

GT 4.0: CSF

GT 4.0: CSF

Table of Contents

1. Key Concepts	1
2. 4.0.0 Release Notes	2
1. Component Overview	2
2. Feature Summary	2
3. Technology Dependencies	2
4. Supported Platforms	3
5. Backward Compatibility Summary	3
6. For More Information	3
3. 4.0.1 Release Notes	4
1. Introduction	4
2. Changes Summary	4
3. Bug Fixes	4
4. Known Problems	4
5. For More Information	4
4. 4.0.2 Release Notes	5
1. Introduction	5
2. Changes Summary	5
3. Bug Fixes	5
4. Known Problems	5
5. For More Information	5
5. 4.0.3 Release Notes	6
1. Introduction	6
2. Changes Summary	6
3. Bug Fixes	6
4. Known Problems	6
5. For More Information	6
6. 4.0.4 Release Notes	7
1. Introduction	7
2. Changes Summary	7
3. Bug Fixes	7
4. Known Problems	7
5. For More Information	7
7. System Administrator's Guide	8
1. Introduction	8
2. Building and installing	8
3. Configuring	8
4. Deploying	10
5. Testing	10
6. Security considerations	10
7. Troubleshooting	10
8. Fact Sheet	12
1. Brief component overview	12
2. Summary of features	12
3. Backward compatibility summary	12
4. Technology dependencies	13
5. Tested platforms	13
6. Associated standards	13
7. For More Information	13

Chapter 1. Key Concepts

See the [key concepts](#)¹ from Execution Management.

¹ <http://www.globus.org/toolkit/docs/4.0/execution/key>

Chapter 2. GT 4.0: Release Notes: Community Scheduler Framework (CSF, contribution)

1. Component Overview

The Community Scheduler Framework 4.0 (CSF4.0) is a WSRF-compliant Grid level meta-scheduling framework built upon the Globus® Toolkit. CSF provides interface and tools for Grid users to submit jobs, create advanced reservations and define different scheduling policies at the Grid level. Using CSF, Grid users are able to access different resource managers, such as LSF, PBS, Condor and SGE, via a single interface.

2. Feature Summary

Features new in release GT 4.0:

- Nothing new since GT 3.2.

Other Supported Features

- Submit, control and monitor jobs at a grid level.
- Create and manage advanced reservations at the grid level.
- Send job and advanced reservation operations to local resource managers.
- A plug in scheduler interface so that site and user specific scheduling policies can be implemented regardless of the underlying resource manager.
- Create queues of jobs, each with separately definable scheduling policies.

Deprecated Features

- None.

3. Technology Dependencies

CSF depends on the following GT components:

- Java WS Core
- Message-Level Security
- Delegation Service
- GRAM
- MDS - internal libraries

CSF depends on the following 3rd party software. The dependencies only exist if there is a requirement to have CSF use an interface (RM Adapter or GRAM) to a local scheduler:

- [LSF](#)¹
- Some other batch scheduler supported by the GRAM interface.

4. Supported Platforms

Tested Platforms for CSF

- Linux (RedHat 9.0)

5. Backward Compatibility Summary

Protocol changes since GT version 3.2:

- The protocol has been changed to be WSRF compliant. There is no backward compatibility with the previous versions.

API changes since GT version 3.2:

- As the factory pattern has been removed from Globus Toolkit 4, the `CreateServiceObject()` of the `QueuingFactoryService`, `JobFactoryService`, and `ReservationFactoryService` are removed. The user doesn't need to create a service instance as before, but instead the user creates a `Resource` by invoking `createResource()` of the `QueuingService`, `JobService`, or `ReservationService`.

Exception changes since GT version 3.2:

- The exceptions from Globus Toolkit 3 that no longer exist have been changed to `RemoteException`.

Schema changes since GT version 3.2:

- See Gram Schema changes to RSL on the [WS GRAM Fact Sheet](#)².

6. For More Information

Click [here](#)³ for more information about this component.

¹ <http://www.platform.com/Products/Platform.LSF.Family/Platform.LSF/>

² [../execution/wsggram/WGRAMFacts.html](#)

³ [index.html](#)

Chapter 3. GT 4.0.1 Release Notes: Community Scheduler Framework (CSF)

1. Introduction

These release notes are for the incremental release 4.0.1. It includes a summary of changes since 4.0.0, bug fixes since 4.0.0 and any known problems that still exist at the time of the 4.0.1 release. This page is in addition to the top-level 4.0.1 release notes at <http://www.globus.org/toolkit/releasenotes/4.0.1>.

