>/container`, but if there are several containers, this may be actually `<TalendRuntimePath>/container2`, and so on. In this document, `<RuntimeContainerPath>` is used to indicate the related container path.

The standard directory layout of the Talend Runtime container installation is as follows:

- `/bin`: startup scripts
- `/etc`: configuration files
 - `/keystores`: KeyStore files for JobServer SSL connection
- `/data`: working directory
 - `/cache`: OSGi framework bundle cache
 - `/generated-bundles`: temporary folder used by the deployer
 - `/log`: log files
- `/deploy`: hot deploy directory
- `/instances`: directory containing child instances
- `/lib`: contains the bootstrap libraries
 - `/lib/ext`: directory for JRE extensions
 - `/lib/endorsed`: directory for endorsed libraries

-
- /system: OSGi bundles repository, laid out as a Maven 2 repository



The data folder contains all the working and temporary files for Karaf. If you want to restart from a clean state, you can wipe out this directory, which has the same effect as using the [clean option \[6\]](#).



Chapter 3. Starting/Stopping Talend Runtime and its ESB Infrastructure Components

This chapter describes how to start and stop Talend Runtime, the various options that are available, and how to start and stop the infrastructure components of Talend ESB.

3.1. Starting Talend Runtime



After starting Talend Runtime, you need to wait a few seconds for initialization to complete before entering the commands. Karaf starts the non core bundles in the background. So even if the console is already available, the job commands may not.

3.1.1. Starting Talend Runtime with console

From a console window:

1. Change to the installation directory:

```
cd <RuntimeContainerPath>
```

2. Run the Talend Runtime:

- **bin\trun.bat** (Windows)
- **bin\trun** (Linux)



Do NOT close the console or shell in which Karaf was started, as that will terminate Karaf (unless Karaf was started using the **nohup** command).

3.1.2. Starting Talend Runtime in the background

Karaf can be easily started as a background process using the following command:

- **bin\start.bat** (Windows)
- **bin\start** (Linux)

Karaf can be reset to a clean state by simply deleting the `<RuntimeContainerPath>/data` folder, or by adding **clean** after the above **start** command.

3.2. Stopping Talend Runtime

Within the Karaf console, you can perform a shutdown using the **Ctrl+D** keys or the following **logout** command.

Within the Karaf console, to perform a clean shutdown of Karaf, use the **system:shutdown** command, with an optional **-f** (force) flag to avoid a **yes/no** confirmation.

Furthermore, from a command shell, you can run **bin\stop.bat** (Windows) or **bin/stop** (Linux).

For more advanced functionality, it is also possible to delay the shutdown using time arguments. The time arguments can have different formats:

- absolute time format `hh:mm`, in which `hh` is the hour (1 or 2 digits) and `mm` is the minute of the hour (in two digits).

For example, the following command will shutdown Karaf at 10:35am:

```
system:shutdown 10:35
```

- duration format `+m`, in which `m` is the number of minutes to wait. The argument `now` is an alias for `+0`.

For example, the following command will shutdown Karaf in 10 minutes:

```
system:shutdown +10
```

3.3. Starting and stopping Talend ESB infrastructure services

Talend Runtime provides infrastructure services, including the Service Locator, Service Activity Monitoring, Security Token Service for its community version, and additional XKMS, Service Registry, Authorization, and Event Logging, etc. for its subscription versions, which are already installed in the Talend Runtime container.

To start all the infrastructure services, after starting the Talend Runtime container, enter the following command at the console prompt:

```
tesb:start-all
```



The Event Logging feature (subscription feature only) will not be automatically started by the **tesb:start-all** command. To start it, you will have to execute its specific command(s). For more information on how to start Event Logging, see the *Talend ESB Infrastructure Services Configuration Guide*.

You can shutdown all Talend ESB infrastructure services by entering:

```
tesb:stop-all
```



If there is another Talend Runtime container running on the same machine and you forgot to adapt the configuration of the used ports, a port conflict may arise while starting an infrastructure component. In this case **tesb:start-all** will not try to start the remaining infrastructure components. Stop all infrastructure services using **tesb:stop-all** and restart them after adapting the configuration. There are configuration scripts simplifying this task.

For more information on how to start and stop and how to configure each of the infrastructure services, see the *Talend ESB Infrastructure Services Configuration Guide*.



Chapter 4. Using the console commands

To see a list of the available commands in the Talend Runtime container (or Karaf) console press the **tab** at the prompt.

The **tab** key toggles completion anywhere on the line, so if you want to see the commands in the `feature` group, type the first letters and hit **tab**. Depending on the commands, completion may be available on options and arguments too.

To view help on a particular command, type the command followed by **--help** or use the **help** command followed by the name of the command:

```
karaf@trun> feature:list --help
```

or

```
karaf@trun> help feature:list
```

DESCRIPTION

```
feature:list
```

```
Lists all existing features available from the defined repositories
```

SYNTAX

```
feature:list [options]
```

OPTIONS

```
-o, --ordered
```

```
    Display a list using alphabetical order
```

```
--help
```

```
    Display this help message
```

```
-i, --installed
```

```
    Display a list of all installed features only
```

```
-r, --required
```

```
    Display a list of all required features only
```

```
-s, --show-hidden
    Display hidden features
--no-format
    Disable table rendered output
```

This chapter gives a summary of the commands, grouped by scope; help for each command is also available [online](#) at the Apache Karaf website.

4.1. Config Scope

The commands within the **config:** scope are used to modify configuration data of Karaf instances (data stored in the `/etc` folder of the Karaf instance.) Updates take effect immediately without need of restarting Karaf. Note there is also a ConfigMBean that can be used to do this configuration within JMX.

Table 4.1. config Scope commands

Command	Parameters	Description
config:cancel		Cancels the changes to the configuration being edited.
config:delete	[options] pid	Deletes a configuration.
config:edit	[options] pid	Creates or edits a configuration.
config:list	[query]	Lists existing configurations.
config:meta	[options]	Lists meta type information.
config:property-append	[options] name value	Appends the given value to an existing property or creates the property with the specified name and value.
config:property-delete	[options] property	Deletes a property from the configuration being edited.
config:property-list	[options]	Lists properties from the currently edited configuration.
config:property-set	[options] property value	Sets a property in the currently edited configuration.
config:update		Saves and propagates changes from the configuration being edited

4.2. Dev Scope

The `dev:` scope only has one command: `dev:dump-create`. You can use it to create zip archive with diagnostic info as follows:

```
dev:dump-create [options] [name]
```

4.3. Feature Scope

The commands within the `feature:` scope are used to provide support for Karaf features, which are predefined collections of bundles used to implement specific services.

Table 4.2. feature Scope commands

Command	Parameters	Description
feature:info	[options] name [version]	Shows information about selected feature.
feature:install	[options] feature	Installs a feature with the specified name and version.
feature:list	[options]	Lists all existing features available from the defined repositories.
feature:regions	[options] [regions]	Prints information about region digraph. Where regions are regions to provide detailed info for.
feature:repo-add	[options] name/url [version]	Adds a features repository. Where: <ul style="list-style-type: none"> name/url is the shortcut name of the features repository or the full URL, version is the version of the features repository if using features repository name as first argument. It should be empty if using the URL
feature:repo-list	[options]	Displays a list of all defined repositories.
feature:repo-refresh	[Feature name or uri] [Feature version]	Refresh a features repository. Where: <ul style="list-style-type: none"> Feature name or uri is a shortcut name of the feature repository or the full URI Feature version is the version of the feature if using the feature name. Should be empty if using the uri
feature:repo-remove	[options] repository	Removes the specified repository features service.
feature:requirement-add	[options] requirements	Adds provisioning requirements.
feature:requirement-list	[options]	Lists provisioning requirements.
feature:requirement-remove	[options] requirements	Removes provisioning requirements.
feature:start	[options] feature	Starts features with the specified name and version.
feature:stop	[options] feature	Stops features with the specified name and version.
feature:uninstall	[options] features	Uninstalls a feature with the specified name and version.
feature:version-list	[options] feature	Lists all versions of a feature available from the currently available repositories.

4.4. JAAS Scope

The commands within the `jaas: scope` are used for management of [JAAS](#) users and realms.

Table 4.3. jaas Scope commands

Command	Parameters	Description
jaas:cancel		Cancels the modification of a JAAS realm.
jaas:group-add	username group	Makes a user part of a group.
jaas:group-create	group	Creates a group in a realm.
jaas:group-delete	username group	Removes a user from a group.
jaas:group-list		Lists groups in a realm.
jaas:group-role-add	group role	Adds a role to a group.
jaas:group-role-delete	group role	Removes a role from a group.
jaas:pending-list		Lists the pending modification on the active JAAS Realm/Login Module.
jaas:realm-list	[options]	Lists JAAS realms.
jaas:realm-manage	[options]	Manages user and roles of a Jaas Realm.
jaas:role-add	username role	Adds a role to a user.
jaas:role-delete	username role	Deletes a role from a user.
jaas:su	[options] [user]	Substitutes user identity.
jaas:sudo	[options] [command]	Executes a command as another user.
jaas:update		Applies pending modification on the edited JAAS Realm.
jaas:user-add	username password	Adds a user.
jaas:user-delete	username	Deletes a user.
jaas:user-list	[options]	Lists the users of the selected JAAS realm/login module.

4.5. Log Scope

The commands within the `log:` scope are used for management of system logs.

Table 4.4. log Scope commands

Command	Parameters	Description
log:clear		Clears log entries.
log:display	[options] [logger]	Displays log entries.
log:exception-display	[logger]	Displays the last occurred exception from the log.
log:get	[options] [logger]	Shows the currently set log level.
log:log	[options] message	Logs a message.
log:set	level [logger]	Sets the log level.
log:tail	[options] [logger]	Continuously displays log entries. Use Ctrl+C to quit this command.

4.6. Shell Scope

The commands within the `shell:` scope are used to provide terminal window commands in the OSGi shell.

Table 4.5. Shell Scope commands

Command	Parameters	Description
<code>shell:alias</code>	<code>command</code>	Creates an alias to a command.
<code>shell:cat</code>	<code>[options] [paths or urls]</code>	Displays the content of a file or URL. Where <code>paths or urls</code> is a list of file paths or urls to display separated by whitespace (use <code>-</code> for STDIN)
<code>shell:clear</code>		Clears the console buffer.
<code>shell:completion</code>	<code>[mode]</code>	Displays or changes the completion mode on the current console session. Where <code>mode</code> is the completion mode to set. The valid completion modes are: <i>global</i> , <i>first</i> , <i>subshell</i> .
<code>shell:date</code>	<code>[options] [format]</code>	Displays the current time in the given FORMAT.
<code>shell:each</code>	<code>values function</code>	Executes a closure on a list of arguments. Where: <ul style="list-style-type: none"> <code>values</code> is the collection of arguments to iterate on <code>function</code> is the function to execute
<code>shell:echo</code>	<code>[options] [arguments]</code>	Echoes or prints arguments to STDOUT, separated by whitespaces.
<code>edit</code>		
<code>shell:env</code>	<code>variable [value]</code>	Gets/sets the value of a console session variable.
<code>shell:exec</code>	<code>command</code>	Executes system processes.
<code>shell:grep</code>	<code>[options] pattern</code>	Prints lines matching the given pattern, in regular expression.
<code>shell:head</code>	<code>[options] [paths or urls]</code>	Displays the first lines of a file. Where <code>paths or urls</code> is a list of file paths or urls to display separated by whitespaces.
<code>shell:history</code>		Prints commands history.
<code>shell:if</code>	<code>condition ifTrue [ifFalse]</code>	If/Then/Else block.
<code>shell:info</code>		Prints system information.
<code>shell:java</code>	<code>[options] className [arguments]</code>	Executes a Java standard application.
<code>less</code>		
<code>shell:logout</code>		Disconnects shell from current session.
<code>more</code>		
<code>shell:new</code>	<code>class [args]</code>	Creates a new java object.
<code>shell:printf</code>	<code>format arguments</code>	Formats and prints arguments.
<code>shell:sleep</code>	<code>[options] duration</code>	Sleeps for a defined <code>duration</code> then wakes up. The default time unit is millisecond, use <code>-s</code> option to use second instead.
<code>shell:sort</code>	<code>[options] [files]</code>	Writes sorted concatenation of all files to standard output, separated by whitespaces.

Command	Parameters	Description
shell:source	<code>script [args]</code>	Runs a script. Where: <ul style="list-style-type: none"> <code>script</code> is a URI pointing to the script <code>args</code> are arguments for the script
shell:stack-traces-print	<code>[print]</code>	Prints the full stack trace in the console when the execution of a command throws an exception.
shell:tac	<code>[options]</code>	Captures the STDIN and returns it as a string. Optionally writes the content to a file.
shell:tail	<code>[options] [path or url]</code>	Displays the last lines of a file. Where <code>path</code> or <code>url</code> is a file path or url to display.
shell:threads	<code>[options] [id]</code>	Prints the current threads (optionally with stacktraces).
shell:watch	<code>[options] command</code>	Watches & refreshes the output of a command.
shell:wc	<code>[options] [files]</code>	Prints newline, word, and byte counts for each file.
shell:while	<code>condition function</code>	Loops while the condition is true. Where: <ul style="list-style-type: none"> <code>condition</code> is the condition of the loop <code>function</code> is the function to execute

4.7. Miscellaneous Scopes

There are a few scopes that offer just one or two commands each. They are listed in the below table.

Table 4.6. Miscellaneous Scope commands

Command	Parameters	Description
package:exports	<code>[options]</code>	Lists exported packages and the bundles that export them.
package:imports	<code>[options]</code>	Lists imported packages and the bundles that import them.
ssh:ssh	<code>[options] hostname</code> <code>[command]</code>	Connects to a remote SSH server.
wrapper:install	<code>[options]</code>	Installs the container as a system service in the OS.



Chapter 5. Deploying Multiple Karaf Containers

This chapter describes how to adapt the configuration of the container(s) to deploy several of them on the same machine.

5.1. Deploying multiple containers using configuration adaption scripts

In order to avoid conflicts between multiple container instances, there are Karaf configuration adaption scripts which automate the adjustment of potentially conflicting parameters. They are based on the **edit** and **property-set** commands described above.

Those scripts are executed within the Karaf container and thus are supported by all platforms compatible with container. So, even if they have a ".sh" extension, their usage is not limited to Linux based operating systems.

To deploy multiple containers:

1. Make sure that the default container you will copy has never been started before. Or if it has already been started, make sure it is stopped (using the **system:shutdown** or **logout** commands or **Ctrl+D**), and that its `data` directory has been deleted, before copying its directory and files to create containers 1, 2 and 3. Otherwise, the data in the default container (which contains absolute paths relating to the default container) will cause problems in the others.
2. In a Linux terminal or Windows command line window, change directory to `<TalendRuntimePath>`.
3. Copy the directory structure of the default container with all its content to produce the first container by executing the most applicable command:

- **cp -r container container1** (Linux)
- **robocopy /e container container1** (Windows)
- **xcopy /e container container1** (older Windows with no robocopy)

Now, you have a second container in directory `container1`.

Repeat the above command to create the second and the third containers in the same way.

4. Start the first container:

- Under Linux, go to directory `<TalendRuntimePath>/container1` and execute the following command:

```
./bin/trun
```

- Under Windows, go to directory `<TalendRuntimePath>\container1` and execute the following command:

```
.\bin\trun.bat
```

Start the second and the third containers in the same way.

