

# CloverETL Server

## Reference Manual



CloverETL

---

# CloverETL Server: Reference Manual

This Reference Manual refers to CloverETL Server 4.9.x release.

Copyright © 2018 Javlin, a.s. All rights reserved.

Javlin

[www.cloveretl.com](http://www.cloveretl.com)

[www.javlininc.com](http://www.javlininc.com)

## Feedback welcome:

If you have any comments or suggestions for this documentation, please send them by email to [support@cloveretl.com](mailto:support@cloveretl.com).

Consider [How to speed up communication with CloverCARE support](#) before contacting the support.

---

---

# Table of Contents

I. CloverETL Server .....	1
1. What is CloverETL Server? .....	2
2. CloverETL Server Architecture .....	4
3. CloverETL Core .....	5
4. CloverETL Worker .....	6
II. Installation .....	8
5. System Requirements .....	9
6. Installing .....	12
Evaluation Server .....	13
Installation .....	13
Production Server .....	16
Apache Tomcat .....	18
Jetty .....	24
IBM WebSphere .....	28
GlassFish / Sun Java System Application Server .....	32
JBoss Application Server .....	36
JBoss Enterprise Application Platform .....	41
Oracle WebLogic Server .....	47
Activation .....	51
CloverETL Server Activation Alternatives .....	53
IBM InfoSphere MDM Plugin Installation .....	54
Installation into Server .....	54
Possible Issues during Installation .....	56
7. Postinstallation Configuration .....	60
Memory Settings .....	60
Maximum Number of Open Files .....	62
Maximum Number of Processes or Threads .....	62
Firewall Exceptions .....	62
8. Optional Installation Steps .....	64
Adding Libraries to the Worker's Classpath .....	64
Worker Support for SMB 2.x and 3.x .....	64
9. Upgrading Server to Newer Version .....	65
III. Configuration .....	67
10. Introduction .....	68
11. Configuration Sources .....	70
Configuration File on Specified Location .....	70
Alternative Configuration Sources .....	70
Priorities of Configuration Sources .....	71
Specifying the Path to the Configuration File .....	72
12. Setup .....	74
Before You Start .....	74
Using Setup .....	75
Configuration File .....	76
License .....	77
Database .....	78
Worker .....	78
Sandboxes .....	79
Encryption .....	79
E-Mail .....	80
LDAP .....	80
Cluster .....	81
13. System Database Configuration .....	83
Embedded Apache Derby .....	85
MySQL .....	86
Creating Database .....	86

---

CloverETL Server Setup .....	86
DB2 .....	87
Creating Database .....	87
CloverETL Server Setup .....	87
DB2 on AS/400 .....	89
Oracle .....	90
Creating Database .....	90
CloverETL Server Setup .....	90
Microsoft SQL Server .....	91
Creating Database .....	91
CloverETL Server Setup .....	91
PostgreSQL .....	93
Creating Database .....	93
CloverETL Server Setup .....	93
JNDI Configuration and Encryption .....	95
JNDI DB Datasource .....	95
JNDI Datasource Troubleshooting .....	96
Encrypted JNDI .....	96
14. List of Configuration Properties .....	106
General Configuration Properties .....	106
Worker - Configuration Properties .....	114
Worker - JNDI Properties .....	117
JDBC Datasources .....	117
JMS Connections .....	118
Worker - SSL Properties .....	120
Job Execution Properties .....	121
List of all properties .....	121
15. Secure Configuration Properties .....	124
16. Logging .....	127
IV. Administration .....	130
17. Monitoring .....	131
Standalone Server Detail .....	131
Cluster Overview .....	137
Node Detail .....	138
Server Logs .....	139
Using the Monitoring .....	140
Restarting the Worker .....	140
Showing Worker's Command Line Arguments .....	140
Suspending the Server .....	140
Resuming the Server .....	140
Displaying List of Threads of the Server Core .....	140
18. Temp Space Management .....	141
Overview .....	141
Management .....	142
19. Secure Parameters .....	147
20. Users and Groups .....	151
LDAP Authentication .....	152
Users .....	156
Groups .....	159
User Lockout .....	171
21. Sandboxes - Server Side Job Files .....	174
Referencing Files from the ETL Graph or Jobflow .....	176
Sandbox Content Security and Permissions .....	177
Sandbox Content .....	178
Download sandbox as ZIP .....	178
Upload ZIP to sandbox .....	179
Download file in ZIP .....	180
Delete Sandbox .....	181

---

---

Job Config Properties .....	183
WebDAV Access to Sandboxes .....	186
WebDAV Clients .....	186
WebDAV Authentication/Authorization .....	186
22. Server Configuration Migration .....	188
Server Configuration Export .....	189
Server Configuration Import .....	190
23. Diagnostics .....	194
24. Troubleshooting Worker .....	196
V. Using Graphs .....	198
25. Graph/Jobflow Parameters .....	199
Parameters by Execution Type .....	200
Adding Another Graph Parameters .....	202
26. Tasks .....	203
Send an Email .....	204
Placeholders .....	205
Execute Shell Command .....	208
Start a Graph .....	211
Start a Jobflow .....	213
Start a Profiler Job .....	214
Abort job .....	215
Archive Records .....	216
Send a JMS Message .....	218
Execute Groovy Code .....	221
27. Manual Task Execution .....	223
28. Scheduling .....	224
Timetable Setting .....	225
Allocations of Scheduled Task on Nodes .....	230
Scheduling the Tasks - Examples .....	231
Start a graph at specific time .....	231
Start a Jobflow once an hour .....	231
Complex Scheduling .....	231
29. Viewing Job Runs - Executions History .....	232
Filtering and ordering .....	232
Tracking .....	234
Log File .....	235
30. Listeners .....	236
Graph Event Listeners .....	238
Graph Events .....	238
Listener .....	239
Tasks .....	239
Use Cases .....	239
Jobflow Event Listeners .....	244
Jobflow Events .....	244
Listener .....	245
Tasks .....	245
JMS Message Listeners .....	246
Universal Event Listeners .....	251
Evaluation Criteria .....	251
File Event Listeners (remote and local) .....	253
Cluster environment .....	254
Supported filesystems and protocols .....	255
Observed file .....	256
File Events .....	257
Check Interval, Task and Use Cases .....	258
Howtos .....	258
Task Failure Listeners .....	261
Task Choice .....	261

---

---

Task Failed E-mail Template .....	262
31. Recommendations for Transformations Developers .....	263
32. Extensibility - CloverETL Engine Plugins .....	264
33. Troubleshooting .....	265
VI. API .....	266
34. Simple HTTP API .....	267
35. JMX mBean .....	278
JMX Configuration .....	278
Operations .....	280
36. SOAP WebService API .....	281
SOAP WS Client .....	281
SOAP WS API Authentication/Authorization .....	281
37. Launch Services .....	282
Launch Services Authentication .....	287
Sending the Data to Launch Service .....	287
Results of the Graph Execution .....	287
38. Data Services .....	289
Overview .....	289
User Interface .....	290
Endpoints .....	290
List of Data Services .....	291
Detail .....	292
Testing and Documentation .....	293
State and History .....	293
Alerts and Notification .....	294
Configuration .....	296
Catalog of Services .....	297
Built-in Data Service Examples .....	298
HTTPS Connectors .....	298
Using Data Services .....	301
Deploying Data Service .....	301
Publishing and Unpublishing Data Service from Sandbox .....	301
Publishing Data Service Examples .....	302
Changing Data Service to Anonymous .....	303
Running Data Service on HTTPS .....	303
Running Data Service on HTTPS on Cluster .....	304
Monitoring Data Service .....	305
Testing Data Service .....	305
Performance Tuning .....	307
Exporting Data Service Configuration .....	308
Importing Data Service Configuration .....	308
Avoiding Premature Marking of Data Service as Failing .....	308
Looking up Particular Data Service .....	308
Resetting State of Failing Data Service Endpoint .....	309
Custom HTTP Headers .....	310
Data Services on Cluster .....	310
VII. Cluster .....	311
39. Clustering Features .....	312
High Availability .....	312
Scalability .....	312
Transformation Requests .....	313
Parallel Data Processing .....	313
Graph Allocation Examples .....	319
Example of Distributed Execution .....	320
Details of the Example Transformation Design .....	320
Scalability of the Example Transformation .....	323
Remote Edges .....	324
40. Cluster Configuration .....	325

---

---

Mandatory Cluster Properties .....	326
Optional Cluster Properties .....	327
Example of 2 Node Cluster Configuration .....	331
Basic 2-nodes Cluster Configuration .....	331
2-nodes Cluster with Proxied Access to Database .....	332
2-nodes Cluster with Load Balancer .....	333
Example of 3 Node Cluster Configuration .....	335
Basic 3-nodes Cluster Configuration .....	335
Jobs Load Balancing Properties .....	337
Running More Clusters .....	338
Cluster Reliability in Unreliable Network Environment .....	339
NodeA Cannot Establish HTTP Connection to NodeB .....	339
NodeA Cannot Establish TCP Connection (Port 7800 by Default) to NodeB .....	340
NodeB is Killed or It Cannot Connect to the Database .....	340
Auto-Resuming in Unreliable Network .....	341
Long-Term Network Malfunction May Cause Jobs to Hang on .....	341
41. Recommendations for Cluster Deployment .....	343
42. Multiple CloverServer Instances on the same Host .....	344
VIII. Security .....	345
43. Security Recommendations for CloverETL Server .....	346

---

# Part I. CloverETL Server

---

# Chapter 1. What is CloverETL Server?

**CloverETL Server** is an enterprise runtime, monitoring and automation platform for the **CloverETL** data integration suite. It is a Java application built to J2EE standards with HTTP and SOAP Web Services APIs providing an additional automation control for integrating the **CloverETL Server** into existing application portfolios and processes.

**CloverETL Server** provides necessary tools to deploy, monitor, schedule, integrate and automate data integration processes in large scale and complex projects. **CloverETL Server** supports a wide range of application servers including Apache Tomcat, Jetty, IBM WebSphere, Sun Glassfish, JBoss AS and Oracle WebLogic.

## CloverETL Server simplifies the process of:

- *Operation* - **CloverETL Server** allows you to [set up \(p. 74\)](#) and [monitor \(p. 131\)](#) the status of the Server and [jobs \(p. 232\)](#) and [notify you via an email \(p. 204\)](#) if the job fails;
- *Automation* - It allows you to efficiently handle jobflow events via [listeners \(p. 236\)](#) and [schedule \(p. 224\)](#) tasks to be triggered as one-time events or repeatedly, as required;
- *Administration* - It helps you manage [users \(p. 156\)](#) and [groups \(p. 159\)](#) and their privileges, create and [configure sandboxes \(p. 174\)](#) and [export the configuration \(p. 188\)](#) to another instance of the Server. Furthermore, the Server provides [API \(p. 266\)](#) and allows you to create it via [Data Services \(p. 289\)](#);
- *Security* - For better control over **CloverETL Server**, you can set up [user lockout \(p. 171\)](#) and [encrypt sensitive data \(p. 147\)](#).

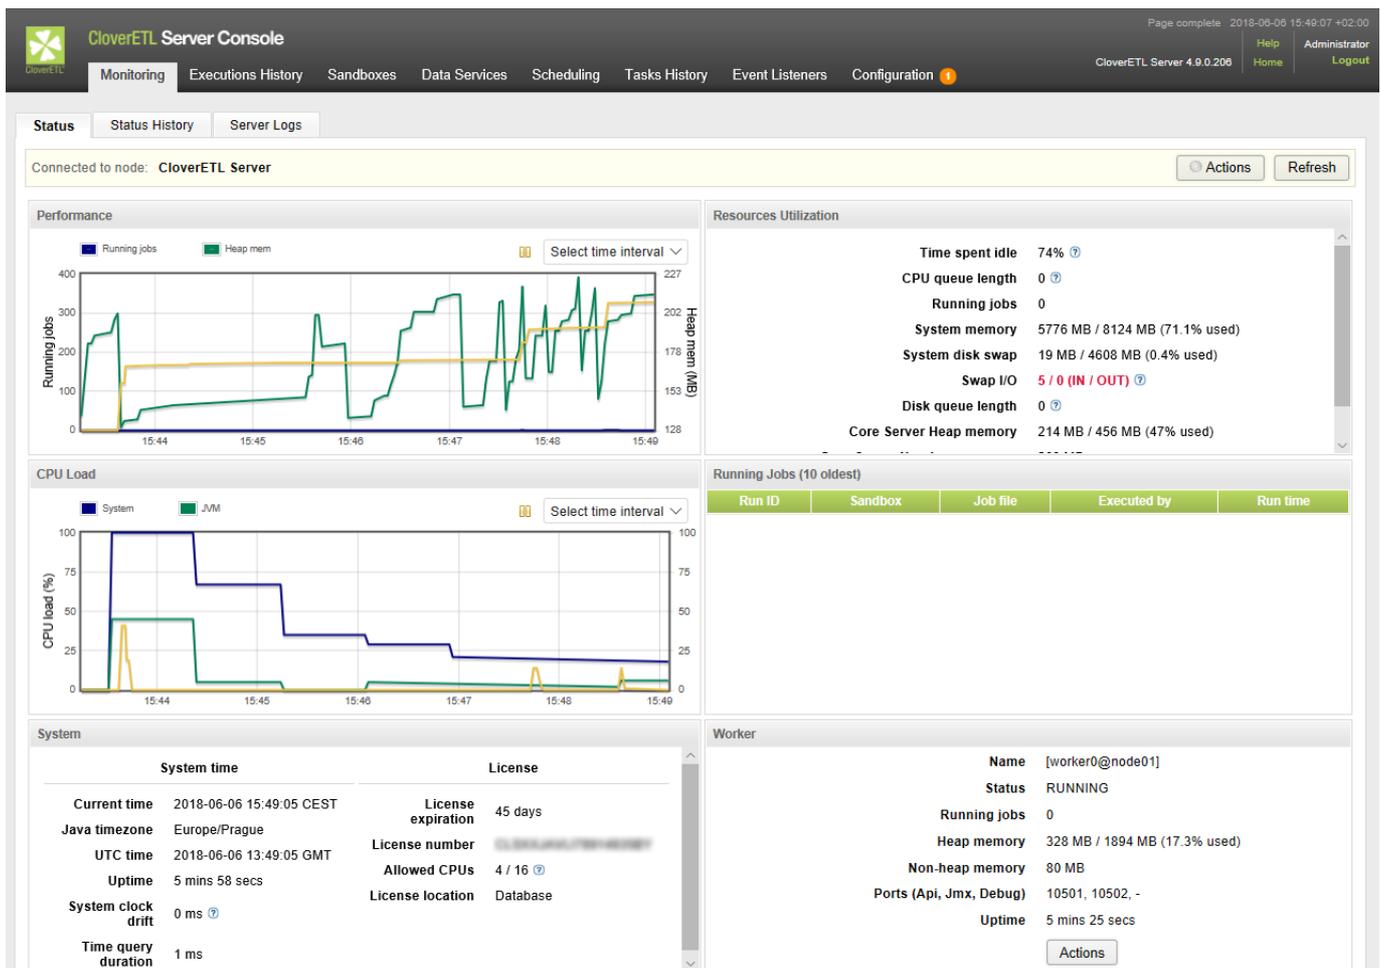


Figure 1.1. CloverETL Server User Interface

To learn more about the architecture of **CloverETL Server**, see [Chapter 2, CloverETL Server Architecture \(p. 4\)](#).

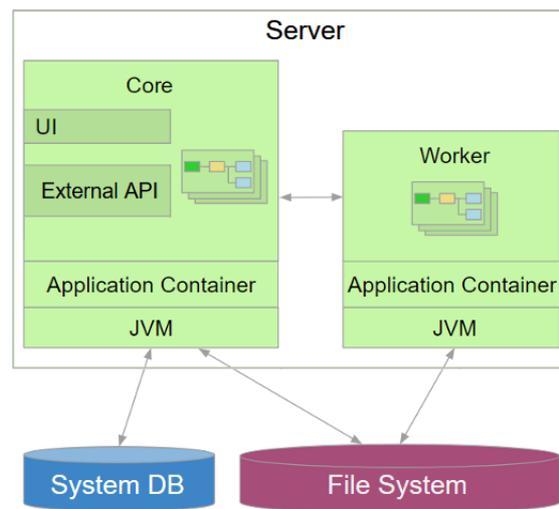
---

## Chapter 2. CloverETL Server Architecture

**CloverETL Server** is a Java application distributed as a web application archive (.war) for an easy deployment on various application servers. It is compatible with Windows and Unix-like operating systems.

Since **CloverETL Server** is a Java application, it requires Java Development Kit (JDK) to run. We **do not recommend** using Java Runtime Environment (JRE) only, since compilation of some CloverETL Server's transformations requires JDK to function properly.

The Server requires some space on the file system to store persistent data (transformation graphs) and temporary data (temporary files, debugging data, etc.). It also requires an external relational database to save run records, permission, users' data, etc.



*Figure 2.1. System Architecture*

The **CloverETL Server** architecture consists of **Core** and **Worker**.

### CloverETL Core

**CloverETL Server's** Core manages [users and groups](#) (p. 151) checks permissions, [schedules](#) (p. 224) execution and provides [management](#) (p. 74) and [monitoring](#) (p. 131) UI. It provides APIs for other applications: [Data Service API](#) (p. 289) [HTTP API](#) (p. 267) and [Web Service API](#) (p. 281). For more information, see Chapter 3, [CloverETL Core](#) (p. 5).

### CloverETL Worker

Worker executes jobs: graphs, jobflows and profiler jobs. The purpose of Worker is to provide a sandboxed execution environment. Worker runs as a separate process. For more information, see Chapter 4, [CloverETL Worker](#) (p. 6).

### Dependencies on External Services

The Server requires a database to store its configuration, user accounts, executions history, etc. It comes bundled with an Apache Derby database to ease the evaluation. To use CloverETL Server in production environment, a [relational database](#) (p. 10) is needed.

The Server needs a connection to an [SMTP server](#) (p. 80) to be able to send you notification emails.

Users and groups' data can be stored in the database or read from an [LDAP](#) (p. 152) server.

### Server Core - Worker Communication

Server Core receives Worker's stdout and stderr. The processes communicate via TCP connections.

---

## Chapter 3. CloverETL Core

**CloverETL Core** is a central point of **CloverETL Server**. It manages and monitors **Workers** that run the jobs. CloverETL Core launches scheduled jobs, checks user permissions, sends notifications and forwards communication between **Designer** and **Worker**.

**CloverETL Core** is the visible part of the **Server**. It has a web-based user interface that allows you to schedule jobs, monitor them, set up user permissions, listeners, notifications, etc.

**CloverETL Core** connects to the system database and stores its configuration and service records in it. The system database is required. If it is configured, the Core connects to an SMTP server to send notification emails or to an LDAP server to authenticate users against an existing LDAP database.

---

## Chapter 4. CloverETL Worker

Worker is a standalone JVM running separately from the main Server. This provides an isolation of the Server Core from executed jobs (e.g. graphs, jobflows, etc.). Therefore, an issue caused by a job in Worker will not affect the Server Core.

Worker does not require any additional installation - it is started and managed by the Server. Worker runs on the same host as the Server Core, i.e. it is not used for parallel or distributed processes. In case of a cluster, **each** cluster node has its **own** Worker.

Worker is a relatively light-weight and simple executor of jobs. It handles job execution requests from the Server Core, but does not perform any high-level job management or scheduling. It communicates with the Server Core via an API for more complex activities, e.g. to request execution of other jobs, check file permissions, etc.

### Configuration

#### General configuration

Worker is started by the Server Core as a standalone JVM process. The default configuration of Worker can be changed in the [Setup](#) (p. 78):

- Heap memory limits
- Port ranges
- Additional command line arguments (e.g. to tweak garbage collector settings)

The settings are stored in the usual Server configuration file. Worker is configured via special [configuration properties](#) (p. 114).

A full [command line](#) (p. 140) of Worker is available in the Monitoring section.

#### Cluster specific configuration

Cluster should use a single [portRange](#) (p. 114): all nodes should have identical value of portRange. That is the preferred configuration, although different ranges for individual nodes are possible.

### Management

The Server manages the runtime of Worker, i.e. it is able to start, stop, restart Worker, etc. Users don't need to manually install and start Worker.

Status of Worker and actions are available in the [Monitoring Worker](#) (p. 133) section.

In case of problems with Worker, see Chapter 24, [Troubleshooting Worker](#) (p. 196).

### Job Execution

By default, all jobs are executed in Worker; yet the Server Core still keeps the capability to execute jobs. It is possible to set specific jobs or whole sandboxes to run in the Server Core via the worker\_execution (p. 184) property on the job or sandbox. It is also possible to disable Worker completely (p. 114), in which case all jobs will be executed in the Server Core.

Executing jobs in the Server Core should be an exception. To see where the job was executed, look in the run details in **Execution History** - in the Executor field. Jobs started in Worker also log a message in their log, e.g. `Job is executed on Worker:[worker0@node01:10500]`.

### Job Configuration

The following areas of Worker configuration affect job execution:

- JNDI

Graphs running in Worker cannot use JNDI as defined in the application container of the Server Core, because Worker is a separate JVM process. Worker provides its own [JNDI configuration](#) (p. 117).

- Classpath

The classpath is not shared between the Server Core and Worker. If you need to add a library to the Worker classpath, e.g. a JDBC driver, follow the instructions in [Adding Libraries to the Worker's Classpath](#) (p. 64).

---

## **Part II. Installation**

---

---

# Chapter 5. System Requirements

---

## Hardware Requirements

The following table shows hardware requirements for both Basic and Corporate **CloverETL Server** edition and for running **CloverETL Server** in a cluster.

Table 5.1. Hardware requirements of CloverETL Server

	Basic Edition	Corporate Edition	Cluster
RAM	4 GB (recommended 16 GB)	8 GB (recommended 64 GB)	8 GB (recommended 64 GB)
Processors	up to 4 cores	16 cores	8 cores <sup>a</sup>
Disk space (installation)	1 GB	1 GB	1 GB
Disk space (tempspace)	> 25 GB <sup>b</sup>	> 25 GB <sup>b</sup>	> 25 GB <sup>b</sup>
Disk space (data)	> 50 GB <sup>b</sup>	> 50 GB <sup>b</sup>	> 50 GB <sup>b</sup>
Disk space (shared) <sup>c</sup>	-	-	> 50 GB <sup>b</sup>

<sup>a</sup> This may vary depending on total number of nodes and cores in license.

<sup>b</sup> Minimum value, the disk space depends on data.

<sup>c</sup> Disk space for shared sandboxes is required only for CloverETL Cluster.

---

## Software Requirements

### Operating system

CloverETL Server is compatible with Windows and Unix-based systems, as well as with other systems supporting Java (Mac OS X, IBM System, etc.).

### Java Virtual Machine

- Oracle JDK 7/8 64 bit
- IBM SDK 7 (for IBM WebSphere only)

Since it is a Java application, it requires Java Development Kit (JDK) to run. We **do not recommend** using Java Runtime Environment (JRE) only, since compilation of some CloverETL Server's transformations requires JDK to function properly.

### Application Server

- [Apache Tomcat 7, 8, 8.5 or 9](#) (p. 18)
- [Jetty 9.1](#) (p. 24)
- [IBM WebSphere 8.5](#) (p. 28)
- [Glassfish 3.1](#) (p. 32)
- [JBoss 6 or 7](#) (p. 36)
- [Oracle WebLogic 11g \(10.3.6\) or 12c \(12.1.2 or 12.1.3\) 32/64 bit](#) (p. 47)

Table 5.2. CloverETL Server Compatibility Matrix

Application Server	CloverETL 3.3	CloverETL 3.4 and 3.5	CloverETL 4.0	CloverETL 4.1 - 4.7		CloverETL 4.8		CloverETL 4.9 and newer	
	Java 6 and 7	Java 6 and 7	Java 7	Java 7	Java 8	Java 7	Java 8	Java 7	Java 8
Tomcat 6 <sup>a</sup>	✓	✓	✓	✓	✓	✗	✗	✗	✗
Tomcat 7	✗	✗	✓	✓	✓	✓	✓	✓	✓
Tomcat 8	✗	✗	✗	✓	✓	✓	✓	✓	✓
Tomcat 8.5	✗	✗	✗	✗	✗	✓ <sup>b</sup>	✓ <sup>b</sup>	✓	✓
Tomcat 9	✗	✗	✗	✗	✗	✗	✓	✗	✓
Pivotal tc Server Standard (3.1.9, Tomcat 8)	✗	✗	✗	✗	✓	✗	✓	✗	✓
Pivotal tc Server Standard (3.2.9, Tomcat 8.5)	✗	✗	✗	✗	✗	✗	✓ <sup>b</sup>	✗	✓
Jetty 6	✓	✓	✗	✗	✗	✗	✗	✗	✗
Jetty 9	✗	✗	✓	✓	✓	✓	✓	✓	✓
WebLogic 11g (10.3.6)	✓	✓	✓	✓	✗	✓	✗	✓	✗
WebLogic 12c (12.1.2)	✗	✓	✓	✓	✗	✓	✗	✓	✗
WebLogic 12c (12.1.3)	✗	✗	✗	✓	✓	✓	✓	✓	✓
JBoss AS 5	✓	✓	✗	✗	✗	✗	✗	✗	✗
JBoss AS 6	✓	✓	✓	✓	✗	✓	✗	✓	✗
JBoss AS 7	✗	✗	✓	✓ <sup>c</sup>	✓ <sup>d</sup>	✓ <sup>c</sup>	✓ <sup>d</sup>	✓ <sup>c</sup>	✓ <sup>d</sup>
Glassfish 2	✓	✓	✗	✗	✗	✗	✗	✗	✗
Glassfish 3	✗	✗	✓	✓	✗	✓	✗	✓	✗
WebSphere 7	✓	✓	✗	✗	✗	✗	✗	✗	✗
WebSphere 8.5	✗	✗	✓	✓	✗	✓	✗	✓	✗

<sup>a</sup> Please note that support for Apache Tomcat 6.0.x has ended on 31 December 2016. For more information, see [End of life for Apache Tomcat 6.0.x](#).

<sup>b</sup> Since 4.8.2

<sup>c</sup> EAP 6.2

<sup>d</sup> EAP 6.4



## Note

We support Java 8 on particular supported application server only if the application server itself officially supports Java 8.

## Database servers

We support the following database servers. The officially supported versions, we are testing against, are in parentheses.

- [MySQL \(5.6.12\)](#) (p. 86)
- [DB2 \(10.5.1\)](#) (p. 87)
- [Oracle \(11.2.0.2.0\)](#) (p. 90)
- [MS SQL Server 2014 \(12.0.5557.0\)](#) (p. 91)
- [PostgreSQL \(9.2.4\)](#) (p. 93)

The evaluation version uses an embedded Apache Derby database for evaluation purposes. We **do not support** using the Derby database in production environment. Instead you can choose one of several supported database servers.

---

## Chapter 6. Installing

[Evaluation Server](#) (p. 12)

[Production Server](#) (p. 12)

This chapter describes two different Server installations - [Evaluation Server](#) (p. 13) and [Production Server](#) (p. 16) - and provide instructions on installing the CloverETL Server License.

### Evaluation Server

The [Evaluation Server](#) (p. 13) consists of **CloverETL Server** bundled with the Tomcat application container. The Server performs basic configuration during the first startup and requires no additional database server. This option is **recommended only for basic evaluation** of **CloverETL Server's** functions.

However with further configuration, it is possible to evaluate other CloverETL Server features and even make the Evaluation Server **ready for production environment**. This process requires a connection to an external, dedicated database and subsequent configuration of services (e.g. SMTP, LDAP, etc.).



#### Important

The Apache Derby DB, bundled with the Evaluation Server, is **not** supported for production environment. Please use one of the supported database systems.

### Production Server

In case of [Production Server](#) (p. 16), the **CloverETL Server** is installed on one of the several compatible application containers. This process requires additional configuration (e.g. memory allocation, database connection, etc.) but allows you to choose an application container and external database according to your preference.

## Evaluation Server

The default installation of **CloverETL Server** uses embedded Apache Derby DB; therefore, it does not require any external database server. Furthermore, it does **not require** any subsequent **configuration**, as **CloverETL Server** configures itself during the first startup. Database tables and some necessary records are automatically created on the first startup with an empty database.

By performing a subsequent configuration, you can evaluate other **CloverETL Server** features (e.g. sending emails, LDAP authentication, clustering, etc.). This way, you can also prepare the Evaluation Server for production environment. However, note that the embedded Apache Derby database is **not** supported for production environment. Therefore, before the subsequent configuration, choose one of the supported external dedicated databases.

If the **CloverETL Server** must be evaluated on application containers other than Tomcat, or you prefer a different database system, proceed with a common installation of [Production Server](#) (p. 16)



### Note

Default login credentials for CloverETL Server Console are:

Username: **clover**

Password: **clover**

## Installation

1. Make sure you have a compatible Java version:



### Important

**CloverETL Server 4.1** and higher requires Oracle JDK or JRE v. **1.7.x** or **higher**. We recommend JDK 1.8.x.

- You can check your installed Java version by typing the following command to the command prompt or terminal:

```
java -version
```

- Alternatively, for **macOS** and **Windows** platforms, see [How to find Java version in Windows or Mac](#).

2. Download and extract the CloverETL Evaluation Server.

- Go to [CloverETL User Login Page](#).
- Using your credentials, log into your account, navigate to the download section and download the CloverETL Evaluation Server Bundle.
- Extract the `.zip` archive. (For example, the name of an archive containing CloverETL Server v 4.9.0 bundled with Tomcat v 8.0.47 will be `CloverETLServer.4.9.0.Tomcat-8.0.47.zip`.)



### Note

It is recommended to place the extracted content on a path that does not contain space character(s).

`C:\Program Files` or `/home/user/some dir` ❌

C:\Users\Username or /home/user/some\_dir ✔

### 3. Set the JAVA\_HOME or JRE\_HOME Environment Variables:

- **Unix-like systems:**

- Using a text editor, open the `setenv.sh` file located in the `[Tomcat_home]/bin/` directory.
- Define the path at the beginning of the file (the path may differ):

```
export JAVA_HOME=/opt/jdk1.8.0_161
```

```
export JAVA_HOME=/opt/jdk1.8.0_121
if [ -n "$JRE_HOME" -a -n "$JAVA_HOME" ]; then
  echo "Using JAVA_HOME instead of JRE_HOME"
  unset JRE_HOME
fi
```

Figure 6.1. `setenv.sh` edited in Linux.

- **Windows system:**

- Using a text editor, open the `setenv.bat` file located in the `[Tomcat_home]\bin\` directory.
- Define the path on the second line of the file (the path may differ):

```
set "JAVA_HOME=C:\java\jdk1.8.0"
```

```
1 @echo off
2 set "JAVA_HOME=C:\java\jdk1.8.0"
3 IF DEFINED JAVA_HOME IF DEFINED JRE_HOME (
4   echo "Using JAVA_HOME instead of JRE_HOME"
5   set JRE_HOME=%JAVA_HOME%
```

Figure 6.2. `setenv.bat` edited in Windows.

### 4. Run Tomcat.

- **Unix-like systems:**

Run the `[Tomcat_home]/bin/startup.sh` file.

- **Windows system:**

Run the `[Tomcat_home]\bin\startup.bat` file.

### 5. Check whether CloverETL Server is running.

- Open a new tab in your browser and type <http://localhost:8083/clover/> in the address bar.
- Use the **default administrator credentials** to access the web GUI: username: **clover**, password: **clover**.



## Note

If you access the web GUI of the CloverETL Server before the product activation, you will be asked to install the CloverETL Server license key.

➔ **Continue with:** [Activation](#) (p. 51)



### Tip

To safely stop the server, run the `[Tomcat_home]/bin/shutdown.sh` or `[Tomcat_home]\bin\shutdown.bat` file for Unix-like or Windows system respectively.

6. **CloverETL Server** is now installed and prepared for basic evaluation. There are couple of sandboxes with various demo transformations installed.

## Production Server

This section describes in detail the installation of CloverETL Server on various application containers and its subsequent configuration required for production environment. For simple evaluation of CloverETL Server features, use Evaluation Server (p. 13) (note that CloverETL Evaluation Server can also be configured for production use).

**CloverETL Server** for production environment is shipped as a *Web application archive* (WAR file) and uses an external, dedicated database, so standard methods for deploying a web application on your application server may be used. However, each application server has specific behavior and features. Detailed information about their installation and configuration can be found in the following chapters.

### List of Suitable Containers

- [Apache Tomcat](#) (p. 18)
- [Jetty](#) (p. 24)
- [IBM WebSphere](#) (p. 28)
- [GlassFish / Sun Java System Application Server](#) (p. 32)
- [JBoss Application Server](#) (p. 36)
- [JBoss Enterprise Application Platform](#) (p. 41)
- [Oracle WebLogic Server](#) (p. 47)

In case of problems during the installation see [Possible Issues during Installation](#) (p. 56).



### Important

**CloverETL Server 4.1** and higher requires Oracle JDK or JRE v. **1.7.x** or **higher**. We recommend JDK 1.8.x.

## Installation and Configuration Procedure

To create a fully working instance of Production CloverETL Server, you should:

### Install an application server

**CloverETL Server** is compatible with several application containers. Following subsections offer detailed instructions on installation of the respective application servers and their subsequent configuration.

### Set up limits on a number of opened files, memory allocation and firewall exceptions

**CloverETL Server's** graph transformations and evaluations may require more memory than the default limit set in the database as well as higher number of simultaneously opened files. Moreover, some components require firewall exceptions to be set. These instructions provide recommendation on adjusting both the [Memory Settings](#) (p. 60) and the [Maximum Number of Open Files](#)(p. 62) as well as [Firewall Exceptions](#) (p. 62).

### Install CloverETL Server into the application server

**CloverETL Server** is provided as a web archive ( .war) file for an easy deployment.

### Create a database dedicated to CloverETL Server

Unlike the Evaluation Server, the Production Server requires that you have created a dedicated database for **CloverETL Server**. In the configuration phase of this manual, you will be guided to Chapter 13, [System Database Configuration](#) (p. 83) with instructions on how to properly configure the properties file of various databases.

### Set up a connection to the database

The **CloverETL Server** Console GUI lets you configure a number of items including database connection, license file, etc. Optionally, you can set up password encryption in configuration files for higher security. For details, see Chapter 12, [Setup](#) (p. 74).

### **Install a license**

To be able to execute graphs, you need to install a valid license. There are three options for [Activation](#) (p. 51).

### **Perform additional Server configuration**

#### **Set up a master password for secure parameters**

When handling sensitive information (e.g. passwords), it is advised to define secure graph parameters. This action requires a master password (see Chapter 19, [Secure Parameters](#) (p. 147)).

#### **Set up SMTP server connection**

**CloverETL Server** lets you configure an SMTP connection for reporting events on the Server via emails.

#### **Configure temp space**

**CloverETL Server** works with temporary directories and files. To ensure components work correctly, you should configure the Temp space location on the file system. For details, see Chapter 18, [Temp Space Management](#) (p. 141).

#### **Configure sandboxes**

Lastly, you should set the content security and user's permissions for sandboxes. For details and instructions, see Chapter 21, [Sandboxes - Server Side Job Files](#) (p. 174).

## Apache Tomcat

[Installation of Apache Tomcat](#) (p. 18)

[Apache Tomcat as a Windows Service](#) (p. 19)

[Apache Tomcat on IBM AS/400 \(iSeries\)](#) (p. 21)

[Installation of CloverETL Server](#) (p. 21)

[Configuration of CloverETL Server on Apache Tomcat](#) (p. 22)



### Important

See [Application Server](#) (p. 9) in system requirements for currently supported **Apache Tomcat** versions and required **Java** versions.

If you encounter any problems during the installation, the [Possible Issues during Installation](#) (p. 56) section may provide a solution.

Please note that since 31 December, 2016, the Apache Tomcat team **has ended Apache Tomcat 6.0.x support**. For more information, see [End of life for Apache Tomcat 6.0.x](#).

## Installation of Apache Tomcat

1. Download the binary distribution: [Tomcat 7](#), [Tomcat 8](#), [Tomcat 8.5](#) or [Tomcat 9](#).

**CloverETL Server** is developed and tested with Apache Tomcat 7.0.x, 8.0.x, 8.5.x or 9.0.x containers. Running the Server with other versions may result in unpredictable behavior.

2. Extract the downloaded archive (zip or tar.gz).
3. Set up JAVA\_HOME to point to the correct Java version:

- **Unix-like systems:**

Setup the path in `/etc/profile` or `/etc/bash.bashrc`:

```
export JAVA_HOME=/path/to/JDK
```

(for example: `export JAVA_HOME=/usr/jdk1.8.0_161`)

- **Windows system:**

Under **System Variables** in **Advanced System Settings**, create a new variable named `JAVA_HOME`. The value should contain the path to the JDK installation directory (e.g. `C:\Java\jdk1.8.0_161`).

4. Run Tomcat:

- **Unix-like systems:**

Run the `[Tomcat_home]/bin/startup.sh` file.

- **Windows system:**

Run the `[Tomcat_home]\bin\startup.bat` file.

5. Check whether Tomcat is running.

- Open a new tab in your browser and type <http://localhost:8080/> in the address bar.