For release notes about 4.0 (including feature summary, technology dependencies, etc) go to the [CSF 4.0 Release Notes](#)¹.

2. Changes Summary

The only change that has occurred for the Community Scheduler Framework (CSF) is that it can now submit GRAM jobs to the Queuing service.

3. Bug Fixes

No bug was fixed for the Community Scheduler Framework (CSF).

4. Known Problems

There are no known problems for the Community Scheduler Framework (CSF) at the time of the 4.0.1 release.

5. For More Information

Click [here](#)² for more information about this component.

¹ http://www.globus.org/toolkit/docs/4.0/contributions/csf/CSF_Release_Notes.html

² [index.html](#)

Chapter 4. GT 4.0.2 Release Notes: Community Scheduler Framework (CSF)

1. Introduction

These release notes are for the incremental release 4.0.2. It includes a summary of changes since 4.0.1, bug fixes since 4.0.1 and any known problems that still exist at the time of the 4.0.2 release. This page is in addition to the top-level 4.0.2 release notes at <http://www.globus.org/toolkit/releasenotes/4.0.2>.

For release notes about 4.0 (including feature summary, technology dependencies, etc) go to the [CSF 4.0 Release Notes](#)¹.

2. Changes Summary

There are no changes for the Community Scheduler Framework (CSF) since the 4.0.1 release.

3. Bug Fixes

No bug was fixed for the Community Scheduler Framework (CSF).

4. Known Problems

There are no known problems for the Community Scheduler Framework (CSF) at the time of the 4.0.2 release.

5. For More Information

Click [here](#)² for more information about this component.

¹ http://www.globus.org/toolkit/docs/4.0/contributions/csf/CSF_Release_Notes.html

² [index.html](#)

Chapter 5. GT 4.0.3 Release Notes: Community Scheduler Framework (CSF)

1. Introduction

These release notes are for the incremental release 4.0.3. It includes a summary of changes since 4.0.2, bug fixes since 4.0.2 and any known problems that still exist at the time of the 4.0.3 release. This page is in addition to the top-level 4.0.3 release notes at <http://www.globus.org/toolkit/releasenotes/4.0.3>.

For release notes about 4.0 (including feature summary, technology dependencies, etc) go to the [CSF 4.0 Release Notes](#)¹.

2. Changes Summary

There are no changes for the Community Scheduler Framework (CSF) since the 4.0.2 release.

3. Bug Fixes

No bug has been fixed for the Community Scheduler Framework (CSF) since GT 4.0.2.

4. Known Problems

There are no known problems for the Community Scheduler Framework (CSF) at the time of the 4.0.3 release.

5. For More Information

Click [here](#)² for more information about this component.

¹ http://www.globus.org/toolkit/docs/4.0/contributions/csf/CSF_Release_Notes.html

² [index.html](#)

Chapter 6. GT 4.0.4 Release Notes: Community Scheduler Framework (CSF)

1. Introduction

These release notes are for the incremental release 4.0.4. It includes a summary of changes since 4.0.3, bug fixes since 4.0.3 and any known problems that still exist at the time of the 4.0.4 release. This page is in addition to the top-level 4.0.4 release notes at <http://www.globus.org/toolkit/releasenotes/4.0.4>.

For release notes about 4.0 (including feature summary, technology dependencies, etc) go to the [CSF 4.0 Release Notes](#)¹.

2. Changes Summary

No changes have been made for CSF since 4.0.3.

3. Bug Fixes

No bugs have been fixed for CSF since 4.0.3.

4. Known Problems

There are no known problems with CSF at the time of the 4.0.4 release.

5. For More Information

Click [here](#)² for more information about this component.

¹ http://www.globus.org/toolkit/docs/4.0/contributions/csf/CSF_Release_Notes.html

² [index.html](#)

Chapter 7. GT 4.0: System Administrator's Guide

1. Introduction

This guide contains advanced configuration information for system administrators working with the Community Scheduler Framework (CSF). It provides references to information on procedures typically performed by system administrators, including installation, configuring, deploying, and testing the installation.