5. When the initialization is complete, run the configuration adaption script at the container prompt to update and save the new settings:

- `source scripts/configureC1.sh` for the first container copy
- `source scripts/configureC2.sh` for the second container copy
- `source scripts/configureC3.sh` for the third container copy
- `source scripts/configureC0.sh` resets the parameters to the default values

All necessary parameters adjustments are done using this single script call. Changes performed by the configuration adaption scripts are persistent and reflected in the configuration files in `container/etc`.

6. Shut down and restart the container after applying a configuration adaption script, to make sure the new parameters are used. Most of the parameter changes will be adapted "on the fly", but for the Jobserver parameters in Talend ESB this is not yet possible.

The ports which are configured using the configuration scripts are described in the table below:

Parameter	configureC0.sh	configureC1.sh	configureC2.sh	configureC3.sh
HTTP Port	8040	8041	8042	8043
HTTPS Port	9001	9002	9003	9004
RMI Registry Port	1099	1100	1101	1102
RMI Server Port	44444	44445	44446	44447
SSH Port	8101	8102	8103	8104
Command Port	8000	8010	8020	8030
File Transfer Port	8001	8011	8021	8031
Monitoring Port	8888	8898	8908	8918

5.2. Troubleshooting

If you get the "Port already in use exception" when starting an alternate container, recheck that there is not already a container running using the default parameters.

If you are still getting the error, it may also be that the port is actually in use by an unrelated process, so change the ports in the configuration adaption scripts and rerun these to apply the changes.



Chapter 6. Remote Console

It does not always make sense to manage an instance of Talend ESB using the local console. Talend ESB can be remotely managed using a remote console.

When you start Karaf, it enables a remote console that can be accessed over SSH from any other Karaf console or plain SSH client. The remote console provides all the features of the local console and gives a remote user complete control over the container and services running inside of it.

6.1. Configuring remote instances

The SSH hostname and port number of the remote console is configured in the `<RuntimeContainerPath>/etc/org.apache.karaf.shell.cfg` configuration file with the following defaults values:

```
sshPort=8101
sshHost=0.0.0.0
sshRealm=karaf
hostKey=${karaf.base}/etc/host.key
```

You can change this configuration using the config commands or by editing the above file, but you will need to restart the ssh console in order for it to use the new parameters.

```
# define helper functions
bundle-by-sn = { bm = new java.util.HashMap ; //
  each (bundles) { $bm put ($it symbolicName) $it } ; $bm get $1 }
bundle-id-by-sn = { b = (bundle-by-sn $1) ; //
  if { $b } { $b bundleId } { -1 } }

# edit config
config:edit org.apache.karaf.shell
config:property-set sshPort 8102
```

```
config:update
# force a restart
bundle:restart --force (bundle-id-by-sn org.apache.karaf.shell.ssh)
```

6.2. Connecting and disconnecting remotely

6.2.1. Using the ssh:ssh command

You can connect to a remote Karaf's console using the `ssh:ssh` command.

```
karaf@trun> ssh:ssh -l tadmin -P tadmin -p 8101 hostname
```



The default password is `tadmin` but we recommend changing it. See [Security](#) for more information.

To confirm that you have connected to the correct Karaf instance, type **shell:info** at the `karaf>` prompt. Information about the currently connected instance is returned, as shown.

```
Karaf
  Karaf home           /local/apache-karaf-2.0.0
  Karaf base           /local/apache-karaf-2.0.0
  OSGi Framework      org.eclipse.osgi - 3.5.1.R35x_v20090827
JVM
  Java Virtual Machine Java HotSpot(TM) Server VM version 14.1-b02
  ...
```

6.2.2. Using the Karaf client

The Karaf client allows you to securely connect to a remote Karaf instance without having to launch a Karaf instance locally.

For example, to quickly connect to a Karaf instance running in server mode on the same machine, run the following command from the `karaf-install-dir` directory: **bin/client**. More usually, you would provide a hostname, port, username and password to connect to a remote instance. And, if you were using the client within a larger script, you could append console commands as follows:

```
bin/client -a 8101 -h hostname -u tadmin -p tadmin feature:install wrapper
```

To display the available options for the client, type:

```
> bin/client --help
Apache Karaf client
  -a [port]      specify the port to connect to
  -h [host]      specify the host to connect to
  -u [user]      specify the user name
  -p [password] specify the password
  --help        shows this help message
  -v            raise verbosity
  -r [attempts] retry connection establishment (up to attempts times)
```



```
-d [delay]      intra-retry delay (defaults to 2 seconds)
[commands]     commands to run
If no commands are specified, the client will be put in an interactive mode
```

6.2.3. Using a plain SSH client

You can also connect using a plain SSH client from your *nix system or Windows SSH client like Putty.

```
~$ ssh -p 8101 tadmin@localhost
tadmin@localhost's password:
```

6.2.4. Disconnecting from a remote console

To disconnect from a remote console, press **Ctrl+D**, **shell:logout** or simply **logout** at the Karaf prompt.

6.3. Stopping a remote instance

6.3.1. Using the remote console

If you have connected to a remote console using **ssh:ssh** or the Karaf client, you can stop the remote instance using **system:shutdown**.



Pressing Ctrl+D in a remote console simply closes the remote connection and returns you to the local shell without shutting off the remote instance.

6.3.2. Using the Karaf client

To stop a remote instance using the Karaf client, run the following from the `karaf-install-dir/lib` directory:

```
bin/client -u tadmin -p tadmin -a 8101 -h hostname system:shutdown
```


Chapter 7. Apache CXF and Camel, commands and tuning

These are commands that are related to Apache CXF and Camel functionality.



CXF and Camel are already preconfigured in the Talend Runtime container, but you need to install Camel / CXF features in native Karaf to use these commands.

To view help on a particular command, type the command followed by `--help` or use the `help` command followed by the name of the command:

```
karaf@trun> cxf:list-endpoints --help
DESCRIPTION
  cxf:list-endpoints

  Lists all CXF Endpoints on a Bus.

SYNTAX
  cxf:list-endpoints [options] [bus]

ARGUMENTS
  bus      The CXF bus name where to look for the Endpoints

OPTIONS
  --help      Display this help message
  -f, --fulladdress  Display full address of an endpoint
```

7.1. Commands supplied by CXF

These commands related to CXF functionality. For more information on CXF see <http://cxf.apache.org/>.

Table 7.1. CXF commands

Command	Parameters	Description
cxf:list-busses		Lists all CXF Busses
cxf:list-endpoints	[options] [bus]	Lists all CXF Endpoints on a Bus. Where bus is the optional CXF bus name where to look for the Endpoints
cxf:stop-endpoint	bus endpoint	Stops a CXF Endpoint on a Bus. Where: <ul style="list-style-type: none"> bus is CXF bus name where to look for the Endpoint. endpoint is the Endpoint name to stop.
cxf:start-endpoint	[options] bus endpoint	Starts a CXF Endpoint on a Bus. Where: <ul style="list-style-type: none"> bus is CXF bus name where to look for the Endpoint. endpoint is the Endpoint name to start.

7.2. Commands supplied by Camel

These commands related to Camel functionality. Help for these commands is available at <http://camel.apache.org/karaf.html>. These are for version Camel 2.15.3, which is the current version used by Talend ESB.



Use TAB key for completion on the name parameters below.

Table 7.2. Camel commands

Command	Parameters	Description
camel:component-list	[options] name	Lists all Camel components that are in use in Karaf.
camel:context-list		Lists the Camel contexts available in the current Karaf instance
camel:context-inflight	[options] name	Lists inflight exchanges.
camel:context-info	[options] name	Displays detailed information about a given Camel context
camel:context-resume	context	Resumes a Camel context.
camel:context-start	context	Starts the given Camel context.
camel:context-stop	context	Stops the given Camel context. It becomes unavailable and can not be started again.
camel:context-suspend	context	Suspends a Camel context.
camel:eip-explain	[options] name nameOrId	Explains the EIP in the CamelContext. Where:

Command	Parameters	Description
		<ul style="list-style-type: none"> name is the name of the Camel context nameOrId is the name of the EIP or a node id to refer to a specific node from the routes
camel:endpoint-explain	[options] name	Explains all Camel endpoints available in the CamelContext.
camel:endpoint-list	[options] name	Lists endpoints from all camel contexts available in the current Karaf instance.
camel:rest-registry-list	[options] name	Lists all Camel REST services enlisted in the Rest Registry from a CamelContext.
camel:rest-show	context	Display the Camel REST definition in XML.
camel:route-list	[name]	Displays the list of Camel routes available in the current Karaf instance, where name is the Camel context name where to look for the route.
camel:route-info	[options] name	Provides detail information about a Camel route.
camel:route-profile	route [context]	Displays information about a Camel route. Where: <ul style="list-style-type: none"> route is the Camel route ID. context is the Camel context name.
camel:route-reset-stats	context	Resets route performance stats from a CamelContext.
camel:route-resume	route context	Resumes the given route. Where: <ul style="list-style-type: none"> route is the Camel route ID or a wildcard expression. context is the Camel context name.
camel:route-show	route [context]	Displays the Camel route definition in XML.
camel:route-start	route [context]	Starts the given route. Where: <ul style="list-style-type: none"> route is the Camel route ID or a wildcard expression. context is the Camel context name.
camel:route-stop	route [context]	Stops the given route. Where: <ul style="list-style-type: none"> route is the Camel route ID or a wildcard expression. context is the Camel context name.
camel:route-suspend	route [context]	Suspends the given route. Where: <ul style="list-style-type: none"> route is the Camel route ID or a wildcard expression. context is the Camel context name.

7.3. Configuring CXF workqueues

CXF workqueues are used for queuing incoming work requests using a thread pool.

The `etc/org.apache.cxf.workqueues.cfg` configuration file is used for workqueue configuration (without this file, the configuration would need to be done for each individual bundle using CXF services). This process can significantly optimize the performance of HTTP/CXF Service request handling in the Talend Runtime container.

This mechanism allows you to configure global workqueues for use by all bundles that are created, and normally services would share a thread pool. However an individual service can override this via local configuration if they have specific requirements.

7.3.1. Usage of queues

Workqueue settings may be used in different scenarios:

- For SOAP-based web service providers, the ws-addressing workqueue (*Configuration files*) can be used, and the default workqueue for One-Way (no response) SOAP calls. (Workqueues are not applicable for REST endpoints.)
- For JMS transport, workqueues are active to handle continuations; for regular JMS, there are similar settings, like the number of consumers or the thread pool for execution, which are implemented outside of workqueues.
- SOAP clients can use the "http-conduit" workqueue for asynchronous calls.

7.3.2. Configuration parameters

A configuration file contains the following parameters:

Name	Description	Default value
<code>org.apache.cxf.workqueue.names</code>	One or more names of workqueues, separated by commas	'default'
<code>org.apache.cxf.workqueue.default.highWaterMark</code>	Maximum number of threads to work on the queue	10
<code>org.apache.cxf.workqueue.default.lowWaterMark</code>	Minimum number of threads to work on the queue	5
<code>org.apache.cxf.workqueue.default.initialSize</code>	Initial number of threads to work on the queue	7
<code>org.apache.cxf.workqueue.default.dequeueTimeout</code>	This is the keep alive time for the threadpool executor. This allows excess threads to be terminated when idle for longer than this time, and they can be created again later if needed.	100 (optional)
<code>org.apache.cxf.workqueue.default.queueSize</code>	Maximum number of entries in the queue	100 (optional)

7.3.3. Configuration files

In the Talend Runtime container, the default configuration file is in `etc/org.apache.cxf.workqueues.cfg`.

Note this file is normally not used or edited directly; it gives workqueue default values that each CXF-using bundle can choose to employ when it creates its work queues.

To configure a workqueue, there is a specific corresponding file: `org.apache.cxf.workqueues-n.cfg`, where typically "n" is the same as the workqueue name; for example, `org.apache.cxf.workqueues-http-conduit.cfg` would configure the `http-conduit` workqueue. At the moment, this functionality is for pre-defined work queues; it is not possible to use it for user-defined workqueues.

Here is the list of pre-defined workqueue names:

Name	Description
default	this means using the default values.
http-conduit	On the client side when using the asynchronous methods, the HTTP conduit must wait for and process the response on a background thread. This can control the queue that is used specifically for that purpose, to limit or expand the number of outstanding asynchronous requests.
jms-continuation	This is used by the JMS transport to handle continuations.
local-transport	The local transport being based on PipedInput/OutputStreams requires the use of separate threads; this workqueue can be used to configure the queue used exclusively for the local-transport.
ws-addressing	For decoupled cases, the ws-addressing layer may need to process the request on a background thread. This can control the workqueue it uses.



You can also update these variables (for example `org.apache.cxf.workqueue.default.initialSize`) using the standard Karaf configuration commands.



Chapter 8. Security

You can configure security settings such as usernames, passwords and roles, and enable password encryption.

The security framework also allows runtime deployment of JAAS based configuration.

It is also possible to deploy security providers such as BouncyCastle.

8.1. Managing users and passwords

The default security configuration uses a property file located at `<RuntimeContainerPath>/etc/users.properties` to store authorized users and their passwords. The default user name is `tadmin` and the associated password is `tadmin` too. We strongly encourage you to change the default password by editing the above file before moving Karaf into production.

The users are currently used in three different places in Karaf:

- access to the SSH console
- access to the JMX management layer
- access to the Web console

Those three ways all delegate to the same JAAS based security authentication.

The `users.properties` file contains one or more lines, each line defining a user, its password and the associated roles:

```
user=password[,role][,role]...
```

8.2. Managing roles

JAAS roles can be used by various components. The three management layers (SSH, JMX and WebConsole) all use a global role based authorization system. The default role name is configured in the `etc/system.properties` using the `karaf.local.roles` system property and the default value is `admin`. All users authenticating for the management layer must have this role defined. The syntax for this value is the following:

```
[classname:]principal
```

Where `classname` is the class name of the principal object (defaults to `org.apache.karaf.jaas.modules.RolePrincipal`) and `principal` is the name of the principal of that class (defaults to `admin`). Note that roles can be changed for a given layer using `ConfigAdmin` in the following configurations:

Layer	PID	Value
SSH	<code>org.apache.karaf.shell</code>	<code>sshRole</code>
JMX	<code>org.apache.karaf.management</code>	<code>jmxRole</code>
Web	<code>org.apache.karaf.webconsole</code>	<code>role</code>

8.3. Enabling password encryption

In order to not keep the passwords in plain text, the passwords can be stored encrypted in the configuration file. This can be easily enabled using the following commands:

```
# edit config
config:edit org.apache.karaf.jaas
config:property-set encryption.enabled true
config:update

# force a restart
bundle:restart
```

The passwords will be encrypted automatically in the `etc/users.properties` configuration file the first time the user logs in. Encrypted passwords are prepended with `{CRYPT}` so that are easy to recognize.

8.4. Managing realms

Karaf supports [JAAS](#) with some enhancements to allow JAAS to work nicely in an OSGi environment. This framework also features an OSGi keystore manager with the ability to deploy new keystores or truststores at runtime.

8.4.1. Overview

The Security framework feature of Karaf allows runtime deployment of JAAS based configuration for use in various parts of the application. This includes the remote console login, which uses the karaf realm, but which is configured with a dummy login module by default. These realms can also be used by the NMR, JBI components or the JMX server to authenticate users logging in or sending messages into the bus.

In addition to JAAS realms, you can also deploy keystores and truststores to secure the remote shell console, setting up HTTPS connectors or using certificates for WS-Security.

A very simple XML schema for spring has been defined, allowing the deployment of a new realm or a new keystore very easily.