If the Apache Tomcat Information page appears (see below), the installation was successful:

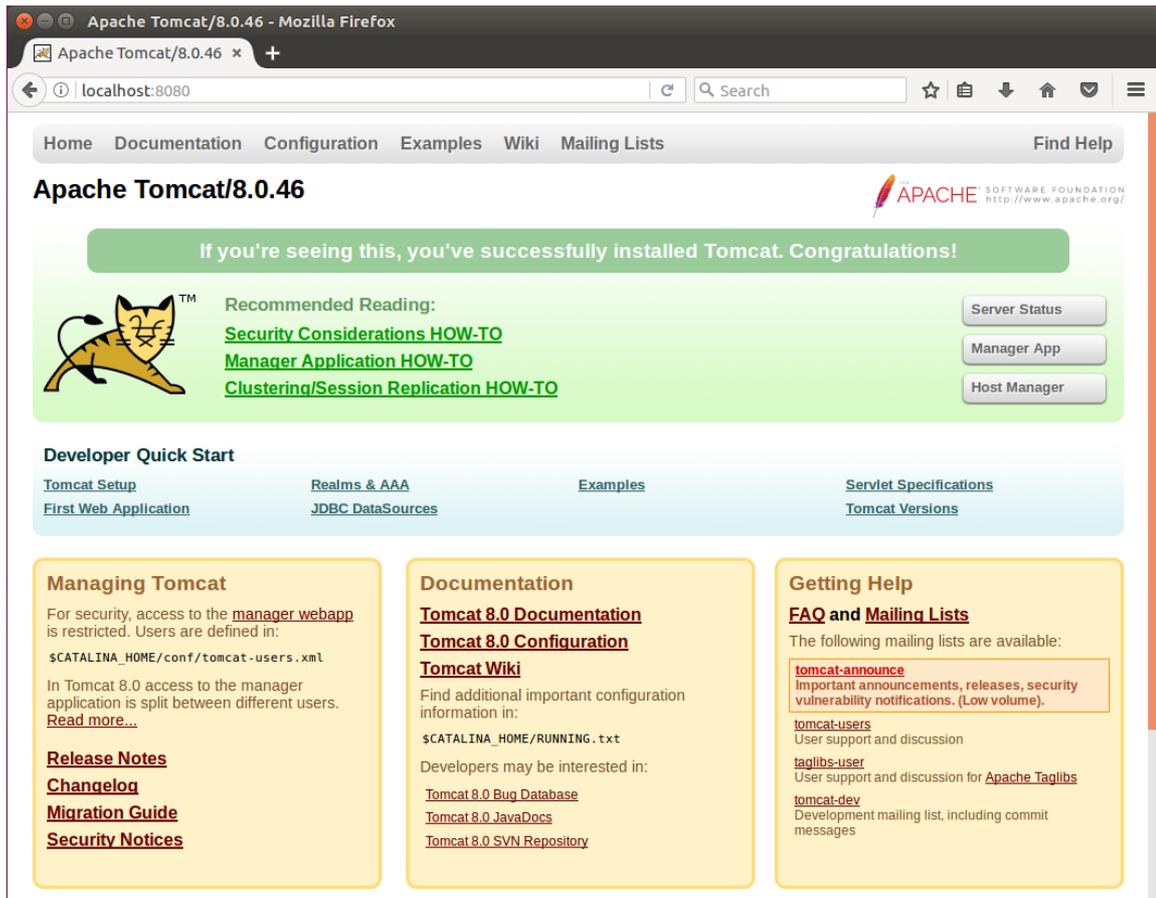


Figure 6.3. Apache Tomcat welcome page



## Tip

For detailed installation instructions, see: [Tomcat 7](#), [Tomcat 8](#), [Tomcat 8.5](#) or [Tomcat 9](#) Setup Guide.



## Note

For the installation on **IBM AS/400**, continue with [Apache Tomcat on IBM AS/400 \(iSeries\)](#) (p. 21).

**Continue with:** [Installation of CloverETL Server](#) (p. 21).

## Apache Tomcat as a Windows Service

1. Download the **32-bit/64-bit Windows Service Installer** file in the **Binary Distributions** section on the [Tomcat 7](#), [Tomcat 8](#), [Tomcat 8.5](#) or [Tomcat 9](#) download page.
2. Use the standard installation wizard to install Apache Tomcat.
3. Check whether Tomcat is running.
  - Type <http://localhost:8080/> in your browser's address bar.
  - If the Apache Tomcat Information page appears, the installation was successful.
4. When Tomcat is installed as a Windows service, **CloverETL** is configured by one of the following options:

## Graphical configuration utility

- Run the [Tomcat\_home]\bin\Tomcat9w.exe file.
- In the **Apache Tomcat Properties** dialog box, select the **Java** tab and set the initial and maximum heap size in **Initial memory pool** and **Maximum memory pool** fields to 512MB and 1024MB respectively. Other configuration parameters can be defined in **Java Options** field, being separated by new line.
- Click on **Apply** and restart the service.



### Note

The **Java** tab allows you to use alternative Java virtual machine by setup of path to `jvm.dll` file.

## Command Prompt tool

- Run the [Tomcat\_home]\bin\Tomcat9.exe file.
- If Tomcat is running, navigate to [Tomcat\_home]\bin and stop the service by typing:

```
.\Tomcat9.exe //SS//Tomcat9
```

in the Command Prompt. (When using different version of Tomcat, change the number in the command to reflect the installed version.)

- Configure the service by typing the command:

```
.\Tomcat9.exe //US//Tomcat9 --JvmMs=512 --JvmMx=1024 --JvmOptions=-Dclover.config.file=C:\path\to\clover-config.pr
```

The parameter `JvmMs` is the initial and `JvmMx` is the maximum heap size in MB; `JvmOptions` are separated by '#' or ' '.



### Important

If you use Java 7, change `-XX:MaxMetaspaceSize` to `-XX:MaxPermSize`.

- Start the service from Windows administration console or by typing the following command in the Command Prompt:

```
.\Tomcat9.exe //TS//Tomcat9
```



### Tip

By default, when Apache Tomcat is run as a Windows service, it is **not available** for Java process monitoring tools (e.g., **JConsole** or **JVisualVM**). However, these tools can still connect to the process via **JMX**. In order to expose Tomcat's Java process via JMX, add the following options to the service settings:

```
-Dcom.sun.management.jmxremote.port=3333
-Dcom.sun.management.jmxremote.ssl=false
-Dcom.sun.management.jmxremote.authenticate=false
```

Once the service is run with these options, you can connect to **port 3333** using JMX and monitor the server.



## Note

The instructions can be applied to Tomcat 7 and Tomcat 8, simply replace the number 9 in the file name with the number 7 or 8.

More information about running Java applications as Windows Service can be found at [Apache Commons](#).

**Continue with:** [Installation of CloverETL Server](#) (p. 21).

## Apache Tomcat on IBM AS/400 (iSeries)

Additional settings are required to run CloverETL Server on the iSeries platform:

1. Declare you are using Java 7.0 32-bit.
2. Run Java with parameter `-Djava.awt.headless=true`.

To configure the settings, modify (or create) the `[Tomcat_home]/bin/setenv.sh` file to contain:

```
JAVA_HOME=/QOpenSys/QIBM/ProdData/JavaVM/jdk70/32bit
JAVA_OPTS="$JAVA_OPTS -Djava.awt.headless=true"
```

**Continue with:** [Installation of CloverETL Server](#) (p. 21)

## Installation of CloverETL Server

1. Check if you meet the prerequisites:
  - Oracle JDK or JRE is installed (See [Java Virtual Machine](#) (p. 9) for required Java version.)
  - JAVA\_HOME or JRE\_HOME environment variable is set (see [Setting up JAVA\\_HOME](#) (p. 18)).
  - A supported version (p. 10) of Apache Tomcat is installed.
2. It is strongly recommended to adjust the default limits for **Memory allocation** (see the [Memory Settings](#) (p. 60) section).

You can set the **initial** and **maximum memory heap size** by adjusting the **Xms** and **Xmx** JVM parameters:

### Unix-like systems:

- Create the `[Tomcat_home]/bin/setenv.sh` file.
- Type or paste in the following lines:

```
export CATALINA_OPTS="$CATALINA_OPTS -XX:MaxMetaspaceSize=512m -Xms128m -Xmx1024m"
export CATALINA_OPTS="$CATALINA_OPTS -Dderby.system.home=$CATALINA_HOME/temp -server"
echo "Using CATALINA_OPTS: $CATALINA_OPTS"
```

### Windows systems:

- Create the `[Tomcat_home]\bin\setenv.bat` file.
- Type or paste in the following lines:

```
set "CATALINA_OPTS=%CATALINA_OPTS% -XX:MaxMetaspaceSize=512m -Xms128m -Xmx1024m"
set "CATALINA_OPTS=%CATALINA_OPTS% -Dderby.system.home=%CATALINA_HOME%/temp -server"
```

```
echo "Using CATALINA_OPTS: %CATALINA_OPTS%"
```



### Important

If you use Java 7, change `-XX:MaxMetaspaceSize` to `-XX:MaxPermSize`.



### Tip

For performance reasons, it is recommended to run the container in the "server mode" by setting the `-server` switch, as seen in the settings above.<sup>1</sup>

Note that on a 64-bit capable JDK, only the Java Hotspot Server VM is supported so the `-server` option is implicit.

3. Go to the download section of your [CloverETL account](#) and download the `clover.war` (web archive) file containing CloverETL Server for Apache Tomcat.
4. Copy `clover.war` to the `[Tomcat_home]/webapps` directory.



### Note

Please note that copying is not an atomic operation. If Tomcat is running, mind the duration of the copying process! Too long copying might cause a failure during deployment as Tomcat tries to deploy an incomplete file. Instead, manipulate the file when the Tomcat is **not** running.

5. Tomcat should automatically detect and deploy the `clover.war` file.
6. Check whether CloverETL Server is running:
  - Run Tomcat.
  - Open a new tab in your browser and type <http://localhost:8080/clover/> in the address bar.
  - Use the **default administrator credentials** to access the web GUI: username: **clover**, password: **clover**.

Continue with: [Configuration of CloverETL Server on Apache Tomcat](#) (p. 22)

## Configuration of CloverETL Server on Apache Tomcat



### Tip

Default installation (without any configuration) is only recommended for evaluation purposes. For production use, at least a dedicated, system database and SMTP server configuration is recommended.

For easy configuration of CloverETL Server, use Setup GUI (p. 74) in which you can configure various properties, including the connection to the database, username and password, path to the license file, private properties, clusters and much more (see Chapter 14, [List of Configuration Properties](#)(p. 106) and Chapter 40, [Cluster Configuration](#) (p. 325)). We recommend you place the file in a specified (p. 70)location and define the path to the file with a system property.

The content of such a file (an example with a PostgreSQL database):

```
jdbc.driverClassName=org.postgresql.Driver
jdbc.url=jdbc:postgresql://127.0.0.1/clover_db?charSet=UTF-8
jdbc.username=yourUsername
jdbc.password=yourPassword
jdbc.dialect=org.hibernate.dialect.PostgreSQLDialect
```

### Properties File in Specified Location

The properties file is loaded from a location specified by a system property or by an environment variable `clover_config_file` or `clover.config.file`.

1. Create the `cloverServer.properties` file in a directory readable by Apache Tomcat. (If you need an example of connection to any of the supported database systems, see Chapter 13, [System Database Configuration](#) (p. 83).)
2. Edit the `[Tomcat_home]/bin/setenv.sh` file (if it does not exist, you can create it).
3. Set the system property by adding the following line into the file:

```
JAVA_OPTS="$JAVA_OPTS -Dclover_config_file=/path/to/cloverServer.properties"
```



### Note

➔ **Continue with:** Chapter 7, [Postinstallation Configuration](#) (p. 60)

## Jetty

---

[Installation of Jetty](#) (p. 24)

[Installation of CloverETL Server](#) (p. 25)

[Configuration of CloverETL Server on Jetty](#) (p. 26)



### Important

See [Application Server](#) (p. 9) in system requirements for currently supported **Jetty** versions and required **Java** versions.

If you encounter any problems during the installation, the [Possible Issues during Installation](#) (p. 56) section may provide a solution.

## Installation of Jetty

1. Download the Jetty release from the [official download page](#).

**CloverETL** Server is currently developed and tested with Jetty 9.2.11 (v20150529) container. Since **CloverETL 4.0**, Jetty 6 is no longer supported (latest tested version is 6.1.22). Running the Server with other versions may result in unpredictable behavior.

2. Extract the downloaded archive (zip or tar.gz).

3. Run Jetty.

- **Unix-like systems:**

- Run `[Jetty_home]/bin/jetty.sh start`

- **Windows system:**

- Run the `[Jetty_home]\java -jar start.jar --exec` command in Windows Command Prompt.

4. Check whether Jetty is running.

- Open a new tab in your browser and type <http://localhost:8080/> in the address bar.



### Note

Since the `clover.war` file is not yet implemented, you may see the **Error 404 - Not found** status code (see the figure below). However, it means that the server is running.

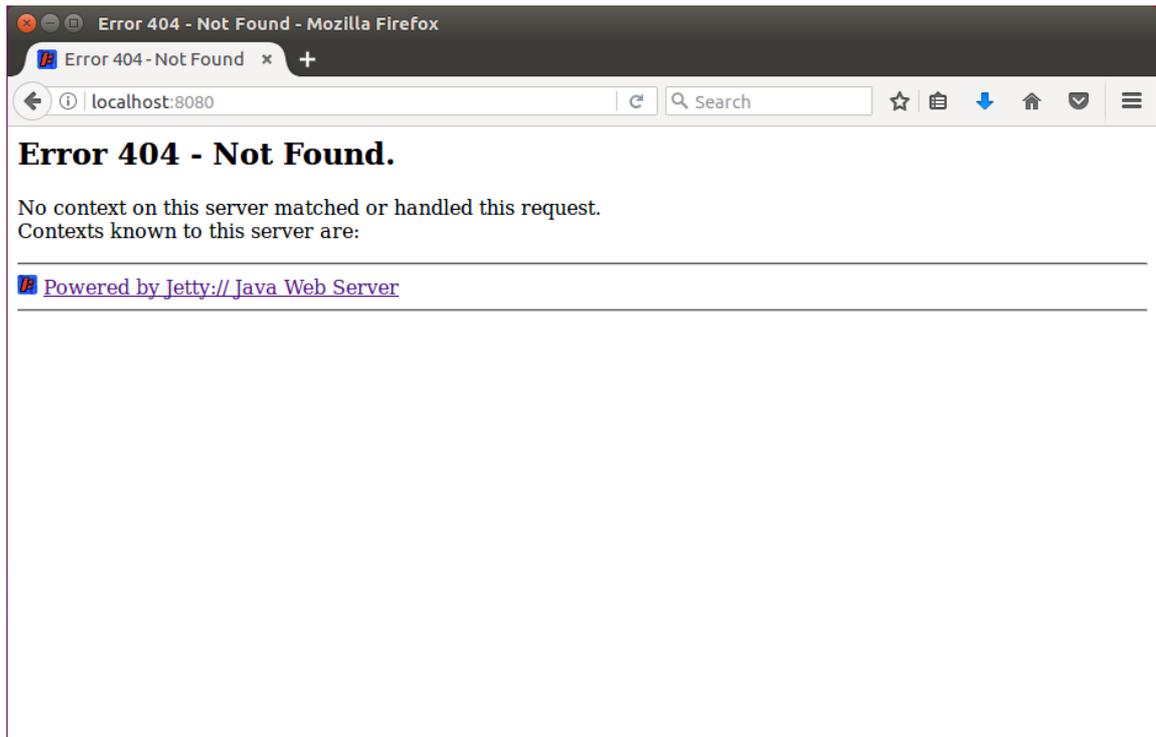


Figure 6.4. Jetty welcome page

## Installation of CloverETL Server

1. Check if you meet the prerequisites:

- Oracle JDK or JRE is installed (see [Java Virtual Machine](#) (p. 9) for the required Java version).
- `JAVA_HOME` or `JRE_HOME` environmental variable is set (see [Setting up JAVA\\_HOME](#) (p. 18)).
- A supported version (p. 10) of Jetty is installed.

2. It is strongly recommended to adjust the default limits for **Memory allocation** (see the [Memory Settings](#) (p. 60) section).

You can set the **initial** and **maximum memory heap size** by adjusting the `Xms` and `Xmx` JVM parameters, and **classloaders memory limit** by adjusting the `XX:MaxMetaspaceSize` parameter:

- **Unix-like systems:**

Edit the `[Jetty_home]/bin/jetty.sh` file.

Type or paste the following line at the end of the file:

```
JAVA_OPTIONS=' $JAVA_OPTIONS -Xms128m -Xmx1024m -XX:MaxMetaspaceSize=256m'
```

- **Windows system:**

Edit the `[Jetty_home]\start.ini` file.

Type or paste the following line at the end of the file:

```
JAVA_OPTIONS=' $JAVA_OPTIONS -Xms128m -Xmx1024m -XX:MaxMetaspaceSize=256m'
```



## Important

If you use Java 7, change `-XX:MaxMetaspaceSize` to `-XX:MaxPermSize`.

3. Go to the download section of your [CloverETL account](#) and download the `clover.war` (web archive) file containing CloverETL Server for Jetty.
4. Copy `clover.war` to the `[Jetty_home]/webapps` directory.
5. Run Jetty.
  - **Unix-like systems:**
    - Run `[Jetty_home]/bin/jetty.sh start`
  - **Windows system:**
    - Run the `[Jetty_home]\java -jar start.jar --exec` command in Windows command prompt.
6. Check whether CloverETL Server is running:
  - Open a new tab in your browser and type <http://localhost:8080/clover/> in the address bar.
  - Use the **default administrator credentials** to access the web GUI: username: **clover**, password: **clover**.

## Configuration of CloverETL Server on Jetty



### Tip

Default installation (without any configuration) is only recommended for evaluation purposes. For production use, at least a dedicated, system database and SMTP server configuration is recommended.

For easy configuration of CloverETL Server, use Setup GUI (p. 74) in which you can configure various properties, including the connection to the database, username and password, path to the license file, private properties, clusters and much more (see Chapter 14, [List of Configuration Properties](#) (p. 106) and Chapter 40, [Cluster Configuration](#) (p. 325)). We recommend you place the file in a specified (p. 70) location and define the path to the file with a system property.

The content of such a file (an example with a PostgreSQL database):

```
jdbc.driverClassName=org.postgresql.Driver
jdbc.url=jdbc:postgresql://127.0.0.1/clover_db?charSet=UTF-8
jdbc.username=yourUsername
jdbc.password=yourPassword
jdbc.dialect=org.hibernate.dialect.PostgreSQLDialect
```



### Note

JDBC Driver must be JDBC 4 compliant and stored in the `[Jetty_home]/lib/ext`.

## Properties file in Specified Location

The properties file is loaded from a location which is specified by the environment/system property `clover_config_file` or `clover.config.file`. This is a recommended way of configuring Jetty.

1. Create the `cloverServer.properties` file in a directory readable by Jetty. (If you need an example of connection to any of the supported database systems, see Chapter 13, [System Database Configuration](#) (p. 83).)

**2. Unix-like systems:**

- a. Edit the `[Jetty_home]/bin/jetty.sh` file.
- b. Set the system property by adding the following line into the file:

```
JAVA_OPTIONS="$JAVA_OPTIONS -Dclover_config_file=/path/to/cloverServer.properties"
```

**Windows system:**

- a. Edit the `[Jetty_home]\start.ini` file.
- b. Set the system property by adding the following line into the file just after the memory settings:

```
JAVA_OPTIONS="$JAVA_OPTIONS -Dclover_config_file=/path/to/cloverServer.properties"
```



**Note**

➔ **Continue with:** Chapter 7, [Postinstallation Configuration](#) (p. 60)

## IBM WebSphere

[Installation of IBM Websphere](#) (p. 28)

[Installation of CloverETL Server](#) (p. 28)

[Configuration of CloverETL Server on IBM WebSphere](#) (p. 30)



### Important

See [Application Server](#) (p. 9) in system requirements for currently supported **IBM WebSphere** versions and required **Java** versions.

In order to ensure reliable function of CloverETL Server always use the latest version of IBM Java SDK. At least SDK 7.0 SR6 (package *IBM WebSphere SDK Java Technology Edition V7.0.6.1*) is recommended. Using older SDKs may lead to deadlocks during execution of specific ETL graphs.

If you encounter any problems during the installation, the [Possible Issues during Installation](#) (p. 56) section may provide a solution.

## Installation of IBM Websphere

1. Create a My IBM account on <https://www.ibm.com>
2. Go to IBM Marketplace and download IBM Installation Manager.
3. Follow the instructions to download IBM Websphere.

CloverETL Server is developed and tested with the IBM WebSphere 8.5 container. Running the Server with other versions may result in unpredictable behavior.

4. Create a profile.

- **Unix-like systems:**

```
Run the [IBM_home]/WebSphere/AppServer/bin/ProfileManagement/pmt.sh.
```

- **Windows system:**

```
Run the [IBM_home]\WebSphere\AppServer\bin\ProfileManagement\pmt.bat.
```



### Important

Make sure the profile name does not contain the keyword "clover", otherwise the CloverETL server won't start properly.

5. Check whether the server is operational.

You will be prompted to run a first-time server management tool that will check whether the installation was successful and the server can be started.

Go to <http://localhost:9060/ibm/console/> where you should be able to see the WebSphere login page.

## Installation of CloverETL Server

1. Check if you meet the prerequisites:

- IBM Java SDK is installed (see [Java Virtual Machine](#) (p. 9) for the required Java version).
- **JAVA\_HOME** or **JRE\_HOME** environmental variable is set (see [Setting up JAVA\\_HOME](#) (p. 18)).

- A supported version (p. 10) of IBM WebSphere is installed.
2. It is strongly recommended to adjust the default limits for **Memory allocation** (see the [Memory Settings](#) (p. 60) section).

You can set the limits in IBM WebSphere's **Integrated Solutions Console** (default URL: <http://localhost:9060/ibm/console/>).

- Go to **Servers** → **Server Types** → **WebSphere application servers** → **[Server\_Name]** (default name: **server1**) → **Java and Process Management** → **Process definition** → **Java Virtual Machine**
- Change the value in the **Maximum heap size** field to 2,048 MB. The default value (256 MB) is insufficient for ETL transformations.

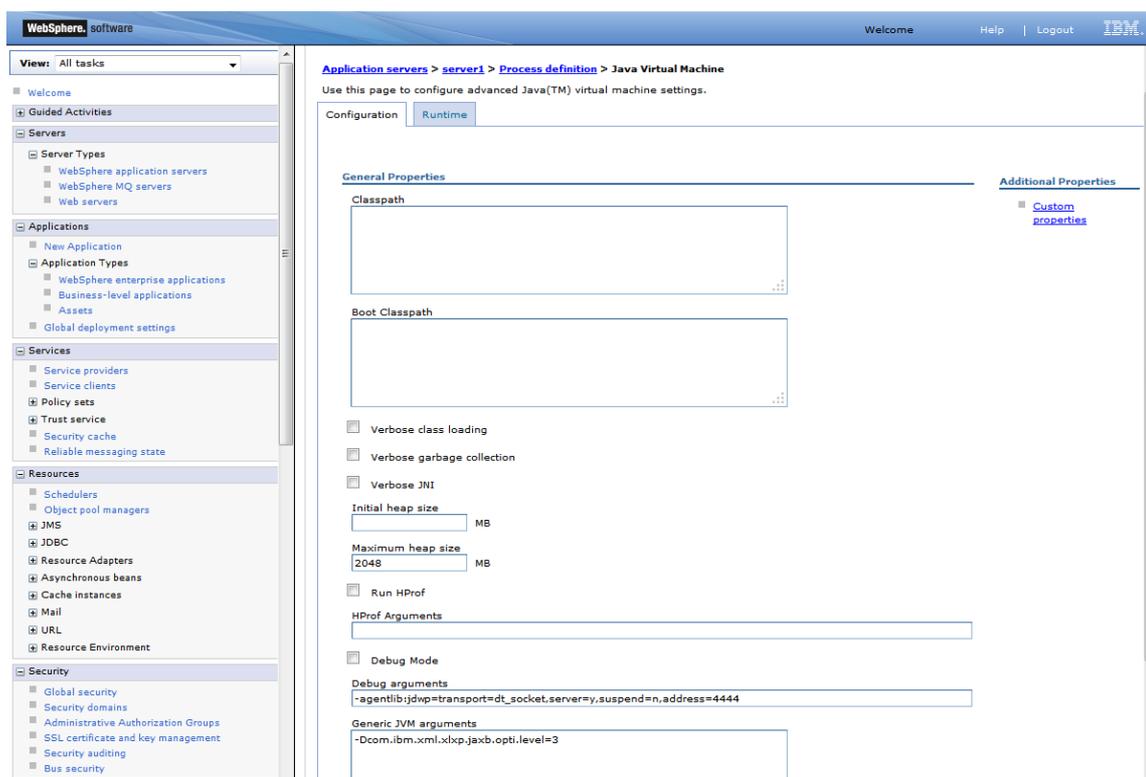


Figure 6.5. Adjusting Maximum heap size limit

- Add the following parameters into the **Generic JVM arguments** field to set the perm space limit and direct memory limit:

```
-XX:MaxPermSize=512M
```

```
-XX:MaxDirectMemorySize=512M
```

- Java runtime settings:

Go to **Servers** → **Server Types** → **WebSphere application servers** → **[Server\_Name]** (default name: **server1**) → **Java SDKs** and select version 1.7 as the default SDK.

- Save the changes to configuration and restart the server so that they take effect.

3. Go to the download section of your [CloverETL account](#) and download the `clover.war` (web archive) file containing CloverETL Server for WebSphere.

4. Deploy the `clover.war` file.

- Go to **Integrated Solutions Console** (default URL: <http://localhost:9060/ibm/console/>).
- Go to **Applications** → **New Application** → **New Enterprise Application**, select the CloverETL Server WAR archive and deploy it to the application server, but do not start it yet.

5. Configure application class loading.

Go to **WebSphere Enterprise Applications** → **clover\_war** (or other name of the Clover application) → **Manage Modules** → **CloverETL** and under **Class loader order** select **Classes loaded with local class loader first (parent last)**.

6. Save the changes to the Server configuration and start the `clover_war` application.

7. Check whether the Server is running.

Provided you set `clover.war` as the application running with "clover" context path, use the following URL (notice the port number has changed):

<http://localhost:9080/clover>



### Note

Please note that some **CloverETL** features using third party libraries do not work properly on IBM WebSphere.

- Hadoop is guaranteed to run only on Oracle Java 1.6+, but Hadoop developers do make an effort to remove any Oracle/Sun-specific code. See [Hadoop Java Versions](#) on Hadoop Wiki.
- **AddressDoctor5** on IBM WebSphere requires additional JVM parameter `-Xms2048k` to prevent AddressDoctor from crashing JVM. See documentation on AddressDoctor component.

## Configuration of CloverETL Server on IBM WebSphere



### Tip

Default installation (without any configuration) is only recommended for evaluation purposes. For production use, at least a dedicated, system database and SMTP server configuration is recommended.

For easy configuration of CloverETL Server, use a Setup GUI (p. 74) in which you can configure various properties, including the connection to the database, username and password, path to the license file, private properties, clusters and much more (see Chapter 14, [List of Configuration Properties](#)(p. 106) and Chapter 40, [Cluster Configuration](#) (p. 325). We recommend you place the file in a specified (p. 70) location and specify the path to the file with a system property.

The content of such a file (an example with a PostgreSQL database):

```
jdbc.driverClassName=org.postgresql.Driver
jdbc.url=jdbc:postgresql://127.0.0.1/clover_db?charSet=UTF-8
jdbc.username=yourUsername
jdbc.password=yourPassword
jdbc.dialect=org.hibernate.dialect.PostgreSQLDialect
```

### Properties File in Specified Location

The properties file is loaded from a location which is specified by the environment/system property `clover_config_file` or `clover.config.file`.

1. Create the `cloverServer.properties` file in a directory readable by IBM WebSphere. (If you need an example of connection to any of the supported database systems, see Chapter 13, [System Database Configuration](#) (p. 83).)
2. Set a system property (or environment variable) `clover_config_file` pointing to the properties file.
  - a. Go to **Integrated Solutions Console** (default URL:<http://localhost:9060/ibm/console/>).
  - b. Go to **Servers** → **WebSphere application servers** → **[Server\_name]** → **Java and Process Management** → **Process Definition** → **Java Virtual Machine** → **Custom Properties**.
  - c. Create a system property named `clover_config_file` whose value is a full path to the properties file (e.g. `cloverServer.properties`) on your file system.
3. Restart IBM WebSphere for changes to take effect.



### Note

➡ **Continue with:** Chapter 7, [Postinstallation Configuration](#) (p. 60)

## GlassFish / Sun Java System Application Server

---

[Installation of GlassFish](#) (p. 32)

[Installation of CloverETL Server](#) (p. 33)

[Configuration of CloverETL Server on GlassFish](#) (p. 34)



### Important

See [Application Server](#) (p. 9) in system requirements for currently supported **GlassFish** versions and required **Java** versions.

GlassFish 3.1.2 contains a bug causing **Launch Services** to work improperly (see <https://java.net/jira/browse/GLASSFISH-18444>). We recommend version 3.1.2.2.

If you encounter any problems during the installation, the [Possible Issues during Installation](#) (p. 56) section may provide a solution.

### Installation of GlassFish

1. Choose and download the GlassFish release from the [official download page](#).

CloverETL Server is developed and tested with the GlassFish 3.1.2.2 container. Running the Server with other versions may result in unpredictable behavior.

2. Extract the downloaded archive, or run the `.exe` file which will guide you through the setup.

3. Run GlassFish.

- Run `[GlassFish_home]/bin/asadmin start-domain` and enter a new password for administrator.

4. Check whether GlassFish is running.

- Open a new tab in your browser and type <http://localhost:8080/> in the address bar.

If the GlassFish welcome page appears (see below), the server is successfully installed.

- The Admin Console is, by default, accessible at <http://localhost:4848/>.

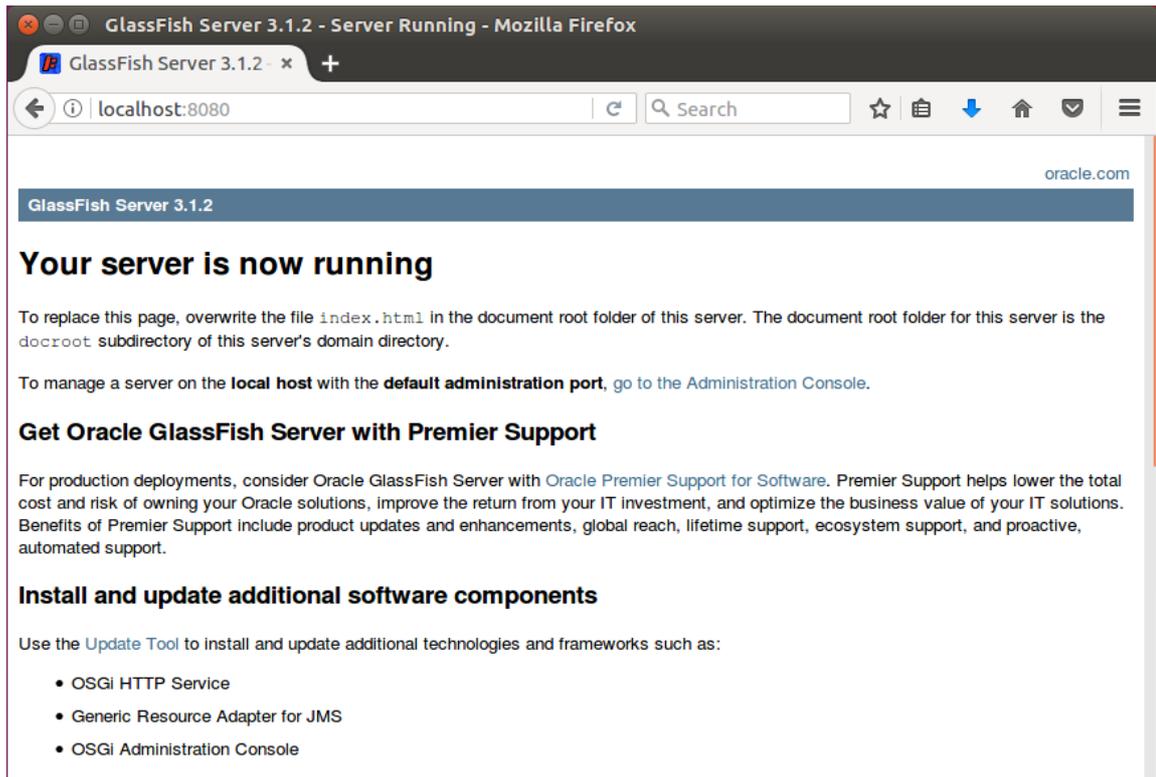


Figure 6.6. Glassfish welcome page

## Installation of CloverETL Server

1. Check if you meet the prerequisites:

- Oracle JDK or JRE is installed (see [Java Virtual Machine](#) (p. 9) for required java version).
- `JAVA_HOME` or `JRE_HOME` environment variable is set (see [Setting up JAVA\\_HOME](#) (p. 18)).
- A supported version (p. 10) of GlassFish 3 is installed.

2. It is strongly recommended to adjust the default limits for **Memory allocation** (see the [Memory Settings](#) (p. 60) section).

You can set the **initial** and **maximum memory heap size** by adjusting the "Xms" and "Xmx" JVM parameters and **perm space** by adjusting the "XX:MaxPermSize" parameter:

- Edit the `[GlassFish_home]/glassfish/domains/domain1/config/domain.xml` file.

Change/add the following sub-elements in the `<java-config>` section:

```
<jvm-options>-XX:MaxPermSize=384m</jvm-options>
<jvm-options>-XX:PermSize=256m</jvm-options>
<jvm-options>-Xms512m</jvm-options>
<jvm-options>-Xmx2g</jvm-options>
```

- Restart GlassFish.

3. Go to the download section of your [CloverETL account](#) and download the `clover.war` (web archive) file containing CloverETL Server for GlassFish 3.

4. Deploy the `clover.war` file.

- Open the **GlassFish Administration Console** (default URL: <http://localhost:4848/>).
- Go to **Applications** and click **Deploy ...**
- Upload the `clover.war` file or select the file from a filesystem if it is present on the machine running GlassFish.
- Make sure the **Web Application** is selected in the **Type** field.  
Type "clover" in both the **Application Name** and the **Context Root** fields.
- Click **OK**.

## Configuration of CloverETL Server on GlassFish



### Tip

Default installation (without any configuration) is only recommended for evaluation purposes. For production use, at least a dedicated, system database and SMTP server configuration is recommended.

For easy configuration of CloverETL Server, use a Setup GUI (p. 74) in which you can configure various properties, including the connection to the database, username and password, path to the license file, private properties, clusters and much more (see Chapter 14, [List of Configuration Properties](#) (p. 106) and Chapter 40, [Cluster Configuration](#) (p. 325)). We recommend you place the file in a specified (p. 70) location and specify the path to the file with a system property.

The content of such a file (an example with a PostgreSQL database):

```
datasource.type=JDBC
jdbc.driverClassName=org.postgresql.Driver
jdbc.url=jdbc:postgresql://127.0.0.1/clover_db?charSet=UTF-8
jdbc.username=yourUsername
jdbc.password=yourPassword
jdbc.dialect=org.hibernate.dialect.PostgreSQLDialect
```

### Properties file in Specified Location

The properties file is loaded from a location which is specified by the environment/system property `clover_config_file` or `clover.config.file`.

1. Create the `cloverServer.properties` file in a directory readable by GlassFish. (If you need an example of connection to any of the supported database systems, see Chapter 13, [System Database Configuration](#) (p. 83).)
2. Set a system property `clover.config.file` pointing to the config properties file:
  - a. Go to **GlassFish Administration Console** (default URL: <http://localhost:4848/>).
  - b. Go to **Configuration** → **System Properties**
  - c. Create a system property named `clover.config.file` whose value is a full path to the file on your file system (e.g.: `/home/clover/cloverServer.properties`).
3. Copy the **JDBC driver** .jar file for a selected database system into `[GlassFish_home]/glassfish/domains/[domain-name]/lib`
4. Restart GlassFish.



**Note**

➔ **Continue with:** Chapter 7, [Postinstallation Configuration](#) (p. 60)

## JBoss Application Server

---

[Installation of JDK](#) (p. 36)

[Installation of JBoss AS](#) (p. 36)

[Installation of CloverETL Server](#) (p. 37)

[Configuration of CloverETL Server on JBoss AS](#) (p. 39)



### Important

See [Application Server](#) (p. 9) in system requirements for currently supported **JBoss AS** versions and required **Java** versions.

If you encounter any problems during the installation, the [Possible Issues during Installation](#) (p. 56) section may provide a solution.

### Installation of JDK

1. Install Oracle JDK (see [Java Virtual Machine](#) (p. 9) for the required Java version).
2. Set up `JAVA_HOME` to point to the correct Java version

- **Unix-like systems:**

Setup the path in `/etc/profile` or `/etc/bash.bashrc`:

- **Windows system:**

Run `sysdm.cpl`. Switch to **Advanced** tab. Open **Environment variables ...** and add new environment variable.

### Installation of JBoss AS

1. Download the JBoss AS release from the [official download page](#).

CloverETL Server is developed and tested with JBoss AS 6 and 7 containers. Running the Server with other versions may result in unpredictable behavior.

2. Extract the downloaded archive (`zip` or `tar.gz`).
3. Run JBoss AS.

- **Unix-like systems:**

- JBoss AS v 6.x.x

```
Run [JBoss_AS_home]/bin/run.sh.
```

- JBoss AS v 7.x.x

```
Run [JBoss_AS_home]/bin/standalone.sh.
```

- **Windows system:**

- JBoss AS v 6.x.x

```
Run [JBoss_AS_home]\bin\run.bat.
```

- JBoss AS v 7.x.x

```
Run [JBoss_AS_home]\bin\standalone.bat.
```

## 4. Check whether JBoss AS is running.

- Open a new tab in your browser and type <http://localhost:8080/> in the address bar.