Important

This information is in addition to the basic Globus Toolkit prerequisite, overview, installation, security configuration instructions in the [GT 4.0 System Administrator's Guide](#)¹. Read through this guide before continuing!

2. Building and installing

If you are going to submit jobs to a local batch scheduler, make sure that the scheduler is installed according to the instructions provided with the scheduling software. For LSF, you will need to install and run the `gabd` service that is used by CSF's RM Adapter to communicate with LSF clusters. The `gabd` is part of the LSF WebGUI package, or can be acquired by contacting Platform Computing's support department.

Make sure that there is a full installation of the Globus Toolkit as detailed in the [GT 4.0 System Administrator's Guide](#)². If you are going to run jobs through the WS GRAM interface, make sure and follow all the instructions for setting up WS GRAM, including running the test cases.

CSF is distributed in a gzipped tarfile that can be installed using `gpt-build`. Since CSF is built from java source, you must make sure that you have an environment set up properly so that you can build GT 4.0 services.

Set up your GT 4.0 build environment as follows:

- Make sure `JAVA_HOME` and `ANT_HOME` are set to the correct installation locations.
- Add `JAVA_HOME/bin` and `ANT_HOME/bin` to your `PATH` environment variable.
- Set `GLOBUS_LOCATION` to the top level directory of your GT 4.0 installation.
- Source either `globus-user-env.csh` or `globus-user-env.sh` depending on your shell.

In order to install CSF, do the following as the Globus container user (i.e. as the user who owns the Globus installation in `GLOBUS_LOCATION`) in the directory containing the CSF contribution (under `contrib/` in the installer directory):

```
globus$ gpt-build csf-4.0-src.tar.gz
globus$ gpt-postinstall
```

3. Configuring

In order to make CSF operational, you must configure the following files:

¹ ../../admin/docbook/

² ../../admin/docbook/

- `resourcemanager-config.xml`
- `metascheduler-config.xml`

The files contain some commented out settings as examples.

3.1. Configuring the Resource Manager Factory Service

Edit `$GLOBUS_LOCATION/etc/metascheduler/resourcemanager-config.xml` and specify a "cluster" section for each resource manager you will be accessing through the RM Adapter.

You must define the following elements:

- `name`: the name by which this cluster can be referenced.
- `type`: the type of resource manager. Currently only supports type "LSF".
- `host`: the host where the resource manager can be contacted. For LSF, this is the host which is running the `gabd`.
- `port`: the port number where the resource manager can be contacted. For LSF, this is specified in `$LSF_EN-VDIR/ga.conf`.

3.2. Configuring CSF

Edit `$GLOBUS_LOCATION/etc/metascheduler/metascheduler-config.xml` and specify:

- `GISHandle`: the endpoint of the local Index Service.
- `registryHandle`: the endpoint of the local container registry.



Note

Do not use 'localhost' or 127.0.0.1 within the endpoint. Use the actual host IP address or fully qualified host name.

3.3. Configuring the Queuing Service

Edit `$GLOBUS_LOCATION/etc/metascheduler/metascheduler-config.xml` in the section "queueConfig". By default, there are no queues configured in CSF. Any job submission to a queue will fail. Each queue has its own configuration section, in which you can specify:

- `name`: this attribute on the `queueConfig` element specifies the name by which this queue is referenced.
- `plugin`: the name of the class that implements the `com.platform.metascheduler.impl.schedPlugin` interface. The default plugin (`schedPluginDefault`) is always loaded for a queue regardless whether it is defined in the queue configuration. If the plugin doesn't exist or does not implement the `schedPlugin` interface, it will not be loaded. You can specify as many plugins for a queue as you wish.
- `scheduleInterval`: the interval in seconds between scheduling cycles. Its value is an integer between 5 and 600. This parameter is optional, and if not defined will default to 30 seconds. This is a parameter of `schedPluginDefault`.

- `throttle`: this is a parameter for the throttle scheduler plugin (`com.platform.metascheduler.impl.schedThrottle`), and sets the maximum number of jobs that can be dispatched to a back end resource manager in each scheduling cycle. The value is an integer greater than 0.

3.4. Support for multiple GT 4.0 hosting environments

CSF supports the management of jobs and reservations across resource managers hosted in multiple GT 4.0 containers which are part of the same Virtual Organization (VO) (i.e. they trust the same CAs). The multiple container support allows CSF services running in one container to send jobs to the Resource Manager Adapters in another container, etc. In order to support this behaviour, there needs to be a central Index Service for storing information about the different RM Adapters in the VO. This Index Service must be hosted in a container that doesn't host the CSF services.