8.4.2. Schema

To override or deploy a new realm, you can use the following XSD which is supported by a Spring namespace handler and can thus be defined in a Spring xml configuration file.

You can find the schema at the following location: <http://karaf.apache.org/xmlns/jaas/v1.1.0>.

Here are two examples using this schema:

```
<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0"
  xmlns:jaas="http://karaf.apache.org/xmlns/jaas/v1.0.0"
  xmlns:ext="http://aries.apache.org/blueprint/xmlns/blueprint-ext/
    v1.0.0">

  <!-- Bean that allows the ${karaf.base} property to be resolved -->
  <ext:property-placeholder
    placeholder-prefix="${ " placeholder-suffix=" }"/>

  <jaas:config name="myrealm">
    <jaas:module className="org.apache.karaf.jaas.modules.properties.
      PropertiesLoginModule"
      flags="required">
      users = ${karaf.base}/etc/users.properties
    </jaas:module>
  </jaas:config>

</blueprint>
```

```
<jaas:keystore xmlns:jaas="http://karaf.apache.org/xmlns/jaas/v1.1.0"
  name="ks"
  rank="1"
  path="classpath:privatestore.jks"
  keystorePassword="keyStorePassword"
  keyPasswords="myalias=myAliasPassword">
</jaas:keystore>
```

The `id` attribute is the blueprint id of the bean, but it will be used by default as the name of the realm if no name attribute is specified. Additional attributes on the `config` elements are a `rank`, which is an integer. When the `LoginContext` looks for a realm for authenticating a given user, the realms registered in the OSGi registry are matched against the required name. If more than one realm is found, the one with the highest rank will be used, thus allowing the override of some realms with new values. The last attribute is `publish` which can be set to false to not publish the realm in the OSGi registry, thereby disabling the use of this realm.

Each realm can contain one or more module definitions. Each module identifies a `LoginModule` and the `className` attribute must be set to the class name of the login module to use. Note that this login module must be available from the bundle classloader, so either it has to be defined in the bundle itself, or the needed package needs to be correctly imported. The `flags` attribute can take one of four values that are explained on the [JAAS documentation](#). The content of the `module` element is parsed as a properties file and will be used to further configure the login module.

Deploying such a code will lead to a `JaasRealm` object in the OSGi registry, which will then be used when using the JAAS login module.

8.4.2.1. Configuration override and use of the rank attribute

The `rank` attribute on the `config` element is tied to the ranking of the underlying OSGi service. When the JAAS framework performs an authentication, it will use the realm name to find a matching JAAS configuration. If multiple configurations are used, the one with the highest `rank` attribute will be used. So if you want to override the default security configuration in Karaf (which is used by the ssh shell, web console and JMX layer), you need to deploy a JAAS configuration with the name `name="karaf"` and `rank="1"`.

```
<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0"
  xmlns:jaas="http://karaf.apache.org/xmlns/jaas/v1.1.0"
  xmlns:ext="http://aries.apache.org/blueprint/xmlns/blueprint-ext/
    v1.0.0">

  <!-- Bean that allows the ${karaf.base} property to be resolved -->
  <ext:property-placeholder
    placeholder-prefix="$[" placeholder-suffix="]"/>

  <jaas:config name="karaf" rank="1">
    <jaas:module className="org.apache.karaf.jaas.modules.properties.
      PropertiesLoginModule" flags="required">
      users = ${karaf.base}/etc/users.properties
      ...
    </jaas:module>
  </jaas:config>

</blueprint>
```

8.4.3. Architecture

Due to constraints in the JAAS specification, one class has to be available for all bundles. This class is called `ProxyLoginModule` and is a `LoginModule` that acts as a proxy for an OSGi defined `LoginModule`. If you plan to integrate this feature into another OSGi runtime, this class must be made available from the system classloader and the related package part of the boot delegation classpath (or be deployed as a fragment attached to the system bundle).

The xml schema defined above allows the use of a simple xml (leveraging spring xml extensibility) to configure and register a JAAS configuration for a given realm. This configuration will be made available into the OSGi registry as a `JaasRealm` and the OSGi specific configuration will look for such services. Then the proxy login module will be able to use the information provided by the realm to actually load the class from the bundle containing the real login module.

8.4.4. Available realms

Karaf comes with several login modules that can be used to integrate into your environment.

8.4.4.1. PropertiesLoginModule

This login module is the one configured by default. It uses a properties text file to load the users, passwords and roles from. This file uses the [properties file format](#). The format of the properties are as follows, each line defining a user, its password and the associated roles: `user=password[,role][,role]...`

```
<jaas:config name="karaf">
  <jaas:module className=
    "org.apache.karaf.jaas.modules.properties.PropertiesLoginModule"
    flags="required">
    users = ${karaf.base}/etc/users.properties
  </jaas:module>
</jaas:config>
```

8.4.4.2. OsgiConfigLoginModule

The OsgiConfigLoginModule uses the OSGi ConfigurationAdmin service to provide the users, passwords and roles. Instead of `users` for the PropertiesLoginModule, this configuration uses a `pid` value for the process ID of the configuration containing user definitions.

8.4.4.3. JDBCLoginModule

The JDBCLoginModule uses a database to load the users, passwords and roles from, provided a data source (normal or XA). The data source and the queries for password and role retrieval are configurable, with the use of the following parameters.

Name	Description
<code>datasource</code>	The datasource as on OSGi ldap filter or as JNDI name
<code>query.password</code>	The SQL query that retrieves the password of the user
<code>query.role</code>	The SQL query that retrieves the roles of the user

Passing a data source as an OSGi ldap filter

To use an OSGi ldap filter, the prefix `osgi:` needs to be provided. See the example below:

```
<jaas:config name="karaf">
  <jaas:module
    className="org.apache.karaf.jaas.modules.jdbc.JDBCLoginModule"
    flags="required">
    datasource = osgi:javax.sql.DataSource/  \
      (osgi.jndi.service.name=jdbc/karafdb)
    query.password = SELECT PASSWORD FROM USERS WHERE USERNAME=?
    query.role = SELECT ROLE FROM ROLES WHERE USERNAME=?
  </jaas:module>
</jaas:config>
```

Passing a data source as a JNDI name

To use a JNDI name, the prefix `jndi:` needs to be provided. The example below assumes the use of Aries JNDI to expose services via JNDI.

```
<jaas:config name="karaf">
  <jaas:module
    className="org.apache.karaf.jaas.modules.jdbc.JDBCLoginModule"
```

```

    flags="required">
    datasource = jndi:aries:services/javax.sql.DataSource/  \\
      (osgi.jndi.service.name=jdbc/karafdb)
      query.password = SELECT PASSWORD FROM USERS WHERE USERNAME=?
      query.role = SELECT ROLE FROM ROLES WHERE USERNAME=?
  </jaas:module>
</jaas:config>

```

8.4.4.4. LDAPLoginModule

The LDAPLoginModule uses a LDAP to load the users and roles, bind the users on the LDAP to check passwords. The LDAPLoginModule supports the following parameters:

Name	Description
connection.url	The LDAP connection URL, e.g. ldap://hostname
connection.username	Admin username to connect to the LDAP. This parameter is optional, if it's not provided, the LDAP connection will be anonymous.
connection.password	Admin password to connect to the LDAP. Only used if the connection.username is specified.
user.base.dn	The LDAP base DN used to looking for user, e.g. ou=user,dc=apache,dc=org
user.filter	The LDAP filter used to looking for user, e.g. (uid=%u) where %u will be replaced by the username.
user.search.subtree	If "true", the user lookup will be recursive (SUBTREE). If "false", the user lookup will be performed only at the first level (ONELEVEL).
role.base.dn	The LDAP base DN used to looking for roles, e.g. ou=role,dc=apache,dc=org
role.filter	The LDAP filter used to looking for user's role, e.g. (member:=uid=%u)
role.name.attribute	The LDAP role attribute containing the role string used by Karaf, e.g. cn
role.search.subtree	If "true", the role lookup will be recursive (SUBTREE). If "false", the role lookup will be performed only at the first level (ONELEVEL).
authentication	Define the authentication backend used on the LDAP server. The default is simple.
initial.context.factory	Define the initial context factory used to connect to the LDAP server. The default is com.sun.jndi.ldap.LdapCtxFactory
ssl	If "true" or if the protocol on the connection.url is ldaps, an SSL connection will be used
ssl.provider	The provider name to use for SSL
ssl.protocol	The protocol name to use for SSL (SSL for example)
ssl.algorithm	The algorithm to use for the KeyManagerFactory and TrustManagerFactory (PKIX for example)
ssl.keystore	The key store name to use for SSL. The key store must be deployed using a jaas:keystore configuration.
ssl.keyalias	The key alias to use for SSL
ssl.truststore	The trust store name to use for SSL. The trust store must be deployed using a jaas:keystore configuration.

An example of LDAPLoginModule usage follows:

```

<jaas:config name="karaf">
  <jaas:module

```

```

    className="org.apache.karaf.jaas.modules.ldap.LDAPLoginModule"
    flags="required">
    connection.url = ldap://localhost:389
    user.base.dn = ou=user,dc=apache,dc=org
    user.filter = (cn=%u)
    user.search.subtree = true
    role.base.dn = ou=group,dc=apache,dc=org
    role.filter = (member:=uid=%u)
    role.name.attribute = cn
    role.search.subtree = true
    authentication = simple
  </jaas:module>
</jaas:config>

```

If you want to use an SSL connection, the following configuration can be used as an example:

```

<ext:property-placeholder />

<jaas:config name="karaf" rank="1">
  <jaas:module
    className="org.apache.karaf.jaas.modules.ldap.LDAPLoginModule"
    flags="required">
    connection.url = ldaps://localhost:10636
    user.base.dn = ou=users,ou=system
    user.filter = (uid=%u)
    user.search.subtree = true
    role.base.dn = ou=groups,ou=system
    role.filter = (uniqueMember=uid=%u)
    role.name.attribute = cn
    role.search.subtree = true
    authentication = simple
    ssl.protocol=SSL
    ssl.truststore=ks
    ssl.algorithm=PKIX
  </jaas:module>
</jaas:config>

<jaas:keystore name="ks"
  path="file:///${karaf.home}/etc/trusted.ks"
  keystorePassword="secret" />

```

8.5. Deploying security providers

Some applications require specific security providers to be available, such as BouncyCastle. The JVM imposes some restrictions about the use of such provider JAR files, namely, that they need to be signed and be available on the boot classpath. One way to deploy such providers is to put them in the JRE folder at `$JAVA_HOME/jre/lib/ext` and modify each provider's security policy configuration (`$JAVA_HOME/jre/lib/security/java.security`) in order to register them. While this approach works fine, it has a global effect and requires that all servers are configured accordingly.

However Talend ESB offers a simple way to configure additional security providers:

- put your provider jar in `[karaf-install-dir]/lib/ext`
- modify the `[karaf-install-dir]/etc/config.properties` configuration file to add the following property:

```
org.apache.karaf.security.providers = xxx,yyy
```

The value of this property is a comma separated list of the provider class names to register. For example:

```
org.apache.karaf.security.providers = \\
  org.bouncycastle.jce.provider.BouncyCastleProvider
```

In addition, you may want to provide access to the classes from those providers from the system bundle so that all bundles can access those. It can be done by modifying the `org.osgi.framework.bootdelegation` property in the same configuration file:

```
org.osgi.framework.bootdelegation = ...,org.bouncycastle*
```




Chapter 9. HTTP Configuration

HTTP is a request and response protocol used to enable communications between clients and servers.

Basically, a client send a request to a server on a particular port. The server listening to that port is waiting for the client's request. And when it receives the request from the client, the server sends back the corresponding response.

To secure this communication, a SSL protocol can be used on top of HTTP to provide security. This will allow the data to be encrypted and safely transfered through a secure HTTP protocol: HTTPS.

9.1. Server HTTP Configuration

The Talend Runtime provides support for HTTP and HTTPS by default with the help of the pax web component. For more information, see the documentation of Pax Web on <http://team.ops4j.org/wiki/display/paxweb/Documentation>.

HTTP / HTTPS configuration for Talend Runtime is done in the following configuration file: `<TalendRuntimePath>/container/etc/org.ops4j.pax.web.cfg`.

9.1.1. Basic configuration

To enable HTTP and HTTPS in Talend Runtime, configure the properties of the `org.ops4j.pax.web.cfg` file as follows:

Property	Default	Description
<code>org.osgi.service.http.port</code>	8080	This property specifies the port used for servlets and resources accessible via HTTP.

Property	Default	Description
		The default value for this property is 8080. You can specify a value of 0 (zero), if you wish to allow Pax Web to automatically determine a free port to use for HTTP access.
<code>org.osgi.service.http.port.secure</code>	8443	This property specifies the port used for servlets and resources accessible via HTTPS. The default value for this property is 8443. You can specify a value of 0 (zero), if you wish to allow Pax Web to automatically determine a free port to use for HTTPS access.
<code>org.osgi.service.http.enabled</code>	true	This property specifies if the HTTP is enabled or disabled. If the value is set to "true", the support for HTTP access is enabled. If the value is set to "false", the support for HTTP access is disabled. The default value is "true".
<code>org.osgi.service.http.secure.enabled</code>	false	This property specifies if the HTTPS is enabled. If the value is set to "true", the support for HTTPS access is enabled. If the value is set to "false", the support for HTTPS access is disabled. The default value is "false".

9.1.2. SSL configuration

To encrypt the communication and secure the identification of a server, you can use a HTTPS protocol. HTTPS is based on SSL, which supports the encryption of messages sent via HTTP.

To secure a communication, HTTPS uses key pairs containing one public key and one private key. Data is encrypted with one key and can only be decrypted with the other key of the key pair. This establishes trust and privacy in message transfers.

To authenticate the Talend Runtime, you need to configure its private key in the `org.ops4j.pax.web.cfg`.

Property	Default	Description
<code>org.ops4j.pax.web.ssl.keystore</code>		Path to the keystore file. See http://team.ops4j.org/wiki/display/paxweb/SSL+Configuration for details.
<code>org.ops4j.pax.web.ssl.keystore.type</code>	JKS	This property specifies the keystore type. By default, the value is JKS.
<code>org.ops4j.pax.web.ssl.password</code>		Password used for keystore integrity check.
<code>org.ops4j.pax.web.ssl.keypassword</code>		Password used for keystore.

9.1.3. Advanced configuration

For a complete list of all advanced configuration properties, see <http://team.ops4j.org/wiki/display/paxweb/Configuration/>.

9.1.4. Default configuration

The Talend Runtime is deployed with the following configuration of the <TalendRuntimePath>/container/etc/org.ops4j.pax.web.cfg file.

```
org.osgi.service.http.port=8040
org.osgi.service.http.port.secure=9001
org.osgi.service.http.secure.enabled=true
org.ops4j.pax.web.ssl.keystore=./etc/keystore.jks
org.ops4j.pax.web.ssl.password=password
org.ops4j.pax.web.ssl.keypassword=password
```



The Certificates deployed with the Talend Runtime by default must not be used for production, but only for demo purposes.



And the key password corresponds to the password generated by the user when he/she generated the key.

9.1.5. Enabling client authentication for SSL

To exchange certificates and allow only "trusted" clients to use the Talend Runtime container HTTP service, you need to follow the following instructions.