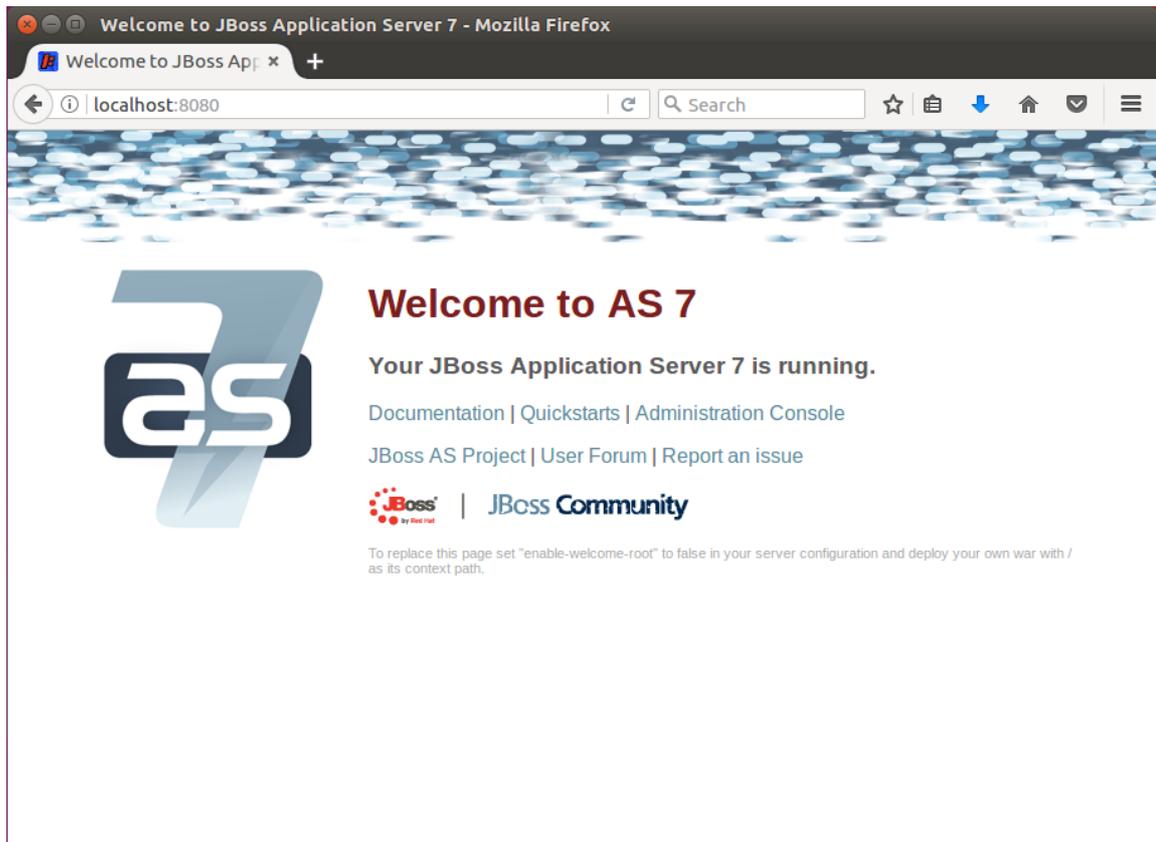


Figure 6.7. JBoss AS 7 welcome page

## Installation of CloverETL Server

### 1. Check if you meet the prerequisites:

- Oracle JDK or JRE is installed (see [Java Virtual Machine](#) (p. 9) for the required Java version).
- **JAVA\_HOME** or **JRE\_HOME** environment variable is set (see [Setting up JAVA\\_HOME](#) (p. 18)).
- A supported version (p. 10) of JBoss AS is installed.

### 2. It is strongly recommended to adjust the default limits for **Memory allocation** (see the [Memory Settings](#) (p. 60) section).

You can set the **initial** and **maximum memory heap size** by adjusting the "Xms" and "Xmx" JVM parameters and **classloaders memory limit** by adjusting the "XX:MaxMetaspaceSize" parameter:

- **Unix-like systems:**

- JBoss AS v 6.x.x

Edit the `[JBoss_AS_home]/bin/run.conf`.

- JBoss AS v 7.x.x

Edit the `[JBoss_AS_home]/bin/standalone.conf`.

Edit the values of the following attributes:

```
-XX:MaxMetaspaceSize=512m -Xms128m -Xmx2048m"
```

- **Windows system:**

- JBoss AS v 6.x.x

Edit the [JBoss\_AS\_home]\bin\run.conf.bat.

- JBoss AS v 7.x.x

Edit the [JBoss\_AS\_home]\bin\standalone.conf.bat.

Edit the values of the following attributes:

```
-XX:MaxMetaspaceSize=512m -Xms128m -Xmx2048m"
```

3. Go to the download section of your [CloverETL account](#) and download the clover.war (web archive) file containing CloverETL Server for JBoss AS.

4. Create a separate JBoss server configuration.

It may be useful to use a specific JBoss server configuration, when it is necessary to run CloverETL:

- isolated from other JBoss applications
- with a different set of services
- with different libraries on the classpath than other applications

See the JBoss manual for details about the JBoss server configuration: [JBoss Server Configurations](#), [Start the Server With Alternate Configuration](#)

5. Configure database connection.

As CloverETL Server's embedded Derby database does not work under JBoss AS, a database connection has to be always configured. We used MySQL accessed via JNDI-bound datasource in this example:

- Create datasource deployment file [JBoss\_AS\_home]/server/[serverConfiguration]/deploy/mysql-ds.xml

```
<datasources>
  <local-tx-datasource>
    <jndi-name>CloverETLServerDS</jndi-name>
    <connection-url>jdbc:mysql://localhost:3306/cloverServerDB</connection-url>
    <driver-class>com.mysql.jdbc.Driver</driver-class>
    <user-name>clover</user-name>
    <password>SecretPasword</password>
  </local-tx-datasource>
</datasources>
```



## Note

Special characters in the XML file have to be typed in as XML entities. For instance, ampersand "&" as "&amp;", etc.

"CloverETLServerDS" is the name under which the datasource will be accessible. The thing to do here is to set database connection parameters (connection-url, driver-class, user-name and password) to the database. The database has to be empty before the first execution, the server creates its tables on its own.

- Put the JDBC driver JAR file for your database to the application server classpath. In this example we copied the file `mysql-connector-java-5.1.5-bin.jar` to `[JBoss_AS_home]/server/[serverConfiguration]/lib`
6. Configure CloverETL Server according to the description in the [next section](#) (p. 39).
  7. Deploy the WAR file:
    - JBoss AS 6.x.x  
Copy `clover.war` to `[JBoss_AS_home]/server/[serverConfiguration]/deploy`
    - JBoss AS 7.x.x  
Copy `clover.war` to `[JBoss_AS_home]/standalone/deployments`
  8. Start JBoss AS:
    - JBoss AS 6.x.x  
Start JBoss AS via `[JBoss_AS_home]/bin/run.sh` (or `run.bat` on Windows OS).  
  
If you want to run JBoss with a specific server configuration, it has to be specified as a parameter: `[JBoss_AS_home]/bin/run.sh -c [serverConfiguration]` If the `serverConfiguration` is not specified, the "default" is used.
    - JBoss AS 7.x.x  
Start JBoss AS via `[JBoss_AS_home]/bin/standalone.sh` (or `standalone.bat` on Windows OS).

## Configuration of CloverETL Server on JBoss AS



### Note

Default installation (without any configuration) does not work under JBoss AS. In order to be able to use the **CloverETL Server**, a working database connection is required.

For an easy configuration of CloverETL Server, use a Setup GUI (p. 74) in which you can configure various properties, including the connection to the database, username and password, path to the license file, private properties, clusters and much more (see Chapter 14, [List of Configuration Properties](#) (p. 106) and Chapter 40, [Cluster Configuration](#) (p. 325)). We recommend you place the file in a specified (p. 70) location and specify the path to the file with a system property.

Content of such a file (example with MySQL database):

```
datasource.type=JNDI
datasource.jndiName=java:/CloverETLServerDS
jdbc.dialect=org.hibernate.dialect.MySQLDialect
```

<code>datasource.type</code>	Indicates the server will use JNDI-bound datasource created in steps above.
<code>datasource.jndiName</code>	Specifies where can the datasource be found in JNDI.
<code>jdbc.dialect</code>	Set the dialect according to your database server (Part III, <a href="#">"Configuration"</a> (p. 67)).

### Properties File in Specified Location

The properties file is loaded from a location which is specified by the environment/system property `clover_config_file` or `clover.config.file`.

1. Create the `cloverServer.properties` file in a directory readable by JBoss AS. (If you need an example of connection to any of the supported database systems, see Chapter 13, [System Database Configuration](#) (p. 83).)
2. Set the system property (or environment variable) `clover.config.file` pointing to the configuration properties file.

It should contain a full path to the `cloverServer.properties` file created in the previous step.

The simplest way is by setting a Java parameter:

**Unix-like systems:**

- a. Edit the file `[JBoss_AS_home]/bin/run.sh` (JBoss AS v 6.x.x) or `standalone.sh` (JBoss AS v 7.x.x).
- b. Add the following line:

```
export JAVA_OPTS="$JAVA_OPTS -Dclover.config.file=/absolute/path/to/cloverServer.properties"
```

**Windows system:**

- a. Edit the file `[JBoss_AS_home]\bin\run.conf.bat` (JBoss AS v 6.x.x) or `\standalone.conf.bat` (JBoss AS v 7.x.x).
- b. Add the following line to the section where options are passed to the JVM:

```
set JAVA_OPTS=%JAVA_OPTS% -Dclover.config.file=\absolute\path\to\cloverServer.properties
```



**Important**

Do not override other settings in the `JAVA_OPTS` property - i.e. memory settings described above.

3. Restart JBoss AS for the changes to take effect.
4. Check the CloverETL Server application is running:

By default, the Server's console is accessible at <http://localhost:8080/clover>.



**Note**

➔ **Continue with:** Chapter 7, [Postinstallation Configuration](#) (p. 60)

**External Documentation Sources**

[JBoss AS Documentation](#)

## JBoss Enterprise Application Platform

---

[Installation of JBoss EAP](#) (p. 41)

[Installation of CloverETL Server](#) (p. 42)

[Configuration of CloverETL Server on JBoss EAP](#) (p. 45)



### Important

See [Application Server](#) (p. 9) in system requirements for currently supported **JBoss EAP** versions and required **Java** versions.

If you encounter any problems during the installation, the [Possible Issues during Installation](#) (p. 56) section may provide a solution.

### Installation of JBoss EAP

1. Using your credentials, log into the customer portal on the [official Red Hat page](#) and download a compatible JBoss EAP version.
2. Extract the downloaded archive (alternatively, you can download and run the installer and follow the instructions).
3. Run JBoss EAP.
  - **Unix-like systems:**  

```
Run [JBoss_EAP_home]/bin/standalone.sh.
```
  - **Windows system:**  

```
Run [JBoss_EAP_home]\bin\standalone.bat.
```
4. Check whether JBoss EAP is running.
  - Open a new tab in your browser and type <http://localhost:8080/> in the address bar.  
You should see the JBoss EAP welcome page (otherwise, please consult the [JBoss EAP guide](#)):

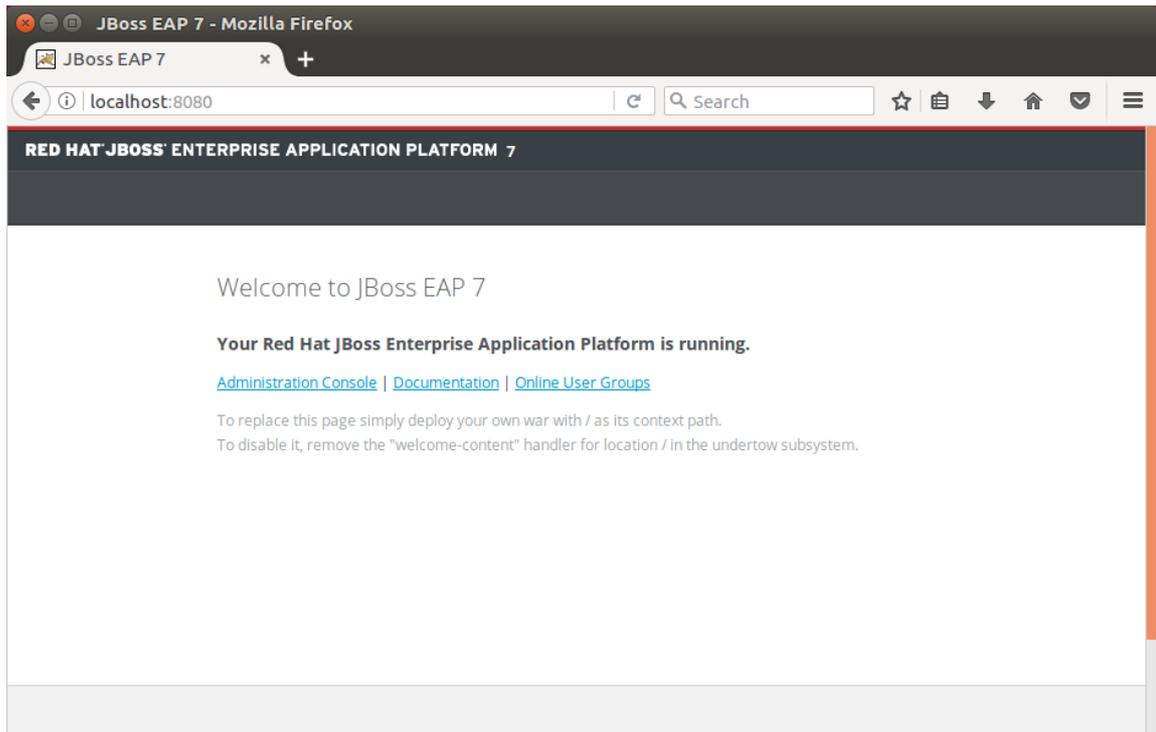


Figure 6.8. JBoss EAP 7 welcome page

- To shutdown JBoss EAP 7, run the following script with command:

```
jboss-cli.sh --connect command=:shutdown
```

## Installation of CloverETL Server

1. Check if you meet the prerequisites:

- Oracle JDK or JRE is installed (see [Java Virtual Machine](#) (p. 9) for the required Java version).
- `JAVA_HOME` or `JRE_HOME` environment variable is set (see [Setting up JAVA\\_HOME](#) (p. 18)).
- A supported version (p. 10) of JBoss EAP (JBoss AS 7) is installed.

CloverETL Server is developed and tested with the JBoss EAP 6.2 and 6.4 (JBoss 7) containers. Running the Server with other versions may result in unpredictable behavior.

2. It is strongly recommended to adjust the default limits for **Memory allocation** (see the [Memory Settings](#) (p. 60) section).

You can set the **initial** and **maximum memory heap size** by adjusting the "Xms" and "Xmx" JVM parameters and **classloaders memory limit** by adjusting the "XX:MaxMetaspaceSize" parameter:

For JBoss EAP standalone mode, follow these steps:

- **Unix-like systems:**
  - Edit the `[JBoss_EAP_home]/bin/standalone.conf` file.
  - Add the following line:

```
JAVA_OPTS="$JAVA_OPTS -XX:MaxMetaspaceSize=512m -Xms128m -Xmx2048m"
```

- **Windows systems:**

- Edit the `[JBoss_EAP_home]\bin\standalone.conf.bat` file.

```
JAVA_OPTS="$JAVA_OPTS -XX:MaxMetaspaceSize=512m -Xms128m -Xmx2048m"
```



## Important

If you use Java 7, change `-XX:MaxMetaspaceSize` to `-XX:MaxPermSize`.

3. Go to the download section of your [CloverETL account](#) and download the `clover.war` (web archive) file containing CloverETL Server for JBoss EAP.
4. Configure the database connection.

By default, CloverETL Server uses an embedded Derby database; however, such setup is not supported for production use.

You can use the database connection provided by JNDI-bound datasource deployed by JBoss EAP. In order to define the datasource, edit the file:

```
[JBoss_EAP_home]/standalone/configuration/standalone.xml
```

and add the definition of the datasource into the section `<subsystem xmlns="urn:jboss:domain:datasources:1.1">` under the element `<datasources>`. Here is an example of datasource connecting to a MySQL database:

```
<datasource jndi-name="java:jboss/datasources/CloverETLServerDS"
  pool-name="CloverETLServerDS-Pool" enabled="true">
  <connection-url>jdbc:mysql://localhost:3307/cloverServerDB</connection-url>
  <driver>com.mysql</driver>
  <transaction-isolation>TRANSACTION_READ_COMMITTED</transaction-isolation>
  <pool>
    <min-pool-size>5</min-pool-size>
    <max-pool-size>50</max-pool-size>
    <prefill>true</prefill>
  </pool>
  <security>
    <user-name>clover</user-name>
    <password>SecretPassword</password>
  </security>
  <statement>
    <prepared-statement-cache-size>32</prepared-statement-cache-size>
    <share-prepared-statements>true</share-prepared-statements>
  </statement>
</datasource>
<drivers>
  <driver name="com.mysql" module="mysql.driver">
    <driver-class>com.mysql.jdbc.Driver</driver-class>
  </driver>
</drivers>
```

5. The datasource definition references a module (`mysql.driver`) with the MySQL JDBC driver. Take the following steps to add the module:



## Note

Under JBoss EAP, there are more options to set up **CloverETL Server's** database: along with JNDI-bound data source, it is possible to use the embedded Derby database or other supported database system specified in the **CloverETL** configuration file.

In order to be able to connect to the database, you need to define a global module so that the driver is available for the CloverETL web application - copying the driver to the `lib/ext` directory of the Server will **not** work. Such module is created and deployed in few steps (the example is for MySQL and module's name is `mysql.driver`):

- a. Create a directory `[JBoss_EAP_home]/modules/mysql/driver/main` (note that the directory path corresponds to module name `mysql.driver`)
- b. Copy the driver `mysql-connector-java-5.1.5-bin.jar` to the directory and create a file `module.xml` there with the following content:

```
<?xml version="1.0" encoding="UTF-8"?>
<module xmlns="urn:jboss:module:1.1" name="mysql.driver">
  <resources>
    <resource-root path="mysql-connector-java-5.1.5-bin.jar" />
  </resources>
  <dependencies>
    <module name="javax.api" />
  </dependencies>
</module>
```

- c. Add the module to global server modules: in case of the standalone JBoss EAP server they are defined in `[JBoss_EAP_home]/standalone/configuration/standalone.xml`. Add the module to the EE domain subsystem section:

```
<subsystem xmlns="urn:jboss:domain:ee:1.1">
  <global-modules>
    <module name="mysql.driver" slot="main" />
  </global-modules>
  <spec-descriptor-property-replacement>false</spec-descriptor-property-replacement>
  <jboss-descriptor-property-replacement>true</jboss-descriptor-property-replacement>
</subsystem>
```

6. Configure CloverETL Server according to a description in the [next section](#) (p. 45).

7. Deploy WAR file.

Copy the `clover.war` file to `[JBoss_EAP_home]/standalone/deployments`.

8. To start the JBoss platform:

- **Unix-like systems:**

Run `[JBoss_EAP_home]/bin/standalone.sh`.

- **Windows system:**

Run `[JBoss_EAP_home]\bin\standalone.bat`.

It may take a couple of minutes for all applications to start.

9. Check the JBoss response and CloverETL Server response.

- By default, the JBoss administration console is accessible at <http://localhost:8080/>. The default username/password is `admin/admin`

- By default, CloverETL Server is accessible at <http://localhost:8080/clover>.

## Configuration of CloverETL Server on JBoss EAP



### Tip

The default installation (without any configuration) is only recommended for evaluation purposes. For production use, at least a dedicated, system database and SMTP server configuration is recommended.

For an easy configuration of CloverETL Server, use a Setup GUI (p. 74) in which you can configure various properties, including the connection to the database, username and password, path to the license file, private properties, clusters and much more (see Chapter 14, [List of Configuration Properties](#) (p. 106) and Chapter 40, [Cluster Configuration](#) (p. 325)). We recommend you place the file in a specified (p. 70) location and specify the path to the file with a system property.

### Properties File in Specified Location

The properties file is loaded from a location which is specified by the environment/system property `clover_config_file` or `clover.config.file`.

1. • Create the `cloverServer.properties` file in a directory readable by JBoss EAP. (If you need an example of connection to any of the supported database systems, see Chapter 13, [System Database Configuration](#) (p. 83).):

```
datasource.type=JNDI
datasource.jndiName=java:jboss/datasources/CloverETLServerDS
jdbc.dialect=org.hibernate.dialect.MySQLDialect
license.file=/home/clover/config/license.dat
```

Do not forget to set a correct JDBC dialect according to your database server (Part III, [“Configuration”](#) (p. 67)). You can set the path to the license file, too.

- Alternatively, you can set "JDBC" `datasource.type` and configure the database connection to be managed directly by CloverETL Server (provided that you have deployed proper JDBC driver module to the Server):

```
datasource.type=JDBC
jdbc.url=jdbc:mysql://localhost:3306/cloverServerDB
jdbc.dialect=org.hibernate.dialect.MySQLDialect
jdbc.driverClassName=com.mysql.jdbc.Driver
jdbc.username=clover
jdbc.password=SecretPassword
license.file=/home/clover/config/license.dat
```

2. Set the `clover.config.file` system property (or environment variable).

It should contain the full path to the `cloverServer.properties` file created in the previous step.

The simplest way to set the system property is to edit the configuration file `[JBoss_EAP_home]/standalone/configuration/standalone.xml`, and to add the following snippet just under the `<extensions>` section:

```
<system-properties>
  <property name="clover.config.file" value="C:/jboss-eap-6.2/cloverServer.properties" />
</system-properties>
```

3. Restart JBoss EAP for the changes to take effect.
4. Check the CloverETL Server application is running:

By default, the Server's console is accessible at <http://localhost:8080/clover>.



### Note

By default, JBoss EAP has enabled HTTP session replication. This requires session serialization that is not supported by CloverETL Server and produces lots of harmless errors in JBoss's console, see below for example:

```
10:56:38,248 ERROR [org.infinispan.transaction.TransactionCoordinator] (http-/127.0.0.1:8080-2)
ISPN000188: Error while processing a commit in a two-phase transaction:
java.lang.UnsupportedOperationException: Serialization of HTTP session objects is not supported
by CloverETL Server - disable the session passivation/replication for this web application.
    at com.cloveretl.server.web.gui.e.writeExternal(Unknown Source) [cs.jar:]
    at org.jboss.marshalling.river.RiverMarshaller.doWriteObject(RiverMarshaller.java:874)
```

To eliminate these errors, disable the session replication. Edit [jboss-home]/standalone/configuration/standalone.xml and comment out the whole `<cache-container name="web" aliases="standard-session-cache">` block under `<subsystem xmlns="urn:jboss:domain:infinispan:1.5">` section.



### Note

➔ **Continue with:** Chapter 7, [Postinstallation Configuration](#) (p. 60)

## Oracle WebLogic Server

---

[Installation of Oracle Weblogic](#) (p. 47)

[Installation of CloverETL Server](#) (p. 48)

[Configuration of CloverETL Server on WebLogic](#) (p. 49)



### Important

See [Application Server](#) (p. 9) in system requirements for currently supported **Oracle WebLogic** versions and required **Java** versions.

If you encounter any problems during the installation, the [Possible Issues during Installation](#) (p. 56) section may provide a solution.

## Installation of Oracle Weblogic

1. Create an Oracle account on <http://www.oracle.com>.
2. Go to the [download page](#), accept the license agreement, download a compatible version of Weblogic server and extract the archive.
3. Set up a domain (the following steps are similar in Windows, simply run the `.cmd` files instead `.sh`).

- Run `[Weblogic_home]/configure.sh`.

After all files are unpacked and the environment is set, you can configure your server domain. Alternatively, you can create and configure it in more details in the following step:

- Run `[Weblogic_home]/oracle_common/bin/config.sh`.

In the installer, you can create and configure your domain, administrator password and other parameters.

4. Start the server.
  - **Unix-like systems:** Run `[Weblogic_home]/user_projects/your_domain_name/startWeblogic.sh`.
  - **Windows system:** Run `[Weblogic_home]\user_projects\your_domain_name\startWeblogic.cmd`.

5. Launch the Administration Console (default URL: `http://localhost:7001/console/`).

You should see the following welcome page:

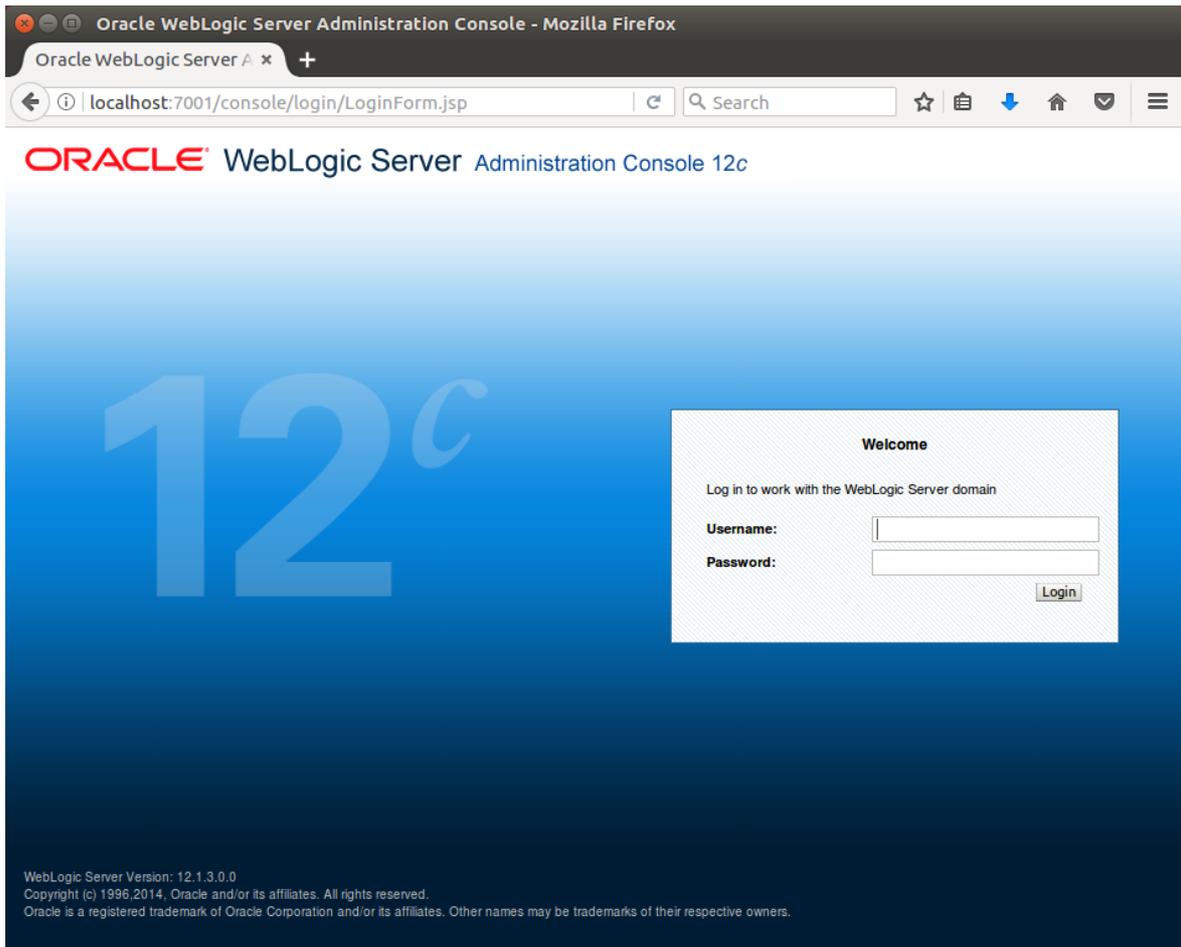


Figure 6.9. WebLogic welcome page

## Installation of CloverETL Server

1. Check if you meet the prerequisites:

- Oracle JDK or JRE is installed (see [Java Virtual Machine](#) (p. 9) for the required Java version).
- `JAVA_HOME` or `JRE_HOME` environment variable is set (see [Setting up JAVA\\_HOME](#) (p. 18)).
- A supported version (p. 10) of Oracle WebLogic Server is installed.

CloverETL Server is developed and tested with the WebLogic Server 11g (10.3.6) and WebLogic Server 12c (12.1.2) containers. Running the Server with other versions may result in unpredictable behavior.

WebLogic has to be running and a domain has to be configured. You can check it by connecting to **Administration Console**: `http://localhost:7001/console/`. **Username** and **password** are specified during installation.

2. It is strongly recommended to adjust the default limits for **Memory allocation** (see the [Memory Settings](#) (p. 60) section).

You can set the **initial** and **maximum memory heap size** by adjusting the "Xms" and "Xmx" JVM parameters and **classloaders memory limit** by adjusting the "XX:MaxMetaspaceSize" parameter:

- **Unix-like systems:**

Edit the start script and add:

```
export JAVA_OPTIONS='$JAVA_OPTIONS -Xms512m -Xmx2048m -XX:MaxMetaspaceSize=512m'
```

- **Windows system:**

See [WebLogic Server Performance and Tuning](#).



### Important

If you use Java 7, change `-XX:MaxMetaspaceSize` to `-XX:MaxPermSize`.

3. Go to the download section of your [CloverETL account](#) and download the `clover.war` (web archive) file containing CloverETL Server for Oracle WebLogic Server.

4. Change HTTP Basic Authentication configuration

- When WebLogic finds an "Authentication" header in an HTTP request, it tries to find a user in its own realm. This behavior has to be disabled so CloverETL could authenticate users itself.
- Edit the configuration file `[domainHome]/config/config.xml` and add:

```
<enforce-valid-basic-auth-credentials>false</enforce-valid-basic-auth-credentials>
```

into `<security-configuration>` element (just before the end tag).

5. Deploy `clover.war` (or an application directory).

Use the **WebLogic Server Administration Console**. See the [Oracle Fusion Middleware Administrator's Guide](#) for details.

6. Configure a license and other properties. See [Configuration of CloverETL Server on WebLogic](#)(p. 49) for details.

7. Check whether CloverETL Server is running.

- Web-app is started automatically after deployment, so you can check whether it is up and running.
- By default, CloverETL Server is accessible at <http://host:7001/clover>.

## Configuration of CloverETL Server on WebLogic



### Tip

The default installation (without any configuration) is only recommended for evaluation purposes. For production use, at least a dedicated, system database and SMTP server configuration is recommended.

For an easy configuration of CloverETL Server, use a [Setup GUI \(p. 74\)](#) in which you can configure various properties, including the connection to the database, username and password, path to the license file, private properties, clusters and much more (see Chapter 14, [List of Configuration Properties](#)(p. 106) and Chapter 40, [Cluster Configuration](#) (p. 325)). We recommend you place the file in a [specified \(p. 70\)](#) location and specify the path to the file with a system property.

The content of such a file (example with a PostgreSQL database):

```
datasource.type=JDBC
jdbc.driverClassName=org.postgresql.Driver
jdbc.url=jdbc:postgresql://127.0.0.1/clover_db?charSet=UTF-8
jdbc.username=yourUsername
jdbc.password=yourPassword
```

```
jdbc.dialect=org.hibernate.dialect.PostgreSQLDialect
```

### Properties File in Specified Location

1. Create the `cloverServer.properties` file in a directory readable by WebLogic. (If you need an example of connection to any of the supported database systems, see Chapter 13, [System Database Configuration](#) (p. 83).)

The configuration file should contain DB datasource configuration, SMTP connection configuration, etc. For details, see Part III, “[Configuration](#)” (p. 67).

2. Set the `clover_config_file` system property (or environment variable) pointing to the configuration properties file.

- Set the `JAVA_OPTIONS` variable in the WebLogic domain start script `[domainHome]/startWebLogic.sh`

```
JAVA_OPTIONS="${JAVA_OPTIONS} -Dclover_config_file=/path/to/clover-config.properties
```

3. Restart WebLogic for changes to take effect.



### Important

When CloverETL Server is deployed on WebLogic and JNDI Datasource pointing to Oracle DB is used, there must be an extra configuration property in the configuration file:

```
quartz.driverDelegateClass=org.quartz.impl.jdbcjobstore.oracle.weblogic.WebLogicOracleDelegate
```



### Note

➔ **Continue with:** Chapter 7, [Postinstallation Configuration](#) (p. 60)

## Activation

To be able to execute graphs, CloverETL Server requires a valid license. You can install and run CloverETL Server without any license, but no graph will be executed.

There are three ways of installing the license. They work on all application servers and can be used at the same time, but **only the most recent valid license is used**.

We recommend using the first and easiest option (for other options, see [CloverETL Server Activation Alternatives](#) (p. 53)):

### CloverETL Server Activation using Web Form

If the CloverETL Server has been started without assigning any license, click the **Activate server** link on the welcome page. You will be redirected to the **Add New License** form where you can upload the license file using the **Browse** button, or simply copy the license from the file and paste it into the **License key text** field.

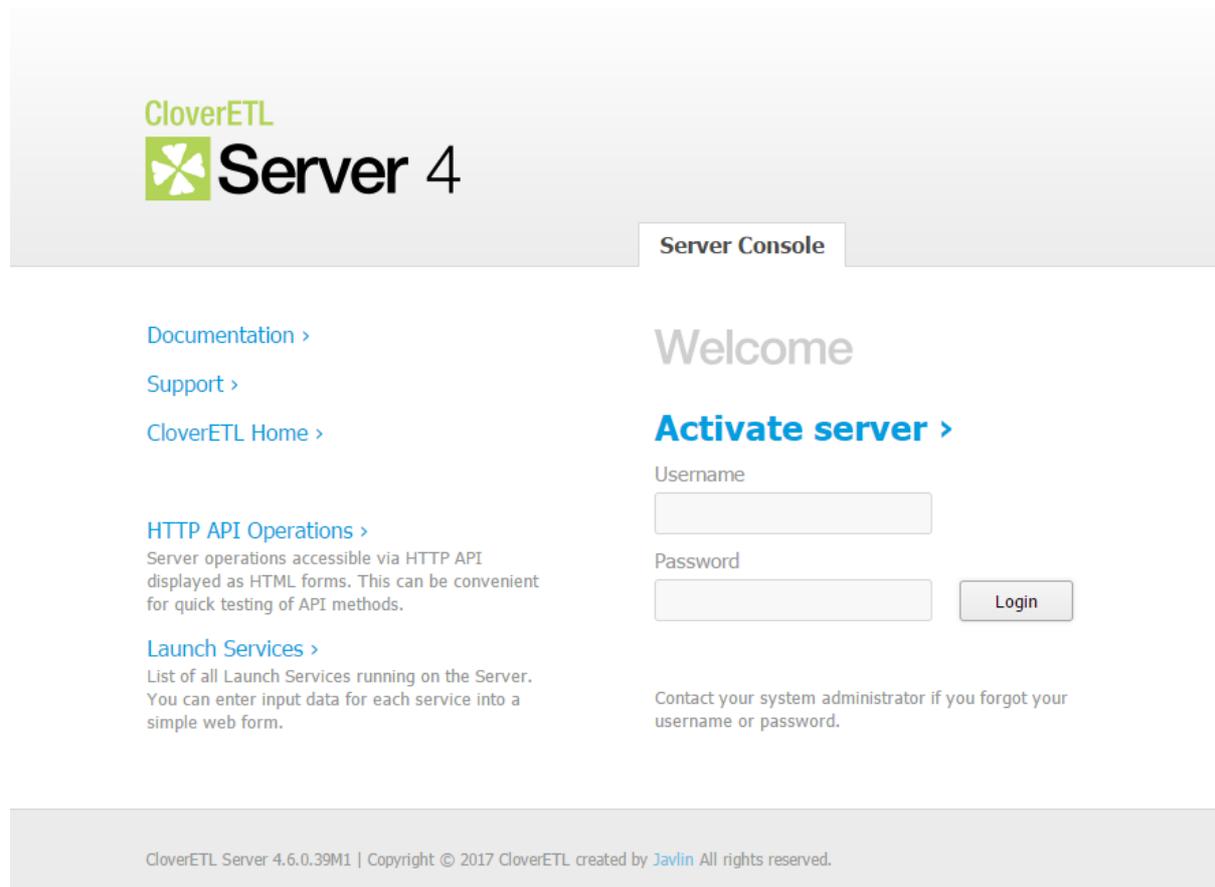
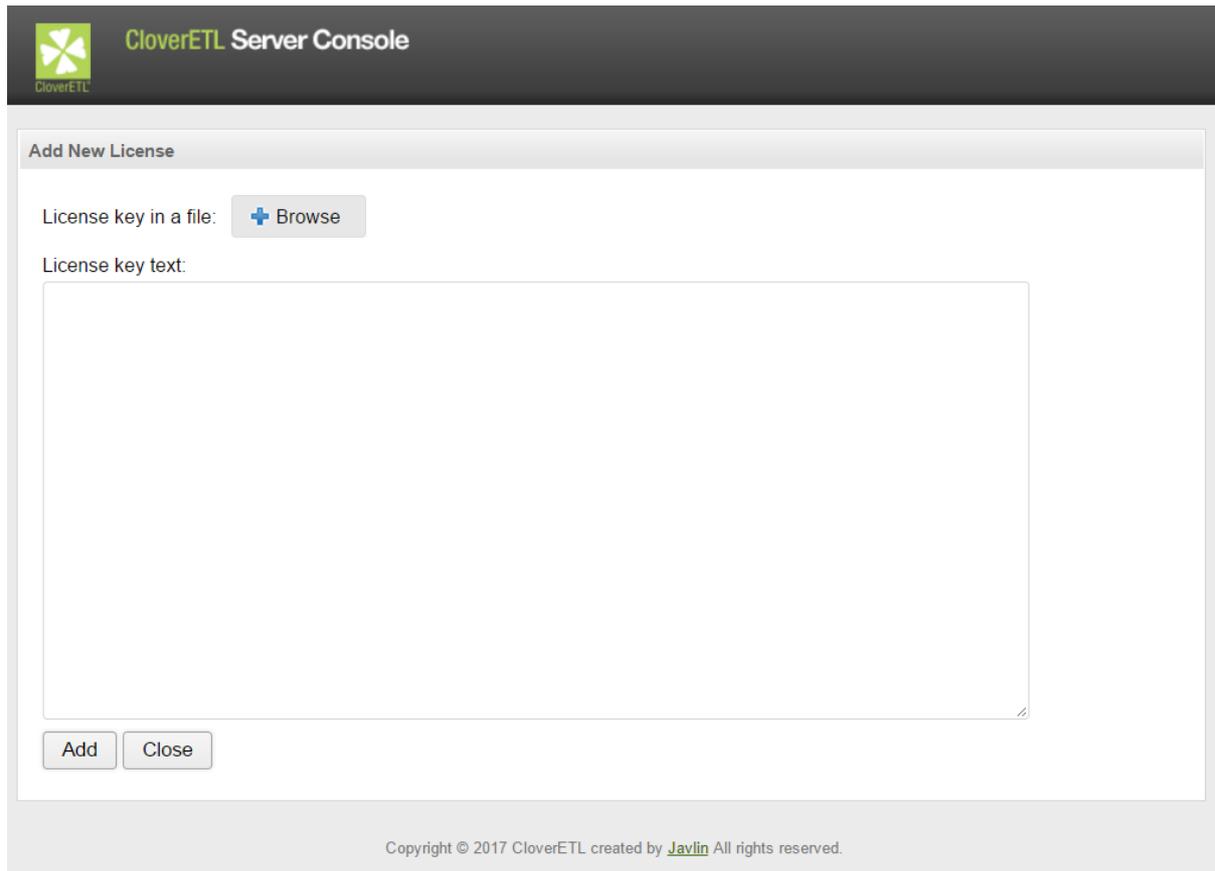


Figure 6.10. Login page of CloverETL Server without license

After clicking the **Update** button, the license is validated and saved to the database table `clover_licenses`. If the license is valid, a table with license's description appears. To proceed to CloverETL Server console click **Continue to server console**.