In order to set up this configuration, the following steps should be taken:

- On the hosts running the CSF services for the VO (1 or more hosts), edit `$GLOBUS_LOCATION/etc/metascheduler/metascheduler-config.xml` and set the endpoint of the central Index Service for the VO in:
 - The `ReservationConfig` section, in the `CommunityGISHandle` element.
 - The `queueConfig` section for each queue you want to use remote RM Adapters, in the `communityGisHandle` parameter.
- On the host running the central Index Service, edit `$GLOBUS_LOCATION/etc/globus_wsrf_mds_index/hierarchy.xml`, and configure the location of the Index Services running in the containers running CSF as downstream elements.

4. Deploying

There are no further deployment instructions other than what is detail in the [GT 4.0 System Administrator's Guide](#)³.

5. Testing

There are currently no tests for Platform CSF.

6. Security considerations

No special security considerations exist at this time.

7. Troubleshooting

7.1. Cannot create a queue

If you can't create a queue, check out if the queue is configured in `$GLOBUS_LOCATION/etc/metascheduler/metascheduler-config.xml`.

³ ../../admin/docbook/

7.2. RM Adapter cannot connect to resource manager

If the container log shows messages that the RM Adapter can't connect to the back end resource manager (i.e. LSF), check that LSF daemons are started and that the LSF `gabd` is also running. Also check that the host name and port number in `$GLOBUS_LOCATION/etc/metascheduler/resourcemanager-config.xml` corresponds to the settings in `ga.conf`.

Chapter 8. GT 4.0 Component Fact Sheet: Community Scheduler Framework (CSF, Contribution)

1. Brief component overview

The Community Scheduler Framework 4.0 (CSF4.0) is a WSRF-compliant Grid level meta-scheduling framework built upon the Globus® Toolkit. CSF provides interface and tools for Grid users to submit jobs, create advanced reservations and define different scheduling policies at the Grid level. Using CSF, Grid users are able to access different resource managers, such as LSF, PBS, Condor and SGE, via a single interface.

2. Summary of features

Features new in release GT 4.0:

- Nothing new since GT 3.2.

Other Supported Features

- Submit, control and monitor jobs at a grid level.
- Create and manage advanced reservations at the grid level.
- Send job and advanced reservation operations to local resource managers.
- A plug in scheduler interface so that site and user specific scheduling policies can be implemented regardless of the underlying resource manager.
- Create queues of jobs, each with separately definable scheduling policies.

Deprecated Features

- None.

3. Backward compatibility summary

Protocol changes since GT version 3.2:

- The protocol has been changed to be WSRF compliant. There is no backward compatibility with the previous versions.

API changes since GT version 3.2:

- As the factory pattern has been removed from Globus Toolkit 4, the `CreateServiceObject()` of the `QueuingFactoryService`, `JobFactoryService`, and `ReservationFactoryService` are removed. The user doesn't need to create a service instance as before, but instead the user creates a Resource by invoking `createResource()` of the `QueuingService`, `JobService`, or `ReservationService`.

Exception changes since GT version 3.2:

- The exceptions from Globus Toolkit 3 that no longer exist have been changed to RemoteException.

Schema changes since GT version 3.2:

- See Gram Schema changes to RSL on the [WS GRAM Fact Sheet](#)¹.

4. Technology dependencies

CSF depends on the following GT components:

- Java WS Core
- Message-Level Security
- Delegation Service
- GRAM
- MDS - internal libraries

CSF depends on the following 3rd party software. The dependencies only exist if there is a requirement to have CSF use an interface (RM Adapter or GRAM) to a local scheduler:

- [LSF](#)²
- Some other batch scheduler supported by the GRAM interface.

5. Tested platforms

Tested Platforms for CSF

- Linux (RedHat 9.0)

6. Associated standards

CSF does not support any standards other than those supported by the Globus Toolkit itself.

7. For More Information

Click [here](#)³ for more information about this component.

¹ ../../execution/wsgram/WSGRAMFacts.html

² <http://www.platform.com/Products/Platform.LSF.Family/Platform.LSF/>

³ index.html