1. Enable the HTTP client auth support in the Karaf-based Talend Runtime container.

When you install the HTTP feature, the container leverages Pax-Web to provide HTTP OSGi service:

```
karaf@trun> feature:install http
```

2. Add a custom etc/org.ops4j.pax.web.cfg file with the following content:

```
org.osgi.service.http.port=8181

org.osgi.service.http.port.secure=8443
org.osgi.service.http.secure.enabled=true
org.ops4j.pax.web.ssl.keystore=./etc/keystores/keystore.jks
org.ops4j.pax.web.ssl.password=password
org.ops4j.pax.web.ssl.keypassword=password
#org.ops4j.pax.web.ssl.clientauthwanted=false
org.ops4j.pax.web.ssl.clientauthneeded=true
```

The `clientauthwanted` and `clientauthneeded` properties are valid for Karaf 2.2.x which uses Pax Web 1.0.x. For more information about the version of Karaf your Talend Runtime container is based on, see the Talend Installation and Upgrade Guide or the Release Notes.

Thanks to the `clientauthneeded` property, the client is "forced" to be trusted.

9.1.5.1. Create the trusted client certificate

You are going to use a keytool (provided with the JDK) to manipulate the keys and certificates.

1. Create two key pairs:
 - one for the server side (use for SSL),
 - one as an example of the client side (use for "trust", should be performed for each client, on the client side).

```
mkdir -p etc/keystores
cd etc/keystores
keytool -genkey -keyalg RSA -validity 365 -alias serverkey -keypass
password -storepass password -keystore keystore.jks
keytool -genkey -keyalg RSA -validity 365 -alias clientkey -keypass
password -storepass password -keystore client.jks
```

These key are self-signed. In a production system, you should use a Certificate Authority (CA).

2. Export the client certificate to be imported in the server keystore:

```
keytool -export -rfc -keystore clientKeystore.jks -storepass password -
alias clientkey -file client.cer
keytool -import -trustcacerts -keystore keystore.jdk -storepass
password -alias clientkey -file client.cer
```

3. Check that the client certificate is trusted in our keystore:

```
keytool -list -v -keystore keystore.jks
...
Alias name: clientkey
Creation date: Dec 12, 2012
Entry type: trustedCertEntry
...
```

4. You can now remove the `client.cer` certificate.

9.1.5.2. Start the container and test with the WebConsole

1. Start the Talend Runtime container:

- `bin/trun` for Linux
- `bin/trun.bat` for Windows

2. Install the WebConsole feature:

```
karaf@trun> feature:install webconsole
```

If you try to access to the WebConsole (using a simple browser) using `https://localhost:8443/system/console`, you get the following message:

```
An error occurred during a connection to localhost:8443.
SSL peer cannot verify your certificate.
(Error code: ssl_error_bad_cert_alert)
```

Which is normal as the browser does not have any trusted certificate.

3. Add the client certificate in the browser.

Firefox supports the import of PKCS12 keystore. So, you are going to "transform" the JKS keystore into a PKCS12 keystore:

```
keytool -importkeystore -srckeystore clientKeystore.jks -srcstoretype
JKS -destkeystore client.pfx -deststoretype PKCS12
```

```

Enter destination keystore password:
Re-enter new password:
Enter source keystore password:
Entry for alias clientkey successfully imported.
Import command completed: 1 entries successfully imported, 0 entries
failed or cancelled

```

Now, you can import the client certificate in Firefox. To do so, in the **Tools** menu, click the **Options** entry, and click on the **Advanced** tab.

You can go in **Certificates** tab and click on **View Certificates** button.

In the **Your Certificates** tab, you can click on the **Import...** button and choose the `client.pfx` keystore file.

4. If you try to access `https://localhost:8443/system/console` again, you will have access as a trusted client and use it.

9.1.6. Configuring jetty for SSL

To turn off having pax-web to directly create the connector, change the `etc/org.ops4j.pax.web.cfg` file as follows:

```

#org.osgi.service.http.port.secure=9001
#org.osgi.service.http.secure.enabled=true
....
org.ops4j.pax.web.config.file=${karaf.base}/etc/jetty.xml

```

In `etc/jetty.xml`, remove the connector already defined there ("`org.eclipse.jetty.server.nio.BlockingChannelConnector`") and replace it with the following one:

```

<Call name="addConnector">
  <Arg>
    <New
      class="org.eclipse.jetty.server.ssl.SslSelectChannelConnector">
        <Set name="port">9001</Set>
        <Set name="maxIdleTime">30000</Set>
        <Set name="keystore">./etc/keystores/keystore.jks</Set>
        <Set name="password">password</Set>
        <Set name="keyPassword">password</Set>
      </New>
    </Arg>
  </Call>

```

Those settings puts the connector on port 9001 to use the **SslSelectChannelConnector** which provides working continuation support.

9.2. Client HTTP Configuration

This section discusses configuring HTTPS using the OSGi Configuration Admin Service.

For additional information, see CXF SSL configuration in <http://cxf.apache.org/docs/client-http-transport-including-ssl-support.html>, which has some sample configurations supported by CXF.

9.2.1. OSGi configuration files

If a Web service is deployed in Talend Runtime as an OSGi bundle, it is now possible to configure SSL via the OSGi Configuration Admin Service. This may be done using the following files (located at `<TalendRuntimePath>/container/etc/`):

- a generic configuration `org.apache.cxf.http.conduits-common.cfg`, used by all HTTPS endpoints in the container,
- endpoint-specific files: `org.apache.cxf.http.conduits-<endpoint_name>.cfg`,
- there may also be additional endpoint-specific configuration files.

Note that instead of the tree-structured XML configurations, these configurations are flat property files. The properties are named after their XML equivalents and are described in detail in [HTTP Conduit OSGi Configuration Parameters](#).

The non-OSGi style of SSL configuration is supported and if present will take precedence over the OSGi configuration.

9.2.2. HTTP Conduit OSGi Configuration Parameters

The configuration files described in this section are located in `<TalendRuntimePath>/container/etc/org.apache.cxf.http.conduits-<endpoint_name>.cfg` in the Talend Runtime.

As an example of the syntax involved, here are the contents of the general `org.apache.cxf.http.conduits-common.cfg` configuration file:

```
url = https://localhost.*
tlsClientParameters.disableCNCheck = true
tlsClientParameters.trustManagers.keyStore.type = JKS
tlsClientParameters.trustManagers.keyStore.password = password
tlsClientParameters.trustManagers.keyStore.file = ./etc/keystore.jks
tlsClientParameters.keyManagers.keyStore.type = JKS
tlsClientParameters.keyManagers.keyStore.password = password
tlsClientParameters.keyManagers.keyStore.file = ./etc/keystore.jks
tlsClientParameters.cipherSuitesFilter.include = *_EXPORT_.*,*_EXPORT\
1024_.*,*_WITH_DES_.*,*_WITH_AES_.*,*_WITH_NULL_.*,*_DH_anon_.*
```

9.2.2.1. The url parameter

The `url` parameter is one of the main parameters. In the configuration files, the `url` parameter defines the list of matching client endpoints for which the contained configuration parameters are applied. (The client endpoint address is retrieved using `HTTPConduit.getAddress()`).

`url` may be a full endpoint address or may be a regular expression containing wild cards - for example:

- `" .* "` matches all endpoints,
- `https.*` matches all client addresses starting with "https".
- `https://localhost.*` matches all client addresses starting with "https" only at localhost.

Note the `org.apache.cxf.http.conduits-common.cfg` file above restricts access to local servers, as the configuration points to a keystore with certificates suitable for the Talend samples. This keystore does not contain the required root certificates from the public certification authorities for public servers, such as Salesforce servers, which are in the standard Java keystore.

All parameters contained in all matching configuration files are collected:

1. in the order defined by the `order` parameter (see the table in [The order parameter](#)),
2. then by an exact match,
3. then by a configuration with a matching conduit bean name.

If a parameter is defined in multiple matching configuration files, then the last parameter definition found is the one that is used.

9.2.2.2. The order parameter

This parameter defines the order in which the parameters in configuration files are applied. Each file has a unique value of `order`. For example:

`abc.cfg`:

```
order = 1
url = .*
client.ReceiveTimeout = 60000;
```

`xyz.cfg`:

```
order = 2
url = .*busy.*
client.ReceiveTimeout = 120000;
```

If the endpoint address contains "busy", then both config files match as applicable, according to the rules in [The url parameter](#).

In this case, `client.ReceiveTimeout` will have the longer timeout value 120000 because the order parameters stipulate that `xyz.cfg` is applied after `abc.cfg`.

9.2.2.3. Configuration properties

In this table, we look at the complete list of possible properties:

Property	Default	Description
<code>url</code>		The endpoint URL - either defined as exact string or as regular expression pattern (see The url parameter)
<code>order</code>	50	Defines the order in which parameters are applied (see The order parameter).
<code>name</code>		If name is defined and is equal to the conduit bean name, <code>HTTPConduit.getBeanName()</code> , the parameter definitions have highest priority, overwriting and extending other matching configurations.
<code>tlsClientParameters.secureSocketProtocol</code>	TLS	Protocol Name. Most common examples are "SSL", "TLS" or "TLSv1".
<code>tlsClientParameters.sslCacheTimeout</code>	JDK default	Sets the SSL Session Cache timeout value for client sessions handled by CXF.
<code>tlsClientParameters.jsseProvider</code>		JSSE provider name.

Property	Default	Description
tlsClientParameters.disableCNCheck	false	Indicates whether that the hostname given in the HTTPS URL will be checked against the service's Common Name (CN) given in its certificate during SOAP client requests - it fails if there is a mismatch. If set to true (not recommended for production use), such checks will be bypassed. That will allow you, for example, to use a URL such as localhost during development.
tlsClientParameters.useHttpsURLConnectionDefaultHostnameVerifier	false	This attribute specifies if <code>HttpsURLConnection.getDefaultHostnameVerifier()</code> should be used to create HTTPS connections. If 'true', the 'disableCNCheck' configuration parameter is ignored.
tlsClientParameters.useHttpsURLConnectionDefaultSslSocketFactory	false	Specifies if <code>HttpsURLConnection.getDefaultSSLContextFactory()</code> should be used to create HTTPS connections. If 'true', 'jsseProvider', 'secureSocketProtocol', 'trustManagers', 'keyManagers', 'secureRandom', 'cipherSuites' and 'cipherSuitesFilter' configuration parameters are ignored.
tlsClientParameters.certConstraints.SubjectDNConstraints.combinator		SubjectDN certificate constraints specification as combinator.
tlsClientParameters.certConstraints.SubjectDNConstraints.RegularExpression		SubjectDN certificate constraints specification as regular expression.
tlsClientParameters.certConstraints.IssuerDNConstraints.combinator		IssuerDN certificate constraints specification as combinator.
tlsClientParameters.certConstraints.IssuerDNConstraints.RegularExpression		IssuerDN certificate constraints specification as regular expression.
tlsClientParameters.secureRandomParameters.algorithm	JVM default	algorithm parameter of the SecureRandom specification.
tlsClientParameters.secureRandomParameters.provider	JVM default	provider parameter of the SecureRandom specification.
tlsClientParameters.cipherSuitesFilter.include		filters the supported CipherSuites, list of CipherSuites that will be supported and used if available.
tlsClientParameters.cipherSuitesFilter.exclude		filters the supported CipherSuites, list of CipherSuites that will be excluded.
tlsClientParameters.cipherSuites	default sslContext cipher suites	CipherSuites that will be supported.
tlsClientParameters.trustManagers.provider	JVM default	Provider of the trust manager.
tlsClientParameters.trustManagers.factoryAlgorithm	JVM default	factory algorithm of the trust manager.
tlsClientParameters.trustManagers.keyPassword	JVM default	Key password of the trust manager.

Property	Default	Description
tlsClientParameters. trustManagers.keyStore.type	JVM default	Keystore type of the trust manager.
tlsClientParameters. trustManagers.keyStore.password	JVM default	Keystore password of the trust manager.
tlsClientParameters. trustManagers.keyStore.provider	JVM default	Keystore provider of the trust manager.
tlsClientParameters. trustManagers.keyStore.url	JVM default	Trust Managers URL to hold X509 certificates.
tlsClientParameters. trustManagers.keyStore.file	JVM default	Trust Managers file to hold X509 certificates.
tlsClientParameters. trustManagers.keyStore.resource	JVM default	Trust Managers resource to hold X509 certificates.
tlsClientParameters. keyManagers.provider	JVM default	Provider of the key manager.
tlsClientParameters. keyManagers.factoryAlgorithm	JVM default	factory algorithm of the key manager.
tlsClientParameters. keyManagers.keyPassword	JVM default	Key password of the key manager.
tlsClientParameters. keyManagers.keyStore.type	JVM default	Keystore type of the key manager.
tlsClientParameters. keyManagers.keyStore.password	JVM default	Keystore password of the key manager.
tlsClientParameters. keyManagers.keyStore.provider	JVM default	Keystore provider of the key manager.
tlsClientParameters. keyManagers.keyStore.url	JVM default	Key managers URL to hold X509 certificates.
tlsClientParameters. keyManagers.keyStore.file	JVM default	Key managers file to hold X509 certificates.
tlsClientParameters. keyManagers.keyStore.resource	JVM default	Key managers resource to hold X509 certificates.
authorization.UserName		Specifies the UserName parameter for configuring the basic authentication method that the endpoint uses preemptively.
authorization.Password		Specifies the Password parameter for configuring the basic authentication method that the endpoint uses preemptively.
authorization.Authorization		Corresponds to the authentication specified in the SPNEGO/Kerberos login.conf.
authorization.AuthorizationType		Authorization type: "Basic", "Digest" or "Negotiation"
proxyAuthorization.UserName		Specifies the UserName parameter for configuring basic authentication against outgoing HTTP proxy servers.
proxyAuthorization.Password		Specifies the Password parameter for configuring basic authentication against outgoing HTTP proxy servers.
proxyAuthorization.Authorization		Proxy authorization type: "Basic", "Digest" or "Negotiation"

Property	Default	Description
proxyAuthorization.AuthorizationType		Corresponds to the proxy authentication specified in the SPNEGO/Kerberos login.conf.
client.ConnectionTimeout	30000	Specifies the amount of time, in milliseconds, that the client will attempt to establish a connection before it times out. 0 specifies that the client will continue to attempt to open a connection indefinitely.
client.ReceiveTimeout	60000	Specifies the amount of time, in milliseconds, that the client will wait for a response before it times out. 0 specifies that the client will wait indefinitely.
client.AutoRedirect	false	Specifies if the client will automatically follow a server issued redirection. The default is false.
client.MaxRetransmits	-1	Specifies the maximum number of times a client will retransmit a request to satisfy a redirect. The default of -1 specifies that unlimited retransmissions are allowed.
client.AllowChunking	true	Specifies whether the client will send requests using chunking. The default is true which specifies that the client will use chunking when sending requests. Chunking cannot be used if either <ul style="list-style-type: none"> • http-conf:basicAuthSupplier is configured to provide credentials preemptively or • AutoRedirect is set to true. In both cases the value of AllowChunking is ignored and chunking is disallowed. See note about chunking below.
client.ChunkingThreshold	4000	Specifies the threshold at which CXF will switch from non-chunking to chunking. By default, messages less than 4K are buffered and sent non-chunked. Once this threshold is reached, the message is chunked.
client.Connection	Keep-Alive	Specifies whether a particular connection is to be kept open or closed after each request/response dialog. There are two valid values: <ul style="list-style-type: none"> • Keep-Alive specifies that the client wants to keep its connection open after the initial request/response sequence. If the server honors it, the connection is kept open until the consumer closes it. • close specifies that the connection to the server is closed after each request/response sequence.
client.DecoupledEndpoint		Specifies the URL of a decoupled endpoint for the receipt of responses over a separate server-client connection.