You can skip adding a license by clicking the **Close** button.



The screenshot shows the 'Add New License' form in the CloverETL Server Console. The form has a dark header with the CloverETL logo and the text 'CloverETL Server Console'. Below the header, the form is titled 'Add New License'. It contains two main input areas: 'License key in a file:' with a '+ Browse' button, and 'License key text:' with a large text area. At the bottom of the form, there are two buttons: 'Add' and 'Close'. A copyright notice is visible at the bottom of the console window: 'Copyright © 2017 CloverETL created by Javlin All rights reserved.'

Figure 6.11. Add new license form

## Updating CloverETL Server License in the Configuration Section

If the license has been already installed, you can still change it by using form in the Server web GUI.

- Go to **server web GUI** → **Configuration** → **Setup** → **License**
- Click **Update license**.

You can paste a license text into a **License key** text area or use the **Browse** button to search for a license file in the filesystem. To skip adding a license, click the **Close** button.

After clicking the **Update** button, the license is saved to the database table `clover_licenses` and reloaded.

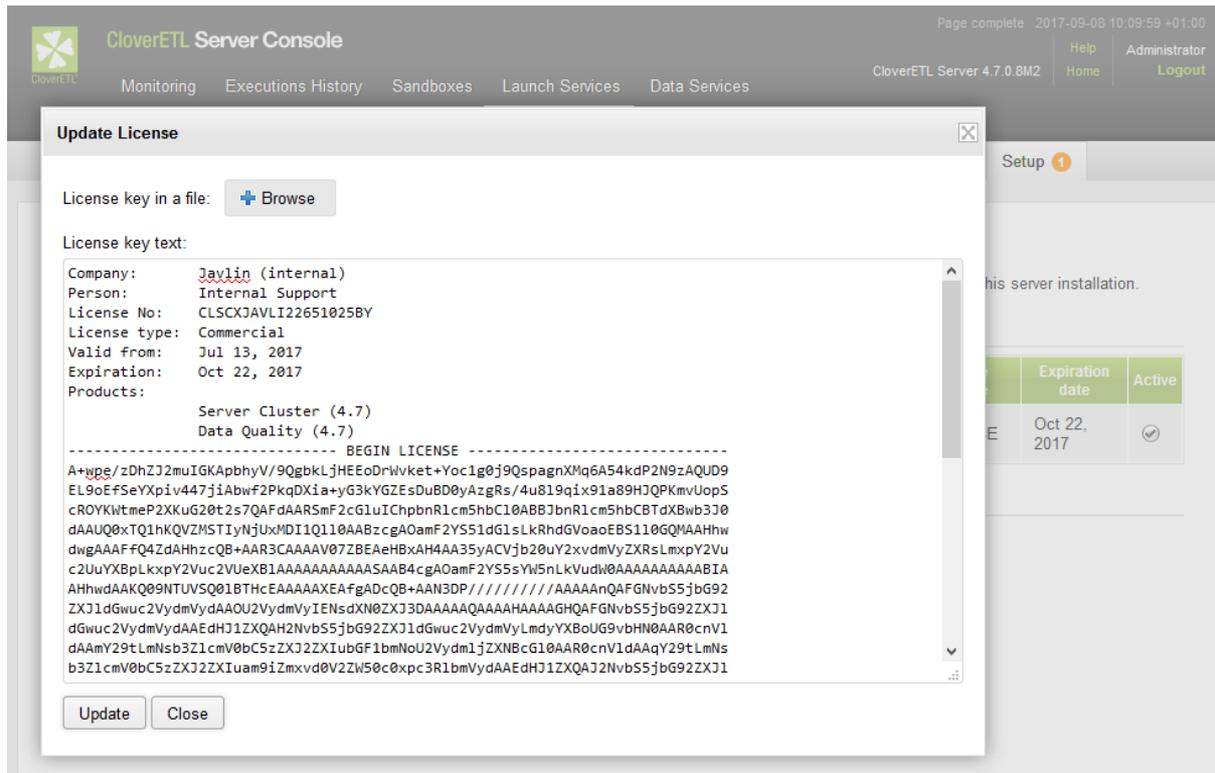


Figure 6.12. Update license form



### Tip

The CloverETL license can be **changed** at any time by replacing `license.dat` file. Afterwards, you have to let CloverETL Server know the license has changed.

- Go to **server web GUI** → **Configuration** → **Setup** → **License**
- Click **Reload license**.
- Alternatively, you can restart the CloverETL Server application.



### Note

The license in the database is common for all nodes in the cluster. Reloading of the license occurs on each node in the cluster.

➔ **Continue with:** Chapter 12, [Setup](#) (p. 74)

## CloverETL Server Activation Alternatives

If, for any reason, you decide to not use the recommended way of installing the server license, you can choose one of the following options:

### Activation Using `license.file` Property

1. Download the `license.dat` file.
2. Set the CloverETL Server `license.file` parameter to the full path of the `license.dat` file.

For list of properties, see Chapter 14, [List of Configuration Properties](#) (p. 106).

- Restart the application server.

## Separate License WAR

A simple approach, but it may be used only for a standalone Server running on Apache Tomcat.

- Download the `clover-license.war` web archive file.
- Copy `clover-license.war` to the `[tomcat_home]/webapps` directory.
- The war file should be detected and deployed automatically without restarting Tomcat.
- Check whether the license web-app is running on:

`http://[host]:[port]/clover-license/` (Note: `clover-license` contextPath is mandatory and cannot be changed)



### Note



Continue with: Chapter 12, [Setup](#) (p. 74)

---

## IBM InfoSphere MDM Plugin Installation

### Downloading

**IBM InfoSphere MDM Components for CloverETL Server** are downloaded as a ZIP file containing the extension. The ZIP file is available for download under your account on [www.cloveretl.com](http://www.cloveretl.com) in **CloverETL Server** download area, under the **Utilities** section as the `ibm-mdm-connectors.${version}.zip` file.

### Requirements

Requirements of **IBM InfoSphere MDM Components**:

- supported OS are Microsoft Windows 32 bit, Microsoft Windows 64 bit, Linux 64 bit, and Mac OS X Cocoa
- at least 512MB of RAM
- installed CloverETL Server

The support for 32 bit Linux was removed in 4.5.0.

### Installation into Server

The following steps are needed to install **IBM InfoSphere MDM Components** into **CloverETL Server**:

- Install **CloverETL Server**, see its documentation for details.
- Download the ZIP file with **IBM InfoSphere MDM Components** for the Server and store it on the system where **CloverETL Server** is installed. For the download instructions, see [Downloading](#) (p. 54).
- The ZIP file contains a **CloverETL** plugin. Your Server installation needs to be configured to find and load the plugin from the ZIP file. This is done by setting the `engine.plugins.additional.src` Server configuration property to the absolute path of the ZIP file, e.g. `engine.plugins.additional.src=c:/Server/ibm-mdm-connectors.4.9.0.zip` (in case the Server is configured via property file).

Details for setting the configuration property depend on your Server installation specifics, application server used, etc. See **CloverETL Server** documentation for details. Typically the property would be set similarly to how you set-up the properties for connection to the Server's database. Updating the configuration property usually requires restart of the Server.

4. To verify that the plugin was loaded successfully, log into the Server's **Reporting Console** and look in the **Configuration** → **CloverETL Info** → **Plugins** page. In the list of plugins, you should see `cloveretl.engine.initiate`.

## Troubleshooting

---

If you get an `Unknown component` or `Unknown connection` error when running a graph with IBM InfoSphere MDM components, it means that the **IBM InfoSphere MDM Components** plugin was not loaded by the Server successfully. Please check the above steps to install the plugin, especially the path to the ZIP file.

---

## Possible Issues during Installation

Since CloverETL Server is considered a universal JEE application running on various application servers, databases and JVM implementations, problems may occur during the installation. These can be solved with a proper configuration of the Server environment. This section contains tips for the configuration.

[Memory Issues on Derby](#) (p. 56)

[JAVA\\_HOME or JRE\\_HOME Environment Variables Are Not Defined](#) (p. 56)

[Apache Tomcat Context Parameters Do Not Have Any Effect](#) (p. 57)

[Tomcat Log File catalina.out Is Missing on Windows](#) (p. 57)

[Timeouts Waiting for JVM](#) (p. 57)

[clover.war as Default Context on WebSphere \(Windows OS\)](#) (p. 58)

[Tomcat 6.0 on Linux - Default DB](#) (p. 58)

[Derby.system.home Cannot be Accessed](#) (p. 58)

[Environment Variables and More than one CloverETL Server Instances Running on Single Machine](#) (p. 59)

[Special Characters and Slashes in Path](#) (p. 59)

[File System Permissions](#) (p. 59)

[JMS API and JMS Third-Party Libraries](#) (p. 59)

[Using an Unsupported JDBC Connector for MySQL](#) (p. 59)

---

## Memory Issues on Derby

If your Server suddenly starts consuming too much resources (CPU, memory) despite having been working well before, it might be caused by a running internal Derby DB. Typically, causes are incorrect/incomplete shutdown of Apache Tomcat and parallel (re)start of Apache Tomcat.

Solution: move to a standard (standalone) database.

How to fix this? Redeploy CloverETL Server:

1. Stop Apache Tomcat and verify there are no other instances running. If so, kill them.
2. Backup the configuration file, if you configured any.
3. Delete the `webapps/clover` directory.
4. Start the Apache Tomcat server. It will automatically redeploy CloverETL Server.
5. Verify you can connect from Designer and from web.
6. Shutdown Apache Tomcat.
7. Restore the configuration file and point it to your regular database.
8. Start Apache Tomcat.

---

## JAVA\_HOME or JRE\_HOME Environment Variables Are Not Defined

If you are getting this error message during an attempt to start your application server (mostly Tomcat), perform the following actions.

### Linux:

This command will help you set a path to the variable on the server.

```
[clover@server /] export JAVA_HOME=/usr/local/jdk1.x.x
```

As a final step, restart the application server.

#### Windows OS:

Set `JAVA_HOME` to your JDK installation directory, e.g. `C:\Program Files\java\jdk1.8.0`.



### Important

Some CloverETL functions requires JDK to work correctly, therefore we do not recommend having only JRE installed.

---

## Apache Tomcat Context Parameters Do Not Have Any Effect

Tomcat may sometimes ignore some context parameters. It may cause strange CloverETL Server behavior, since it appears as configured, but only partially. Some parameters are accepted, some are ignored. This issue is rare, however it may occur in some environments. Such behavior is consistent, so restart has no effect. It's possibly related to Tomcat issues: [Bug #47516](#) and [Bug #50700](#) To avoid this, please use a properties file instead of context parameters to configure CloverETL Server.

---

## Tomcat Log File `catalina.out` Is Missing on Windows

Tomcat start batch files for Windows aren't configured to create the `catalina.out` file which contains the standard output of the application. The `catalina.out` file may be vital when Tomcat isn't started in the console and an issue occurs. Or even when Tomcat is executed in the console, it may be closed automatically just after the error message appears in it.

Please follow these steps to enable `catalina.out` creation:

- Modify `[Tomcat_home]/bin/catalina.bat` and add a parameter `/B` to the lines where the `_EXECJAVA` variable is set. There should be two such lines:

```
set _EXECJAVA=start /B [the rest of the line]
```

Parameter `/B` causes, that "start" command doesn't open a new console window, but runs the command in its own console window.

- Create a new startup file, e.g. `[Tomcat_home]/bin/startupLog.bat`, containing a single line:

```
catalina.bat start > ..\logs\catalina.out 2<&1
```

It executes Tomcat in the usual way, but the standard output isn't put to the console, but to the `catalina.out` file.

Then use the new startup file instead of `[Tomcat_home]/bin/startup.bat`.

---

## Timeouts Waiting for JVM

If you get the Jetty application server successfully running but cannot start CloverETL Server, it might be because of the wrapper waiting for JVM too long (it is considered a low-memory issue). Examine `[Jetty_home]\logs\jetty-service.log` for the following line:

```
Startup failed: Timed out waiting for signal from JVM.
```

If it is there, edit `[Jetty_home]\bin\jetty-service.conf` and add these lines:

```
wrapper.startup.timeout=60
wrapper.shutdown.timeout=60
```

If that does not help either, try setting 120 for both values. Default timeouts are 30.

## clover.war as Default Context on WebSphere (Windows OS)

If you are deploying `clover.war` on the IBM WebSphere server without the context path specified, be sure to check whether it is the only application running in the context root. If you cannot start CloverETL Server on WebSphere, check the log and look for a following message:

```
com.ibm.ws.webcontainer.exception.WebAppNotLoadedException:
Failed to load webapp: Failed to load webapp: Context root /* is already bound.
Cannot start application CloverETL
```

If you can see it, then this is the case. The easiest way to fix the issue is to stop all other (sample) applications and leave only `clover.war` running on the server. That should guarantee the server will be available in the context root from now on (e.g. `http://localhost:9080/`).

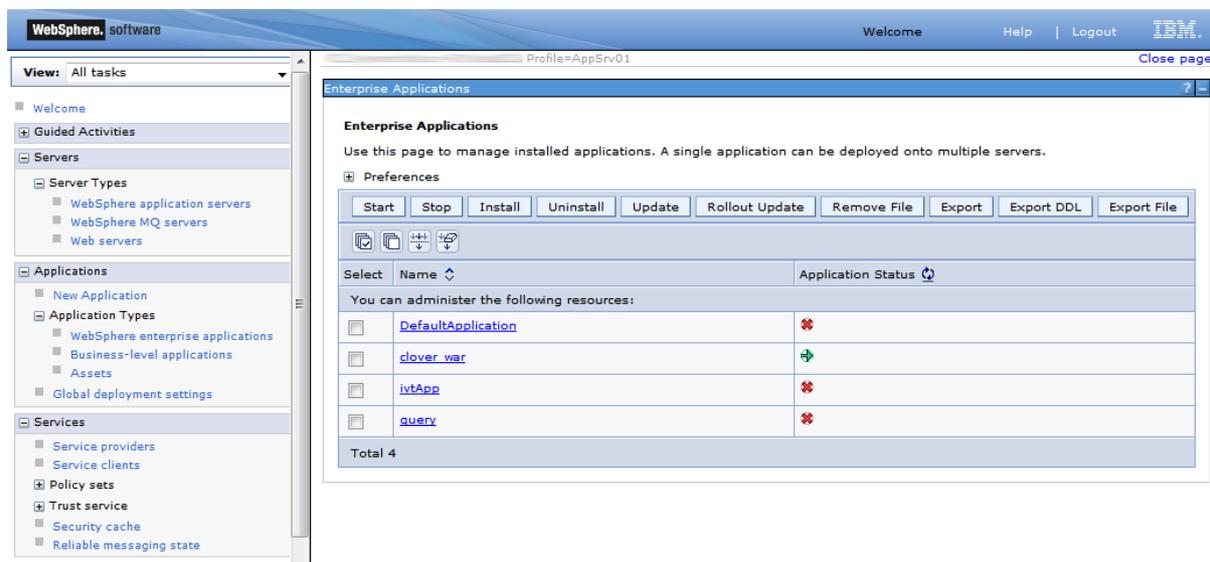


Figure 6.13. CloverETL Server as the only running application on IBM WebSphere

## Tomcat 6.0 on Linux - Default DB

When using the internal (default) database on Linux, your CloverETL Server might fail on the first start for no obvious reasons. Chances are that the `/var/lib/tomcat6/databases` directory was not created (because of access rights in parent folders).

Solution: Create the directory yourself and try restarting the Server. This simple fix was successfully tested with CloverETL Server deployed as a WAR file via the Tomcat web administration tool.

## Derby.system.home Cannot be Accessed

If the Server cannot start and the following message is in the log:

```
java.sql.SQLException: Failed to start database 'databases/cloverserver'
```

then see the next exception for details. After that, check settings of the `derby.system.home` system property. It may point to an unaccessible directory, or files may be locked by another process. We suggest you set a specific directory as the system property.

## Environment Variables and More than one CloverETL Server Instances Running on Single Machine

If you are setting environment variables like `clover_license_file` or `clover_config_file`, remember you should not be running more than one CloverETL Server. Therefore, if you ever need to run more instances at once, use other ways of setting parameters (see Part III, “[Configuration](#)” (p. 67) for description of all possibilities). The reason is the environment variables are shared by all applications in use causing them to share configurations and fail unexpectedly. Instead of the environment variables, you can use system properties (passed to the application container process using parameter with `-D` prefix: `-Dclover_config_file`).

## Special Characters and Slashes in Path

When working with servers, be sure to follow the folder naming rules. Do not use any special characters in the server path, e.g. spaces, accents, diacritics are not recommended. It can produce issues which are hard to find. If you are experiencing weird errors and cannot trace the source of them, install the application server in a safe destination like:

```
C:\JBoss6\
```

Similarly, use slashes but never backslashes in paths inside the `*.properties` files, e.g. when pointing to the CloverETL Server license file. If you incorrectly use a backslash, it will be considered an escape character and the server may not work properly. This is an example of a correct path:

```
license.file=C:/CoverETL/Server/license.dat
```

## File System Permissions

The application server must be executed by an OS user with proper read/write permissions on file system. Problem may occur, if `app-server` is executed by a root user for the first time, so log and other temp files are created by root user. When the same `app-server` is executed by another user, it will fail because it cannot write to root's files.

## JMS API and JMS Third-Party Libraries

Missing JMS libraries do not cause failure of the server startup, but it is an issue of deployment on an application server, thus it is still related to this chapter.

`clover.war` itself does not contain `jms.jar`, so it has to be on an application server's classpath. Most of the application servers have `jms.jar` by default, but Tomcat, for example, does not. So if the JMS features are needed, the `jms.jar` has to be added explicitly.

If the "JMS Task" feature is used, there must be third-party libraries on a Server's classpath as well. The same approach is recommended for JMS Reader/Writer components, even if these components allow to specify external libraries. It is due to common memory leak in these libraries which causes "OutOfMemoryError: PermGen space".

## Using an Unsupported JDBC Connector for MySQL

CloverETL Server requires MySQL 5 up to version 5.5 included. Using an unsupported JDBC connector for MySQL might cause an exception, for example:

```
could not execute query
```

```
You have an error in your SQL syntax; check the manual that corresponds to your MySQL server version for the right syntax to use near 'OPTION SQL_SELECT_LIMIT=DEFAULT' at line 1
```

---

## Chapter 7. Postinstallation Configuration

[Memory Settings](#) (p. 60)

[Maximum Number of Open Files](#) (p. 62)

[Maximum Number of Processes or Threads](#) (p. 62)

[Firewall Exceptions](#) (p. 62)

---

### Memory Settings

Current implementation of Java Virtual Machine allows only a global configuration of memory for the JVM system process. Thus the whole application server, together with WARs and EARs running on it, share one memory space.

Default JVM memory settings is **too low** for running an application container with CloverETL Server. Some application servers, like IBM WebSphere, increase JVM defaults themselves, however they **still may be too low**.

The **optimal memory limits** depend on many conditions, i.e. transformations which CloverETL should execute. Please note that the maximum limit isn't the amount of permanently allocated memory, but limit which can't be exceeded. If the limit is exhausted, the `OutOfMemoryError` is raised.

#### JVM Memory Areas

JVM memory consists of several areas: **heap memory**, **PermGen space**, **direct memory** and **stack memory**. Since JVM memory is not just HEAP memory, you should not set the HEAP limit too high; in case it consumes whole RAM, JVM won't be able to allocate direct memory and stack for new threads.

*Table 7.1. JVM Memory Structure*

Type	Description
Heap memory	Heap is an area of memory used by JVM for dynamic memory allocation. Required heap memory size depends on various factors (e.g. complexity of graphs, number of graphs running in parallel, type of component, etc.), see the respective server container's installation guide in this documentation. (Note that current heap memory usage can be observed in CloverETL Server Console (p. 131).)
PermGen Space	Permanent Generation - separate memory space containing class definitions and related metadata. (PermGen was removed from Java 8.)
Direct Memory	Memory used by graph edges and buffers for I/O operations.
Stack Memory	Stack Memory contains local, method specific variables and references to other objects in the method. Each thread has its own stack; therefore, the memory usage depends on the number of components running in parallel.

#### Configuring Memory

You can set the minimum and maximum memory heap size by adjusting the "Xms" and "Xmx" JVM parameters. There are more ways to change the settings depending on the used application container.

#### Recommended Server Core and Worker Heap Memory Configuration

Optimal distribution of main memory between Server Core and Worker depends on nature of executed tasks. The recommended defaults of Server Core heap size and Worker heap size for different RAM sizes are in the table below.

Heap limit is *not* a limit of the full memory used by JVM. JVM uses memory in addition to the heap size for other memory spaces, e.g. direct memory. We recommend to set the heap limit to no more than 80% of system memory size, to leave space for the operating system and other JVM memory spaces.

Table 7.2.

RAM Size	Server Core Heap	Worker Heap	Remaining for OS (estimated)
4 GB	1 GB	1 GB	1 GB
8 GB	2-3 GB	2-3 GB	2 GB
16 GB	4 - 8 GB	4 - 8 GB	4 GB
32 GB	4 - 8 GB	16 - 20 GB	8 GB
64 GB	4 - 8 GB	42 - 52 GB	8 GB

### Memory Configuration in Different Java Versions

In Java 7 and earlier, the memory space for loading classes (so called "PermGen space") is separated from the heap memory, and can be set by the JVM parameter `-XX:MaxPermSize`. By default, it is just 64 MB which is not enough for enterprise applications. Again, suitable memory limit depends on various criteria, but 512MB should be enough in most cases. If the PermGen space maximum is too low, `OutOfMemoryError: PermGen space` may occur.

In Java 8, the memory space for loading classes (so called "Metaspace") is separated from heap, and can be set by the JVM parameter `-XX:MaxMetaspaceSize`. The default maximum Metaspace size is unlimited.

Please see the specific container section for details on memory settings.

#### Metaspace

We recommend you to put limit on metaspace memory. Add `-XX:MaxMetaspaceSize=size` to command line parameters of Server Core or Worker. Replace `size` with a suitably high limit.

See <https://docs.oracle.com/javase/8/docs/technotes/tools/unix/java.html>

#### Direct memory

To avoid excessive usage of direct memory, we add `-Djdk.nio.maxCachedBufferSize=262144` to command line of Worker. We recommend you to add this system property also to the command line of Server Core. This system property is available since java 1.8.0\_102.

In Apache Tomcat, add this system property to `JAVA_OPTS` environment variable, which is configured in `bin/setenv.sh` file.

### Codecache Size

Some CloverETL Server installations can occasionally run into performance issue: JVM is running more than hundred times slower. The issue can be caused by a full code cache ([Java SE Embedded: Developer's Guide - Codecache Tuning](#)). The reserved code cache size is platform dependent and can be too small for CloverETL Server. It is highly recommended to increase the code cache size using the following JVM argument:

```
-XX:ReservedCodeCacheSize=256m
```

## Maximum Number of Open Files

When using resource-demanding components, such as **FastSort**, or when running a large number of graphs concurrently, you may reach the system limit on simultaneously open files. This is usually indicated by the `java.io.IOException: Too many open files` exception.

The default limit is fairly low in many Linux distributions (e.g. 4096 in Ubuntu). Such a limit can be easily exceeded, considering that one FastSort component can open up to 1,000 files when sorting 10 million records. Furthermore, some application containers recommend increasing the limit themselves (8,192 for IBM WebSphere).

Therefore, it is recommended to increase the limit for production systems. Reasonable limits vary from 10,000 to about 100,000 depending on the expected load of **CloverETL Server** and the complexity of your graphs.

The current limit can be displayed in most UNIX-like systems using the `ulimit -Sn` command.

The exact way of increasing the limit is OS-specific and is beyond the scope of this manual.

## Maximum Number of Processes or Threads

If you run graphs with many subgraphs containing many components, you may reach the limit on number of threads per user or per system. In this case, you can find `java.lang.OutOfMemoryError: unable to create new native thread` in graph's log.

The current limit on number of processes/threads per user can be displayed in most UNIX-like systems using the `ulimit -Su` command. Note that the documentation on `ulimit` may not distinguish between processes and threads. The limit on number of threads per system can be displayed using the `sysctl kernel.threads-max` command. The exact way of increasing the limit is OS-specific and is beyond the scope of this manual.

## Firewall Exceptions

In order to function properly, CloverETL Server requires an outside communication. The table below describes both incoming and outgoing communication of CloverETL Server. Please, configure your firewall exceptions accordingly.

Table 7.3. Firewall Exceptions

Traffic	Communication	Description & Components
<b>Incoming</b>	HTTP(S)	Communication between Designer and Server
	JMX	Tracking and debugging information
<b>Outgoing (depending on an actual usage)</b>	JDBC	Connection to databases (DBInputTable, DBOutputTable, DBExecute)
	JMX	Receiving and sending JMS messages (JMSReader, JMSWriter, JMS Listener)
	HTTP(S)	Requesting and receiving responses from servers (Readers, WebserviceClient, HTTPConnector)
	SMTP	Sending data converted into emails (EmailSender)
	IMAP/ POP3	Receiving emails (EmailReader)
	FTP/ SFTP/ FTPS:	Remote file reading and writing (readers, writers)



**Note**

➔ **Continue with:** Chapter 13, [System Database Configuration](#) (p. 83)

---

## Chapter 8. Optional Installation Steps

This chapter describes optional installation steps for items not specified in the previous sections.

---

### Adding Libraries to the Worker's Classpath

Worker may need additional libraries, e.g. a JDBC driver library or the Bouncy Castle cryptographic library. There are two ways to add the libraries.

#### Directory in Application Server

Create a `worker-lib` directory in the `${user.data.home}/CloverETL/` directory and place the libraries there.

This approach works in Apache Tomcat, Glassfish and JBoss.

#### Configuration Property

Create a directory containing the libraries and set the `worker.classpath` configuration property to the path to this directory.

---

### Worker Support for SMB 2.x and 3.x

Worker supports the SMB 2.x and 3.x protocol. It utilizes the [SMBJ library](#) dependent on [Bouncy Castle](#).

Before you start using SMB 2.x/3.x & Bouncy Castle:

1. Go to the official [Latest Bouncy Castle Java Releases](#) page.
2. Locate the section "SIGNED JAR FILES" and download the latest release.

Download a required `.jar` file. The filename consists of the name `bcprov-jdk15on`, followed by a version number, for example:

```
bcprov-jdk15on-158.jar
```

3. a. **Recommended** - after you download the `.jar` file, you can point to the directory containing the file by adding the following line:

```
-Dcloveretl.smb2.bouncycastle.jar.file=path/to/bcprov-jdk15on.jar
```

to the **JVM arguments** field in the Worker (p. 78) tab of Setup GUI.

- b. **Alternatively** - create a `worker-lib` directory in the `${user.data.home}/CloverETL` directory and move the `jar` file there;
4. Now in case of using the recommended method, you only need to restart Worker. When using the alternative method, restart **CloverETL Server** for the changes to take effect.

---

# Chapter 9. Upgrading Server to Newer Version

## General Notes on Upgrade

- An upgrade of CloverETL Server requires down time; plan a maintenance window.
- A successful upgrade requires about 30 minutes; rollback requires 30 minutes.
- Perform the steps below in a development/testing environment first before moving onto a production environment.

## Upgrade Prerequisites

- Having a new CloverETL Server web application archive (`clover.war` appropriate for the application server used) & license files available.
- Having [release notes](#) for the particular CloverETL version available (and all versions between current and intended version to be upgraded to).
- Having the graphs and jobs updated and tested with regards to [Known Issues & Compatibility](#) for the particular CloverETL version.
- Having the CloverETL Server configuration properties file externalized from default location, see Chapter 11, [Configuration Sources](#) (p. 70).
- Standalone database schema where CloverETL Server stores configuration, see Chapter 13, [System Database Configuration](#) (p. 83).
- Having a separate sandbox with a test graph that can be run at any time to verify that CloverETL Server runs correctly and allows for running jobs.

## Upgrade Instructions

1. Suspend all sandboxes, wait for running graphs to finish processing.
2. Shutdown the CloverETL Server application (or all servers, if they run in a cluster mode).
3. Backup the existing CloverETL database schema (if any changes to the database schema are necessary, the new server will automatically make them when you start it for the first time).
4. Backup the existing CloverETL web application archive (`clover.war`) & license files (on all nodes).
5. Backup the existing CloverETL sandboxes (on all nodes).
6. Re-deploy the CloverETL Server web application. Instructions how to do that are application server dependent - see [Production Server](#) (p. 16) for installation details on all supported application servers. After the re-deployment, your new server will be configured based on the previous version's configuration.
7. Replace old license files by the valid one (or you can later use the web GUI form to upload new license). The license file is shipped as a text containing a unique set of characters. If you:
  - received the new license as a file (`*.dat`), then simply use it as new license file.
  - have been sent the license text, e.g. inside an email, then copy the license contents (i.e. all text between `Company` and `END LICENSE`) into a new file called `clover-license.dat`. Next, overwrite the old license file with the new one or upload it in the web GUI.

For details on license installation, see [Activation](#) (p. 51).

8. Start the CloverETL Server application (on all nodes).

9. Review that contents of all tabs in the CloverETL Server Console, especially scheduling and event listeners looks OK.
10. Update graphs to be compatible with the particular version of CloverETL Server (this should be prepared and tested in advance).
11. Resume the test sandbox and run a test graph to verify functionality.
12. Resume all sandboxes.

## Upgrade from 4.8.x or earlier to 4.9.x or later



### Significant architectural change in version 4.9.0

By default since 4.9.0, jobs (graphs, jobflow, data services) are executed in a standalone JVM called Worker (p. 6).

To run jobs in the Server Core (i.e. in the same way as in earlier versions of **CloverETL Server**), you can disable execution in Worker for particular jobs or sandboxes (p. 184) or disable Worker completely (p. 114).

**Note** that Launch Services and Profiler are only available in the Server Core. These features require no additional configuration.

### Configuration Changes

Worker is the executor of jobs, all jobs run in the Worker by default. It runs in a separate process (JVM), so it requires configuration in addition to the Server Core.

The Worker's configuration relates to memory size, classpath, command line options and JNDI. For an overview of Worker related configuration, see Introduction to Worker configuration (p. 68) The introduction provides an overview of the new Worker specific configuration with links to details.

### Changes to Jobs

If your jobs (ETL graphs, jobflows, etc.) use JNDI for database or JMS connections, you need to configure JNDI on Worker (see JNDI in Worker (p. 117)). You might also need to update your jobs to use the new JNDI resources configured in Worker - they might be available on new JNDI paths. If you do not need these JNDI resources on the Server Core anymore, consider removing them from the Server Core.

### Rollback Instructions

1. Shutdown the CloverETL Server application.
2. Restore the CloverETL Server web application (`clover.war`) & license files (on all nodes).
3. Restore the CloverETL Server database schema.
4. Restore the CloverETL sandboxes (on all nodes).
5. Start the CloverETL Server application (on all nodes).
6. Resume the test sandbox and run a test graph to verify functionality.
7. Resume all sandboxes.



### Important

**Evaluation Version** - a mere upgrade of your license is not sufficient. When moving from an evaluation to production server, you should not use the default configuration and database. Instead, take some time to configure CloverETL Server so that it best fits your production environment.

---

# Part III. Configuration

---

---

## Chapter 10. Introduction

This part describes in detail the configuration options for **CloverETL Server** used in a production environment. In the following chapters, you will find information on setting required properties and parameters, description of CloverETL Server's Setup GUI elements, parameters for specific database configuration, list of properties used in general configuration, instructions on encrypting confidential properties and log files setting.



### Note

We recommend the default installation (without any configuration) only for evaluation purposes. For production use, you should configure a dedicated, system database and set up an SMTP server for sending notifications.

## CloverETL Server Configuration Procedure

For initial configuration after the installation of **CloverETL Server**, follow these steps:

### Choose a configuration source

Choose a source of configuration data for **CloverETL Server**. There are several options for configuration sources; however, we recommend using a property file on specified location. For more information, see Chapter 11, [Configuration Sources](#) (p. 70).

### Set up a database dedicated to CloverETL Server

Now, you should set up a CloverETL Server's database and configure a connection to the database. Choose a supported [database system](#) (p. 10) and read Chapter 13, [System Database Configuration](#) (p. 83) for information and examples on how to create a database, add user/role for Clover, grant it required rights/privileges, etc.

Once you have set up the database, configure CloverETL Server's connection to the database using the [Setup GUI](#) (p. 74).

### Activate the server with a license

After you have set up the database, configured the connection to it and specified the source for configuration data, you can activate the Server with your license. While it is possible to activate the Server immediately after installation, we do not recommend this, since after the activation, the license information is stored in the database. To activate the Server, follow the information in the [Activation](#) (p. 51) section.

### Configure the server

Finally, you can configure the server features. The Setup with a user friendly GUI allows you to configure the basic, most important features including, encryption of sensitive data, SMTP for email notifications, etc. For more information, see Chapter 12, [Setup](#) (p. 74).

### Configure Worker

Worker is the executor of jobs, all jobs run in the Worker by default. It runs in a separate process (JVM), so it requires configuration in addition to the Server Core.

Basic configuration of the Worker can be done in the [Worker](#) (p. 78) page of the Setup (p. 74) For all configuration properties of Worker see [Worker - Configuration Properties](#) (p. 114). See the Chapter 24, [Troubleshooting Worker](#) (p. 196) section for useful tips on solving issues.

The following are the typical areas that need to be configured for Worker:

- **Heap memory size** - required heap size for Server Core and Worker depends on the nature of executed jobs. In general, Worker should have higher heap allocated, as it runs the jobs which represent the bulk of memory consumption.

See our recommendations (p. 60) for heap sizes of Worker and Server Core.

Heap size of Worker can be easily configured in the Worker (p. 78) tab of Setup (p. 74) or via the `worker.maxHeapSize` (p. 115) configuration property.

- **Classpath** - the Worker's classpath is separate from Server Core (i.e. application container classpath). Any libraries needed by jobs executed on Worker need to be added on the Worker's classpath.

See the `worker.classpath` (p. 114) configuration property for more details.

- **Command line options** - Worker is started as a separate process with its own JVM. If you need to set JVM command line options, e.g. for garbage collector tweaking, better diagnostics, etc., then you need to set them on the Worker's JVM. See Additional Diagnostic Tools (p. 196) section for useful options for troubleshooting and debugging Worker.

Command line options of Worker can be easily customized in the Worker (p. 78) tab of Setup (p. 74) or via the `worker.jvmOptions` (p. 115) configuration property.

- **JNDI** - Worker has its own JNDI pool separate from the application container JNDI pool. If your jobs use JNDI resources (to obtain JDBC or JMS connections), you have to configure the Worker's JNDI pool and its resources.

See [JNDI in worker](#) (p. 117) section for more details.

### Encrypt the configuration

As the last step, we strongly recommend you to encrypt the configuration file to protect your sensitive data.

For more information, see the [Encryption](#) (p. 79) section.