Property	Default	Description
		Warning: You must configure both the client and server to use WS-Addressing for the decoupled endpoint to work.
client.ProxyServer		Specifies the URL of the proxy server through which requests are routed.
client.ProxyServerPort		Specifies the port number of the proxy server through which requests are routed.
client.ProxyServerType	HTTP	Specifies the type of proxy server used to route requests. Valid values are: HTTP (default), SOCKS
client.NonProxyHosts		a (possibly empty) list of hosts which should be connected directly and not through the proxy server; it may contain wild card expressions.

9.2.3. Chunking

There are a number of parameters related to chunking, so we discuss this here in more detail. There are two ways of putting the body of a message into an HTTP stream:

1. Standard scheme: this is used by most browsers. It consists in specifying a Content-Length header in the HTTP headers. This allows the receiver to know how much data is coming and when to stop reading.

The problem with this approach is that the length needs to be pre-determined. The data cannot be streamed as generated as the length needs to be calculated upfront.

Thus, if chunking is turned off, we need to buffer the data in a byte buffer (or temp file if it is too large) so that the Content-Length can be calculated.

2. Chunked scheme: with this mode, the data is sent to the receiver in chunks. Each chunk is preceded by a hexadecimal chunk size. When a chunk size is 0, the receiver knows that all the data has been received.

This mode allows better streaming, as we just need to buffer a small amount, up to 8K by default. When the buffer fills, the chunk is written out.

Some parameters in *Configuration properties* allow you to specify the details of the chunking.

In general, the Chunked scheme will perform better as the streaming can take place directly. However, there are some problems with chunking:

- Many proxy servers don't understand it, especially older proxy servers. Many proxy servers want the Content-Length up front so they can allocate a buffer to store the request before passing it onto the real server.
- Some of the older Web Services stacks also have problems with Chunking - specifically, older versions of .NET.



If you are getting strange errors (generally not SOAP faults, but other HTTP-type errors) when trying to interact with a service, try turning off chunking to see if that helps.



Chapter 10. Logging System

Karaf provides a powerful logging system based on OPS4j Pax Logging. In addition to being a standard OSGi Log service, it supports the following APIs:

- Apache Commons Logging
- SLF4J
- Apache Log4j
- Java Util Logging

Karaf also comes with a set of console commands that can be used to display, view and change the log levels.

10.1. Configuration

10.1.1. Configuration file

The configuration of the logging system uses a [standard Log4j configuration file](#) at the following location:
<RuntimeContainerPath>/etc/org.ops4j.pax.logging.cfg

You can edit this file at runtime and any change will be reloaded and be effective immediately.

10.1.2. Configuring the appenders

The default logging configuration defines three appenders:

- the `stdout` console appender is disabled by default. If you plan to run Karaf in server mode only (i.e. with the locale console disabled), you can turn on this appender on by adding it to the list of configured appenders using the `log4j.rootLogger` property
- the `out` appender is the one enabled by default. It logs events to a number of rotating log files of a fixed size. You can easily change the parameters to control the number of files using `maxBackupIndex` and their size `maxFileSize`.
- the `sift` appender can be used instead to provide a per-bundle log file. The default configuration uses the bundle symbolic name as the file name to log to

10.1.3. Changing the log levels

The default logging configuration sets the logging levels so that the log file will provide enough information to monitor the behavior of the runtime and provide clues about what caused a problem. However, the default configuration will not provide enough information to debug most problems. The most useful logger to change when trying to debug an issue with Karaf is the root logger. You will want to set its logging level to `DEBUG` in the `org.ops4j.pax.logging.cfg` file.

```
log4j.rootLogger=DEBUG, out, osgi:VmLogAppender
...
```

When debugging a problem in Karaf you may want to change the level of logging information that is displayed on the console. The example below shows how to set the root logger to `DEBUG` but limiting the information displayed on the console to `>WARN`.

```
log4j.rootLogger=DEBUG, out, stdout, osgi:VmLogAppender
log4j.appender.stdout.threshold=WARN
...
```

10.2. Console Log Commands

The log scope comes with several commands -- see [Log Scope](#) for a full list.

For example, if you want to debug something, you might want to run the following commands:

```
>log:set DEBUG
... do something ...
>log:display
```

Note that the log levels set using the `log:set` commands are not persistent and will be lost upon restart. To configure those in a persistent way, you should edit the configuration file mentioned above using the `config` commands or directly using a text editor of your choice. The log commands have a separate configuration file: `<RuntimeContainerPath>/etc/org.apache.karaf.log.cfg`

10.3. Advanced Configuration

The logging backend uses `Log4j`, but offers a number of additional features.

10.3.1. Filters

Appender filters can be added using the following syntax:

```
log4j.appender.[appender-name].filter.[filter-name]=[filter-class]
log4j.appender.[appender-name].filter.[filter-name].[option]=[value]
```

Below is a real example:

```
log4j.appender.out.filter.fl=org.apache.log4j.varia.LevelRangeFilter
log4j.appender.out.filter.fl.LevelMax=FATAL
log4j.appender.out.filter.fl.LevelMin=DEBUG
```

10.3.2. Nested appenders

Nested appenders can be added using the following syntax:

```
log4j.appender.[appender-name].appenders= //
    [comma-separated-list-of-appender-names]
```

Below is a real example:

```
log4j.appender.async=org.apache.log4j.AsyncAppender
log4j.appender.async.appenders=jms

log4j.appender.jms=org.apache.log4j.net.JMSAppender
...
```

10.3.3. Error handlers

Error handlers can be added using the following syntax:

```
log4j.appender.[appender-name].errorhandler=[error-handler-class]
log4j.appender.[appender-name].errorhandler.root-ref=[true|false]
log4j.appender.[appender-name].errorhandler.logger-ref=[logger-ref]
log4j.appender.[appender-name].errorhandler.appender-ref=[appender-ref]
```

10.3.4. OSGi specific MDC attributes

Pax-Logging provides the following attributes by default:

- `bundle.id`: the id of the bundle from which the class is loaded
- `bundle.name`: the symbolic-name of the bundle
- `bundle.version`: the version of the bundle

10.3.5. MDC sifting appender

An MDC sifting appender is available to split the log events based on MDC attributes. Below is a configuration example for this appender:

```
log4j.appender.sift=org.apache.log4j.sift.MDCSiftingAppender
log4j.appender.sift.key=bundle.name
log4j.appender.sift.default=karaf
log4j.appender.sift.appender=org.apache.log4j.FileAppender
log4j.appender.sift.appender.layout=org.apache.log4j.PatternLayout
log4j.appender.sift.appender.layout.ConversionPattern=  \\  
    %d{ABSOLUTE} | %-5.5p | %-16.16t | %-32.32c{1} | %-32.32C %4L | %m%n
log4j.appender.sift.appender.file=${karaf.data}/log/${bundle.name}.log
log4j.appender.sift.appender.append=true
```

10.3.6. Enhanced OSGi stack trace renderer

This renderer is configured by default in Karaf and will give additional informations when printing stack traces. For each line of the stack trace, it will display OSGi specific informations related to the class on that line: the bundle id, the bundle symbolic name and the bundle version. This information can greatly help diagnosing problems in some cases. The information is appended at the end of each line in the following format id:name:version as shown below:

```
java.lang.IllegalArgumentException: Command not found: *:foo
  at org.apache.felix.gogo.runtime.shell.Closure.execute
    (Closure.java:225)[21:org.apache.karaf.shell.console:2.1.0]
  at org.apache.felix.gogo.runtime.shell.Closure.executeStatement
    (Closure.java:162)[21:org.apache.karaf.shell.console:2.1.0]
  at org.apache.felix.gogo.runtime.shell.Pipe.run
    (Pipe.java:101)[21:org.apache.karaf.shell.console:2.1.0]
  at org.apache.felix.gogo.runtime.shell.Closure.execute
    (Closure.java:79)[21:org.apache.karaf.shell.console:2.1.0]
  at org.apache.felix.gogo.runtime.shell.CommandSessionImpl.execute
    (CommandSessionImpl.java:71)[21:org.apache.karaf.shell.console:2.1
.0]
  at org.apache.karaf.shell.console.jline.Console.run
    (Console.java:169)[21:org.apache.karaf.shell.console:2.1.0]
  at java.lang.Thread.run(Thread.java:637)[:1.6.0_20]
```

10.3.7. Using your own appenders

If you plan to use your own appenders, you need to create an OSGi bundle and attach it as a fragment to the bundle with a symbolic name of `org.ops4j.pax.logging.pax-logging-service`. This way, the underlying logging system will be able to see and use your appenders. So for example you write a log4j appender:

```
class MyAppender extends AppenderSkeleton {
    ...
}
```

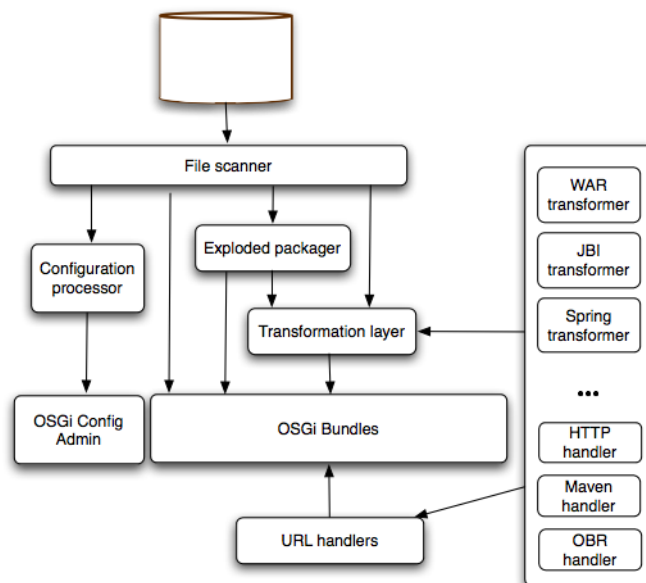
Then you need to package the appender in a jar with a Manifest like this:

```
Manifest:
Bundle-SymbolicName: org.mydomain.myappender
Fragment-Host: org.ops4j.pax.logging.pax-logging-service
...
```

Now you can use the appender in your log4j config file like shown in the config examples above.

Chapter 11. Deployer

The Karaf deployer is used for deploying bundles or groups of bundles (called features) into the Karaf container. The following diagram describes the architecture of the deployer.



11.1. Features deployer

To be able to hot deploy features from the deploy folder, you can simply drop a Feature descriptor on that folder. A bundle will be created and its installation (automatic) will trigger the installation of all features contained in the descriptor. Removing the file from the deploy folder will uninstall the features. If you want to install a single feature, you can do so by writing a feature descriptor like the following:

```
<features>
  <repository>mvn:org.apache.servicemix.nmr/apache-servicemix-nmr/  \\
    1.0.0/xml/features</repository>
  <feature name="nmr-only">
    <feature>nmr</feature>
  </feature>
</features>
```

For more informations about features, see [Provisioning](#).

11.2. Spring deployer

Karaf includes a deployer that is able to deploy plain blueprint or spring-dm configuration files. The deployer will transform on the fly any spring configuration file dropped into the deploy folder into a valid OSGi bundle. The generated OSGi manifest will contain the following headers:

```
Manifest-Version: 2
Bundle-SymbolicName: [name of the file]
Bundle-Version: [version of the file]
Spring-Context: *;publish-context:=false;create-asynchronously:=true
Import-Package: [required packages]
DynamicImport-Package: *
```

The name and version of the file are extracted using a heuristic that will match common patterns. For example `my-config-1.0.1.xml` will lead to `name = my-config` and `version = 1.0.1`. The default imported packages are extracted from the Spring file definition and includes all classes referenced directly. If you need to customize the generated manifest, you can do so by including an xml element in your Spring configuration:

```
<spring:beans ...>
  <manifest>
    Require-Bundle= my-bundle
  </manifest>
  ...
</spring:beans>
```

11.3. Wrap deployer

The wrap deployer allows you to hot deploy non-OSGi jar files ("classical" jar files) from the deploy folder. It's a standard deployer (you don't need to install additional Karaf features):

```
karaf@trun(>) la|grep -i wrap
 35 | Active   | 24 | 4.0.0.SNAPSHOT          | Apache Karaf :: Deployer ::
    Wrap Non OSGi Jar
 92 | Active   | 5  | 2.4.1                   | OPS4J Pax Url - wrap:
```

Karaf wrap deployer looks for jar files in the deploy folder. The jar files is considered as non-OSGi if the MANIFEST doesn't contain the `Bundle-SymbolicName` and `Bundle-Version` attributes, or if there is no MANIFEST at all. The non-OSGi jar file is transformed into an OSGi bundle. The deployer tries to populate the `Bundle-SymbolicName` and `Bundle-Version` extracted from the jar file path. For example, if you simply copy `commons-lang-2.3.jar` (which is not an OSGi bundle) into the deploy folder, you will see:

```
karaf@trun> la|grep -i commons-lang
```

```
[ 41] [Active ] [      ] [ 60] commons-lang (2.3)
```

If you take a look on the commons-lang headers, you can see that the bundle exports all packages with optional resolution and that Bundle-SymbolicName and Bundle-Version have been populated:

```
karaf@trun> bundle:headers 41

commons-lang (41)
-----
Specification-Title = Commons Lang
Tool = Bnd-0.0.357
Specification-Version = 2.3
Specification-Vendor = Apache Software Foundation
Implementation-Version = 2.3
Generated-By-Ops4j-Pax-From = wrap:file:/home/onofreje/workspace/karaf/
  assembly/target/apache-karaf-2.99.99-SNAPSHOT/deploy/commons-lang-2.3
  .jar$ Bundle-SymbolicName=commons-lang&Bundle-Version=2.3
Implementation-Vendor-Id = org.apache
Created-By = 1.6.0_21 (Sun Microsystems Inc.)
Implementation-Title = Commons Lang
Manifest-Version = 1.0
Bnd-LastModified = 1297248243231
X-Compile-Target-JDK = 1.1
Originally-Created-By = 1.3.1_09-85 ("Apple Computer, Inc.")
Ant-Version = Apache Ant 1.6.5
Package = org.apache.commons.lang
X-Compile-Source-JDK = 1.3
Extension-Name = commons-lang
Implementation-Vendor = Apache Software Foundation

Bundle-Name = commons-lang
Bundle-SymbolicName = commons-lang
Bundle-Version = 2.3
Bundle-ManifestVersion = 2

Import-Package =
  org.apache.commons.lang;resolution:=optional,
  org.apache.commons.lang.builder;resolution:=optional,
  org.apache.commons.lang.enum;resolution:=optional,
  org.apache.commons.lang.enums;resolution:=optional,
  org.apache.commons.lang.exception;resolution:=optional,
  org.apache.commons.lang.math;resolution:=optional,
  org.apache.commons.lang.mutable;resolution:=optional,
  org.apache.commons.lang.text;resolution:=optional,
  org.apache.commons.lang.time;resolution:=optional
Export-Package =
  org.apache.commons.lang;uses:="org.apache.commons.lang.builder,
  org.apache.commons.lang.math,org.apache.commons.lang.exception",
  org.apache.commons.lang.builder;
  uses:="org.apache.commons.lang.math,org.apache.commons.lang",
  org.apache.commons.lang.enum;uses:=org.apache.commons.lang,
  org.apache.commons.lang.enums;uses:=org.apache.commons.lang,
  org.apache.commons.lang.exception;uses:=org.apache.commons.lang,
  org.apache.commons.lang.math;uses:=org.apache.commons.lang,
  org.apache.commons.lang.mutable;uses:="org.apache.commons.lang,
  org.apache.commons.lang.math",
  org.apache.commons.lang.text;uses:=org.apache.commons.lang,
  org.apache.commons.lang.time;uses:=org.apache.commons.lang
```

You may set the manifest headers by specifying them as URL parameters. The token '\$' is used to separate the hierarchy path and query parts of the URL, with standard URL syntax for query parameters of key=value pairs joined by '&'. On the command line you must use single quotes around the URL to prevent the shell from attempting variable expansion, and if present, you must escape exclamation marks with a backslash.

```
install -s 'wrap:mvn:jboss/jbossall-client/4.2.3.GA/  \\
  $Bundle-SymbolicName=jbossall-client&Bundle-Version=4.2.3.GA&  \\
  Export-Package=org.jboss.remoting;version="4.2.3.GA",!*'
```

If defined in a features.xml file, you'll need to escape any ampersands and quotes, or use a CDATA tag:

```
<bundle>wrap:mvn:jboss/jbossall-client/4.3.2.GA/  \\
  $Bundle-SymbolicName=jbossall-client  \\
  &Bundle-Version=4.3.2.GA  \\
  &Export-Package=org.jboss.remoting;version="4.3.2.GA",!*
</bundle>
```

11.4. War deployer

See [Web Applications](#) for information on web application (war) deployment.



Chapter 12. Servlet Context

A servlet context defines a set of methods which allows a servlet to communicate with its servlet container. Karaf and CXF provide servlet context custom configuration, for building services. For example, you can deploy into a servlet container, using a servlet transport, `CXFServlet`. The following section explains servlet context configuration in more detail.

Please read the [Servlet Transport](#) on the Apache CXF website for additional information about this servlet context.

12.1. Configuration

First we look at the configuration files. The configuration of the servlet context for the OSGi HTTP Service is specified in a file at the following location: `<RuntimeContainerPath>/etc/org.apache.cxf.osgi.cfg`

The configuration of the port for the OSGi HTTP Service is specified in: `<RuntimeContainerPath>/etc/org.ops4j.pax.web.cfg`

You can edit these files at runtime and any change will be reloaded and be effective immediately.

12.2. Configuring the context

When editing the files, the `org.apache.cxf.osgi.cfg` file specified prefix for the services is:

```
org.apache.cxf.servlet.context=/services
```

The `org.ops4j.pax.web.cfg` file specified port for the services is:

```
org.osgi.service.http.port=8040
```

12.2.1. Relative endpoint address

The CXFServlet uses a relative address for the endpoint rather than a full http address. For example, given an implementation class called `GreeterImpl` with endpoints `greeter` and `greeterRest`, the relative endpoint addresses would be configured as:

```
<jaxws:endpoint id="greeter"
  implementor="org.apache.hello_soap_http.GreeterImpl"
  address="/Greeter1"/>

<jaxrs:server id="greeterRest"
  serviceClass="org.apache.hello_soap_http.GreeterImpl"
  address="/GreeterRest"/>
```

The cumulative result of these changes is that the endpoint address for the servlet will be: `http://{server}:8040/services/Greeter1` and `http://{server}:8040/services/GreeterRest`



Chapter 13. Provisioning

Karaf provides a simple, yet flexible, way to provision applications or "features". Such a mechanism is mainly provided by a set of commands available in the features shell. The provisioning system uses xml "repositories" that define a set of features.

13.1. Example: Deploying a sample feature

In the rest of this chapter, we describe in detail the process involved in provisioning. But as a quick demonstration, we'll run a sample Apache Camel feature already present in the Talend ESB distribution. In the console, run the following commands:

```
feature:repo-add mvn:org.apache.camel/camel-example-osgi/2.10.0/xml/
features
feature:install camel-example-osgi
```

The example installed uses Camel to start a timer every 2 seconds and output a message on the console. These **feature:repo-add** and **feature:install** commands download the Camel features descriptor and install this example. The output is as follows:

```
>>>> SpringDSL set body:  Fri Jan 07 11:59:51 CET 2011
>>>> SpringDSL set body:  Fri Jan 07 11:59:53 CET 2011
>>>> SpringDSL set body:  Fri Jan 07 11:59:55 CET 2011
```

13.1.1. Stopping and uninstalling the sample application

To stop this demo, run the following command:

```
feature:uninstall camel-example-osgi
```

13.2. Repositories

So, first we look at feature repositories. The features are described in a features XML descriptor, a "features repository". This features XML has a schema, so for more information on this schema, see **Features XML Schema** in the **Provisioning** section of the latest version of the *Karaf User Guide* (<http://karaf.apache.org/>). We recommend using this XML schema. It will allow Karaf to validate your repository before parsing. You may also verify your descriptor before adding it to Karaf by simply validation, even from IDE level. Here is an example of such a repository:

```
<features xmlns="http://karaf.apache.org/xmlns/features/v1.0.0">
  <feature name="spring" version="3.0.4.RELEASE">
    <bundle>mvn:org.apache.servicemix.bundles/  \ \
      org.apache.servicemix.bundles.aopalliance/1.0_1</bundle>
    <bundle>mvn:org.springframework/spring-core/3.0.4.RELEASE</bundle>
    <bundle>mvn:org.springframework/spring-beans/3.0.4.RELEASE</bundle>
    <bundle>mvn:org.springframework/spring-aop/3.0.4.RELEASE</bundle>
    <bundle>mvn:org.springframework/spring-context/3.0.4.RELEASE</bundle>
    <bundle>
      mvn:org.springframework/spring-context-support/3.0.4.RELEASE
    </bundle>
  </feature>
</features>
```

A repository includes a list of feature elements, each one representing an application that can be installed. The feature is identified by its name which must be unique amongst all the repositories used and consists of a set of bundles that need to be installed along with some optional dependencies on other features and some optional configurations for the Configuration Admin OSGi service.

References to features define in other repositories are allow and can be achieved by adding a list of repository.

```
<features xmlns="http://karaf.apache.org/xmlns/features/v1.0.0">
  <repository>mvn:org.apache.servicemix.nmr/apache-servicemix-nmr/  \ \
    1.3.0/xml/features</repository>
  <repository>mvn:org.apache.camel.karaf/apache-camel/2.5.0/xml/features
  </repository>
  <repository>mvn:org.apache.karaf/apache-karaf/2.1.2/xml/features
  </repository>
  ...
</features>
```

Be careful when you define them as there is a risk of 'cycling' dependencies. Note: By default, all the features defined in a repository are not installed at the launch of Apache Karaf (see [Service configuration](#) for more information).

13.2.1. Bundles

The main information provided by a feature is the set of OSGi bundles that defines the application. Such bundles are URLs pointing to the actual bundle jars. For example, one would write the following definition:

```
<bundle>http://repo1.maven.org/maven2/org/apache/servicemix/nmr/  \ \
  org.apache.servicemix.nmr.api/1.0.0-m2/  \ \
  org.apache.servicemix.nmr.api-1.0.0-m2.jar</bundle>
```


Doing this will make sure the above bundle is installed while installing the feature. However, Karaf provides several URL handlers, in addition to the usual ones (file, http, etc...). One of these is the Maven URL handler, which allow reusing Maven repositories to point to the bundles.

13.2.1.1. Maven URL Handler

The equivalent of the above bundle would be:

```
<bundle>
  mvn:org.apache.servicemix.nmr/org.apache.servicemix.nmr.api/1.0.0-m2
</bundle>
```

In addition to being less verbose, the Maven URL handlers can also resolve snapshots and can use a local copy of the jar if one is available in your Maven local repository.

The org.ops4j.pax.url.mvn bundle resolves mvn URLs. This flexible tool can be configured through the configuration service. For example, to find the current repositories type: **karaf@trun:/> config:list** and the following will display:

```
Pid:                org.ops4j.pax.url.mvn
BundleLocation:    mvn:org.ops4j.pax.url/pax-url-mvn/0.3.3
Properties:
  service.pid = org.ops4j.pax.url.mvn
  org.ops4j.pax.url.mvn.defaultRepositories = file:/opt/development/  \ \
    karaf/assembly/target/apache-felix-karaf-1.2.0-SNAPSHOT/  \ \
    system@snapshots
  org.ops4j.pax.url.mvn.repositories = http://repol.maven.org/maven2,
    http://svn.apache.org/repos/asf/servicemix/m2-repo
  below = list of repositories and even before the local repository
```

The repositories checked are controlled by these configuration properties. For example, org.ops4j.pax.url.mvn.repositories is a comma separated list of repository URLs specifying those remote repositories to be checked. So, to replace the defaults with a new repository at http://www.example.org/repo on the local machine:

```
karaf@trun:/> config:edit org.ops4j.pax.url.mvn
karaf@trun:/> config:property-list
  service.pid = org.ops4j.pax.url.mvn
  org.ops4j.pax.url.mvn.defaultRepositories = file:/opt/development/karaf/
    assembly/target/apache-felix-karaf-1.2.0-SNAPSHOT/system@snapshots
  org.ops4j.pax.url.mvn.repositories = http://repol.maven.org/maven2,
    http://svn.apache.org/repos/asf/servicemix/m2-repo
  below = list of repositories and even before the local repository
karaf@trun:/> config:property-set org.ops4j.pax.url.mvn.repositories
  http://www.example.org/repo
karaf@trun:/> config:update
```

By default, snapshots are disabled. To enable an URL for snapshots append @snapshots. For example: http://www.example.org/repo@snapshots. Repositories on the local are supported through file:/ URLs.

13.2.1.2. Bundle start-level

By default, the bundles deployed through the feature mechanism will have a start-level equals to the value defined in the configuration file config.properties with the variable karaf.startlevel.bundle=60. This value can be changed using the xml attribute start-level.

```
<feature name='my-project' version='1.0.0'>
  <feature version='2.4.0'>camel-spring</feature>
  <bundle start-level='80'>mvn:com.mycompany.myproject/  \\
    myproject-dao</bundle>
  <bundle start-level='85'>mvn:com.mycompany.myproject/  \\
    myproject-service</bundle>
  <bundle start-level='85'>mvn:com.mycompany.myproject/  \\
    myproject-camel-routing</bundle>
</feature>
```

The advantage to define the start-level of a bundle is that you can deploy all your bundles including those of the project with the 'infrastructure' bundles required (e.g : camel, activemq) at the same time and you will have the guaranty when you use Spring Dynamic Module (to register service through OSGi service layer), Blueprint that by example Spring context will not be created without all the required services installed.

13.2.1.3. Bundle stop/start

The OSGi specification allows to install a bundle without starting it. To use this functionality, simply add the following attribute in your <bundle> definition

```
<feature name='my-project' version='1.0.0'>
  <feature version='2.4.0'>camel-spring</feature>
  <bundle start-level='80' start='false'>mvn:com.mycompany.myproject/  \\
    myproject-dao</bundle>
  <bundle start-level='85' start='false'>mvn:com.mycompany.myproject/  \\
    myproject-service</bundle>
  <bundle start-level='85' start='false'>mvn:com.mycompany.myproject/  \\
    myproject-camel-routing</bundle>
</feature>
```

13.2.1.4. Bundle dependency

A bundle can be flagged as being a dependency. Such information can be used by resolvers to compute the final list of bundles to be installed.

13.2.2. Dependent Features

Dependent features are useful when a given feature depends on another feature to be installed. Such a dependency can be expressed easily in the feature definition:

```
<feature name="jbi">
  <feature>nmr</feature>
  ...
</feature>
```

The effect of such a dependency is to automatically install the required nmr feature when the jbi feature will be installed. A version range can be specified on the feature dependency:

```
<feature name="spring-dm">
```

```
<feature version="[2.5.6,4)">spring</feature>
...
</feature>
```

In such a case, if no matching feature is already installed, the feature with the highest version available in the range will be installed. If a single version is specified, this version will be chosen. If nothing is specified, the highest available will be installed.

13.2.3. Configurations

The configuration section allows to deploy configuration for the OSGi Configuration Admin service along a set of bundles. Here is an example of such a configuration:

```
<config name="com.foo.bar">
  myProperty = myValue
</config>
```

The name attribute of the configuration element will be used as the ManagedService PID for the configuration set in the Configuration Admin service. When using a ManagedServiceFactory, the name attribute is `servicePid_aliasId_`, where `servicePid` is the PID of the ManagedServiceFactory and `aliasId` is a label used to uniquely identify a particular service (an alias to the factory generated service PID). Deploying such a configuration has the same effect than dropping a file named `com.foo.bar.cfg` into the `etc` folder.

The content of the configuration element is set of properties parsed using the [standard java property mechanism](#). Such configuration as usually used with Spring-DM or Blueprint support for the Configuration Admin service, as in the following example, but using plain OSGi APIs will of course work the same way:

```
<bean ...>
  <property name="propertyName" value="\${myProperty}" />
</bean>
<osgix:cm-properties id="cmProps" persistent-id="com.foo.bar">
  <prop key="myProperty">myValue</prop>
</osgix:cm-properties>
<ctx:property-placeholder properties-ref="cmProps" />
```

There may also be cases where you want to make the properties from multiple configuration files available to your bundle context. This is something you may want to do if you have a multi-bundle application where there are application properties used by multiple bundles, and each bundle has its own specific properties. In that case, `<ctx:property-placeholder>` won't work as it was designed to make only one configuration file available to a bundle context. To make more than one configuration file available to your bundle-context you would do something like this:

```
<beans:bean id="myBundleConfigurer" class=
  "org.springframework.beans.factory.config.PropertyPlaceholderConfig">
  <beans:property name="ignoreUnresolvablePlaceholders" value="true"/>
  <beans:property name="propertiesArray">
    <osgix:cm-properties id="myAppProps" persistent-id="myApp.props"/>
    <osgix:cm-properties id="myBundleProps"
      persistent-id="my.bundle.props"/>
  </beans:property>
</beans:bean>
```

In this example, we are using SpringDM with `osgi` as the primary namespace. Instead of using `ctx:context-placeholder` we are using the "PropertyPlaceholderConfig" class. Then we are passing in a beans array and inside of that array is where we set our `osgix:cm-properties` elements. This element "returns" a properties bean.

For more informations about using the Configuration Admin service in Spring-DM, see the [Spring-DM documentation](#).

13.2.4. Configuration files

In certain cases it is needed not only to provide configurations for the configuration admin service but to add additional configuration files e.g. a configuration file for jetty (jetty.xml). It even might be help full to deploy a configuration file instead of a configuration for the config admin service since. To achieve this the attribute `finalname` shows the final destination of the configfile, while the value references the Maven artifact to deploy.

```
<configfile
  finalname="/etc/jetty.xml">mvn:org.apache.karaf/apache-karaf/  \\
  ${project.version}/xml/jettyconfig</configfile>
```

13.2.5. Feature resolver

The resolver attribute on a feature can be set to force the use of a given resolver instead of the default resolution process. A resolver will be use to obtain the list of bundles to actually install for a given feature. The default resolver will simply return the list of bundles provided in the feature description. The obr resolver can be installed and used instead of the standard one. In that case, the resolver will use the OBR service to determine the list of bundles to install (bundles flagged as dependency will only be used as possible candidates to solve various constraints).

13.3. Commands

13.3.1. Repository management

The following commands can be used to manage the list of descriptors known by Karaf. They use URLs pointing to features descriptors. These URLs can use any protocol known to Apache Karaf, the most common ones being http, file and mvn.