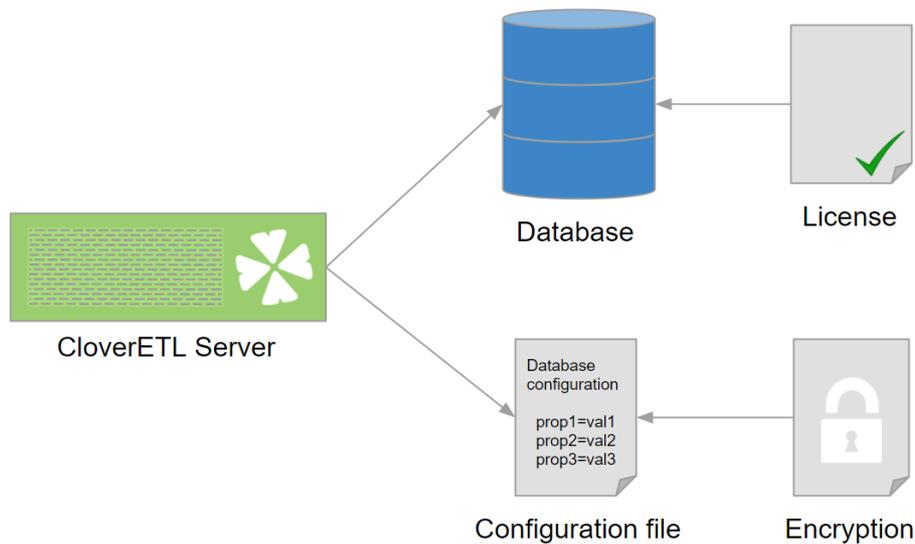


Figure 10.1. CloverETL Server's System Database Configuration

---

# Chapter 11. Configuration Sources

[Configuration File on Specified Location](#) (p. 70)

[Alternative Configuration Sources](#) (p. 70)

[Priorities of Configuration Sources](#) (p. 71)

[Specifying the Path to the Configuration File](#) (p. 72)

Once installed, CloverETL Server requires the configuration of essential (database connection, license, sandboxes) and optional (SMTP, LDAP, etc.) features. The configuration is done by specifying [configuration properties](#) (p. 106) in a [property-key]=[property-value] format.



## Important

The configuration might contain sensitive data; therefore, **CloverETL Server** enables you to encrypt the configuration properties (for more information, see Chapter 15, [Secure Configuration Properties](#) (p. 124)).

CloverETL can load the configuration properties from several sources. We **recommend** the easiest, most convenient way:

---

## Configuration File on Specified Location

The `cloverServer.properties` configuration file is a text file which contains all **CloverETL** settings. You can edit the file either manually or using a much **simpler and intuitive** [Setup GUI](#) (p. 74).

**Example of configuration file's content for PostgreSQL database:**

```
# Modify the url, username and password for your environment.
jdbc.driverClassName=org.postgresql.Driver
jdbc.url=jdbc:postgresql://hostname/clover_db?charset=UTF-8
jdbc.username=user
jdbc.password=pass
jdbc.dialect=org.hibernate.dialect.PostgreSQLDialect
```

The path to the file can be specified by system property, context parameter or environment variable.

We **recommend** specifying the path with the `clover.config.file` system property.

**Example for Apache Tomcat:**

Edit `bin/setenv.sh` (or `bin\setenv.bat`) and add `-Dclover.config.file=/absolute/path/to/cloverServer.properties` to `CATALINA_OPTS`.

For more information and examples on the supported application containers, see [Specifying the Path to the Configuration File](#) (p. 72).

## Alternative Configuration Sources

There are other sources of configuration properties, as well. Each source containing the configuration data has a different priority. If a property isn't set, application's default setting is used.



## Warning

Combining configuration sources could lead to a confusing configuration which would make maintenance much more difficult.

- **Environment Variables**

Environment variables are variables configured by means of your operating system (e.g. `$PATH` is an environment variable).

Set an environment variable with the `clover.` prefix, i.e. (`clover.config.file`).

Some operating systems may not use a dot (`.`) character, so underlines (`_`) may be used instead of dots. So the `clover_config_file` name works as well.

- **System Properties**

System properties are configured by means of JVM, i.e. with the `-D` argument (`-Dclover.config.file`).

Set a system property with the `clover.` prefix, i.e. (`clover.config.file`).

Underlines (`_`) may be used instead of dots (`.`) so the `clover_config_file` name works as well.

- **Configuration File on Default Location**

A text file containing configured CloverETL properties. By default, CloverETL searches for the file on the `[AppServerDir]/cloverServer.properties` path.

- **Modification of Context Parameters in web.xml**

This way **isn't recommended**, since it requires a modification of the WAR file, but it may be useful when none of the approaches above are possible.

Unzip `clover.war` and modify the `WEB-INF/web.xml` file. Add the following piece of code into the file:

```
<context-param>
  <param-name>[property-name]</param-name>
  <param-value>[property-value]</param-value>
</context-param>
```

- **Context Parameters (Available on Apache Tomcat)**

Some application servers allow you to set context parameters without modification of the WAR file.

This way of configuration is possible, but it is **not recommended**, as Apache Tomcat may ignore some context parameters in some environments. Using the configuration file is almost as convenient and much more reliable.

**Example for Apache Tomcat:**

On Tomcat, it is possible to specify context parameters in a context configuration file `[Tomcat_home]/conf/Catalina/localhost/clover.xml` which is created automatically just after deployment of the CloverETL Server web application.

You can specify a property by adding this element:

```
<Parameter name="[propertyName]" value="[propertyValue]" override="false" />
```

(Note: by setting the `override` attribute to `false`), the context parameter does not override the default setting associated with the owning host.)

---

## Priorities of Configuration Sources

Configuration sources have the following priorities (from the highest to lowest):

1. **Context parameters**

Context parameters are specified in an application server or directly in a `web.xml` file (**not recommended**).

## 2. External configuration file

The path to the external configuration file can be specified in several ways. CloverETL Server attempts to find the file in this order (only one of them is loaded):

- a. the path specified with a `config.file` context parameter;
- b. the path specified with a `clover_config_file` or `clover.config.file` system property (**recommended**);
- c. the path specified with a `clover_config_file` or `clover.config.file` environment variable;
- d. the default location (`[AppServerDir]/cloverServer.properties`).

## 3. System properties

## 4. Environment variables

## 5. Default values

---

# Specifying the Path to the Configuration File

[Setup \(p. 74\)](#) uses the configuration file to save the Server's settings. The path to the file is specified by the `clover.config.file` system property. Each application server has a different way to configure the path:

- **Apache Tomcat**

Edit `bin/setenv.sh` (or `bin/setenv.bat`) and add `-Dclover.config.file=/absolute/path/to/cloverServer.properties` to `CATALINA_OPTS`.

See also [Apache Tomcat](#) (p. 18).

- **Jetty**

Edit `bin/jetty.sh` and add `-Dclover.config.file=/absolute/path/to/cloverServer.properties` to `JAVA_OPTS`.

See also [Jetty](#) (p. 24).

- **GlassFish**

Add the `clover.config.file` property in application server's GUI (accessible on <http://localhost:4848>).

The property can be added under **Configuration** → **System Properties**.

See also [GlassFish / Sun Java System Application Server](#) (p. 32).

- **JBoss Application Server**

Edit the file `[JBoss_AS_home]/bin/run.sh` (JBoss AS v 6.x.x) or `standalone.sh` (JBoss AS v 7.x.x) and add the following line:

```
export JAVA_OPTS="$JAVA_OPTS -Dclover.config.file=/absolute/path/to/cloverServer.properties"
```

See also [JBoss Application Server](#) (p. 36).

- **JBoss Enterprise Application Platform**

Edit the configuration file `/standalone/configuration/standalone.xml` and add the following snippet just under the `<extensions>` section:

```
<system-properties>
  <property name="clover.config.file" value="C:/jboss-eap-6.2/cloverServer.properties" />
</system-properties>
```

See also [JBoss Enterprise Application Platform](#) (p. 41).

- **IBM WebSphere**

1. Go to **Integrated Solutions Console** (default URL: <http://localhost:9060/ibm/console/>).
2. Go to **Servers** → **WebSphere application servers** → **[Server\_name]** → **Java and Process Management** → **Process Definition** → **Java Virtual Machine** → **Custom Properties**.
3. Create a system property named `clover_config_file` whose value is a full path to the properties file (e.g. `cloverServer.properties`) on your file system.

See also [IBM WebSphere](#) (p. 28).

- **Weblogic**

Set `JAVA_OPTIONS` variable in the WebLogic domain start script `[domainHome]/startWebLogic.sh`

```
JAVA_OPTIONS="${JAVA_OPTIONS} -Dclover_config_file=/path/to/clover-config.properties
```

See also [Oracle WebLogic Server](#) (p. 47).



## Note

➔ **Continue with:** Chapter 13, [System Database Configuration](#) (p. 83)

---

# Chapter 12. Setup

[Before You Start](#) (p. 74)

[Using Setup](#) (p. 75)

While it is possible to configure **CloverETL Server** by modifying the configuration file (p. 76) in a text editor, the **Setup** with a user-friendly GUI offers a much easier way of configuring basic properties according to your preferences and requirements. Setup is accessible from **Server Console** under **Configuration** → **Setup**.

It lets you configure:

- License (p. 77)
- Database Connection (p. 78)
- Worker (p. 78)
- Sandbox Paths (p. 79)
- Encryption (p. 79)
- E-mail (p. 80)
- LDAP Connection (p. 80)
- Cluster Configuration (p. 81)

**Remember** that you should create a database for **CloverETL Server** and add a user/role for Clover with appropriate rights **before** you set up the connection to the database in the server's Setup GUI.

**CloverETL Server requires a working database connection** for storing license information. Therefore, it allows you to access the Setup and configure the connection **prior** the Server **activation** - simply log in the Server Console and click the **Close** button. Otherwise, you would have to activate the server again, after switching from Derby to a new system database.



## Important

To access the Setup section, you need the [Server Setup permission](#) (p. 170).



## Tip

To keep your settings and data in case of a database migration (e.g. from evaluation to production environment), see Chapter 22, [Server Configuration Migration](#) (p. 188).

---

## Before You Start

Before you start using the Setup, you have to specify the path to the configuration file where the Setup saves the settings, and add required libraries to the classpath.

### 1. Specify the Path to the Configuration File

Setup uses the configuration file to save the server's settings. The path to the file is specified by the `clover.config.file` system property. Each application server has a different way to configure it; for more information, see the [Specifying the Path to the Configuration File](#) (p. 72) section.

### 2. Add Libraries to the Classpath

As the next step, you should place the libraries required for further configuration on the application server's classpath (in most cases, this is done by placing the files into a specific directory). You usually need a JDBC driver for the connection to the database or a `.jar` file with an encryption provider.

For example, in case of Apache Tomcat, copy the libraries to the `TOMCAT/lib/` folder. Then, restart the application server.

### 3. Configure Particular Items

Now, you can use the Setup to configure the rest of the Server's features (see the following section). The configuration is then saved into a file defined in the `clover.config.file` property.

If you wish to encrypt the sensitive data in you configuration file, set up the **Encryption** first.

### 4. Restart the Server (if needed)

Some changes require you to restart the Server. These changes are indicated by the  icon. Other changes (e.g. License, Sandboxes, E-mail, LDAP) are applied immediately and do not require a restart.

---

## Using Setup

Each setup page consists of a menu with setup tabs on the left, a main configuration part in the middle and a configuration status and text on the right side.

**The main configuration part contains several buttons:**

- **Save** saves changes made to the configuration file. The changes in the configuration must be valid.
- **Save Anyway** saves the configuration even if it is invalid. For example, a database connection is considered invalid if a required library is missing.
- **Validate** validates the configuration on a current tab. If you see the **Save** button disabled, use **Validate** to validate the configuration first.
- **Discard Changes** discards unsaved changes and returns to currently used values.

**The following icons can appear in the Server GUI:**

-  configured tab
-  inactive tab
-  error
-  warning
-  restart required
-  pending changes which have not been saved yet

If an error/warning icon appears, a status message on the right side of the Setup GUI will provide relevant details.

If you start the Server without configuration, you will see decorators pointing to the Setup. The decorators mark problems which require your attention. The displayed number corresponds to the number of items.

The screenshot displays the CloverETL Server Console interface. The top navigation bar includes 'Monitoring', 'Executions History', 'Sandboxes', 'Scheduling', 'Tasks History', 'Event Listeners', 'Configuration', and 'Setup'. The 'Setup' tab is active, showing the 'Server Configuration File' section. The configuration file content is as follows:

```

profiler.reporting.console=http://localhost:8083/profiler
server.env.min_heap_memory=600

##Following properties are primarily used in the sandboxes root path specification.
sandboxes.home=${catalina.home}/sandboxes
sandboxes.home.local=${catalina.home}/sandboxes-local
sandboxes.home.partitioned=${catalina.home}/sandboxes-partitioned

## Uncomment lines bellow to enable user lockout after number of failed logins
## Number of failed login attempts after which a next failed login attempt will lock the user
## 0 means feature is switched off
## default suggested value is 50
#security.lockout.login.attempts=50
## Periods are specified in seconds
## Period of time during which the failed login attempts are counted
## Default is 60s (1 min)
#security.lockout.reset.period=60
## Period of time after which a successful login attempt will unlock previously locked user
## Default is 300s (5 min)
#security.lockout.unlock.period=300
## Comma separated list of emails which will be notified when user is locked out.
#security.lockout.notification.email=

## Uncomment lines bellow to enable cluster mode
#cluster.enabled=true
#cluster.node.id=node01
#cluster.jgroups.bind_address=localhost
#cluster.jgroups.start_port=7800
#cluster.jgroups.tcpping.initial_hosts=localhost[7800]
#cluster.http.url=http://localhost:8083/clover

# Uncomment and tweak one of the following sections to use a

```

The right sidebar shows 'Current Status' as 'OK' and a status message: 'The configuration file is C:/Users/Jan/Desktop/CloverETLServer.4.9.0M1.Tomcat-8.0.47/cloverconf/clover.properties set by system property clover\_config\_file.' Below this is an 'Additional Information' section stating: 'This section is to specify advanced configuration options not presented in other sections. Use caution for there is no validation of most of entered settings.'

At the bottom of the configuration editor, there are four buttons: 'Discard Changes', 'Save', 'Download', and 'Update Status'.

Figure 12.1. Setup GUI with decorators

## Configuration File

The **Configuration** tab displays the content of the configuration file. For basic setting, you do not have to edit the content of the file manually. Instead, use the particular Setup tab to configure the corresponding subsystem. For advanced setting, see Chapter 14, [List of Configuration Properties](#) (p. 106).

### Tab specific buttons:

- **Download** allows you to download the properties file.
- **Update Status** updates the content of the Configuration File tab (e.g. in case the file has been manually edited).

### Server Configuration File

Configuration file of the CloverETL Server that stores basic settings. Most of the settings defined in the other Setup sections are stored in this file.

```
# Cluster Settings
cluster.enabled = true
cluster.node.id = Node01
cluster.group.name = cloverCluster
cluster.http.url = http://centos64-7-0.javlin.eu:8080/clover
cluster.jgroups.bind_address = 127.0.0.1
cluster.jgroups.start_port = 7800
# Database
datasource.type = JDBC
jdbc.driverClassName = org.postgresql.Driver
jdbc.dialect = org.hibernate.dialect.PostgreSQLDialect
jdbc.url = jdbc:postgresql://localhost:5432/clover410glassfish
jdbc.username = clover
jdbc.password = conf#9k2SKSP#J6xUTTswWvF1RA==
# Encryption Settings
security.config_properties.encryptor.algorithm = PBewithSHA1AndDESede
# SMTP Settings
clover.smtp.transport.protocol = smtp
clover.smtp.host = localhost
clover.smtp.port = 25
clover.smtp.timeout = 5000
clover.smtp.authentication = true
clover.smtp.username = clover
clover.smtp.password = conf#EXGVQUu86a0j9jrnPLirgmHq3WmyKdId
# LDAP Settings
security.authentication.allowed_domains = clover
```

Discard Changes Save Download Update Status

Figure 12.2. Example of the Server Configuration file

## License

The **License** tab lets you **add/update**, **show** or **reload** the license. The type of the license determines what CloverETL Server's features are activated.



### Important

The license is **stored in the database**; therefore, you should configure the database prior activating the Server with a license. Otherwise, you will have to specify the license again.

Similarly, if you change the database, you will be prompted to re-activate the Server.

### Tab specific buttons:

- **Update License** allows you to add a new, or update an existing license.
- **Show License** opens the window with an active license.
- **Reload License** reloads the license from the database.

### License

CloverETL Server license grants you the ability to use the server and specifies capabilities of this server installation.

Available Licenses

License number	Company name	Products	License storage	Expiration date	Active
Javlin-Internal-License	Javlin	Data Quality(4.1.x); Server(4.1.x)	DATABASE	Jul 21, 2016	<input checked="" type="checkbox"/>

Show detail

Update License Show License Reload License

Figure 12.3. The License tab

## Database

The **Database** tab lets you configure the connection to the database. You can connect via:

- **JDBC**

Under JDBC connection, choose your **Database** from the first drop-down list. This will enable you to choose a **Database URL** template from the second drop-down list. In this template (e.g. `jdbc:postgresql://host:5432/dbname` for PostgreSQL), replace the `host` and `dbname` keywords with proper values.

Next, enter the **User name** and **Password** for your database (default User name/Password: `clover/clover`).



### Note

An Apache Derby JDBC 4 compliant driver is bundled with CloverETL Server. If you use another database system, add a JDBC 4 compliant driver on the classpath.

- **JNDI**

With JNDI, you can access the Datasource on an application server level. Select your **Database** and choose a suitable item from a JNDI tree For more information, see [JNDI DB Datasource](#) (p. 95).

The figure shows two side-by-side screenshots of the 'Database Setup' configuration page. The left screenshot is for a JDBC connection, and the right is for a JNDI connection.

**Left Screenshot (JDBC):**

- Database Setup:** Database is used to store CloverETL Server's configuration (users, event listeners, etc.) and runtime state information (running jobs, task logs, etc.).
- Connection Type:**  JDBC Connection,  JNDI Data Source.
- Connection Specification:**
  - Database: PostgreSQL (dropdown)
  - Database URL: jdbc:postgresql://localhost:5432/clover410glassfish (dropdown)
  - User name: clover (text input)
  - Password: [masked] (password input)
- Buttons: Discard Changes, Save, Validate.

**Right Screenshot (JNDI):**

- Database Setup:** Database is used to store CloverETL Server's configuration (users, event listeners, etc.) and runtime state information (running jobs, task logs, etc.).
- Connection Type:**  JDBC Connection,  JNDI Data Source.
- Connection Specification:**
  - Database: PostgreSQL (dropdown)
  - JNDI data source name: java:comp/env/jdbc/clover\_postgresql (text input)
  - JNDI Initial Context Settings:** (expandable section)
  - JNDI View:**
    - java:comp/env
      - jdbc
        - clover\_postgresql (org.apache.tomcat.dbcp.dbcp.BasicDataSource)
        - clover\_mysql (org.apache.tomcat.dbcp.dbcp.BasicDataSource)
- Buttons: Discard Changes, Save, Validate.

Figure 12.4. Database connection configuration for JDBC (left) and JNDI (right)

## Worker

The **Worker** tab lets you configure Worker features.

You can change **Initial** and **Maximum heap size** for Worker, include **JVM arguments** and set **Port range** for communication with Worker.

**Worker**

This subsystem of CloverETL Server is responsible for running jobs

**Memory Settings**

Initial heap size (MB)  ?

Maximum heap size (MB)  ?

**JVM arguments**

**Connection**

Port range  -  ?

Figure 12.5. The Worker tab

## Sandboxes

The **Sandboxes** tab lets you configure a path to shared (p. 316) local (p. 316) and partitioned (p. 317) sandboxes.

**Sandboxes**

Sandboxes store project's files such as graphs, metadata, input and output data, etc.

**Sandbox Location**

Sandboxes home  ?  
Resolved path: /home/clover/CloverETL/sandboxes

**Sandboxes**

Sandboxes store project's files such as graphs, metadata, input and output data, etc.

**Sandbox Location**

Shared sandboxes home  ?  
Resolved path: /home/clover/CloverETL/sandboxes

Local sandboxes home  ?  
Resolved path: /home/clover/CloverETL/sandboxes-local

Partitioned sandboxes home  ?  
Resolved path: /home/clover/CloverETL/sandboxes-partitioned

Figure 12.6. Sandbox path configuration with clustering disabled (left) and enabled (right)

## Encryption

To secure your sensitive data in the configuration file, you can use the Encryption feature. In the **Encryption** tab, choose a desired **Encryption provider** and **Encryption algorithm**; among the default algorithms, the *PBEWithSHA1AndDESede* is the strongest.

However, since the default algorithms are generally weaker, we **recommend** using [Bouncy Castle](#) - a free custom JCE (Java Cryptography Extension) provider offering a **higher strength of encryption**.

Note that in case you use a custom provider, the libraries have to be added on the `classpath` (in the same way as database libraries).

### Tab specific buttons:

- **Validate** allows you to check validity of the encryption.
- **Save & Encrypt** saves the configuration and encrypt the passwords.

Figure 12.7. Encryption configuration

## E-Mail

The **E-mail** tab lets you configure a connection to an SMTP server so the CloverETL Server can send an email (p. 204) reporting the server status / events on the server.

To make sure the configuration of **Outgoing SMTP Server** is correct, you can check that the SMTP server can be reached by sending a **test email** from the dialog.

### Tab specific buttons:

- **Send e-mail** Allows you to test your SMTP setting by sending an email.

Figure 12.8. E-mail configuration

## LDAP

The **LDAP** tab lets you use an existing LDAP database for user authentication.

Specify a connection to the LDAP server and define a pattern for user DN. **Login Test** allows you to validate the login using any user matching the pattern.

For more information, see [LDAP Authentication](#) (p. 152).

### Tab specific buttons:

- **Validate Login** allows you to test your LDAP authentication setup.

**LDAP Setup**

CloverETL Server can authenticate users against an LDAP directory.

Enable LDAP authentication

**Authentication Policy**

Use LDAP for user authentication only  
 Use LDAP for user authentication and user synchronization

**Connection Specification**

Context factory

LDAP host

IP port

Use encryption (SSL)

Referral processing

**User Authentication**

User DN pattern  ⓘ  
Example: uid=\${username},ou=employees,dc=company,dc=com or just \${username}

**Login Test**

User name

Password

Figure 12.9. LDAP configuration

## Cluster

The **Cluster** tab lets you configure clustering features.

In case your license does not allow clustering, the `Enable clustering` checkbox is grayed out and the note `The license does not allow clustering.` appears at the top.

For more information, see Part VII, “[Cluster](#)” (p. 311).

**Cluster Node Setup**

CloverETL Server can distribute workload within a cluster of servers for higher performance and reliability.

Enable clustering

**Cluster Identification**

Cluster group name  ?

**Node Identification**

Cluster Node ID  ?

This node URL  ?

Bind address  ?

IP port

Figure 12.10. Cluster configuration

---

## Chapter 13. System Database Configuration

The CloverETL Server license, as well as user's information, event listeners and other services, are saved in a database. For stability and performance reasons, the default Apache Derby database is **not** supported for production environment; therefore, you should choose one of the supported DB systems.



### Important

Since CloverETL Server stores important data in a database, you should create a system database and set up a working connection **before** you activate the Server with license and configure it.

For details on how to set up a connection to an external system database, see the list of examples below. The examples contain details on creating databases in DB systems supported by CloverETL Server and configuring a working connection between the database and the Server.

It is possible to specify common **JDBC** DB connection properties (see below) or a **JNDI** location of DB Datasource.

### Clustered Deployment

In a clustered deployment, at least one node in the cluster must have a DB connection configured. Other nodes may have their own direct connection (to the same DB), or may use another node as a proxy for persistent operations; however, the scheduler is active only on nodes with a direct connection. For more information about the feature, see Part VII, “[Cluster](#)” (p. 311).

### Setting up a CloverETL Server's System Database

#### 1. Create a database

- Choose one of the supported database systems and create a database dedicated to CloverETL Server. Add a user/role for Clover and grant it required rights/privileges.

#### 2. Configure common JDBC connection properties

- Some JDBC connection properties are common for all supported database systems. If you use a **properties file** for configuration, specify these properties:

<b>jdbc.driverClassName</b>	Class name for JDBC driver name.
<b>jdbc.url</b>	JDBC URL used by CloverETL Server to store data.
<b>jdbc.username</b>	JDBC database username.
<b>jdbc.password</b>	JDBC database password.
<b>jdbc.dialect</b>	Hibernate dialect to use in Object-relational mapping (ORM).

#### 3. Add a JDBC 4 compliant driver on the classpath.

- As the last step, add a JDBC 4 compliant driver on the classpath. A JDBC Driver which doesn't meet JDBC 4 won't work properly.

Below is a list of examples of individual database systems configurations.

### Examples of Database Configurations

- [Embedded Apache Derby](#) (p. 85)
- [MySQL](#) (p. 86)

- [DB2](#) (p. 87)
- [Oracle](#) (p. 90)
- [Microsoft SQL Server](#) (p. 91)
- [PostgreSQL](#) (p. 93)
- [JNDI DB Datasource](#) (p. 95)

For officially supported versions of particular database systems, see [Database servers](#) (p. 10).

## Embedded Apache Derby

The Apache Derby embedded DB is used with a default CloverETL Server installation. It uses the working directory as a storage for data persistence by default. This may be a problem on some systems. In case of any problems with connecting to Derby DB, we recommend you configure a connection to external DB or at least specify the Derby home directory:

Configure the `derby.system.home` system property to set path which is accessible for application server. You can specify this system property with this JVM execution parameter:

```
-Dderby.system.home=[derby_DB_files_root]
```

Example of a properties file configuration:

```
jdbc.driverClassName=org.apache.derby.jdbc.EmbeddedDriver
jdbc.url=jdbc:derby:databases/cloverDb;create=true
jdbc.username=user
jdbc.password=password
jdbc.dialect=com.cloveretl.server.dbschema.DerbyDialect
```

Take a closer look at the `jdbc.url` parameter. The `databases/cloverDb` part means a subdirectory for DB data. This subdirectory will be created in the directory which is set as `derby.system.home` (or in the working directory if `derby.system.home` is not set). You may change the default value `databases/cloverDb`.

A Derby JDBC 4 compliant driver is bundled with CloverETL Server, thus there is no need to add it on the classpath.



### Note

➔ Continue with: [Encrypted JNDI](#) (p. 96) or [Activation](#) (p. 51)

## MySQL

[Creating Database](#) (p. 86)

[CloverETL Server Setup](#) (p. 86)

CloverETL Server supports MySQL 5, up to version 5.6 included.

### Creating Database

---

The following steps will create a `clover_db` database and the `clover` user with `clover` password .

1. Create database `clover_db`, set charset and collate.

```
CREATE SCHEMA clover_db CHARACTER SET utf8 COLLATE utf8_unicode_ci;
```

2. Use `clover_db` as the current database.

```
USE clover_db;
```

3. Create a new user with password and host.

```
CREATE USER 'clover'@'%' IDENTIFIED BY 'clover';
```

4. Add all privileges to user 'clover' in DB `clover_db`.

```
GRANT ALL ON clover_db.* TO 'clover'@'%';
```

5. Reload privileges.

```
FLUSH privileges;
```

### CloverETL Server Setup

---

Example of a properties file configuration:

```
jdbc.driverClassName=com.mysql.jdbc.Driver
jdbc.url=jdbc:mysql://127.0.0.1:3306/clover_db?useUnicode=true&characterEncoding=utf8
jdbc.username=clover
jdbc.password=clover
jdbc.dialect=org.hibernate.dialect.MySQLDialect
```

Add a JDBC 4 compliant driver on the classpath. A JDBC Driver which doesn't meet JDBC 4 won't work properly.



#### Note

➔ Continue with: [Encrypted JNDI](#) (p. 96) or [Activation](#) (p. 51)

## DB2

[Creating Database](#) (p. 87)

[CloverETL Server Setup](#) (p. 87)

[Troubleshooting](#) (p. 88)

[DB2 on AS/400](#) (p. 89)

### Creating Database

---

1. Create a dedicated user for the Clover database and set a password (UNIX/Linux).

```
useradd clover
```

```
passwd clover
```

2. Create a new database.

```
db2 "CREATE DATABASE cloverdb PAGESIZE 32768 RESTRICTIVE"
```

3. Activate the database.

```
db2 activate db cloverdb
```

4. Connect to the database.

```
db2 connect to cloverdb
```

5. Grant the user DBADM authority (DBADM authority is an administrative authority for a specific database. The database administrator possesses the privileges that are required to create objects and issue database commands. By default, DATAACCESS and ACCESSCTRL authority are also granted).

```
db2 "GRANT DBADM ON DATABASE TO USER clover"
```

6. Disconnect from database

```
db2 connect reset
```

### CloverETL Server Setup

---

Example of a properties file configuration:

```
jdbc.driverClassName=com.ibm.db2.jcc.DB2Driver  
jdbc.url= jdbc:db2://localhost:50000/clover  
jdbc.username=clover  
jdbc.password=clover  
jdbc.dialect=org.hibernate.dialect.DB2Dialect
```

Add a JDBC 4 compliant driver on the classpath. A JDBC driver which doesn't meet JDBC 4 specifications won't work properly.

## Troubleshooting

### Wrong pagesize

The *cloverdb* database has to be created with suitable `PAGESIZE`. DB2 has several possible values for this property: 4096, 8192, 16384 or 32768.

CloverETL Server should work on DB with `PAGESIZE` set to 16384 or 32768. If the `PAGESIZE` value is not set properly, there should be an error message in the log file after failed CloverETL Server startup.

#### The error indicating wrong pagesize:

```
ERROR:
DB2 SQL Error: SQLCODE=-286, SQLSTATE=42727, SQLERRMC=16384;
ROOT, DRIVER=3.50.152
```

`SQLERRMC` contains suitable value for `PAGESIZE`.

#### Solution:

You can create a database with a proper page size using the `PAGESIZE` command, e.g.:

```
CREATE DB clover PAGESIZE 32768;
```

### The table is in the reorg pending state

On rare occasions, the `ALTER TABLE` commands may cause tables to remain in "reorg pending state". This behavior is specific for DB2. The `ALTER TABLE DDL` commands are executed only during the first start of a new CloverETL Server version.

#### The issue may return the following error messages:

```
Operation not allowed for reason code "7" on table "DB2INST2.RUN_RECORD"..
SQLCODE=-668, SQLSTATE=57016
```

```
DB2 SQL Error: SQLCODE=-668, SQLSTATE=57016, SQLERRMC=7;DB2INST2.RUN_RECORD, DRIVER=3.5
```

In this case, the "RUN\_RECORD" table is in the "reorg pending state" and "DB2INST2" is the DB instance name.

#### Solution:

Go to DB2 console and execute the following command (for table `run_record`):

```
reorg table run_record
```

#### DB2 console output should look like this:

```
db2 => connect to clover1
Database Connection Information

Database server          = DB2/LINUX 9.7.0
SQL authorization ID    = DB2INST2
Local database alias    = CLOVER1

db2 => reorg table run_record
DB20000I The REORG command completed successfully.
db2 => disconnect clover1
DB20000I The SQL DISCONNECT command completed successfully.
```

"clover1" is DB name

### **DB2 does not allow ALTER TABLE which trims DB column length.**

This problem depends on the DB2 configuration and has been experienced only on some AS400s, so far. CloverETL Server applies a set of DP patches during the first installation after the application upgrade. Some of these patches may apply column modifications which trim the length of the text columns. These changes never truncate any data, however DB2 does not allow this since it "may" truncate some data. DB2 refuses these changes even in empty DB table.

#### **Solution:**

Disable the DB2 warning for data truncation, restart CloverETL Server which applies patches, then enable DB2 warning again.

### **DB2 on AS/400**

---

The connection on AS/400 might be slightly different.

Example of a properties file configuration:

```
jdbc.driverClassName=com.ibm.as400.access.AS400JDBCdriver
jdbc.username=user
jdbc.password=password
jdbc.url=jdbc:as400://host/cloversrv;libraries=cloversrv;date format=iso
jdbc.dialect=org.hibernate.dialect.DB2400Dialect
```

Use credentials of your OS user for `jdbc.username` and `jdbc.password`.

`cloversrv` in `jdbc.url` above is the name of the DB schema.

#### **You can create the schema in AS/400 console:**

1. execute command STRSQL (**SQL console**)
2. execute command

```
CREATE COLLECTION cloversrv IN ASP 1
```

`cloversrv` is the name of the DB schema and it may be at most 10 characters long

Proper JDBC driver must be in the application server classpath.

Use `jt400ntv.jar` JDBC driver found in `/QIBM/ProdData/Java400` on the server.

Add a JDBC 4 compliant driver on the classpath. A JDBC driver which doesn't meet JDBC 4 specifications won't work properly.



#### **Note**



**Continue with:** [Encrypted JNDI](#) (p. 96) or [Activation](#) (p. 51)

## Oracle

[Creating Database](#) (p. 90)

[CloverETL Server Setup](#) (p. 90)

### Creating Database

---

Run the following script to create a role (cloverRole), user (cloverUser) with password (cloverPassword) and tablespace for CloverETL Server:

```
-- Create a new role and grant it privileges.
CREATE ROLE cloverRole NOT IDENTIFIED;

GRANT CREATE SESSION TO cloverRole;
GRANT ALTER SESSION TO cloverRole;
GRANT CREATE TABLE TO cloverRole;
GRANT CREATE SEQUENCE TO cloverRole;
GRANT CREATE TRIGGER TO cloverRole;

-- Create a new database user with password.
CREATE USER cloverUser IDENTIFIED BY cloverPassword;

-- Set quota on tablespace.
GRANT UNLIMITED TABLESPACE TO cloverUser;

-- Connect a new role to a new user.
GRANT cloverRole TO cloverUser;
```

### CloverETL Server Setup

---

Example of a properties file configuration:

```
jdbc.driverClassName=oracle.jdbc.OracleDriver
jdbc.url=jdbc:oracle:thin:@host:1521:db
jdbc.username=cloverUser
jdbc.password=cloverPassword
jdbc.dialect=org.hibernate.dialect.Oracle10gDialect
```

Add a JDBC 4 compliant driver on the classpath. A JDBC driver which doesn't meet the JDBC 4 specifications won't work properly.

**These are the privileges which have to be granted to a schema used by CloverETL Server:**

```
CONNECT
CREATE SESSION
CREATE/ALTER/DROP TABLE
CREATE/ALTER/DROP SEQUENCE

QUOTA UNLIMITED ON <user_tablespace>;
QUOTA UNLIMITED ON <temp_tablespace>;
```



#### Note

➔ Continue with: [Encrypted JNDI](#) (p. 96) or [Activation](#) (p. 51)

## Microsoft SQL Server

[Creating Database](#) (p. 91)

[CloverETL Server Setup](#) (p. 91)

### Creating Database

---

It is advised to use SQL Server Authentication instead of Windows Authentication. To enable it, select the server instance in Microsoft SQL Server Management Studio: go to **Properties** → **Security** → **Server authentication** and select the SQL Server and Windows Authentication mode. The server instance needs to be restarted.



#### Note

Make sure you have:

- TCP/IP Enabled in **SQL Server Network Configuration** → **Protocols**
- TCP Port set to **1433** in **TCP/IP Properties** → **IP Addresses** → **IPAll**

1. Create a new database

```
CREATE DATABASE clover_db;
```

2. Enable Read Committed Snapshot Isolation on the new database

```
ALTER DATABASE clover_db SET READ_COMMITTED_SNAPSHOT ON;
```

3. Create a new login role.

```
CREATE LOGIN clover with PASSWORD = 'clover', DEFAULT_DATABASE = clover_db;
```

4. Connect to the database.

```
USE clover_db;
```

5. Create a new database user.

```
CREATE USER clover FOR LOGIN clover;
```

6. Add a database role membership db\_owner (Members of the db\_owner fixed database role can perform all configuration and maintenance activities on the database, and can also drop the database).

```
EXEC sp_addrolemember 'db_owner','clover';
```

### CloverETL Server Setup

---

Using MS SQL requires configuration of the database server:

1. Run **Microsoft SQL Server Management Studio** tool;

2. Create a new user under **Security/Logins**;
3. Under **Databases** create a new database (e.g. 'clover') for CloverETL Server, select the user from the previous step as owner of the database;
4. Under database **Options**, set the **Is Read Committed Snapshot On** option to **True**.

Example of a properties file configuration:

```
jdbc.driverClassName=com.microsoft.sqlserver.jdbc.SQLServerDriver
jdbc.url=jdbc:sqlserver://localhost:1433;instance=SQLSERVERINSTANCE;database=clover_db
jdbc.username=clover
jdbc.password=clover
jdbc.dialect=org.hibernate.dialect.SQLServerDialect
```

Add a JDBC 4 compliant driver on the application server classpath. A JDBC driver that does not meet the JDBC 4 specifications will not work properly.



## Note

➔ Continue with: [Encrypted JNDI](#) (p. 96) or [Activation](#) (p. 51)

---

## PostgreSQL

[Creating Database](#) (p. 93)

[CloverETL Server Setup](#) (p. 93)

---

### Creating Database

Advanced users can create their own table space

We are going to create a database for Clover to use a 'user group' role which will own the database and a user role which we will add to the user group. This user role will be then used by the Server to access the database.