- **feature:repo-add**: Add a list of repository URLs to the features service
- **feature:repo-remove**: Remove a list of repository URLs from the features service
- **feature:repo-list**: Display the repository URLs currently associated with the features service.
- **feature:repo-refresh**: Reload the repositories to obtain a fresh list of features

Karaf maintains a persistent list of these repositories so that if you add one URL and restart Karaf, the features will still be available. The **repo-refresh** command is mostly used when developing features descriptors: when changing the descriptor, it can be handy to reload it in the Kernel without having to restart it or to remove then add again this URL.

13.3.2. Features management

Common feature: scope commands used in features management include **feature:install**, **feature:uninstall**, and **feature:list**. See [Feature Scope](#) for more information on these commands.

13.3.3. Examples

To install features using mvn handler:

```
feature:repo-add mvn:org.apache.servicemix.nmr/apache-servicemix-nmr/1.0.0-m2/xml/features
feature:install nmr
```

To use a file handler to deploy a features file (note the path is relative to the Apache Karaf installation directory):

```
feature:repo-add file:base/features/features.xml
```

To deploy bundles from file system without using Maven: As we can use file:// as protocol handler to deploy bundles, you can use the following syntax to deploy bundles when they are located in a directory which is not available using Maven:

```
<features xmlns="http://karaf.apache.org/xmlns/features/v1.0.0">
  <feature name="spring-web" version="2.5.6.SEC01">
    <bundle>file:base/bundles/spring-web-2.5.6.SEC01.jar</bundle>
  </feature>
</features>
```

Note the path again is relative to Apache Karaf installation directory.

13.4. Service configuration

A simple configuration file located in [FELIX:karaf]/etc/org.apache.karaf.features.cfg can be modified to customize the behavior when starting the Kernel for the first time. This configuration file contains two properties:

- featuresBoot: a comma separated list of features to install at startup
- featuresRepositories: a comma separated list of feature repositories to load at startup

This configuration file is of interest if you plan to distribute a Apache Karaf distribution which includes pre-installed features.

Chapter 14. Web Applications

Karaf provides an ability to deploying WAR-based web applications within the Jetty server contained in the Karaf instance.

14.1. Installing WAR support

The following steps will install the "war" feature (support for deploying WAR files with Servlet and JSPs into a Jetty server) into your Karaf instance.

1. List the available features:

```
karaf@trun> feature:list
State      Name
...
[uninstalled] [2.2.9] obr          karaf-2.2.9
[uninstalled] [2.2.9] config       karaf-2.2.9
[uninstalled] [2.2.9] http         karaf-2.2.9
[uninstalled] [2.2.9] war          karaf-2.2.9
[uninstalled] [2.2.9] webconsole  karaf-2.2.9
[installed  ] [2.2.9] ssh          karaf-2.2.9
...
```

2. Install the war feature (and the sub-features it requires):

```
karaf@trun> feature:install war
```

Note: you can use the `-v` or `--verbose` switch to see exactly what Karaf does, and this will show the war, http and jetty features being installed, along with dependent bundles.

3. Verify the features were installed:

```
karaf@trun> feature:list
State      Name
...
[installed ] [2.2.9] http karaf-2.2.9
[installed ] [2.2.9] war  karaf-2.2.9
...
```

- Verify the installed bundles were started

```
karaf@trun> list
START LEVEL 100
ID   State      Level  Name
...
[ 32] [Active ] [ ] [ 60] geronimo-servlet_2.5_spec (1.1.2)
[ 33] [Active ] [ ] [ 60] Apache ServiceMix :: Bundles ::
      jetty (6.1.22.2)
[ 34] [Active ] [ ] [ 60] OPS4J Pax Web - API (1.0.0)
[ 35] [Active ] [ ] [ 60] OPS4J Pax Web - Service SPI (1.0.0)
[ 36] [Active ] [ ] [ 60] OPS4J Pax Web - Runtime (1.0.0)
[ 37] [Active ] [ ] [ 60] OPS4J Pax Web - Jetty (1.0.0)
[ 38] [Active ] [ ] [ 60] OPS4J Pax Web - Jsp Support (1.0.0)
[ 39] [Active ] [ ] [ 60] OPS4J Pax Web - Extender - WAR (1.0.0)
[ 40] [Active ] [ ] [ 60] OPS4J Pax Web - Extender - Whiteboard
      (1.0.0)
[ 42] [Active ] [ ] [ 60] OPS4J Pax Web - FileInstall Deployer
      (1.0.0)
[ 41] [Active ] [ ] [ 60] OPS4J Pax Url - war:, war-i: (1.2.4)
...
```

Note: the version numbers of the packages may be different in your installation.

- The Jetty server should now be listening on `http://localhost:8181/`, but with no published applications available.

```
HTTP ERROR: 404
NOT_FOUND
RequestURI=/
Powered by jetty://
```

14.2. Deploying a WAR to the installed web feature

The following steps will describe how to install a simple WAR file (with JSPs or Servlets) to the just installed web feature.

- To deploy a WAR (JSP or Servlet) to Jetty, update its MANIFEST.MF to include the required OSGi headers as described here: <http://team.ops4j.org/wiki/display/paxweb/WAR+Extender>
- Copy the updated WAR (archive or extracted files) to the deploy directory.

If you want to deploy a sample web application into Karaf, you could use the following command:

```
karaf@trun> bundle:install -s webbundle:http://tomcat.apache.org/
tomcat-5.5-doc/appdev/sample/sample.war?Bundle-SymbolicName=tomcat-
sample&Webapp-Context=/sample
```


Then open your web browser and point to `http://localhost:8181/sample/index.html`.



Chapter 15. Monitoring and Administration using JMX

Apache Karaf provides a large set of MBeans that allow you to fully monitor and administrate Karaf using any JMX client (for example, JConsole provided in the Oracle or IBM JDK). They provide more or less the same actions that you can do using the Karaf shell commands. The list of MBeans available:

Table 15.1. Karaf Management MBeans

MBean (org.apache.karaf:type value)	Description
admin	administrates the child instances
bundles	manipulates the OSGi bundles
config	manipulates the Karaf configuration files (in the /etc folder) and the ConfigAdmin layer
dev	provides information and administration of the OSGi framework
diagnostic	used to create information files (dumps) about Karaf activity
features	manipulate the Karaf features
log	manipulate the logging layer
packages	manipulate the PackageAdmin layer and get information about exported and imported packages
services	to get information about the OSGi services
system	to shutdown the Karaf container itself
web	to get information about the Web bundles (installed with the war feature)
obr	to manipulate the OBR layer (installed with the obr feature)

For MBean registration, the Karaf container exposes a MBeanServer as an OSGi service. This server is created when Karaf boots and can be used to register your own MBeans. The process is really simple as it only requires that

the bean/pojo of the project extend the `javax.management.StandardMBean` class and implement an interface where the name of the class must contain the extension MBean. Below shows an example based on Karaf WebMBean:

1. Create the interface definition.

```
/**
 * Describe the web MBean.
 */
public interface WebMBean {
    ...
}
```

2. Create the implementation. The class must extend the `javax.management.StandardMBean` class and implement the interface.

```
/**
 * Implementation of the web MBean.
 */
public class Web extends StandardMBean implements WebMBean {

    private WebContainerService webContainerService;

    public Web() throws NotCompliantMBeanException {
        super(WebMBean.class);
    }
    ...
}
```

3. Create the Blueprint file. The blueprint file will contain definition of the beans to be instantiated and OSGI services to look up or register. Basically, 2 things should be done : lookup to retrieve the MBeanServer and next registers its Mbean using the MBeanRegister of Karaf.

```
<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0">

    <!-- Reference to the MBean Server -->
    <reference id="mbeanServer"
interface="javax.management.MBeanServer">
        <reference-listener ref="mbeanRegistrar"
            bind-method="registerMBeanServer"
            unbind-method="unregisterMBeanServer"/>
    </reference>

    <!-- Web MBean -->
    <bean id="webMBean"

class="org.apache.karaf.management.mbeans.web.internal.WebMBeanImpl">
        <property name="bundleContext" ref="blueprintBundleContext"/>
    </bean>

    <!-- MBean Registrar -->
    <bean id="mbeanRegistrar"
        class="org.apache.karaf.management.MBeanRegistrar" init-
method="init"
        destroy-method="destroy">
        <property name="mbeans">
            <map>
                <entry value="org.apache.karaf:type=web,name=
${karaf.name}"
                    key-ref="webMBean"/>
            </map>
        </property>
    </bean>
</blueprint>
```

```
        </property>  
    </bean>  
</blueprint>
```




Chapter 16. Installing the Talend Runtime container as a service

The Talend Runtime container is based on Apache Karaf. Karaf Wrapper (for service wrapper) makes it possible to install the Talend Runtime container as a Windows Service. Likewise, the scripts shipped with Karaf also make it very easy to install the Talend Runtime container as a daemon process on Unix systems.

To install Talend Runtime container as a service, you first have to install the wrapper, which is an optional feature, and you can then install the service.

The Wrapper correctly handles "user log outs" under Windows, service dependencies, and the ability to run services which interact with the desktop.

16.1. Supported platforms

The following platforms are supported by the Wrapper:

- AIX
- FreeBSD
- HP-UX, 32-bit and 64-bit versions
- SGI Irix
- Linux kernels 2.2.x, 2.4.x, 2.6.x. Known to work with Debian, Ubuntu, and Red Hat, but should work with any distribution. Currently supported on both 32-bit and 64-bit x86, Itanium, and PPC systems.
- Macintosh OS X
- Sun OS, Solaris 9 and 10. Currently supported on both 32-bit and 64-bit sparc, and x86 systems.

- Windows - Windows 2000, XP, 2003, Vista, 2008 and Windows 7. Currently supported on both 32-bit and 64-bit x86 and Itanium systems. Also known to run on Windows 98 and ME, however due the lack of support for services in the OS, the Wrapper can be run only in console mode.

16.2. Installing the wrapper

First, to install the wrapper, simply:

1. Browse to the `bin` folder of the Talend Runtime container directory, then launch:
 - `trun.bat` in Administrator mode on Windows
 - `trun` as root user on Linux
2. To install the wrapper feature, simply type:
 - **karaf@trun> feature:install wrapper** on Windows.
 - **trun@root> feature:install wrapper** on Linux.

Once installed, wrapper feature will provide **wrapper:install** new command in the `trun`:

```
trun@root> wrapper:install --help
DESCRIPTION
  wrapper:install

  Install the container as a system service in the OS.

SYNTAX
  wrapper:install [options]

OPTIONS
  -d, --display
    The display name of the service.
  --help
    Display this help message

  -s, --start-type
    Mode in which the service is installed. AUTO_START or
    DEMAND_START (Default: AUTO_START)
    (defaults to AUTO_START)

  -n, --name
    The service name that will be used when installing the
    service. (Default: Karaf)
    (defaults to karaf)

  -D, --description
    The description of the service.
    (defaults to null)
```

3. To set up the installation of the service, type in the following command:
 - **karaf@trun> wrapper:install** on Windows.
 - **trun@root> wrapper:install** on Linux.



In Windows you may get the following message when attempting to install the service:

```
wrapper | OpenSCManager failed - Access is denied. (0x5)
```

In this case launch the Command Prompt in Administrator mode. Right-click **Command Prompt** in the **Start** menu and select **Run as administrator**.

For instance, to register Talend Runtime container as a service (depending on the running OS), in automatic start mode, simply type:

- For Windows:

```
karaf@trun> wrapper:install -s AUTO_START -n TALEND-ESB-CONTAINER -d  
Talend-ESB-Container -D "Talend ESB Container Service"
```

- For Linux:

```
trun@root> wrapper:install -s AUTO_START -n TALEND-ESB-CONTAINER -d  
Talend-ESB-Container -D "Talend ESB Container Service"
```

Here is an example of **wrapper:install** command executing on Windows (Note: this also contains instructions on how to remove the service):

```
karaf@trun> wrapper:install -s AUTO_START -n TALEND-ESB-CONTAINER  
-d Talend-ESB-Container -D "Talend ESB Container Service"
```

```
Creating file: C:\work\6.0.1-release\Talend-ESB-V6.0.1\container\bin\  
TALEND-ESB-CONTAINER-wrapper.exe  
Creating file: C:\work\6.0.1-release\Talend-ESB-V6.0.1\container\etc\  
TALEND-ESB-CONTAINER-wrapper.conf  
Creating file: C:\work\6.0.1-release\Talend-ESB-V6.0.1\container\bin\  
TALEND-ESB-CONTAINER-service.bat  
Creating file: C:\work\6.0.1-release\Talend-ESB-V6.0.1\container\lib\  
wrapper.dll  
Creating file: C:\work\6.0.1-release\Talend-ESB-V6.0.1\container\lib\  
karaf-wrapper.jar  
Creating file: C:\work\6.0.1-release\Talend-ESB-V6.0.1\container\lib\  
karaf-wrapper-main.jar
```

Setup complete. You may wish to tweak the JVM properties in the wrapper configuration file:

```
C:\work\6.0.1-release\Talend-ESB-V6.0.1\container\etc\TALEND-  
ESB-CONTAINER-wrapper.conf  
before installing and starting the service.
```

To install the service, run:

```
C:> C:\work\6.0.1-release\Talend-ESB-V6.0.1\container\bin\  
TALEND-ESB-CONTAINER-service.bat install
```

Once installed, to start the service run:

```
C:> net start "TALEND-ESB-CONTAINER"
```

Once running, to stop the service run:

```
C:> net stop "TALEND-ESB-CONTAINER"
```

Once stopped, to remove the installed the service run:

```
C:> C:\work\6.0.1-release\Talend-ESB-V6.0.1\container\bin\  
TALEND-ESB-CONTAINER-service.bat remove
```

Here is an example of **wrapper:install** command executing on Linux:

```
trun@root> wrapper:install -s AUTO_START -n TALEND-ESB-CONTAINER \
    -d Talend-ESB-Container -D "Talend ESB Container Service"
Creating file: /home/onofreje/6.0.1-release/Talend-ESB-V6.0.1/container/
bin/KARAF-wrapper
Creating file: /home/onofreje/6.0.1-release/Talend-ESB-V6.0.1/container/
bin/KARAF-service
Creating file: /home/onofreje/6.0.1-release/Talend-ESB-V6.0.1/container/
etc/KARAF-wrapper.conf
Creating file: /home/onofreje/6.0.1-release/Talend-ESB-V6.0.1/container/
lib/libwrapper.so
Creating file: /home/onofreje/6.0.1-release/Talend-ESB-V6.0.1/container/
lib/karaf-wrapper.jar
Creating file: /home/onofreje/6.0.1-release/Talend-ESB-V6.0.1/container/
lib/karaf-wrapper-main.jar

Setup complete. You may want to tweak the JVM properties in the wrapper
configuration file:
    /home/onofreje/apache-karaf-2.1.3/etc/KARAF-wrapper.conf
before installing and starting the service.
```

16.3. Installing the service

16.3.1. On Windows

1. Open a CMD window in Administrator mode.
2. Browse to the bin folder of the Talend Runtime installation directory, then type in the following command:

```
TALEND-ESB-CONTAINER-service install
```

Here is an example of the installation of Talend Runtime container as a service on Windows:

```
C:\Builds\Talend-Runtime\bin>TALEND-ESB-CONTAINER-service.bat install
wrapper | Talend ESB Container installed.
```

The Talend Runtime service is created and can be viewed by selecting Control Panel > Administrative Tools > Services in the Start menu of Windows.

3. You can then run the **net start "TALEND-CONTAINER"** and **net stop "TALEND-ESB-CONTAINER"** commands to manage the service.