Database name: clover\_db

UserGroup: cloveretl

User: clover

Password: clover

1. Optionally, you can create a new tablespace
2. Connect as postgres (default admin) to the default DB postgres and execute the following commands:

```
CREATE ROLE cloveretl NOSUPERUSER NOCREATEDB NOCREATEROLE NOINHERIT NOLOGIN;  
CREATE ROLE clover NOSUPERUSER NOCREATEDB NOCREATEROLE INHERIT LOGIN ENCRYPTED PASSWORD 'clover';  
GRANT cloveretl TO clover;  
CREATE DATABASE clover_db;  
GRANT ALL ON DATABASE clover_db TO cloveretl;  
REVOKE ALL ON DATABASE clover_db FROM public;
```

To separate the database into its own tablespace, create a tablespace before creating the database.

and use the following command to create the database:

```
CREATE DATABASE clover_db WITH OWNER cloveretl TABLESPACE tablespace_name;
```

For more information, see the [PostgreSQL documentation](#).

---

### CloverETL Server Setup

Example of a properties file configuration:

```
jdbc.driverClassName=org.postgresql.Driver  
jdbc.url=jdbc:postgresql://localhost/clover_db?charSet=UTF-8  
jdbc.username=clover  
jdbc.password=clover  
jdbc.dialect=org.hibernate.dialect.PostgreSQLDialect
```

Add a JDBC 4 compliant driver on the classpath. A JDBC driver which doesn't meet the JDBC 4 specifications won't work properly.

The JDBC driver for PostgreSQL can be downloaded from the [official PostgreSQL page](#).

**Example for Apache Tomcat:** place the libraries into the TOMCAT/lib directory.

See also [PostgreSQL documentation on jdbc URL](#).



### Note

➔ Continue with: [Encrypted JNDI](#) (p. 96) or [Activation](#) (p. 51)

---

## JNDI Configuration and Encryption

[JNDI DB Datasource](#) (p. 95)

[JNDI Datasource Troubleshooting](#) (p. 96)

[Encrypted JNDI](#) (p. 96)

---

### JNDI DB Datasource

CloverETL Server can connect to a database using JNDI Datasource which is configured in an application server.

#### Example for Apache Tomcat and PostgreSQL database:

- **JNDI Datasource Definition**

First you need to **define a JNDI Datasource** in an application server. The following context resource configuration may be added to the [Tomcat\_home]/conf/server.xml file to the <Host> element.

**Note:** Do not put the code into the <GlobalNamingResources> element, since the resource would not be visible by the CloverETL webapp.

```
<Context path="/clover">
  <Resource name="jdbc/clover_server"
    auth="Container"
    type="javax.sql.DataSource"
    factory="org.apache.tomcat.jdbc.pool.DataSourceFactory"
    driverClassName="org.postgresql.Driver"
    url="jdbc:postgresql://127.0.0.1:5432/clover_db"
    username="clover"
    password=""
    maxTotal="20"
    maxIdle="10"
    maxWaitMillis="-1"/>
</Context>
```

- **JNDI Connection Configuration**

Now that the Datasource is defined, you should **configure the connection**.

The following parameters may be set in the same way as other parameters (in the properties file or the Tomcat context file). You can also set the parameters in the [Database \(p. 78\)](#) tab of the [Setup \(p. 74\)](#) GUI.

```
datasource.type=JNDI # type of Datasource; must be set, because the default value
datasource.jndiName=jdbc/clover_server # JNDI location of DB Datasource; the default value is java:c
jdbc.dialect=org.hibernate.dialect.PostgreSQLDialect # Set the dialect according to DB which DataSource is connect
# The correct dialect can be found in the examples of DB conf
```

Since the DB connection contains sensitive information (e.g. username, password, etc.), CloverETL provides the [JNDI Encryption \(p. 96\)](#) feature.



#### Tip

The resource configuration may also be added to the context file [Tomcat\_home]/conf/Catalina/localhost/clover.xml.



#### Important

Special characters typed in the context file have to be specified as XML entities, e.g. ampersand "&" as "&#38;", etc.

For a detailed list of parameters which can be set up in the configuration file, see Chapter 14, [List of Configuration Properties](#) (p. 106).

## JNDI Datasource Troubleshooting

---

### JNDI Datasource in Oracle WebLogic

CloverETL's default Quartz configuration does not work with default JNDI Datasource from WebLogic. Proceed with one of the following options:

1. Configure the Quartz JDBC delegate manually **before** the server is started;
2. Disable JDBC type wrapping in the WebLogic's Datasource configuration.

### Apache Tomcat's DBCP JNDI pool

The default JNDI pool **DBCP** in **Apache Tomcat** does not handle connections efficiently. With the **DBCP** JNDI pool, low performance can be seen if `DBOutputTable` with returning statement is used.

Therefore, `tomcat-jdbc-pool` is used instead by adding the `factory="org.apache.tomcat.jdbc.pool.DataSourceFactory"` attribute to the definition of the JNDI resource. See [The Tomcat JDBC Connection Pool](#)

## Encrypted JNDI

---

The encryption feature allows you to protect your sensitive data defined in the Datasource definition (e.g. username, password, etc.), which are by default stored in plain text. The configuration differs between particular application servers.

[Encrypted JNDI on Tomcat](#) (p. 97)

[Encrypted JNDI on Jetty 9 \(9.2.6\)](#) (p. 98)

[Encrypted JNDI on JBoss 6.0.0](#) (p. 99)

[Encrypted JNDI on JBoss 7](#) (p. 101)

[Encrypted JNDI on Glassfish 3 \(3.1.2.2\)](#) (p. 103)

[Encrypted JNDI on WebSphere 8.5.5.0](#) (p. 104)

[Encrypted JNDI on WebLogic](#) (p. 105)

## Encrypted JNDI on Tomcat

You need `secure-cfg-tool` to encrypt the passwords. Use the version of `secure-cfg-tool` corresponding to the version of CloverETL Server. Usage of the tool is described in Chapter 15, [Secure Configuration Properties](#) (p. 124).

Use `encrypt.sh` or `encrypt.bat` for password encryption. Place the encrypted password into a configuration file, and put `cloveretl-secure-jndi-resource-{version}.jar` and `jaspyt-1.9.0.jar` files on the classpath of the application server. The `.jar` files can be found in the `tomcat-secure-jndi-resource` directory packed in `secure-cfg-tool`.

The `tomcat-secure-jndi-resource` directory contains a useful README file with further details on encrypted JNDI.

### Example of encrypted JNDI connection for PostgreSQL

Encrypt the password:

1. `./encrypt.sh -a PBEWithSHA1AndDESede`
2. The configuration is placed in `${CATALINA_HOME}/conf/context.xml`. Note that the encryption algorithm `PBEWithSHA1AndDESede` is not default.

```
<Resource name="jdbc/clover_server"
  auth="Container"
  factory="com.cloveretl.secure.tomcatresource.SecureDataSourceFactory"
  secureAlgorithm="PBEWithSHA1AndDESede"
  type="javax.sql.DataSource"
  driverClassName="org.postgresql.Driver"
  url="jdbc:postgresql://127.0.0.1:5432/clover_db?charset=UTF-8"
  username="conf#rPz5F0o7HPn4dFTRV5Ourg=="
  password="conf#4KlNp8/FVDR+rTwx0dEqWA=="
  maxTotal="20"
  maxIdle="10"
  maxWaitMillis="-1"/>
```

If you use other JCE (e.g. Bouncy Castle), it has to be added to the classpath of the application server (`${CATALINA_HOME}/lib`). The `encrypt` command requires the path to directory with JCE, too.

```
./encrypt.sh -l ~/lib/ -c
org.bouncycastle.jce.provider.BouncyCastleProvider -a
PBEWITHSHA256AND256BITAES-CBC-BC
```

```
<Resource name="jdbc/clover_server"
  auth="Container"
  factory="com.cloveretl.secure.tomcatresource.SecureDataSourceFactory"
  secureProvider="org.bouncycastle.jce.provider.BouncyCastleProvider"
  secureAlgorithm="PBEWITHSHA256AND256BITAES-CBC-BC"
  type="javax.sql.DataSource"
  driverClassName="org.postgresql.Driver"
  url="jdbc:postgresql://127.0.0.1:5432/clover_db?charset=UTF-8"
  username="conf#Ws9IuHKO9h7hMjP1lr31VxdI1A9LKIaYfGEUmLet9rA="
  password="conf#Cj1v59Z5nCBHaktn6Ubgst4Iz69JLQ/q6/32Xwr/IEE="
  maxTotal="20" maxIdle="10"
  maxWaitMillis="-1"/>
```

## Encrypted JNDI on Jetty 9 (9.2.6)



### Note

See the Jetty documentation on [Secure Password Obfuscation](#).

Configuration of a JNDI JDBC connection pool is stored in the plain text file, `$JETTY_HOME/etc/jetty.xml`.

```
<New id="MysqlDB" class="org.eclipse.jetty.plus.jndi.Resource" >
  <Arg></Arg>
  <Arg>jdbc/MysqlDS</Arg>
  <Arg>
    <New class="com.mysql.jdbc.jdbc2.optional.MysqlConnectionPoolDataSource" >
      <Set name="URL">jdbc:mysql://localhost:3306/clover_empty</Set>
      <Set name="User">user</Set>
      <Set name="Password">password</Set>
    </New>
  </Arg>
</New>
```

### Obfuscating the Password

Password can be obfuscated using the `org.eclipse.jetty.util.security.Password` class within `lib/jetty-util-{VERSION}.jar`:

```
java -cp lib/jetty-util-9.2.6.v20141205.jar org.eclipse.jetty.util.security.Password password
```

Command returns obfuscated and hashed password. The obfuscated password will be used to replace the plain password value.

### Replacing the Password

Replace the plain text password with the Call element. Its only argument is a string starting with the OBF: prefix returned by the command mentioned in the previous section.

```
<New id="MysqlDB" class="org.eclipse.jetty.plus.jndi.Resource" >
  <Arg></Arg>
  <Arg>jdbc/MysqlDS</Arg>
  <Arg>
    <New class="com.mysql.jdbc.jdbc2.optional.MysqlConnectionPoolDataSource" >
      <Set name="URL">jdbc:mysql://localhost:3306/clover_empty</Set>
      <Set name="User">user</Set>
      <Set name="Password">
        <Call class="org.eclipse.jetty.util.security.Password" name="deobfuscate">
          <Arg>OBF:1v2jluumlxtvlzejlzerlxtnluvklvlv</Arg>
        </Call>
      </Set>
    </New>
  </Arg>
</New>
```



### Note

Password in the JMS connection can also be obfuscated.

## Encrypted JNDI on JBoss 6.0.0



### Note

See the JBoss documentation on [Encrypting Data Source Passwords](#)

(In the documentation, `client/jboss-logging-spi.jar` is used; however in newer version, the `client/jboss-logging.jar` can be used instead.)

Original data source with an unencrypted password:

```
<datasources>
  <local-tx-datasource>
    <jndi-name>MySQLDS</jndi-name>
    <connection-url>jdbc:mysql://127.0.0.1:3306/clover</connection-url>
    <driver-class>com.mysql.jdbc.Driver</driver-class>
    <user-name>user</user-name>
    <password>password</password>
  </local-tx-datasource>
</datasources>
```

1. Encrypt the data source password:

- **Unix-like systems:**

```
java -cp client/jboss-logging.jar:lib/jbossx.jar org.jboss.resource.security.SecureIdentityLoginModule password
```

- **Windows system:**

```
java -cp client\jboss-logging.jar;lib\jbossx.jar org.jboss.resource.security.SecureIdentityLoginModule password
```

The command will return an encrypted password, e.g. `5dfc52b51bd35553df8592078de921bc`.

2. Create a new application authentication policy in `conf/login-config.xml` within currently used server's profile directory (e.g. `server/default/conf/login-config.xml`).

```
<application-policy name="EncryptDBPassword">
  <authentication>
    <login-module code="org.jboss.resource.security.SecureIdentityLoginModule" flag="required">
      <module-option name="username">user</module-option>
      <module-option name="password">5dfc52b51bd35553df8592078de921bc</module-option>
      <module-option name="managedConnectionFactoryName">jboss.jca:name=MySQLDS,service=LocalTxCM</module-option>
    </login-module>
  </authentication>
</application-policy>
```

3. Replace authentication entries with a reference to the application authentication policy.

```
<security-domain>EncryptDBPassword</security-domain>
```

The final data source looks like this:

```
<datasources>
  <local-tx-datasource>
    <jndi-name>MySQLDS</jndi-name>
    <connection-url>jdbc:mysql://127.0.0.1:3306/clover</connection-url>
    <driver-class>com.mysql.jdbc.Driver</driver-class>
    <security-domain>EncryptDBPassword</security-domain>
  </local-tx-datasource>
</datasources>
```

```
</local-tx-datasource>  
</datasources>
```

The same mechanism can be probably used for JMS, as well.

```
<tx-connection-factory>  
...  
<security-domain-and-application>RealmWithEncryptedPassword</security-domain-and-application>  
...  
</tx-connection-factory>
```

## Encrypted JNDI on JBoss 7

JBoss 7 - JBoss EAP 6.2.0.GA - AS 7.3.0.Final-redhat-14



### Note

For details, see [Using Encrypted DataSource Password in JBoss AS7](#).

Configuration steps are similar to configuring of JBoss 6.

The configuration takes place in a single configuration file, e.g. for standalone profile `JBOSS_HOME/standalone/configuration/standalone.xml`.

Original data source:

```
<datasources>
  <datasource jndi-name="java:/MysqlDS" pool-name="MySQLPool">
    <connection-url>jdbc:mysql://localhost:3306/cluster</connection-url>
    <driver>mysql</driver>
    <pool>
      <max-pool-size>30</max-pool-size>
    </pool>
    <security>
      <user-name>user</user-name>
      <password>password</password>
    </security>
  </datasource>

  <drivers>
    <driver name="mysql" module="com.cloveretl.jdbc">
      <driver-class>com.mysql.jdbc.Driver</driver-class>
    </driver>
  </drivers>
</datasources>
```

1. In `JBOSS_HOME` directory run the cli command:

```
java -cp modules/system/layers/base/org/picketbox/main/picketbox-4.0.19.SP2-redhat-1.jar:client/jboss-logging.jar
```

The command will return an encrypted password, e.g. `5dfc52b51bd35553df8592078de921bc`.

2. Add a new security-domain to `security-domains`, the password value is a result of the command from the previous step.

```
<security-domain name="EncryptDBPassword" cache-type="default">
  <authentication>
    <login-module code="org.picketbox.datasource.security.SecureIdentityLoginModule" flag="required">
      <module-option name="username" value="user"/>
      <module-option name="password" value="5dfc52b51bd35553df8592078de921bc"/>
      <module-option name="managedConnectionFactoryName" value="jboss.jca:service=LocalTxCM,name=MySQLPool"/>
    </login-module>
  </authentication>
</security-domain>
```

3. Replace user and password with a reference to the security domain.

```
<datasources>
  <datasource jndi-name="java:/MysqlDS" pool-name="MySQLPool" enabled="true" use-java-context="true">
    <connection-url>jdbc:mysql://localhost:3306/cluster</connection-url>
    <driver>mysql</driver>
    <pool>
      <max-pool-size>30</max-pool-size>
    </pool>
    <security>
```

```
<security-domain>EncryptDBPassword</security-domain>
</security>
</datasource>

<drivers>
  <driver name="mysql" module="com.cloveretl.jdbc">
    <driver-class>com.mysql.jdbc.Driver</driver-class>
  </driver>
</drivers>
</datasources>
```

It is possible that the same mechanism can also be used for JMS.

## Encrypted JNDI on Glassfish 3 (3.1.2.2)

Configuration of JDBC connection pool is stored in the plain text file \$DOMAIN/config/domain.xml.

```
<jdbc-connection-pool driver-classname="com.mysql.jdbc.Driver" datasource-classname="" res-type="java.sql.Driver"
  <property name="URL" value="jdbc:mysql://localhost:3306/clover_empty"></property>
  <property name="user" value="user"></property>
  <property name="password" value="password"></property>
</jdbc-connection-pool>
```

The password is not encrypted, but can be replaced with so called password alias:

A password alias stores a password in an encrypted form in the domain keystore, providing a clear-text alias name to use instead of the password. In password files and the domain configuration file, use the form `${ALIAS=alias-name}` to refer to the encrypted password.

### Creating a Password Alias

There are two ways to create a password alias: using `create-password-alias` command in a command-line administration utility, or in the web Server Administration Console in the Password Aliases section (**Domain** → **Password Aliases**).

### Replacing the Password with the Password Alias

Replace the password (the attribute value) with a `${ALIAS=password_alias_name}` string, where `password_alias_name` is the name of the alias.

```
<jdbc-connection-pool driver-classname="com.mysql.jdbc.Driver" datasource-classname="" res-type="java.sql.Driver"
  <property name="URL" value="jdbc:mysql://localhost:3306/clover_empty"></property>
  <property name="user" value="user"></property>
  <property name="password" value="${ALIAS=password_alias_name}"></property>
</jdbc-connection-pool>
```



### Note

- Glassfish Administration Server Console mentions a lower case keyword (alias); if it doesn't work, try changing to upper case (ALIAS).
- Password for a JMS connection can be replaced with an alias as well.

### **Encrypted JNDI on WebSphere 8.5.5.0**

In WebSphere, user credentials aren't saved in plain text, but as J2C authentication data. (see [How to Create a WAS JDBC Provider, J2C Authentication Alias, and Data Source for the IBM i](#)).

The same mechanism can also be used for JMS connection (see IBM's instructions on [Configuring an external JMS provider](#)).

## Encrypted JNDI on WebLogic

Password in a JNDI datasource file is encrypted by default when created by admin's web console (Service/Datasource).

Example of datasource file (located in the DOMAIN/config/jdbc/ directory):

```
<?xml version='1.0' encoding='UTF-8'?>
<jdbc-data-source xmlns="http://xmlns.oracle.com/weblogic/jdbc-data-source" xmlns:sec="http://xmlns.oracle.com/web
  <name>MysqlDS</name>
  <jdbc-driver-params>
    <url>jdbc:mysql://127.0.0.1:3306/clover</url>
    <driver-name>com.mysql.jdbc.Driver</driver-name>
    <properties>
      <property>
        <name>user</name>
        <value>user</value>
      </property>
    </properties>
    <password-encrypted>{AES}zIiq6/JutK/wD4CcRPX1pOueIKqc6uRVxAnZzcC3pI=</password-encrypted>
  </jdbc-driver-params>
  <jdbc-connection-pool-params>
    <test-table-name>SQL SELECT 1</test-table-name>
  </jdbc-connection-pool-params>
  <jdbc-data-source-params>
    <jndi-name>jdbc/MysqlDS</jndi-name>
    <global-transactions-protocol>OnePhaseCommit</global-transactions-protocol>
  </jdbc-data-source-params>
</jdbc-data-source>
```

The same mechanism is also used for encrypting password in the JMS connection (see Oracle's instructions on [Configuring an external JMS provider](#)).



### Note

➔ Continue with: [Activation](#) (p. 51)

# Chapter 14. List of Configuration Properties

[General Configuration Properties](#) (p. 106)

[Worker - Configuration Properties](#) (p. 114)

[Worker - JNDI Properties](#) (p. 117)

[Worker - SSL Properties](#) (p. 120)

[Job Execution Properties](#) (p. 121)

Below you can find the *configuration properties* available in CloverETL Server. The essential properties can be configured using the [Setup GUI \(p. 74\)](#). Other properties serve to tweak various features of CloverETL Server. However, these properties have to be configured manually, e.g. by editing the configuration file.

In CloverETL Server UI, you can view the properties and their values in Configuration > CloverETL Info > Server Properties.

Additional properties used for cluster configuration can be found in Chapter 39, [Clustering Features](#) (p. 312).



## Important

*Configuration property* and *system property* are not the same. *Configuration properties* can be configured in **Setup** section or in `cloveretl.properties` file. *System properties* serve to configure the JVM. E.g. in Apache Tomcat, they are configured in `bin/setenv.[bat|sh]` file using `-D` prefix.

---

## General Configuration Properties

[Configuration file](#) (p. 106)

[License](#) (p. 106)

[Engine](#) (p. 106)

[Sandboxes](#) (p. 107)

[Database connection](#) (p. 107)

[Security](#) (p. 108)

[SMTP](#) (p. 109)

[Logging](#) (p. 109)

[Thread Manager](#) (p. 110)

[Archivator](#) (p. 111)

[Properties resolver](#) (p. 111)

[Data Services](#) (p. 112)

[API](#) (p. 112)

[JVM](#) (p. 112)

[Misc](#) (p. 113)

Table 14.1. General configuration

Key	Description	Default Value
<b>Configuration file</b>		
<code>clover.config.file</code>	Absolute path to location of a CloverETL Server configuration file	<code>/absolute/path/to/cloverServer.properties</code>
<b>License</b>		
<code>license.file</code>	Absolute path to location of a CloverETL Server license file ( <code>license.dat</code> )	
<code>license.context_names</code>	A comma-separated list of web-app contexts which may contain license. Each of them has to start with a slash! Works only on Apache Tomcat.	<code>/clover-license/ clover_license</code>
<b>Engine</b>		

Chapter 14. List of Configuration Properties

Key	Description	Default Value
engine.config.file	location of a CloverETL engine configuration properties file	properties file packed with CloverETL
engine.plugins.additional.src	This property may contain an absolute path to some "source" of additional CloverETL engine plugins. These plugins are not a substitute for plugins packed in WAR. "Source" may be a directory or a zip file. Both, a directory and a zip, must contain a subdirectory for each plugin. Changes in the directory or the ZIP file apply only when the server is restarted. For details see Chapter 32, <a href="#">Extensibility - CloverETL Engine Plugins</a> (p. 264).	empty
<b>Sandboxes</b>		
sandboxes.home	This property is primarily intended to be used as a placeholder in the sandbox root path specification. So the sandbox path is specified with the placeholder and it's resolved to the real path just before it's used. The sandbox path may still be specified as an absolute path, but placeholder has some significant advantages: <ul style="list-style-type: none"> <li>* sandbox definition may be exported/imported to another environment with a different directory structure</li> <li>* user creating sandboxes doesn't have to care about physical location on the filesystem</li> <li>* each node in cluster environment may have a different "sandboxes.home" value, so the directory structure doesn't have to be identical</li> </ul> The default value uses content of a "user.data.home" configuration property, which points to the home directory of the user which runs the JVM process. Directory depends on the OS. On Unix-like systems it's typically /home/[username]	\${user.data.home}/CloverETL/sandboxes
sandboxes.access.check.boundaries.enabled	<p>true   false If it is set to false, then the path relative to a sandbox root may point out of the sandbox. No file/folder outside of the sandbox is accessible by the relative path otherwise.</p>	true
<b>Database connection</b>		
datasource.type	Set this explicitly to JNDI if you need CloverETL Server to connect to a DB using JNDI datasource. In such case, "datasource.jndiName" and "jdbc.dialect" parameters must be set properly. Possible values: JNDI   JDBC	JDBC
datasource.jndiName	JNDI location of a DB DataSource. It is applied only if "datasource.type" is set to "JNDI".	java:comp/env/jdbc/clover_server
jdbc.driverClassName	class name for JDBC driver name	
jdbc.url	JDBC URL used by CloverETL Server to store data	
jdbc.username	JDBC database user name	

Chapter 14. List of Configuration Properties

Key	Description	Default Value
jdbc.password	JDBC database password	
jdbc.dialect	hibernate dialect to use in ORM	
quartz.driverDelegateClass	SQL dialect for quartz. Value is automatically derived from "jdbc.dialect" property value.	
<b>Security</b>		
private.properties	List of server properties which are used only by the CloverETL Server code. So these properties are not accessible outside of the ServerFacade. By default, there are all properties which may contain password in the list, so their values are not visible for web GUI users. The values are replaced by a single star "*". Changes in this list may cause unexpected behavior of some server API.	jdbc.password, executor.password, security.ldap.password, clover.smtp.password
security.session.validity	Session validity in milliseconds. When the request of logged-in user/client is detected, validity is automatically prolonged.	14400000
security.session.exchange.limit	Interval for exchange of invalid tokens in milliseconds.	360000
security.default_domain	Domain in which all new users are included. Stored in user's record in the database. Shouldn't be changed unless the "clover" must be white-labelled.	clover
security.basic_authentication.features_list	List of features which are accessible using HTTP and which should be protected by Basic HTTP Authentication. The list has form of semicolon separated items; Each feature is specified by its servlet path.	/request_processor;/simpleHttpApi;/launch;/launchIt;/downloadStorage;/downloadFile;/uploadSandboxFile;/downloadLog;/webdav
security.basic_authentication.realm	Realm string for HTTP Basic Authentication.	CloverETL Server
security.digest_authentication.features_list	List of features which are accessible using HTTP and which should be protected by HTTP Digest Authentication. The list has form of semi-colon separated items. Each feature is specified by its servlet path.  Please keep in mind that HTTP Digest Authentication is feature added to the version 3.1. If you upgraded your older CloverETL Server distribution, users created before the upgrade cannot use the HTTP Digest Authentication until they reset their passwords. So when they reset their passwords (or the admin does it for them), they can use Digest Authentication as well as new users.	
security.digest_authentication.storeA1.enabled	Switch whether the A1 Digest for HTTP Digest Authentication should be generated and stored or not. Since there is no CloverETL Server API using the HTTP Digest Authentication by default, it's	false

Chapter 14. List of  
Configuration Properties

Key	Description	Default Value
	recommended to keep it disabled. This option is not automatically enabled when any feature is specified in the security.digest_authentication.features_list property.	
security.digest_authentication.realm	Realm string for HTTP Digest Authentication. If it is changed, all users have to reset their passwords, otherwise they won't be able to access the server features protected by HTTP digest Authentication.	CloverETL Server
security.digest_authentication.nonce_validity	Interval of validity for HTTP Digest Authentication specified in seconds. When the interval passes, server requires new authentication from the client. Most of the HTTP clients do it automatically.	300
security.lockout.login.attempts	The number of failed login attempts after which a next failed login attempt will lock the user. Set the value to 0 to disable the function. Since 4.8.0M1.	50
security.lockout.reset.period	Period of time in seconds during which the failed login attempts are counted. Since 4.8.0M1.	60
security.lockout.unlock.period	Period of time in seconds after which a successful login attempt will unlock the previously locked user. Since 4.8.0M1.	300
<b>SMTP</b>		
clover.smtp.transport.protocol	SMTP server protocol. Possible values are "smtp" or "smtps".	smtp
clover.smtp.host	SMTP server hostname or IP address	
clover.smtp.port	SMTP server port	
clover.smtp.authentication	true/false If it is false, username and password are ignored.	
clover.smtp.username	SMTP server username	
clover.smtp.password	SMTP server password	
clover.smtp.additional.*	Properties with a "clover.smtp.additional." prefix are automatically added (without the prefix) to the Properties instance passed to the Mailer. May be useful for some protocol specific parameters. The prefix is removed.	
<b>Logging</b>		
logging.project_name	Used in log messages where it is necessary to name the product name.	CloverETL
logging.default_subdir	Name of a default subdirectory for all server logs; it is relative to the path specified by system property "java.io.tmpdir". Don't specify as an absolute path, use properties which are intended for absolute path.	cloverlogs
logging.logger.server_audit.enabled	Enables logging of operations called on ServerFacade and JDBC proxy interfaces. The name of the output file is "server-audit.log". It is stored in the same directory as other CloverETL Server log files by default. The default	false

Chapter 14. List of Configuration Properties

Key	Description	Default Value
	logging level is DEBUG so it logs all operations which may process any change.	
logging.logger.server_integration.enabled	Enables logging of Designer-Server calls. The name of the output file is "server-integration.log". It is stored in the same directory as other CloverETL Server log files by default. The default logging level is INFO. Username is logged, if available. JDBC and CTL debugging is not logged.	true
launch.log_dir	Location, where server should store launch requests logs. See Chapter 37, <a href="#">Launch Services</a> (p. 282) for details.	\${java.io.tmpdir}/[logging.default_subdir]/launch where \${java.io.tmpdir} is system property
graph.logs_path	Location, where server should store Graph run logs. See Chapter 16, <a href="#">Logging</a> (p. 127) for details.	\${java.io.tmpdir}/[logging.default_subdir]/graph where \${java.io.tmpdir} is system property
logging.appender.jobs.pattern_layout	Pattern of the jobs' log messages	%d %-5p %-3X{runId} [%t] %m%n
logging.appender.jobs.encoding	Encoding of the jobs' log files	UTF-8
logging.mem_appender.WORKER.pattern_layout	Format of log that can be seen in <b>Monitoring &gt; Logs &gt; Worker</b> .	
logging.mem_appender.WORKER.size_limit	Size of log that can be seen in <b>Monitoring &gt; Logs &gt; Worker</b> .	
<b>Thread Manager</b>		
threadManager.pool.corePoolSize	Number of threads which are always active (running or idling). Related to a thread pool for processing server events.	4
threadManager.pool.queueCapacity	Max size of the queue (FIFO) which contains tasks waiting for an available thread. Related to a thread pool for processing server events. For queueCapacity=0, there are no waiting tasks, each task is immediately executed in an available thread or in a new thread.	0
threadManager.pool.maxPoolSize	Max number of active threads. If no thread from a core pool is available, the pool creates new threads up to "maxPoolSize" threads. If there are more concurrent	8192

Chapter 14. List of Configuration Properties

Key	Description	Default Value
	tasks then maxPoolSize, thread manager refuses to execute it.	
threadManager.pool.allowCoreThreadTimeOut	Switch for idling threads timeout. If true, the "corePoolSize" is ignored so all idling threads may be time-outed	false
threadManager.pool.keepAliveSeconds	timeout for idling threads in seconds	20
<b>Archivator</b>		
task.archivator.batch_size	Max number of records deleted in one batch. It is used for deleting of archived run records.	50
task.archivator.archive_file_prefix	Prefix of archive files created by the archivator.	cloverArchive_
<b>Properties resolver</b>		
properties_resolver.resolve_server_props.server_props_list_additional	A list of properties from a subset of properties, which values are resolved. The properties' values may use system properties or environment variables as placeholders. The values are resolved during the server startup. If the system property is changed later, the resolved CloverETL Server property value doesn't change. Users may use this property, if some property they need to resolve is missing in the property: properties_resolver.resolve_server_props.server_props_list_default. If the property to resolve is already specified by the property properties_resolver.resolve_server_props.server_props_list_default, don't add it to this property.	
properties_resolver.resolve_server_props.server_props_list_default	A list of properties from a subset of properties, which values are resolved. The properties' values may use system properties or environment variables as placeholders. Values are resolved during the server startup. If the system property is changed later, the resolved CloverETL Server property value doesn't change. Users are discouraged from modification of the property, unless it's necessary. Instead, users may add more properties by modifying property: properties_resolver.resolve_server_props.server_props_list_additional	sandboxes.home, sandboxes.home.local, sandboxes.home.partitioned, cluster.jgroups.bind_address, cluster.jgroups.start_port, cluster.jgroups.external_address, cluster.jgroups.external_port, cluster.jgroups.tcpping.initial_hosts, cluster.group.name, cluster.http.url
properties_resolver.placeholders	server_props_list_default A list of properties from a subset of properties, that may be used as placeholders and shall be resolved if used in paths. The properties can be used if you define a path to the root of a sandbox, or to locations of local or partitioned sandboxes, or path to a script, or path in archiver job. Users are strongly discouraged from modification of the property. The property name changed since CloverETL 4.2, however the obsolete	sandboxes.home, sandboxes.home.local, sandboxes.home.partitioned, user.data.home

Chapter 14. List of  
Configuration Properties

Key	Description	Default Value
	name is also still accepted to maintain backwards compatibility.	
<b>Data Services</b>		
dataservice.invocation.record.max.age	It sets the maximal age in minutes before the record is removed from the database. The default is 1440 min = 24 h.	1440
dataservice.failure.ratio.min.record.count	Used for Data Service failure indication. It represents the minimum number of invocations required to evaluate whether the percentage of failures is over the threshold. Ensures that during periods of low traffic the endpoint does not switch to failing state. 10 by default.	10
<b>API</b>		
launch.http_header_prefix	Prefix of HTTP headers added by launch services to the HTTP response.	X-cloveretl
http.api.enabled	Enables or disables simple HTTP API.  If the HTTP API is disabled, there is no link to HTTP API operations in login page, the HTTP API tab in Launch Service that is accessible under <b>Test</b> button is not visible, and the HTTP API, the <code>/clover/httpapi.jsp</code> and HTTP API servlet are not accessible.  Available since 4.8.0M1. See Chapter 34, <a href="#">Simple HTTP API</a> (p. 267).	true
webDav.method.propfind.maxDepth	Maximum depth for webDAV method PROPFIND. When the depth is not specified, the default is supposed to be infinite (according to the rfc2518), however it's necessary to set some limit, otherwise the webDav client might overload the server filesystem.  Also if the depth value specified by webDAV client in the request is higher than the pre-configured max depth, only the pre-configured maximum is used.	40
<b>JVM</b>		
server.env.min_heap_memory	Sets the required minimal heap memory threshold. If the configuration of CloverETL Server is set to less heap memory, a warning is displayed. Experienced users can change the default value to avoid the warning when running the server on a system with lower memory. The threshold is in megabytes.	900
server.env.min_nonheap_memory	Sets the required minimal non-heap memory threshold. If the configuration of CloverETL server is set to less non-heap memory, a warning is displayed. Experienced users can change the default value to avoid the warning when running the server on a system with lower memory. The threshold is in megabytes.	256

Chapter 14. List of  
Configuration Properties

Key	Description	Default Value
jvm.implementation.check.enabled	Displays warnings when unsupported Java implementation is used.	true
<b>Misc</b>		
temp.default_subdir	Name of a default subdirectory for server tmp files; it is relative to the path specified by system property "java.io.tmpdir".	clovertmp
graph.pass_event_params_to_graph_in_old_style	Since 3.0. It is a switch for backwards compatibility of passing parameters to the graph executed by a graph event. In versions prior to 3.0, all parameters are passed to executed graph. Since 3.0, just specified parameters are passed. Please see <a href="#">Start a Graph</a> (p. 211) for details.	false
cluster.node.invocation.record.info.interval	Sets the interval for synchronization of the Data Services health state between the cluster nodes. The time is in milliseconds.	30000
clover.event.fileCheckMinInterval	Interval of the timer, running file event listener checks (in milliseconds). See <a href="#">File Event Listeners (remote and local)</a> (p. 253) for details.	1000
clover.event.groovyCheckMinInterval		1000
clover.inDevelopment		1000
launchservices.visible	Display Launch Services (p. 282) in the main menu in server UI. Can be true or false.	false

## Worker - Configuration Properties

Table 14.2. Server - Worker configuration

Key	Description	Default Value
worker.initialWorkers	<p>Enable/disable the Worker. To enable Worker, set to 1 (this is the default). <b>To disable Worker and run all jobs in Core Server, set to 0.</b></p> <p>Starting more than one Worker is currently (in 4.9.0) not supported.</p>	1
worker.portRange	<p>Port range used for communication between Server Core and Worker and between Workers on different cluster nodes. Communication between Server Core and Worker is done on localhost. Workers on different cluster nodes communicate directly with each other over these ports - in Cluster setup, this port range should be open in firewall for other Cluster nodes.</p> <p>This property can be easily configured in the Worker (p. 78) tab of Setup (p. 74).</p> <p><code>worker.portRange</code> should contain at least 5 ports for 1 node (depending on other options, a node takes at most 5 ports from the range). We recommend to use <code>portRange</code> of at least 10 ports to avoid possible problems with occupied ports after restart of Worker.</p> <p>If more cluster nodes run on the same machine, make sure that there are enough free ports for Workers of all cluster nodes on the machine. The default configuration of <code>worker.portRange</code> is sufficient for that.</p>	10500-10600
worker.connectTimeout	<p>Timeout for connection initialization between Worker and Server Core, in both directions. The timeout is in milliseconds.</p> <p>This setting can be useful in case of handling communication issues between Server Core and Worker, typically under high load you might want to increase the timeout.</p>	5000
worker.readTimeout	<p>Read timeout for communication requests between Worker and Server Core, in both directions. If a request is not completely served before reaching this limit, the connection is terminated. The timeout is in milliseconds.</p> <p>This setting can be useful in case of handling communication issues between Server Core and Worker, typically under high load you might want to increase the timeout.</p>	600000
worker.classpath	<p>A directory with additional <code>.jar</code> files to be added to the Worker's classpath. The <code>.jar</code> files would typically be libraries used by ETL graphs (e.g. JDBC drivers for database connections) or JDBC drivers used in JNDI connections defined in Worker (see <a href="#">Worker - JNDI Properties</a> (p. 117)).</p> <p>The Worker's classpath is separate from Server Core (i.e. application container classpath). Any libraries needed by</p>	<code>\${user.data.home}/CloverETL/worker-lib</code>

Chapter 14. List of  
Configuration Properties

Key	Description	Default Value
	<p>jobs executed on Worker need to be added on the Worker's classpath.</p> <p>The property has a default value that points to the <code>CloverETL/worker-lib</code> directory in the user's home directory. The directory is automatically scanned and all libraries found are added to the Worker's classpath.</p> <p>The property can contain paths to multiple directories. The separator between the directories can be a colon (on Linux and Mac) or semicolon (Linux, Mac and Windows), e.g.:</p> <pre>worker.classpath=/home/clover/worker-lib;/opt/worker-lib-2</pre> <p>If a directory is added on the Worker's classpath, its subdirectories are automatically added too.</p> <p>Some basic wildcards are supported: <code>directory-*</code> and <code>directory-?</code>.</p>	
worker.maxHeapSize	<p>The maximum Java heap size of Worker in MB, it will be translated to the <code>-Xmx</code> option for the Worker's JVM. Jobs executed in the Worker require heap memory based on their complexity, dataset size, etc.</p> <p>See our recommendations (p. 60) for heap sizes of Worker and Server Core.</p> <p>This property can be easily configured in the Worker (p. 78) tab of Setup (p. 74).</p> <p>Setting to 0 uses Java default heap size (automatically determined by Java). This setting is not recommended for production usage.</p>	0
worker.initHeapSize	<p>The initial Java heap size of Worker in MB, it will be translated to the <code>-Xmx</code> option for the Worker's JVM. We recommend to set this to the same value as <code>worker.maxHeapSize</code></p> <p>This property can be easily configured in the Worker (p. 78) tab of Setup (p. 74).</p> <p>Setting to 0 uses Java default initial heap size (automatically determined by Java). This setting is not recommended for production usage.</p>	0
worker.jvmOptions	<p>Adds Java command line options for the Worker's JVM. This property is useful to tweak the configuration of the Worker's JVM, e.g. to tune garbage collector settings. These command line options override default options of the JVM.</p> <p>E.g. enable G1 garbage collector: <code>-XX:+UseG1GC -XX:MaxGCPauseMillis=200</code>.</p> <p>See Additional Diagnostic Tools(p. 196)section for useful options for troubleshooting and debugging Worker.</p>	

Chapter 14. List of  
Configuration Properties

Key	Description	Default Value
	This property can be easily configured in the Worker (p. 78) tab of Setup (p. 74).	
worker.enableDebug	Remote Java debugging of Worker, enables <a href="#">JDWP</a> . Enabling this allows you to connect a Java debugger remotely to the running Worker process, to debug your Java transformations, investigate issues, etc. The port used by the debugger is determined dynamically and can be seen in the Worker section of the Monitoring page.	false
worker.inheritSystemProperties	<p>Sets whether system Java properties are inherited from the Server Core process to the Worker process. We automatically inherit some system properties to simplify the Worker configuration, see below for a list.</p> <p>This functionality is enabled by default. Use this property to disable this behavior in case some of the inherited properties would cause issues.</p> <p>The following system Java properties are inherited from the Server Core to Worker:</p> <pre># Clover properties com.opensys.cloveretl.addressdoctor.setConfigFile com.opensys.cloveretl.addressdoctor.persistent cloveretl.smb2.bouncycastle.jar.file  # Standard Java properties java.library.path java.io.tmpdir java.rmi.server.hostname XX:MaxPermSize XX:MaxMetaspaceSize  # SSL related properties javax.net.ssl.keyStore, javax.net.ssl.keyStorePassword javax.net.ssl.trustStore, javax.net.ssl.trustStorePassword javax.net.ssl.keyAlias https.protocols  # Proxy configuration *.proxyHost, *.proxyPort *.proxyUser, *.proxyPassword, *.nonProxyHosts socksProxyHost, socksProxyPort, socksProxyVersion java.net.socks.username, java.net.socks.password java.net.useSystemProxies</pre>	true
worker.javaExecutable	<p>Absolute path to the Java binary for Worker process, e.g. /user/local/java/bin/java. Value is automatically determined based on \$JAVA_HOME environment variable.</p> <p>Use this property if you need to use a specific Java binary for running the Worker.</p>	

## Worker - JNDI Properties

The Worker has its own JNDI pool separate from the application container JNDI pool. If your jobs use JNDI resources (to obtain JDBC or JMS connections), you have to configure the Worker's JNDI pool and its resources.

The worker JNDI properties must be configured using the `clover.properties` configuration file. Libraries used by the JNDI resources must be added to the Worker's classpath, see `worker.classpath` (p. 114).

It is possible to define multiple datasources pointing to different databases or JMS queues, see examples below. The datasources are indexed in configuration, their properties have suffix `[0]`, `[1]`, etc. Even a single datasource must have the `[0]` index.

## JDBC Datasources

Worker uses the Apache DBCP2 pool for its JNDI functionality. Any DBCP2 configuration attribute is supported, see [DBCP attributes](#). The only mandatory properties are `jdbcName` and `url`.

See table below for basic JNDI properties.

You can monitor the state of the datasources via JMX. See [Additional Diagnostic Tools](#) (p. 196) for details on how to enable JMX on Worker. Then you can connect to the Worker's JMX interface with tools like `jconsole` and monitor the JNDI datasources, e.g. for the number of currently open connections. The related MBeans are under the `Tomcat/DataSource/localhost//javax.sql.DataSource` path:

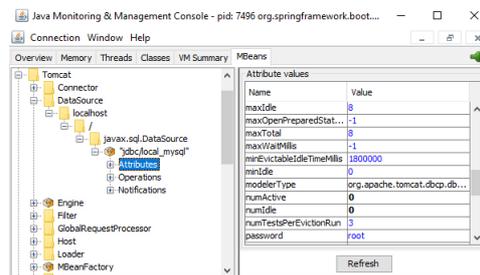


Figure 14.1. MBean for a JNDI datasource in jconsole

Table 14.3. Properties for JDBC JNDI Resources in Worker

Key	Description	Example
worker.jndi.datasource[0].jndiName	The name of the JNDI datasource. Mandatory.	jdbc/database_name
worker.jndi.datasource[0].url	The JDBC connection URL. Mandatory.	jdbc:postgresql:// hostname:5432/ database_name
worker.jndi.datasource[0].username	The user name for a database connection.	clover
worker.jndi.datasource[0].password	The password for a database connection. The password value can be encrypted using the secure configuration tool, see Secure Configuration Properties (p. 124).	clover
worker.jndi.datasource[0].driverClassName	The database driver classname. The database driver must be on the Worker classpath, see worker.classpath (p. 114).	org.postgresql.Driver
worker.jndi.datasource[0].maxIdle	The maximum number of idle database connections in a pool. Set to -1 for no limit.	10
worker.jndi.datasource[0].maxTotal	The maximum number of database connections in a pool. Set to -1 for no limit.	20
worker.jndi.datasource[0].maxWaitMillis	The maximum time Worker waits for a database connection to become available. In milliseconds, set to -1 for no limit.	30000
worker.jndi.datasource[0].dbcpAttribute	Any DBCP2 attribute, e.g. worker.jndi.datasource[0].initSQL. See <a href="#">DBCP attributes</a> .	

The following example shows configuration of two JDBC Datasources.

```
worker.jndi.datasource[0].jndiName=jdbc/postgresql_finance
worker.jndi.datasource[0].url=jdbc:postgresql://finance.example.com:5432/finance
worker.jndi.datasource[0].maxIdle=5
worker.jndi.datasource[0].maxTotal=10
worker.jndi.datasource[0].maxWaitMillis=-1
worker.jndi.datasource[0].username=finance_user
worker.jndi.datasource[0].password=conf#eCf1GD1DtKSJjh9VyD1Rh7IftAbI/vsH
worker.jndi.datasource[0].driverClassName=org.postgresql.Driver

worker.jndi.datasource[1].jndiName=jdbc/MySQLDB
worker.jndi.datasource[1].url=jdbc:mysql://marketing.example.com:3306/marketing?useUnicode=true&characterEncoding=utf8
worker.jndi.datasource[1].maxIdle=10
worker.jndi.datasource[1].maxTotal=20
worker.jndi.datasource[1].maxWaitMillis=-1
worker.jndi.datasource[1].username=marketing_user
worker.jndi.datasource[1].password=conf#JWsMa2okg7Dq2gtLBM84sE==
worker.jndi.datasource[1].driverClassName=com.mysql.jdbc.Driver
```

## JMS Connections

Worker can use any JMS broker to define JMS connections in JNDI. Any JMS broker configuration attribute is supported. The mandatory properties are `jndiName`, `factoryClass`, `typeClass` and `typeInterface`. See table below for basic JNDI properties for JMS resources.

Table 14.4. Properties for JMS JNDI Resources in Worker

Key	Description	Example
worker.jndi.jms[0].jndiName	The name of the JNDI JMS resource. Mandatory.	jms/jms_queue
worker.jndi.jms[0].factory	Factory class for creating the JMS resource. This is JMS broker specific. Mandatory.	org.apache.activemq.jndi.JNDIReferenceFactory
worker.jndi.jms[0].type	Implementation class of the JMS resource. This is JMS broker specific. Mandatory.	org.apache.activemq.command.ActiveMQQueue
worker.jndi.jms[0].jmsProperty	Configuration property for the JMS resource. Any configuration property supported by the JMS broker can be used.	worker.jndi.jms[0].brokerUrl

The following example shows configuration of several JMS resources.

```
worker.jndi.jms[0].jndiName=jms/CloverConnectionFactory
worker.jndi.jms[0].type=org.apache.activemq.ActiveMQConnectionFactory
worker.jndi.jms[0].factory=org.apache.activemq.jndi.JNDIReferenceFactory
worker.jndi.jms[0].brokerUrl=tcp://localhost:61616?jms.prefetchPolicy.queuePrefetch=1
worker.jndi.jms[0].brokerName=LocalActiveMQBroker

worker.jndi.jms[1].jndiName=jms/CloverQueue
worker.jndi.jms[1].type=org.apache.activemq.command.ActiveMQQueue
worker.jndi.jms[1].factory=org.apache.activemq.jndi.JNDIReferenceFactory
worker.jndi.jms[1].physicalName=TestQueue
```

## Worker - SSL Properties

In Cluster, Workers of each node communicate with each other directly for increased performance. This communication is used to transport data of cluster remote edges in clustered jobs between the nodes. For increased security, it is possible to use SSL for the remote edge communication.

SSL communication between Workers needs to be enabled and configured separately from SSL of the application container that runs Server Core. The `worker.ssl.enabled` property is used to enable/disable SSL. If a Cluster node's "self" URL is using HTTPS, we automatically set the property to true. Configuration of SSL consists of setting paths and passwords of KeyStore and TrustStore, see the table below for details.

Note that if the standard SSL related system properties (`javax.net.ssl.keyStore`, `javax.net.ssl.keyStorePassword`, `javax.net.ssl.keyAlias`, `javax.net.ssl.trustStore` and `javax.net.ssl.trustStorePassword`) are used to configure KeyStore/TrustStore for the Server Core, they are propagated to Worker; therefore, their respective `worker.ssl` properties do not need to be configured.

Recommended steps to enable SSL for inter-worker communication are:

- Enable SSL for each cluster node, via the application container settings. Configure TrustStore and KeyStore via the standard `javax.net.ssl.*` properties.
- Set `cluster.http.url` for each node to point to its own HTTPS URL
- Check that communication between Cluster nodes over SSL works and that the nodes can correctly see each other. The Monitoring page of Server Console should show the whole cluster group and its nodes correctly.
- Worker should automatically inherit the above SSL configuration.
- Run a clustered job on Worker

Table 14.5. Properties for SSL communication in Worker

Key	Description	Example
<code>worker.ssl.enabled</code>	Enables or disables an SSL connection for Worker. Note that if the Server runs on HTTPS, SSL is enabled automatically; however, this property has a higher priority.	<code>true/false</code>
<code>worker.ssl.keyStore</code>	Absolute path to the KeyStore file.	<code>path/to/keyStore.file</code>
<code>worker.ssl.keyStorePassword</code>	The KeyStore password.	
<code>worker.ssl.keyAlias</code>	The alias of the key in keyStore. <b>Optional</b> - the property does not have to be specified if there is only one key in the KeyStore.	
<code>worker.ssl.port</code>	The port for SSL communication with Worker. The property is configured automatically and the value is set from <a href="#">worker.portRange</a> (p. 114).	

## Job Execution Properties

Table 14.6. Defaults for job execution configuration - see [Job Config Properties](#) (p. 183) for details

Key	Description	Default Value
executor.tracking_interval	An interval in milliseconds for scanning of a current status of a running graph. The shorter interval, the bigger log file.	2000
executor.log_level	Log level of graph runs. TRACE   DEBUG   INFO   WARN   ERROR	INFO
executor.max_job_tree_depth	Defines maximal depth of the job execution tree, e.g. for recursive job it defines the maximal level of recursion (counting from root job).	32
executor.max_running_concurrently	Amount of graph instances which may exist (or run) concurrently. 0 means no limits.	0
executor.max_graph_instance_age	Specifies how long can a graph instance be idling before it is released from memory. Interval is in milliseconds. 0 means no caching.  This property has been renamed since 2.8. Original name was executor.maxGraphInstanceAge	0
executor.classpath	Classpath for transformation/processor classes used in the graph. Directory [Sandbox_root]/trans/ does not have to be listed here, since it is automatically added to a graph run classpath.	
executor.skip_check_config	Disables check of graph configuration. Increases performance of a graph execution; however, it may be useful during graph development.	true
executor.password	This property is deprecated. The password for decoding encoded DB connection passwords.	
executor.verbose_mode	If true, more descriptive logs of graph runs are generated.	true
executor.use_jmx	If true, the graph executor registers JMX mBean of the running graph.	true
executor.debug_mode	If true, edges with enabled debug store data into files in debug directory.	false

## List of all properties

[clover.event.fileCheckMinInterval](#) (p. 113)  
[clover.event.fileCheckMinInterval](#) (p. 113)  
[clover.inDevelopment](#) (p. 113)  
[clover.smtp.additional.\\*](#) (p. 109)  
[clover.smtp.authentication](#) (p. 109)  
[clover.smtp.host](#) (p. 109)  
[clover.smtp.password](#) (p. 109)  
[clover.smtp.port](#) (p. 109)  
[clover.smtp.transport.protocol](#) (p. 109)  
[clover.smtp.username](#) (p. 109)  
[cluster.node.invocation.record.info.interval](#) (p. 113)

[config.file](#) (p. 106)  
[dataservice.invocation.record.max.age](#) (p. 112)  
[dataservice.failure.ratio.min.record.count](#) (p. 112)  
[datasource.type](#) (p. 107)  
[datasource.jndiName](#) (p. 107)  
[engine.config.file](#) (p. 107)  
[engine.plugins.additional.src](#) (p. 107)  
[executor.classpath](#) (p. 121)  
[executor.debug\\_mode](#) (p. 121)  
[executor.log\\_level](#) (p. 121)  
[executor.max\\_job\\_tree\\_depth](#) (p. 121)  
[executor.max\\_running\\_concurrently](#) (p. 121)  
[executor.max\\_graph\\_instance\\_age](#) (p. 121)  
[executor.password](#) (p. 121)  
[executor.skip\\_check\\_config](#) (p. 121)  
[executor.tracking.interval](#) (p. 121)  
[executor.use\\_jmx](#) (p. 121)  
[executor.verbose\\_mode](#) (p. 121)  
[graph.logs\\_path](#) (p. 110)  
[graph.pass\\_event\\_params\\_to\\_graph\\_in\\_old\\_style](#) (p. 113)  
[http.api.enabled](#) (p. 112)  
[jdbc.dialect](#) (p. 108)  
[jdbc.driverClassName](#) (p. 107)  
[jdbc.password](#) (p. 108)  
[jdbc.url](#) (p. 107)  
[jdbc.username](#) (p. 107)  
[jvm.implementation.check.enabled](#) (p. 113)  
[launch.http\\_header\\_prefix](#) (p. 112)  
[launch.log\\_dir](#) (p. 110)  
[launchservices.visible](#) (p. 113)  
[license.context\\_names](#) (p. 106)  
[license.file](#) (p. 106)  
[logging.appender.jobs.encoding](#) (p. 110)  
[logging-appender-jobs-pattern\\_layout](#) (p. 110)  
[logging.default\\_subdir](#) (p. 109)  
[logging.logger.server\\_audit.enabled](#) (p. 109)  
[logging.logger.server\\_integration.enabled](#) (p. 110)  
[logging.mem\\_appender.WORKER.pattern\\_layout](#) (p. 110)  
[logging.mem\\_appender.WORKER.size\\_limit](#) (p. 110)  
[logging.project\\_name](#) (p. 109)  
[private.properties](#) (p. 108)  
[properties\\_resolver.placeholders.server\\_props\\_list\\_default](#) (p. 111)  
[properties\\_resolver.resolve\\_server\\_props.server\\_props\\_list\\_additional](#) (p. 111)  
[properties\\_resolver.resolve\\_server\\_props.server\\_props\\_list\\_default](#) (p. 111)  
[quartz.driverDelegateClass](#) (p. 108)  
[sandboxes.access.check.boundaries.enabled](#) (p. 107)  
[sandboxes.home](#) (p. 107)  
[security.basic\\_authentication.features\\_list](#) (p. 108)  
[security.basic\\_authentication.realm](#) (p. 108)  
[security.default\\_domain](#) (p. 108)  
[security.digest\\_authentication.features\\_list](#) (p. 108)  
[security.digest\\_authentication.nonce\\_validity](#) (p. 109)  
[security.digest\\_authentication.realm](#) (p. 109)  
[security.digest\\_authentication.storeAll.enabled](#) (p. 108)  
[security.lockout.login.attempts](#) (p. 109)  
[security.lockout.reset.period](#) (p. 109)  
[security.lockout.unlock.period](#) (p. 109)  
[security.session.exchange.limit](#) (p. 108)

[security.session.validity](#) (p. 108)  
[server.env.min\\_heap\\_memory](#) (p. 112)  
[server.env.min\\_nonheap\\_memory](#) (p. 112)  
[task.archivator.archive\\_file\\_prefix](#) (p. 111)  
[task.archivator.batch\\_size](#) (p. 111)  
[temp.default\\_subdir](#) (p. 113)  
[threadManager.pool.allowCoreThreadTimeOut](#) (p. 111)  
[threadManager.pool.corePoolSize](#) (p. 110)  
[threadManager.pool.keepAliveSeconds](#) (p. 111)  
[threadManager.pool.maxPoolSize](#) (p. 110)  
[threadManager.pool.queueCapacity](#) (p. 110)  
[webDav.method.propfind.maxDepth](#) (p. 112)  
[worker.additional.classpath](#) (p. 114)  
[worker.enableDebug](#) (p. 116)  
[worker.inheritSystemProperties](#) (p. 116)  
[worker.initHeapSize](#) (p. 115)  
[worker.javaExecutable](#) (p. 116)  
[worker.jvmOptions](#) (p. 115)  
[worker.maxHeapSize](#) (p. 115)  
[worker.portRange](#) (p. 114)  
[worker.connectTimeout](#) (p. 114)  
[worker.readTimeout](#) (p. 114)  
[worker.initialWorkers](#) (p. 114)  
[worker.jndi.datasource\[0\].url](#) (p. 118)  
[worker.jndi.datasource\[0\].maxIdle](#) (p. 118)  
[worker.jndi.datasource\[0\].maxTotal](#) (p. 118)  
[worker.jndi.datasource\[0\].maxWaitMillis](#) (p. 118)  
[worker.jndi.datasource\[0\].username](#) (p. 118)  
[worker.jndi.datasource\[0\].password](#) (p. 118)  
[worker.jndi.datasource\[0\].driverClassName](#) (p. 118)  
[worker.jndi.datasource\[0\].jndiName](#) (p. 118)  
[worker.jndi.jms\[0\].jndiName](#) (p. 119)  
[worker.jndi.jms\[0\].factory](#) (p. 119)  
[worker.jndi.jms\[0\].type](#) (p. 119)  
[worker.jndi.jms\[0\].jmsProperty](#) (p. 119)  
[worker.ssl.enabled](#) (p. 120)  
[worker.ssl.keyStore](#) (p. 120)  
[worker.ssl.keyStorePassword](#) (p. 120)  
[worker.ssl.keyAlias](#) (p. 120)

---

# Chapter 15. Secure Configuration Properties

Some configuration properties can be confidential (e.g. a password to a database, mail client, etc.) and thus it's desirable to encrypt them. For this purpose, there is a command-line utility *secure-cfg-tool.jar*.

[Basic Utility Usage](#) (p. 124)

[Advanced Usage - Custom Settings](#) (p. 125)

---

## Basic Utility Usage

1. Download the utility archive file (*secure-cfg-tool.zip*) and unzip it.

The utility is available in the download section of your CloverETL account - at the same location as the download of **CloverETL Server**.

2. Execute the script given for your operating system, *encrypt.bat* for MS Windows, *encrypt.sh* for Linux. You will be asked for inserting a value of a configuration property intended to be encrypted.

Example:

```
C:\secure-cfg-tool>encrypt.bat
```

```
*****
```

```
Secure config encryption (use --help or -h option to show help)
```

```
*****
```

```
***** Config settings *****
```

```
Provider: SunJCE
```

```
Algorithm: PBESWithMD5AndDES
```

```
*****
```

```
Enter text to encrypt: mypassword
```

```
Text to encrypt: "mypassword"
```

```
Encrypted text: conf#eCf1GD1DtKSJjh9VyDIRh7IftAbI/vsH
```

```
C:\secure-cfg-tool>
```

If you want to configure the way the values are encrypted, see [Advanced Usage - Custom Settings](#) (p. 125)

3. The encrypted string has *conf#encrypted\_property* format and can be used as a value of a configuration property in the properties file, *clover.xml* file or *web.xml* file (see details about configuration sources in Chapter 11, [Configuration Sources](#) (p. 70)).

**Example of a configuration property file with encrypted password:**

```
jdbc.driverClassName=org.postgresql.Driver
jdbc.url=jdbc:postgresql://127.0.0.1/clover_db?charSet=UTF-8
jdbc.username=yourUsername
jdbc.password=conf#eCf1GD1DtKSJjh9VyDIRh7IftAbI/vsH
jdbc.dialect=org.hibernate.dialect.PostgreSQLDialect
```

Alternatively, you can use the following command:

```
java -jar secure-cfg-tool.jar
```



## Important

Values encrypted by a Secure parameter form (Chapter 19, [Secure Parameters](#) (p. 147)) **cannot** be used as a value of a configuration property.

## Advanced Usage - Custom Settings

The way of encrypting configuration values described above uses default configuration settings (a default provider and algorithm). If you need to customize the settings, use the following parameters of the *secure-cfg-tool.jar* utility.

Table 15.1. Parameters

Parameter	Description	Example
--algorithm, -a	algorithm to encrypt	--algorithm PBEWithMD5AndDES
--file, -f	config file location	-f C:\User\John\cloverServer.properties
--help, -h	show help	--help
--providerclass, -c	custom provider class	-c org.provider.ProviderClass
--providerlocation, -l	path to jar/folder containing a custom provider class (it will be added to the classpath)	--providerlocation C:\User\John\lib\customprovider.jar, -l C:\User\John\lib\
--providers, -p	print available security providers and their algorithms	--providers



## Note

To demonstrate usage of an external provider the Bouncy Castle provider is used.

To find out a list of algorithms, use **-p** or **--providers**

```
C:\secure-cfg-tool>encrypt.bat -p
```

If you want to find out a list of algorithms of an external provider, you must pass the provider's class name and path to jar file(s).

```
C:\secure-cfg-tool>encrypt.bat org.bouncycastle.jce.provider.BouncyCastleProvider -p -c -l C:\User\John\bcprov-jdk15on-152.jar
```

Result might look like this:

```
***** List of available providers and their algorithms *****
```

```
Provider: SunJCE
Provider class: com.sun.crypto.provider.SunJCE
Algorithms:
PBEWithMD5AndDES
PBEWithSHA1AndDESede
PBEWithSHA1AndRC2_40
Provider: BC
Provider class: org.bouncycastle.jce.provider.BouncyCastleProvider
Algorithms:
PBEWITHMD2ANDDES
PBEWITHMD5AND128BITAES-CBC-OPENSSL
PBEWITHMD5AND192BITAES-CBC-OPENSSL
PBEWITHMD5AND256BITAES-CBC-OPENSSL
```

The provider class is displayed on the row starting with *Provider class*, algorithms are strings with *PBE* prefix. Both can be used to configure encryption.

### Configuring the Encryption Process

The algorithm and provider can be passed to the utility in two ways.

- **Using command line arguments**

To change the algorithm, use the argument **-a**. The provider remains default (SunJCE in case of Oracle Java):

```
C:\secure-cfg-tool>encrypt.bat -a PBEWithMD5AndDES
```

To use an external provider, you must specify the provider's class name (the **--providerclass** or **-c** arguments) and add jar(s) to the classpath (the **--providerlocation** or **-l** arguments). Provider location must point to a concrete jar file or directory containing the jar(s) and can be used several times for several paths:

```
C:\secure-cfg-tool>encrypt.bat -a PBEWITHSHA256AND256BITAES-CBC-BC -c  
org.bouncycastle.jce.provider.BouncyCastleProvider -l C:\User\John\bcprov-  
jdk15on-152.jar
```

- **Using configuration file**

A configuration file is a common properties file (text file with key-value pairs):

```
[property-key]=[property-value]
```

See the following example of `secure.config.example.properties` distributed within `secure-cfg-tool.zip`:

```
security.config_properties.encrypted.providerClassName=org.bouncycastle.jce.provider.BouncyCastleProvider  
security.config_properties.encrypted.algorithm=PBEWITHSHA256AND256BITAES-CBC-BC  
security.config_properties.encrypted.provider.location=C:\\User\\libs
```

You must also set the path to the file using the **-f** argument:

```
C:\secure-cfg-tool>encrypt.bat -f path/to/secure.config.example.properties
```



### Note

More jar locations can be set in the `security.config_properties.encrypted.providerLocation` property. The locations are delimited by semicolon.

### Configuring an application server

CloverETL Server application needs to know how the values have been encrypted, therefore the properties must be passed to the server (see details in Part III, “[Configuration](#)” (p. 67)). For example:

```
...  
security.config_properties.encrypted.providerClassName=org.bouncycastle.jce.provider.BouncyCastleProvider  
security.config_properties.encrypted.algorithm=PBEWITHSHA256AND256BITAES-CBC-BC  
...
```



### Important

If a third-party provider is used, its classes must be accessible to the application server. Property `security.config_properties.encrypted.providerLocation` will be ignored.

---

## Chapter 16. Logging

[Main Logs](#) (p. 127)

[Another Useful Logging Settings](#) (p. 127)

[Access Log in Apache Tomcat](#) (p. 128)

[Application Server Logs](#) (p. 128)

[Graph Run Logs](#) (p. 128)

[Server Audit Logs](#) (p. 128)

[Designer-Server Integration Logs](#) (p. 129)

---

### Main Logs

The CloverETL Server uses the log4j library for logging. The WAR file contains the default log4j configuration. The log4j configuration file `log4j.xml` is placed in `WEB-INF/classes` directory.

By default, log files are produced in the directory specified by the `java.io.tmpdir` system property in the `cloverlogs` subdirectory.

The `java.io.tmpdir` system property usually points to a common system temp directory, i.e. `/tmp`. On Apache Tomcat, it is usually the `$TOMCAT_HOME/temp` directory.

The default logging configuration (`log4j.xml` bundled in the `clover.war`) may be changed to another log4j configuration file using system property `log4j.configuration`. If you override the configuration, only the properties from the new file are used.

The `log4j.configuration` should contain the URL of the new log4j configuration file, not a simple file system path, for example:

```
log4j.configuration=file:/home/clover/config/log4j.xml
```



#### Tip

It is better to copy the original file and modify the copy, than to create a new file.

Please note that `log4j.configuration` is not a CloverETL Server configuration property, but a system property, thus it must be set on the JVM command line by `-Dlog4j.configuration` or in other way suitable for the application container. See the [Installation \(p. 12\)](#) chapter for more information on how to set a system property for each application container.

Since such a configuration overrides the default configuration, it may influence Graph run logs. So your own log configuration has to contain following fragment to preserve Graph run logs:

```
<logger name="Tracking" additivity="false">
  <level value="debug"/>
</logger>
```

---

### Another Useful Logging Settings

These system properties allow for logging of HTTP requests/responses to stdout:

Client side:

```
com.sun.xml.ws.transport.http.client.HttpTransportPipe.dump=true (For more information, see the Integrating CloverETL Designer with CloverETL Server chapter of the CloverETL Designer User's Guide.)
```

Server side:

```
com.sun.xml.ws.transport.http.HttpAdapter.dump=true
```

---

## Access Log in Apache Tomcat

If you need to log all requests processed by the server, add the following code to `/$CATALINA_HOME/conf/server.xml`.

```
<Valve className="org.apache.catalina.valves.AccessLogValve" directory="logs"
  prefix="localhost_access_log" suffix=".txt"
  pattern="%h %l %u %t %D %r %s %b" />
```

The format defined above has following meaning

[IP address] [date-time] [processing duration in milliseconds] [method] [URL] [protocol]

The log will look like the next line

```
172.17.30.243 - - [13/Nov/2014:12:53:03 +0000] 2 "POST /clover/sDispatcher/clusterNodeA
```

See also [Valve](#) in documentation on Apache Tomcat.

---

## Application Server Logs

If you use Apache Tomcat, it logs into `/$CATALINA_HOME/logs/catalina.out` file.

---

## Graph Run Logs

Each graph or jobflow run has its own log file – for example, in the Server Console, section "Executions History".

By default, these log files are saved in the subdirectory `cloverLogs/graph` in the directory specified by `java.io.tmpdir` system property.

It's possible to specify a different location for these logs with the CloverETL `graph.logs_path` property. This property does not influence main Server logs.

---

## Server Audit Logs

*Server Audit Log* logs operations called on *ServerFacade* and *JDBC proxy* interfaces.

Audit logging can be enabled by setting (adding) the value of CloverETL property `logging.logger.server_audit.enabled` to `true`. In server GUI, you can change the property value in **Configuration** → **Setup** → **Configuration File**. *Audit logging* is disabled by default.

The name of output file is `server-audit.log`. The file is in the same directory as main server log files. Default log level is `DEBUG`, so all operations which may do any change or another important operations (e.g. login or `openJdbcConnection`) are logged. To enable logging of all operations, change log level to `TRACE` in the `log4j` configuration.

Each logged operation is logged by two messages: entering method and exiting method (if the exception is raised, it's logged instead of output parameters)

- Entering method (marked as "inputParams"). All method's parameters (except for passwords) are printed.
- Exiting method (marked as "outputParams"). Method's return value is printed.
- Exception in method (marked as "EXCEPTION"). Exception's stacktrace is printed.

Message also contains:

- username, if the user is known
- client IP address, if it's known
- cluster node ID
- Interface name and the operation name

Values of transient and lazy initialized (in entity classes) fields and fields with binary content are not printed.

---

## Designer-Server Integration Logs

The logging of Designer-Server integration can be enabled with [logging.logger.server\\_integration.enabled](#) (p. 110) configuration property. The name of the log file is `server-integration.log`.

The log format is date and time, IP address of Designer, user name, operation, result of the operation (success/failure) and duration in milliseconds.

```
2018-03-07 16:42:00,525 10.0.3.2 user=clover, operation=executeGraph SUCCESS duration=576 ms
```

---

## Worker Log

*Worker log* logs operation performed by the Worker. The Worker is configured with a separate `log4j.xml` configuration file.

---

# **Part IV. Administration**

---

# Chapter 17. Monitoring

Monitoring section in the server Web GUI displays useful information about current performance of the standalone CloverETL Server or all cluster nodes if the clustering is enabled.

Monitoring section of the standalone server has slightly different design from cluster environment. In case of standalone server, the server-view is the same as node detail in cluster environment.

The section is refreshed every 15 seconds so the displayed data is up-to-date. The page can also be refreshed manually by the **Refresh** button.

## Standalone Server Detail

Standalone server detail view displays info collected from the standalone server. The info is grouped in several panels. The following ones are displayed by default.

- [Performance](#) (p. 132)
- [Resource Utilization](#) (p. 132)
- 10 longest-running jobs
- [System](#) (p. 133)
- [Worker](#) (p. 133)
- [Status History](#) (p. 134)

You can display the hidden actions with **Actions** button: choose **Actions** → **Show details**.

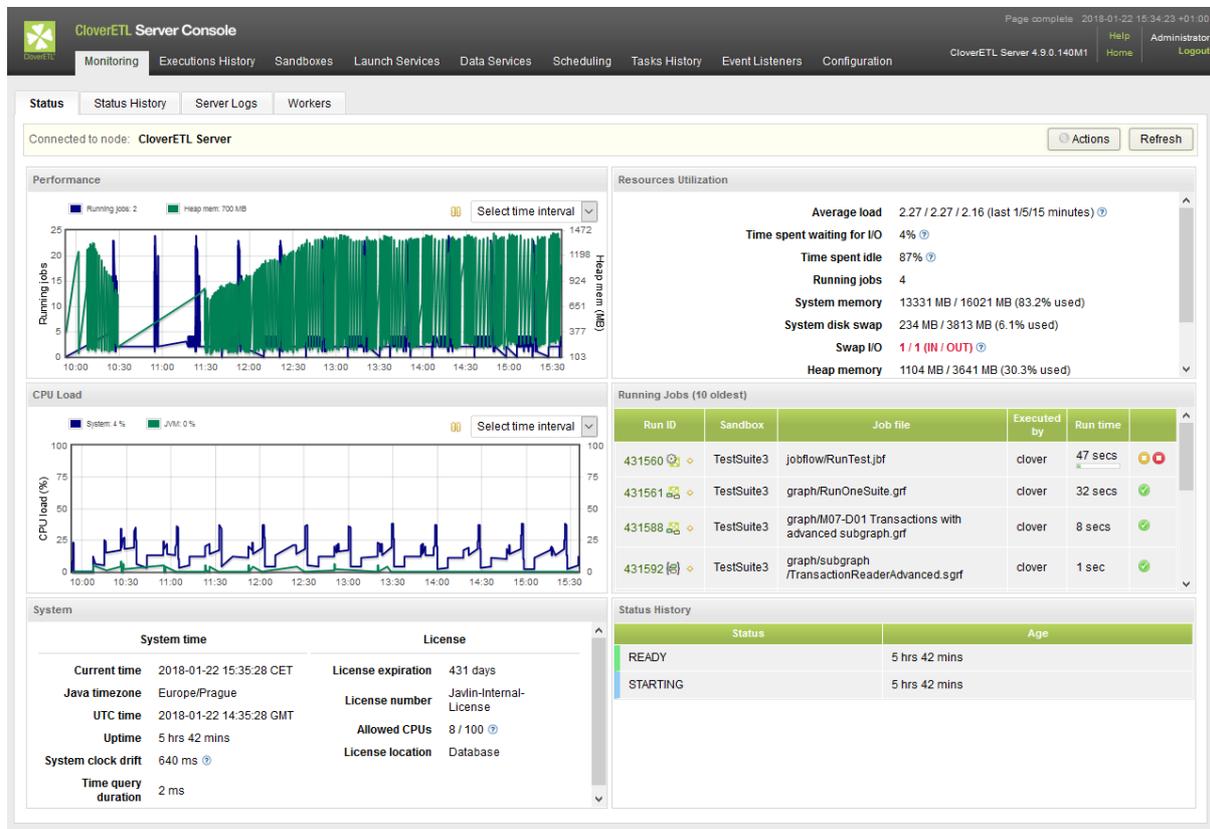


Figure 17.1. Standalone server detail

## Performance

The Performance panel contains a chart with two basic performance statistics: a number of running jobs and an amount of used heap memory. The graph displays values gathered within a specific interval. The interval can be set up with the combo box above the graph or it can be configured by "cluster.node.sendinfo.history.interval" config property. By default, the values are gathered within a couple of last minutes.

Note that the heap memory is constantly oscillating, even in an idle state, since it is periodically managed by JVM garbage collector (i.e. the temporary data required for running CloverETL Server is periodically removed from/allocated to the heap memory).

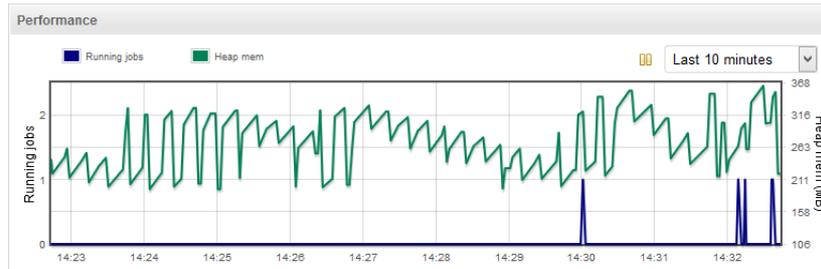


Figure 17.2. Performance

## Resource Utilization

Depending on the operating system, the Resource Utilization panel shows the following statistics:

Table 17.1. Resource Utilization

Resource	Operating System		Description
	UnWindo		
Average load	✓	✗	Shows the number of running/queued tasks per core in the last 1/5/15 minutes.
Time spent waiting for I/O	✓	✗	Shows the percentage of time the server is waiting for input/output data.
Time spent idle	✓	✗	Shows the percentage of time the server is idle (i.e. no jobs are running).
Running jobs	✓	✓	Shows the number of currently running jobs (e.g. graphs, jobflows).
System memory	✓	✓	Shows the <b>total</b> usage of RAM - including operating system and its processes, running programs, etc. (current / maximum (% of max used)).
System disk swap	✓	✓	Shows the usage of the swap memory (current / maximum (% of max used)).
Swap I/O	✓	✗	Amount of memory swapped to/from the disk (amount IN / amount OUT).
Heap memory	✓	✓	Monitors usage of the heap memory (current / maximum (% of max used)). For more information on heap memory, see Table 7.1, " <a href="#">JVM Memory Structure</a> " (p. 60).
Non-heap memory	✓	✓	Monitors usage of the non-heap memory (current / maximum (% of max used)). For more information on non-heap memory, see Table 7.1, " <a href="#">JVM Memory Structure</a> " (p. 60).
Event listeners	✓	✓	Shows the number of active event listeners. For more information, see Chapter 30, <a href="#">Listeners</a> (p. 236).

## CPU Load

The CPU Load panel displays a chart with info about total CPU load and CPU load caused by JVM.