To remove the service, type in the following command in the command window:

```
TALEND-ESB-CONTAINER-service.bat remove
```

16.3.2. On Linux

The way the service is installed depends upon your flavor of Linux:

16.3.2.1. On Redhat/Fedora/CentOS Systems

To install the service:

```
$ ln -s /home/onofreje/6.0.1-release/Talend-ESB-V6.0.1/container/bin/  
TALEND-ESB-CONTAINER-service /etc/init.d/  
$ chkconfig TALEND-ESB-CONTAINER-service --add
```

To start the service when the machine is rebooted:

```
$ chkconfig TALEND-ESB-CONTAINER-service on
```

To disable starting the service when the machine is rebooted:

```
$ chkconfig TALEND-ESB-CONTAINER-service off
```

To start the service:

```
$ service TALEND-ESB-CONTAINER-service start
```

To stop the service:

```
$ service TALEND-ESB-CONTAINER-service stop
```

To uninstall the service:

```
$ chkconfig TALEND-ESB-CONTAINER-service --del  
$ rm /etc/init.d/TALEND-ESB-CONTAINER-service
```

16.3.2.2. On Ubuntu/Debian Systems

To install the service:

```
$ ln -s /home/onofreje/6.0.1-release/Talend-ESB-V6.0.1/container/bin/  
TALEND-ESB-CONTAINER-service /etc/init.d/
```

To start the service when the machine is rebooted:

```
$ update-rc.d TALEND-ESB-CONTAINER-service defaults
```

To disable starting the service when the machine is rebooted:

```
$ update-rc.d -f TALEND-ESB-CONTAINER-service remove
```

To start the service:

```
$ /etc/init.d/TALEND-ESB-CONTAINER-service start
```

To stop the service:

```
$ /etc/init.d/TALEND-ESB-CONTAINER-service stop
```

To uninstall the service:

```
$ rm /etc/init.d/TALEND-ESB-CONTAINER-service
```

16.3.3. Configuration Hints

If you need to append parameters to the "java" invocation (like memory configuration) those can be added in the KARAF-wrapper file using `wrapper.java.additional.n=PARAMETER` where "n" is the number of the additional config (typically you simply look for the last entry and use n+1) and PARAMETER is any JVM parameter you would like to append, such as `-XX:MaxPermSize=1024m`.



Chapter 17. Setting up the Talend ESB Active/Passive configuration

Apache Karaf natively provides a failover mechanism. It uses a kind of master/slave topology where one instance is active and the others are in standby.

Karaf provides failover capability using either a simple lock file or a JDBC locking mechanism.

In both cases, a container-level lock system allows bundles to be preloaded into the slave Karaf instance in order to provide faster failover performance.

So, Talend ESB, based on Apache Karaf, benefits from those mechanisms.

17.1. HA/failover (active/passive)

The Talend ESB failover capability uses a lock system.

This container-level lock system allows bundles installed on the master to be preloaded on the slave, in order to provide faster failover performance.

Two types of lock are supported:

- filesystem lock
- database lock

When a first instance starts, if the lock is available, it takes the lock and become the master.

If a second instance starts, it tries to acquire the lock. As the lock is already hold by the master, the instance becomes a slave, in standby mode (not active). A slave periodically check if the lock has been released or not.

17.1.1. Filesystem lock

The Talend ESB instances share a lock on the filesystem. It means that the filesystem storing the lock has to be accessible to the different instances (using SAN, NFS, and so on).

The configuration of the lock system has to be defined in the `etc/system.properties` file, on each instance (master/slave):

```
karaf.lock=true
karaf.lock.class=org.apache.karaf.main.SimpleFileLock
karaf.lock.dir=<PathToLockFileDirectory>
karaf.lock.delay=10
```

- `karaf.lock` property enables the HA/failover mechanism.
- `karaf.lock.class` property contains the class name providing the lock implementation. Here, the filesystem lock is used.
- `karaf.lock.dir` property contains the location where the lock will be written. All instances have to share the same lock.
- `karaf.lock.delay` property is the interval period (in seconds) to check if the lock has been released or not.

17.1.2. Database lock

It is not always possible and easy to have a shared filesystem between multiple Talend ESB instances.

Instead of sharing a filesystem, Talend ESB supports the sharing of a database.

The master instance holds the lock by locking a table in the database. If the master loses the lock, a waiting slave gains access to the locking table, acquires the lock on the table and starts.

The database lock uses JDBC (Java DataBase Connectivity). To use database locking, you have to:

- copy the JDBC driver in the `lib/ext` folder on each instance. The JDBC driver to use is the one corresponding to the database used for the locking system.
- update `etc/system.properties` file on each instance:

```
karaf.lock=true
karaf.lock.class=org.apache.karaf.main.DefaultJDBCLock
karaf.lock.level=50
karaf.lock.delay=10
karaf.lock.jdbc.url=jdbc:derby://dbserver:1527/sample
karaf.lock.jdbc.driver=org.apache.derby.jdbc.ClientDriver
karaf.lock.jdbc.user=user
karaf.lock.jdbc.password=password
karaf.lock.jdbc.table=KARAF_LOCK
karaf.lock.jdbc.clustername=karaf
karaf.lock.jdbc.timeout=30
```

- `karaf.lock` property enabled the HA/failover mechanism.
- `karaf.lock.class` property contains the class name providing the lock implementation. The `org.apache.karaf.main.DefaultJDBCLock` is the most generic database lock system implementation. Apache Karaf supports lock system for specific databases. For more information, see the

sections: [Lock on Oracle](#), [Lock on Derby](#), [Lock on MySQL](#), [Lock on PostgreSQL](#), and [Lock on Microsoft SQLServer](#).

- `karaf.lock.level` property is the container-level locking. For more information, see the sections: [Lock on Oracle](#), [Lock on Derby](#), [Lock on MySQL](#), [Lock on PostgreSQL](#), and [Lock on Microsoft SQLServer](#).
- `karaf.lock.delay` property is the interval period (in seconds) to check if the lock has been released or not.
- `karaf.lock.jdbc.url` property contains the JDBC URL to the database (derby in this example).
- `karaf.lock.jdbc.driver` property contains the class name of the JDBC driver to use (derby in this example).
- `karaf.lock.jdbc.user` property contains the username to use to connect to the database.
- `karaf.lock.jdbc.password` property contains the password to use to connect to the database.
- `karaf.lock.jdbc.table` property contains the database table to use for the lock.

The Talend ESB will not start if the JDBC driver is not present in the `lib/ext` folder.

The *sample* database will be created automatically if it does not exist.

If the connection to the database is lost, the master instance tries to gracefully shutdown, allowing a slave instance to become the master when the database is back. The former master instance will required a manual restart.

17.1.2.1. Lock on Oracle

Talend ESB supports Oracle database for locking. The lock implementation class name to use is `org.apache.karaf.main.lock.OracleJDBCLock`:

```
karaf.lock=true
karaf.lock.class=org.apache.karaf.main.lock.OracleJDBCLock
karaf.lock.jdbc.url=jdbc:oracle:thin:@hostname:1521:XE
karaf.lock.jdbc.driver=oracle.jdbc.OracleDriver
karaf.lock.jdbc.user=user
karaf.lock.jdbc.password=password
karaf.lock.jdbc.table=KARAF_LOCK
karaf.lock.jdbc.clustername=karaf
karaf.lock.jdbc.timeout=30
```

The Oracle JDBC driver file (`ojdbc*.jar`) has to be copied in the `lib/ext` folder.

The `karaf.lock.jdbc.url` property contains a JDBC URL which requires an active SID. It means that you must manually create the Oracle database instance first before using the lock mechanism.

17.1.2.2. Lock on Derby

Talend ESB supports Apache Derby database for locking. The lock implementation class name to use is `org.apache.karaf.main.lock.DerbyJDBCLock`:

```
karaf.lock=true
karaf.lock.class=org.apache.karaf.main.lock.DerbyJDBCLock
karaf.lock.jdbc.url=jdbc:derby://127.0.0.1:1527/dbname
karaf.lock.jdbc.driver=org.apache.derby.jdbc.ClientDriver
```

```
karaf.lock.jdbc.user=user
karaf.lock.jdbc.password=password
karaf.lock.jdbc.table=KARAF_LOCK
karaf.lock.jdbc.clustertype=karaf
karaf.lock.jdbc.timeout=30
```

The Derby JDBC driver file name has to be copied in the lib/ext folder.

17.1.2.3. Lock on MySQL

Talend ESB supports MySQL database for locking. The lock implementation class name to use is `org.apache.karaf.main.lock.MySQLJDBCLock`:

```
karaf.lock=true
karaf.lock.class=org.apache.karaf.main.lock.MySQLJDBCLock
karaf.lock.jdbc.url=jdbc:mysql://127.0.0.1:3306/dbname
karaf.lock.jdbc.driver=com.mysql.jdbc.Driver
karaf.lock.jdbc.user=user
karaf.lock.jdbc.password=password
karaf.lock.jdbc.table=KARAF_LOCK
karaf.lock.jdbc.clustertype=karaf
karaf.lock.jdbc.timeout=30
```

The MySQL JDBC driver file name has to be copied in lib/ext folder.

17.1.2.4. Lock on PostgreSQL

Talend ESB supports PostgreSQL database for locking. The lock implementation class name to use is `org.apache.karaf.main.lock.PostgreSQLJDBCLock`:

```
karaf.lock=true
karaf.lock.class=org.apache.karaf.main.lock.PostgreSQLJDBCLock
karaf.lock.jdbc.url=jdbc:postgresql://127.0.0.1:1527/dbname
karaf.lock.jdbc.driver=org.postgresql.Driver
karaf.lock.jdbc.user=user
karaf.lock.jdbc.password=password
karaf.lock.jdbc.table=KARAF_LOCK
karaf.lock.jdbc.clustertype=karaf
karaf.lock.jdbc.timeout=0
```

The PostgreSQL JDBC driver file has to be copied in the lib/ext folder.

17.1.2.5. Lock on Microsoft SQLServer

Talend ESB supports Microsoft SQLServer database for locking. The lock implementation class name to use is `org.apache.karaf.main.lock.SQLServerJDBCLock`:

```
karaf.lock=true
karaf.lock.class=org.apache.karaf.main.lock.SQLServerJDBCLock
karaf.lock.level=50
karaf.lock.delay=10
```



```
karaf.lock.jdbc.url=jdbc:jtds:sqlserver://127.0.0.1;databaseName=sample
karaf.lock.jdbc.driver=net.sourceforge.jtds.jdbc.Driver
karaf.lock.jdbc.user=user
karaf.lock.jdbc.password=password
karaf.lock.jdbc.table=KARAF_LOCK
karaf.lock.jdbc.clustername=karaf
karaf.lock.jdbc.timeout=30
```

The JTDS JDBC driver file has to be copied in the `lib/ext` folder.

17.1.3. Container-level locking

Talend ESB supports container-level locking. It allows bundles to be preloaded into the slave instance.

Thanks to that, switching to a slave instance is very fast as the slave instance already contains all required bundles.

The container-level locking is supported in both filesystem and database lock mechanisms.

The container-level locking uses the `karaf.lock.level` property:

```
karaf.lock.level=50
```

The `karaf.lock.level` property tells the Talend ESB instance how far up the boot process to bring the OSGi container.

All bundles with an ID equals or lower to this start level will be started in that Talend ESB instance.

As reminder, the bundles start levels are specified in `etc/startup.properties`, in the `url=level` format.

Level	Behavior
1	A 'cold' standby instance. Core bundles are not loaded into container. Slaves will wait until lock acquired to start server.
<50	A 'hot' standby instance. Core bundles are loaded into the container. Slaves will wait until lock acquired to start user level bundles. The console will be accessible for each slave instance at this level.
>50	This setting is not recommended as user bundles will end up being started.

Using 'hot' standby means that the slave instances are running and bind some ports. So, if you use master and slave instances on the same machine, you have to update the slave configuration to bind the services (ssh, JMX, etc) on different port numbers.



Chapter 18. Troubleshooting Talend ESB

This chapter describes how to fix problems related to JVM memory allocation.

18.1. Memory Allocation Parameters

The Talend ESB start scripts define the JVM memory allocation parameters used for running Talend ESB.

JAVA_MIN_MEM. Determines the start size of the Java heap memory. Higher values may reduce garbage collection and improve performance. Corresponds to the JVM parameter `-Xms`.

JAVA_MAX_MEM. Determines the maximum size of the Java heap memory. Higher values may reduce garbage collection, improve performance and avoid "Out of memory" exceptions. Corresponds to the JVM parameter `-Xmx`.

JAVA_PERM_MEM. Determines the start size of permanent generation. This is a separate heap space that is normally not garbage collected. Java classes are loaded in permanent generation. Using Talend ESB many classes are loaded, specially when using security features and deploying multiple services to the same container. Increase if a "PermGen space" exception is thrown. Corresponds to the JVM parameter `-XX:PermSize`

JAVA_MAX_PERM_MEM. Determines the maximum size of permanent generation. This is a separate heap space that is normally not garbage collected. Increase if a "PermGen space" exception is thrown. Corresponds to the JVM parameter `-XX:MaxPermSize`

Additionally the following parameters are used as default to enable class unloading from permanent generation which reduces the overall permanent generation space usage: `-XX:+CMSClassUnloadingEnabled` and `-XX:+UseConcMarkSweepGC`.



If you set `JAVA_OPTS` yourself from outside or inside the `<RuntimeContainerPath>/bin/setenv` script you overwrite all default settings provided by karaf including the memory settings described above. This could be useful if you need additional options. On Linux/Solaris you can set

DUMP_JAVA_OPTS=true in setenv to find out how karaf would set JAVA_OPTS so you can copy or adapt these settings to fit your requirements.

18.2. On Windows

Adapt the memory allocation parameters in the Karaf start script at the following location: `<RuntimeContainerPath>/bin/setenv.bat`.

In configuring the script variable

```
%Bit_64%
```

you can set different values depending whether your JVM is 64-bit. 64-bit-JVMs need significantly more memory.



There is known issue related 32-bit JVM on 64-bit Windows.

If you set JAVA_OPTS yourself you might get errors during initialization of VM. Because of this problem you would not be able to run tesb container or build examples. In this case, we recommend to set following values for memory parameters: `JAVA_OPTS=-Xmx512m -XX:MaxPermSize=256m`

18.3. On Linux/Solaris

Adapt the memory allocation parameters in the Karaf start script at the following location: `<RuntimeContainerPath>/bin/setenv`.

In configuring the script variable

```
$JAVA_SIXTY_FOUR
```

you can set different values depending whether your JVM is 64-bit. 64-bit-JVMs need significantly more memory.

On Solaris, there is an additional flag `JAVA_64BIT_SOLARIS`, which is set in `<RuntimeContainerPath>/bin/setenv`. The default is `JAVA_64BIT_SOLARIS=true`. This flag is evaluated only for Solaris and determines whether the JVM operates in 64-bit or 32-bit mode. This is necessary, because in Solaris (only), a JVM capable of 64-bit starts normally in 32-bit mode. Set `JAVA_64BIT_SOLARIS=false` if you really want to choose 32-bit mode.

Note that in Redhat Linux Enterprise 5/64 Bit, the memory allocation is different depending on whether you use the container through the wrapper utility or not. If you have installed the service with the wrapper feature, you need to tweak the JVM parameters in the wrapper.conf file, see <http://wrapper.tanukisoftware.com/doc/english/prop-java-maxmemory.html> for details.