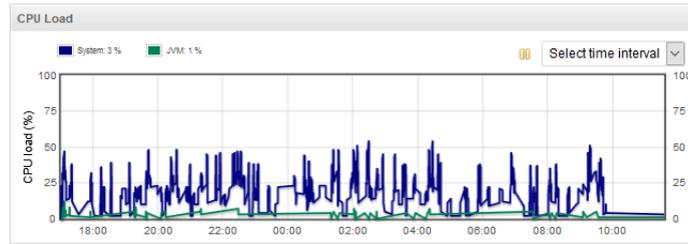


Figure 17.3. CPU Load

## Running Jobs

Running jobs panel lists currently running jobs, 10 oldest runs are displayed.

Run ID	Sandbox	Job file	Executed by	Run time	
205	JobflowExamples	jobflow/01-Automation-BasicFileProcessing.jbf	clover	1 sec	✓
206	JobflowExamples	graph/01-Automation-BasicFileProcessing/FileOperationsProcess.grf	clover	213 ms	✓
207	JobflowExamples	graph/01-Automation-BasicFileProcessing/FileOperationsProcess.grf	clover	96 ms	✓
208	JobflowExamples	graph/01-Automation-BasicFileProcessing/FileOperationsProcess.grf	clover	125 ms	✓

Figure 17.4. Running jobs

## System

System panel contains info about operating system and license.

System time		License	
Current time	2015-11-16 14:23:28 GMT	License expiration	247 days
Java timezone	Europe/London	License number	Javin-Internal-License
UTC time	2015-11-16 14:23:28 GMT	Allowed CPUs	2 / 8
Uptime	3 days 22 hrs	License location	Database
System clock drift	1 ms		
Time query duration	1 ms		

Figure 17.5. System

## Worker

Worker panel contains basic information about Worker and offers several actions for managing Worker.

<b>Name</b>	[worker0@Dyje]
<b>Status</b>	RUNNING
<b>Running jobs</b>	41
<b>Heap memory</b>	400 MB / 1353 MB (29.6% used)
<b>Non-heap memory</b>	74 MB
<b>Ports (Server, Api, Jmx, Ssl, Debug)</b>	10500, 10501, 10502, -, -
<b>Uptime</b>	4 mins 41 secs
<input type="button" value="Actions"/>	

Figure 17.6. Worker

## Status History

Status history panel displays node statuses history since restart of the server.

Status History	
Status	Age
READY	17 mins 23 secs
SUSPENDED	17 mins 33 secs
READY	3 days 22 hrs
STARTING	3 days 22 hrs

Figure 17.7. Status History

## User's Access

User's Access panel lists info about activities on files performed by users. The list displays a timestamp of an event, a username, and a name of the method.

Users' Accesses			
Date	Username	Address	Method
2015-11-16 11:22:21 CET	clover	172.22.22.29	logout
2015-11-16 10:04:08 CET	clover	172.22.22.29	executeGraphAsync
2015-11-16 10:02:26 CET	clover	172.22.22.29	validateClasspath

Figure 17.8. Users' Accesses panel

## Classloader cache

Classloader cache lists all currently cached classloaders. The classloader cache may be empty as classloader caching is disabled by default.

Classloader cache (empty)

Figure 17.9. Classloader cache

## Status

Status panel displays current node status since last server restart. It displays current server status (ready, stopped, ...), exact Java version, exact CloverETL Server version, way of access to database, URLs for synchronous and asynchronous messaging, available heap and non-heap memory, etc.

Status	
Last status check time	2015-11-16 14:44:53
Available heap memory (used / committed / max)	293 MB / 481 MB / 910 MB
Available non-heap memory (used / committed / max)	228 MB / 241 MB / -0 MB
Status	READY
CloverETL version	4.1.1.28/28
Java version	1.8.0_60
Active threads	37
Loaded classes	22183
Running jobs	[]
Active locks	[]
Event listeners	
Number of job pools	0
Number of jobs in the pools	0
Number of transformation definitions in use	0
Nodes accessible by asynchronous messaging	
Nodes accessible by synchronous messaging	
Database access	local PostgreSQL Native Driver PostgreSQL 9.3 JDBC4.1 (build 1102)
Asynchronous messaging info	

Figure 17.10. Status

## Heartbeat

Heartbeat panel displays a list of heartbeat events and their results.

Heartbeat						
Current time	System clock drift	CPU load	Heap mem used	Active threads	Loaded classes	Running jobs
2015-11-16 14:45:05	0 ms	0.03	293 MB	37	22183	[]
2015-11-16 14:45:02	0 ms	0.03	293 MB	37	22183	[]
2015-11-16 14:45:00	0 ms	0.03	280 MB	37	22183	[]
2015-11-16 14:44:58	0 ms	0.03	280 MB	37	22183	[]
2015-11-16 14:44:56	0 ms	0.03	280 MB	37	22183	[]
2015-11-16 14:44:54	0 ms	0.03	266 MB	37	22183	[]

Figure 17.11. Heartbeat

## Threads

Threads panel lists java threads and their states.

Threads			
Thread name	Thread state	Waited time (ms)	Blocked time (ms)
AsyncFileHandlerWriter-1304836502	TIMED_WAITING	-1	-1
ContainerBackgroundProcessor[StandardEngine[Catalina]]	TIMED_WAITING	-1	-1
Finalizer	WAITING	-1	-1
GC Daemon	TIMED_WAITING	-1	-1
Java2D Disposer	WAITING	-1	-1
MultiThreadedHttpConnectionManager cleanup	WAITING	-1	-1
NioBlockingSelector.BlockPoller-1	RUNNABLE	-1	-1
NioBlockingSelector.BlockPoller-2	RUNNABLE	-1	-1

Figure 17.12. Threads

## Quartz

Quartz panel lists scheduled actions: their name, description, start time, end time, time of previous event, time of next event and expected final event.

Quartz						
Name	Description	Start time	End time	Previous event (within this up-time)	Next event	Expected final event
trigger_174	Parsing structured data	2015-11-16 12:00:00		2015-11-16 14:50:00	2015-11-16 15:00:00	
trigger_36	Delete old execution artifacts	2015-11-12 16:15:57		2015-11-16 14:00:00	2015-11-16 15:00:00	
trigger_38	Delete old debug files	2015-11-12 16:15:57		2015-11-16 14:50:00	2015-11-16 15:00:00	

Figure 17.13. Quartz

## Cluster Overview

Cluster overview displays info collected from all cluster nodes. The info is grouped in several panels:

- List of nodes with a toolbar - allows manipulation with selected nodes
- Status history - Displays last 10 status changes for all cluster nodes
- Node detail - Displays several basic performance attributes for selected nodes. It's visible on the right side only when activated by button on the toolbar.
- Running jobs - It's displayed only when there are running jobs.

The screenshot shows the CloverETL Server Console interface. The main content area is titled 'Cluster Overview' and displays the following information:

**Clustering enabled** Cluster group name: **CloverCluster** Connected to node: **Node01** [Actions] [Refresh]

**Cluster Nodes**

Node	Host	Cores	Uptime
<input checked="" type="checkbox"/> Node01	http://10.0.3.15:8083/clover	2	6 hrs 56 mins
<input checked="" type="checkbox"/> Node02	http://10.0.5.3:8084/clover	2	6 hrs 25 mins

**Status History (last 10)**

Node	Status	Detected by	Age
Node02	READY	Node02	6 hrs 25 mins
Node02	STARTING	Node02	6 hrs 25 mins
node01	STOPPED	node01	6 hrs 26 mins
node01	SUSPENDED	node01	6 hrs 27 mins
node01	FORCED_SUSPEND	node01	6 hrs 27 mins
node01	FORCED_SUSPEND	Node01	6 hrs 28 mins
node01	READY	node01	6 hrs 28 mins
node01	STARTING	node01	6 hrs 28 mins
Node01	READY	Node01	6 hrs 56 mins

**Summary**

**Running jobs** 0  
**Nodes** 2 / 4  
**System CPUs** 4 / 32

**Resources Utilization Node01**

**Running jobs** 0  
**System memory** 3387 MB / 3791 MB (89.3% used)  
**System disk swap** 19 MB / 2064 MB (0.9% used)  
**Heap memory** 582 MB / 1820 MB (32% used)  
**Non-heap memory** 201 MB / unknown MB  
**Event listeners** 0 active

**Resources Utilization Node02**

**Running jobs** 0  
**System memory** 3387 MB / 3791 MB (89.3% used)  
**System disk swap** 19 MB / 2064 MB (0.9% used)  
**Heap memory** 281 MB / 1366 MB (20.6% used)  
**Non-heap memory** 175 MB / unknown MB  
**Event listeners** 0 active

Copyright © 2017 CloverETL created by Javlin. All rights reserved.

Figure 17.14. Cluster overview

## Node Detail

Node Detail is similar to the [Standalone Server Detail](#) (p. 131) mentioned above, however it displays detail info about node selected in the tree on the left.

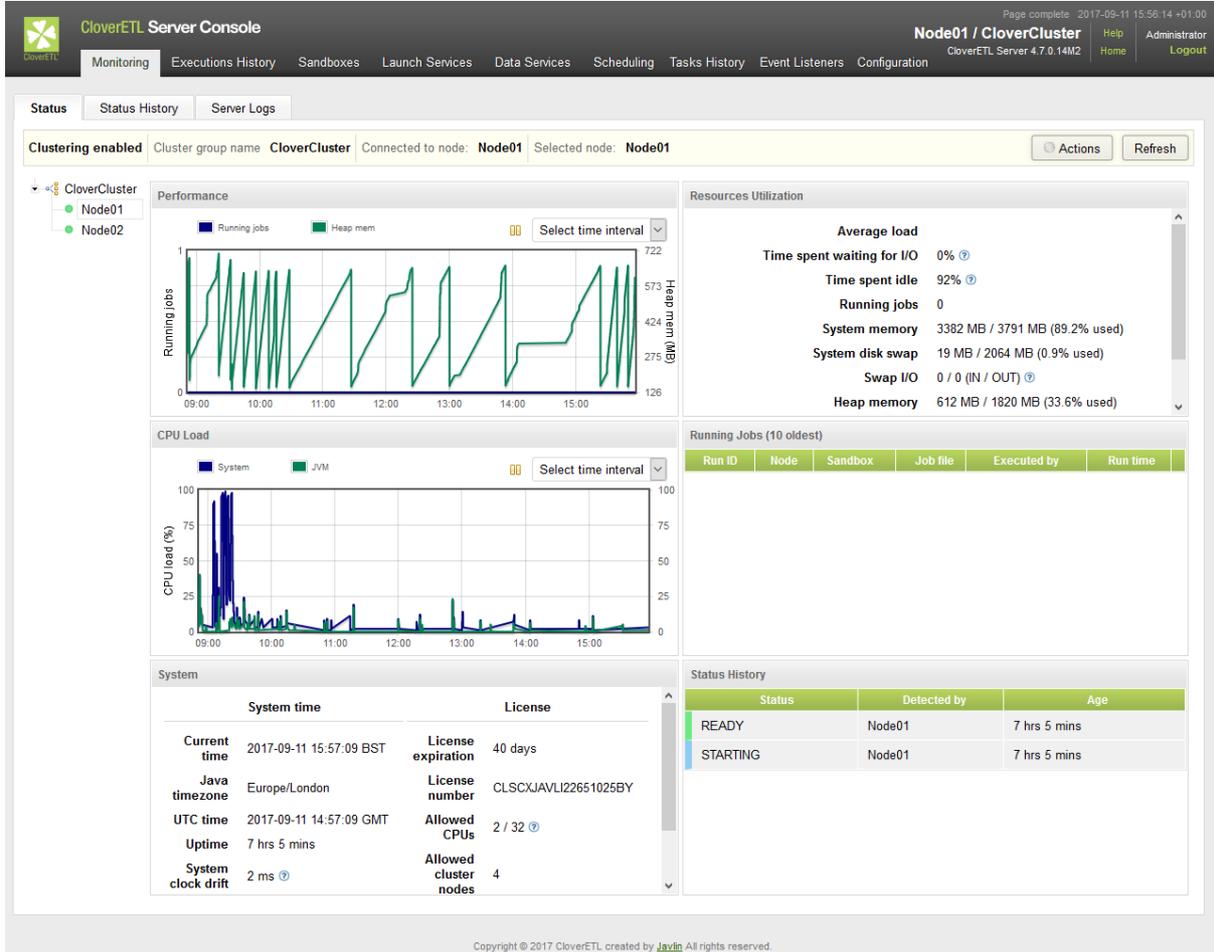


Figure 17.15. Node detail

## Server Logs

**Server Logs** tab allows user to investigate log messages logged on other cluster nodes. Since the log messages are collected in memory, the maximum number of collected messages is relatively low by default, however it's customisable.

There are different "Log types":

- **COMMON** - Ordinary server logs as stored in log files. It contains information on successful and unsuccessful logins, start and end of job execution etc.
- **WORKER** - Worker related log.

The log file is `java.io.tmpdir}/cloverlogs/worker-[node_name].log`.

- **CLUSTER** - Only cluster - related messages are visible in this log. It contains information on job delegation and other types of messages related to cluster communication.
- **LAUNCH\_SERVICES** - Only requests for launch services
- **AUDIT** - Detail logging of operations called on the CloverETL Server core. Since the full logging may affect server performance, it's disabled by default. See [Server Audit Logs](#) (p. 128) for details

The log file is `java.io.tmpdir}/cloverlogs/user-action.log`.

- **USER\_ACTION** - Contains some of user operations, e.g. login, logout, user creation, job execution, file synchronization (upload to server)

The corresponding log file is `java.io.tmpdir}/cloverlogs/user-action.log`.

- **SERVER\_INTERACTION** - Interaction between Designer and Server.

The log file is `java.io.tmpdir}/cloverlogs/server-integration.log`.

The screenshot shows the CloverETL Server Console interface. At the top, there is a navigation menu with options like Monitoring, Executions History, Sandboxes, Launch Services, Data Services, Scheduling, Tasks History, Event Listeners, and Configuration. The main content area is titled "Server Logs" and includes a search filter for Node ID (Node01) and Log type (COMMON). Below the filter, there is a list of log entries with columns for timestamp, log level, and message. The log entries show various initialization and configuration messages, such as "EventsHelper INFO TaskProcessors dictionary map...", "ServiceLogAdvice INFO Designer-server integration log enabled...", and "JGroups asynchronous messaging - initializing...".

Figure 17.16. Server Logs

See also: Chapter 16, [Logging](#) (p. 127).

---

## Using the Monitoring

---

### Restarting the Worker

Switch to **Monitoring** > **Status** tab.

In the worker tile, click **Actions** button and select **Finish jobs and restart** from the menu.

The worker will be restarted.

See also [Worker](#) (p. 133).

---

### Showing Worker's Command Line Arguments

Switch to **Monitoring** > **Status** tab.

In the worker tile, click the **Actions** button and select **Show command line**.

The command line arguments and parameters will be displayed.

See also [Worker](#) (p. 133).

---

### Suspending the Server

Switch to **Monitoring** > **Status** tab.

Click the **Actions** button (in the upper left corner) and select **Suspend** from the menu.

The server will be suspended. If there is a job running, the node will wait until it finishes.

If you need to suspend the node immediately without waiting for jobs to finish, use **Suspend at once** instead of **Suspend**.

---

### Resuming the Server

Switch to **Monitoring** > **Status** tab.

Click the **Actions** button (in the upper right corner) and select **Resume** from the menu.

The server will be resumed.

---

### Displaying List of Threads of the Server Core

Switch to **Monitoring** > **Status** tab.

Click the **Actions** button (in the upper right corner) and select **Show Details**. Several new tiles will appear. Search for the **Threads** tile.

# Chapter 18. Temp Space Management

Many of the components available in the CloverETL Server require temporary files or directories in order to work correctly. *Temp space* is a physical location on the file system where these files or directories are created and maintained. CloverETL Server allows you to configure and manage temp spaces - you can specify their locations, see usage of the filesystem etc.

To access this administration section, you need [Temp Space Management permission](#) (p. 170).

## Overview

The overview of temp spaces defined in CloverETL Server is available under *Configuration > Temp space management > Overview*

The overview panel displays list of temp spaces for each node in the cluster. These properties are displayed for each temp space:

- **Root Path** - location of the temp space with unresolved placeholders (see note below for placeholders)
- **Resolved Path** - location of the temp space with resolved placeholders (see note below for placeholders)
- **Free Space** - remaining space for the temp space
- **Filesystem Size** - all available space for the temp space (actual size of the filesystem where the temp space resides)
- **Filesystem Usage** - size of used space in percentage
- **Available** - the directory exists and is writable
- **Status** - current status of temp space, can be Active or Suspended



### Note

It is possible to use system properties and environment variables as placeholders. See [Using environment variables and system properties](#) (p. 143).

Node	Root Path	Resolved Path	Free Space	Filesystem Size	Filesystem Usage	Available	Status
Node01							
	<code>\$(java.io.tmpdir)/clover_temp_Node01</code>	<code>/home/clover/opt/server/CloverETL.Server.4.7.0M2.Tomcat-8.0.30-4-7-0-M2-R14/temp/clover_temp_Node01</code>	601.5 MB	29.5 GB	98%	<input checked="" type="checkbox"/>	Active
	<code>/tmp</code>	<code>/tmp</code>	601.5 MB	29.5 GB	98%	<input checked="" type="checkbox"/>	Active
Node02							
	<code>\$(java.io.tmpdir)/clover_temp_Node02</code>	<code>/home/clover/opt/server/CloverETL.Server.4.7.0M2.Tomcat-8.0.30/temp/clover_temp_Node02</code>	601.5 MB	29.5 GB	98%	<input checked="" type="checkbox"/>	Active
	<code>/tmp</code>	<code>/tmp</code>	601.5 MB	29.5 GB	98%	<input checked="" type="checkbox"/>	Active

Figure 18.1. Configured temp spaces overview - one default temp space on each cluster node

## Management

Temp space management offers an interface to add, disable, enable and delete a temp space. It is accessible under *Configuration > Temp space management > Edit*.

The screen is divided in two drop-down areas: Global Configuration and Per Node Configuration. The *Global configuration* manages temp spaces of standalone server or in case of a server cluster temp spaces on all its nodes. The *Per Node Configuration* allows to maintain temp spaces on each particular node.

[Initialization](#) (p. 142)

[Adding Temp Space](#) (p. 142)

[Using environment variables and system properties](#) (p. 143)

[Disabling Temp Space](#) (p. 144)

[Enabling Temp Space](#) (p. 145)

[Removing Temp Space](#) (p. 145)

## Initialization

---

When CloverETL Server is starting the system checks temp space configuration: in case no temp space is configured a new default temp space is created in the directory where `java.io.tmpdir` system property points. The directory is named as follows:

- `${java.io.tmpdir}/clover_temp` in case of a standalone server
- `${java.io.tmpdir}/clover_temp_<node_id>` in case of server cluster

## Adding Temp Space

---

In order to define new temp space enter its path into text field under last row in the table and click the **Add** link. If the directory entered does not exist, it will be created.



### Tip

The main point of adding additional temp spaces is to enable higher system throughput - therefore the paths entered should point to directories residing on different physical devices to achieve maximal I/O performance.

The screenshot displays the CloverETL Server Console interface. At the top, there's a navigation bar with 'Configuration' selected. Below it, the 'Temp Space Management' tab is active. The main content area is split into two sections: 'Global Configuration' and 'Detailed Configuration'. The 'Global Configuration' section shows a table with 'Root Path' and 'Operations' columns. A row for '/tmp' is visible with a 'Disable' button. Below it is an 'Add' button. The 'Detailed Configuration' section shows a table with columns: Node, Root Path, Resolved Path, Free Space, Filesystem Size, Filesystem Usage, Available, and Operations. It lists configurations for Node01 and Node02, showing resolved paths and usage percentages.

Figure 18.2. Newly added global temp space.

## Using environment variables and system properties

Environment variables and system properties can be used in the temp space path as a placeholder; they can be arbitrarily combined and resolved paths for each node may differ in accord with its configuration.



### Note

The environment variables have higher priority than system properties of the same name. The path with variables are resolved after system has added new temp space and when the server is starting. In case the variable value has been changed it is necessary to restart the server so that the change takes effect.

### Examples:

- Given that an environment variable `USERNAME` has a value `Filip`, and is used as a placeholder in the path `C:\Users\${USERNAME}\tmp`, the resolved path is `C:\Users\Filip\tmp`.
- Given that Java system property `java.io.tmpdir` has a value `C:\Users\Filip\AppData\Local\Temp` and the property is used as a placeholder in the path `${java.io.tmpdir}\temp_folder`, the resolved path is `C:\Users\Filip\AppData\Local\Temp\temp_folder`.
- Node `node01` has been started with `-Dcustom.tmporary.dir=C:\tmp_node01` parameter. Node `node02` has been started with `-Dcustom.tmporary.dir=C:\tmp_node02` parameter. The declared path is `${custom.tmporary.dir}`. The resolved path is different for each node, `C:\tmp_node01` for `node01` and `C:\tmp_node02` for `node02`.
- When the declared path is `${java.io.tmpdir}\${USERNAME}\tmp_folder`, the resolved path is `C:\tmp\Filip\tmp_folder`.

## Chapter 18. Temp Space Management

Global Configuration							
Root Path							Operations
\$[java.io.tmpdir]/\${USERNAME}/tmp_folder							Disable
\$[java.io.tmpdir]/temp_folder							Disable
C:/Users/\${USERNAME}/tmp							Disable
\${custom.tmporary.dir}							Disable
<input type="text"/>							Add

Detailed Configuration							
Node	Root Path	Resolved Path	Free Space	Filesystem Size	Filesystem Usage	Available	Operations
node01							
	`\${java.io.tmpdir}/clover_temp_node01`	C:\Users\Filip\AppData\Local\Temp\clover_temp_node01	2.4 GB	17.3 GB	86%	<input checked="" type="checkbox"/>	Disable
	C:/Users/\${USERNAME}/tmp	C:\Users\Filip\tmp	2.4 GB	17.3 GB	86%	<input checked="" type="checkbox"/>	Disable
	`\${java.io.tmpdir}/temp_folder`	C:\Users\Filip\AppData\Local\Temp\temp_folder	2.4 GB	17.3 GB	86%	<input checked="" type="checkbox"/>	Disable
	`\${java.io.tmpdir}/\${USERNAME}/tmp_folder`	C:\Users\Filip\AppData\Local\Temp\Filip\tmp_folder	2.4 GB	17.3 GB	86%	<input checked="" type="checkbox"/>	Disable
	`\${custom.tmporary.dir}`	C:\tmp_node01	2.4 GB	17.3 GB	86%	<input checked="" type="checkbox"/>	Disable
	<input type="text"/>						Add
node02							
	`\${java.io.tmpdir}/clover_temp_node02`	C:\Users\Filip\AppData\Local\Temp\clover_temp_node02	24 GB	29.5 GB	18%	<input checked="" type="checkbox"/>	Disable
	C:/Users/\${USERNAME}/tmp	C:\Users\Filip\tmp	24 GB	29.5 GB	18%	<input checked="" type="checkbox"/>	Disable
	`\${java.io.tmpdir}/temp_folder`	C:\Users\Filip\AppData\Local\Temp\temp_folder	24 GB	29.5 GB	18%	<input checked="" type="checkbox"/>	Disable
	`\${java.io.tmpdir}/\${USERNAME}`	C:\Users\Filip\AppData\Local\Temp\Filip\tmp_folder	24 GB	29.5 GB	18%	<input checked="" type="checkbox"/>	Disable
	`\${custom.tmporary.dir}`	C:\tmp_node02	24 GB	29.5 GB	18%	<input checked="" type="checkbox"/>	Disable
	<input type="text"/>						Add

*Figure 18.3. Temp spaces using environment variables and system properties*

### Disabling Temp Space

To disable a temp space click on "Disable" link in the panel. Once the temp space has been disabled, no new temporary files will be created in it, but the files already created may be still used by running jobs. In case there are files left from previous or current job executions a notification is displayed.



#### Note

The system ensures that at least one enabled temp space is available.

## Chapter 18. Temp Space Management

The screenshot shows the CloverETL Server Console interface. At the top, there is a navigation menu with options like Monitoring, Executions History, Sandboxes, Launch Services, Data Services, Scheduling, Tasks History, Event Listeners, Configuration, Users, Groups, Security, Temp Space Management, System Info, CloverETL Info, Export, Import, and Setup. The main content area is divided into two sections: Global Configuration and Detailed Configuration.

**Global Configuration**

Root Path	Operations
<code>\$(java.io.tmpdir)/\${USER}/tmp_folder</code>	Enable Remove
<code>/home/\${USER}/tmp</code>	Disable
<code>/tmp</code>	Disable
<input type="text"/>	Add

**Detailed Configuration**

Node	Root Path	Resolved Path	Free Space	Filesystem Size	Filesystem Usage	Available	Operations
<b>NodeC6</b>							
	<code>\$(java.io.tmpdir)/\${USER}/tmp_folder</code>	<code>/home/clover/opt/apache-tomcat-7.0.54/temp/clover/tmp_folder</code>	2.4 GB	17.3 GB	86%	<input checked="" type="checkbox"/>	Enable Remove
	<code>/home/\${USER}/tmp</code>	<code>/home/clover/tmp</code>	2.4 GB	17.3 GB	86%	<input checked="" type="checkbox"/>	Disable
	<code>/tmp</code>	<code>/tmp</code>	2.4 GB	17.3 GB	86%	<input checked="" type="checkbox"/>	Disable
	<code>\$(java.io.tmpdir)/clover_temp_NodeC6</code>	<code>/home/clover/opt/apache-tomcat-7.0.54/temp/clover_temp_NodeC6</code>	2.4 GB	17.3 GB	86%	<input checked="" type="checkbox"/>	Disable
	<input type="text"/>						Add
<b>NodeC7</b>							
	<code>\$(java.io.tmpdir)/\${USER}/tmp_folder</code>	<code>/home/user1/opt/apache-tomcat-7.0.56/temp/user1/tmp_folder</code>	24 GB	29.5 GB	18%	<input checked="" type="checkbox"/>	Enable Remove
	<code>/home/\${USER}/tmp</code>	<code>/home/user1/tmp</code>	24 GB	29.5 GB	18%	<input checked="" type="checkbox"/>	Disable
	<code>/tmp</code>	<code>/tmp</code>	24 GB	29.5 GB	18%	<input checked="" type="checkbox"/>	Disable
	<code>\$(java.io.tmpdir)/clover_temp_NodeC7</code>	<code>/home/user1/opt/apache-tomcat-7.0.56/temp/clover_temp_NodeC7</code>	24 GB	29.5 GB	18%	<input checked="" type="checkbox"/>	Disable

Figure 18.4. Disable operation reports action performed

### Enabling Temp Space

To enable a temp space click on "Enable" link in the panel. Enabled temp space is active, i.e. available for temporary files and directories creation.

### Removing Temp Space

To remove a temp space click on "Remove" link in the panel. Only the disabled temp space may be removed. If there are any running jobs using the temp space, the system will not allow its removal. In case there are some files left in the temp space directory, it is possible to remove them in the notification panel. The available options are:

- *Remove* - remove temp space from a system, but keep its content
- *Remove and delete* - remove the temp space from a system and its content too
- *Cancel* - do not proceed with operation

## Chapter 18. Temp Space Management

The screenshot shows the CloverETL Server Console interface. The top navigation bar includes 'Monitoring', 'Executions History', 'Sandboxes', 'Launch Services', 'Data Services', 'Scheduling', 'Tasks History', 'Event Listeners', and 'Configuration'. The 'Configuration' section is active, showing 'Temp Space Management'.

The 'Global Configuration' section displays a table of root paths:

Root Path	Operations
<code>\$(java.io.tmpdir)/\${USER}/tmp_folder</code>	Enable, Remove
<code>/home/\${USER}/tmp</code>	Disable
<code>/tmp</code>	Disable
<input type="text"/>	Add

The 'Detailed Configuration' section shows a table of temp spaces across nodes:

Node	Root Path	Resolved Path	Filesystem Size	Filesystem Usage	Available	Operations	
NodeC6	<code>\$(java.io.tmpdir)/\${USER}/tmp_folder</code>	<code>/home/clover/opt/apache-tomcat-7.0.54/temp/clover/tmp_folder</code>	17.3 GB	86%	<input checked="" type="checkbox"/>	Enable, Remove	
	<code>/home/\${USER}/tmp</code>	<code>/home/clover/tmp</code>	17.3 GB	86%	<input checked="" type="checkbox"/>	Disable	
	<code>/tmp</code>	<code>/tmp</code>	2.4 GB	17.3 GB	86%	<input checked="" type="checkbox"/>	Disable
	<code>\$(java.io.tmpdir)/clover_temp_NodeC6</code>	<code>/home/clover/opt/apache-tomcat-7.0.54/temp/clover_temp_NodeC6</code>	2.4 GB	17.3 GB	86%	<input checked="" type="checkbox"/>	Disable
<input type="text"/>						Add	
NodeC7	<code>\$(java.io.tmpdir)/\${USER}/tmp_folder</code>	<code>/home/user1/opt/apache-tomcat-7.0.56/temp/user1/tmp_folder</code>	24 GB	29.5 GB	18%	<input checked="" type="checkbox"/>	Enable, Remove
	<code>/home/\${USER}/tmp</code>	<code>/home/user1/tmp</code>	24 GB	29.5 GB	18%	<input checked="" type="checkbox"/>	Disable
	<code>/tmp</code>	<code>/tmp</code>	24 GB	29.5 GB	18%	<input checked="" type="checkbox"/>	Disable
	<code>\$(java.io.tmpdir)/clover_temp_NodeC7</code>	<code>/home/user1/opt/apache-tomcat-7.0.56/temp/clover_temp_NodeC7</code>	24 GB	29.5 GB	18%	<input checked="" type="checkbox"/>	Disable
<input type="text"/>						Add	

A modal dialog titled 'Remove Temp Space' is overlaid on the table, asking 'Do you want to remove the temp space?' with 'Remove' and 'Cancel' buttons.

Copyright © 2014 CloverETL created by Javin All rights reserved.

Figure 18.5. Remove operation asks for confirmation in case there are data present in the temp space

# Chapter 19. Secure Parameters

[Secure parameters configuration](#) (p. 149)

[Installing Bouncy Castle JCE provider](#) (p. 150)

Transformation graphs in **CloverETL Server** environment allow you to define secure graph parameters. Secure graph parameters are regular graph parameters, either internal or external (in a \*.prm file), but the values of the graph parameters are not stored in plain text on the file system - encrypted values are persisted instead. This allows you to use graph parameters to handle sensitive information, typically credentials such as passwords to databases.

Secure parameters are only available in **CloverETL Server** environment, including working with **CloverETL Server Projects** in **CloverETL Designer**.

The encryption algorithm must be initialized with a **master password**. The master password has to be manually set after server installation in *Configuration > Security > Secure Parameters > Master password*. Secure parameters cannot be used before the master password is set.

The maximum length of the master password is 255 characters; there are no other restrictions or complexity requirements.

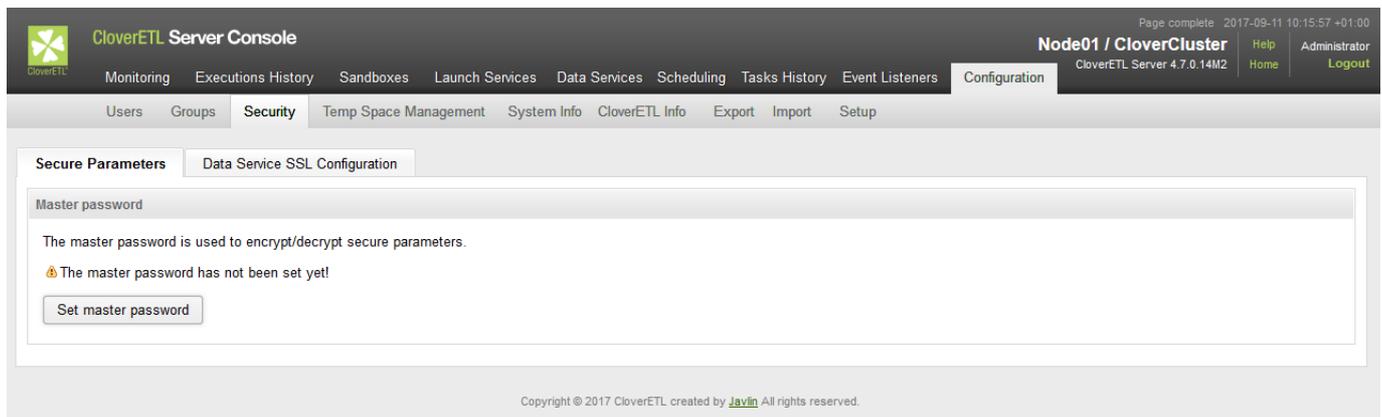


Figure 19.1. Master password initialization

After setting the master password secure parameters are fully available in **Graph parameter editor** in **CloverETL Designer**. When setting value of a secure parameter, it will be automatically encrypted using the master password. Secure parameters are automatically decrypted by server in graph runtime. A parameter value can also be encrypted in the **CloverETL Server Console** in the *Configuration > Security > Secure Parameters* page - use the **Encrypt text** section.

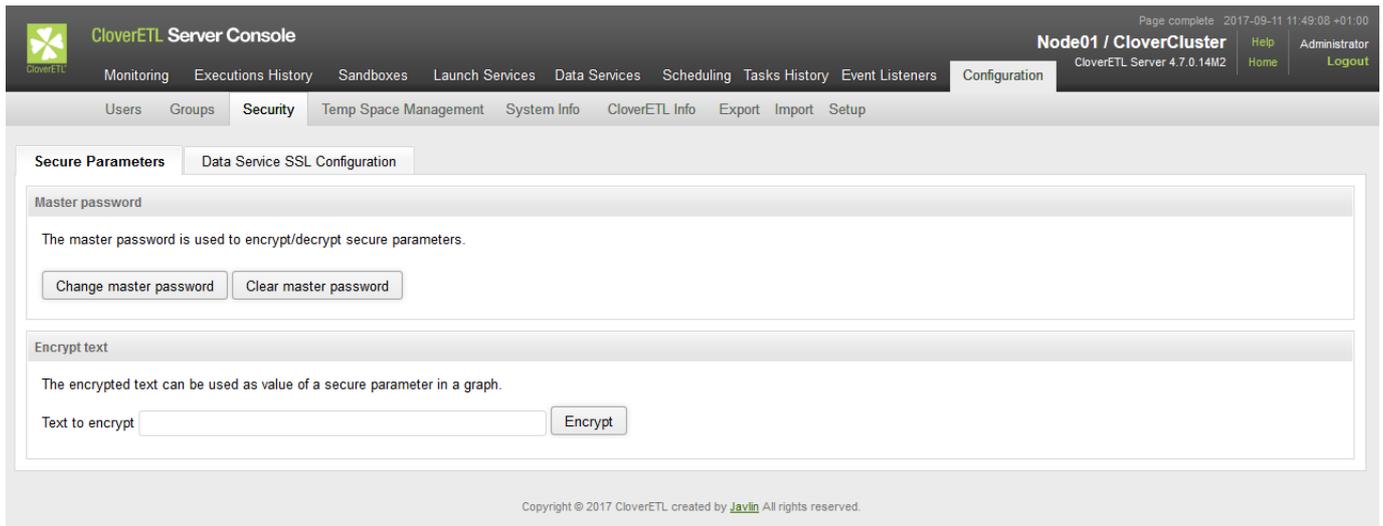


Figure 19.2. Graph parameters tab with initialized master password

If you change the master password, the secure parameters encrypted using the old master password cannot be decrypted correctly anymore. In that case existing secure parameters need to be encrypted again with the new master password. That can be accomplished simply by setting their value (non-encrypted) again in the **Graph parameter editor**. Similar master password inconsistency issue can occur if you move a transformation graph with some secure parameters to another server with a different master password. So it is highly recommended to use the identical master password for all your **CloverETL Server** installations.

See documentation of secure parameters in **CloverETL Designer** manual for further details.

## Secure parameters configuration

Encryption of secure parameters can be further customized via server configuration parameters.

Table 19.1. Secure parameters configuration parameters

Property name	Default value	Description
security.job_parameters.encryptor.algorithm	PBEWithMD5AndDES	<p>The algorithm to be used for encryption. This algorithm has to be supported by your JCE provider (if you specify a custom one, or the default JVM provider if you don't). The name of algorithm should start with <i>PBE</i> prefix.</p> <p>The list of available algorithms depends on your JCE provider, e.g. for the default <i>SunJCE</i> provider you can find them on <i>SunJCEProvider</i> or for the <i>Bouncy Castle</i> provider on <i>Bouncy Castle Specifications</i> (section <i>Algorithms/PBE</i>).</p>
security.job_parameters.encryptor.master_password_encryption.password	clover	The password used to encrypt values persisted in the database table <i>secure_param_passwd</i> (the master password is persisted there).
security.job_parameters.encryptor.providerClassName	Empty string. The default JVM provider is used (e.g. for Oracle Java the <i>SunJCE</i> provider is used)	The name of the security provider to be asked for the encryption algorithm. It must implement <i>java.security.Provider</i> interface. For example set to <i>org.bouncycastle.jce.provider.BouncyCastleProvider</i> for the <i>Bouncy Castle</i> JCE provider, see below.

## Installing Bouncy Castle JCE provider

Algorithms provided by JVM could be too weak to satisfy an adequate security. Therefore it is recommended to install a third-party JCE provider. Following example demonstrates installation of one concrete provider, *Bouncy Castle* JCE provider. Another provider would be installed similarly.

1. Download Bouncy Castle provider jar (e.g. `bcprov-jdk15on-150.jar`) from [http://bouncycastle.org/latest\\_releases.html](http://bouncycastle.org/latest_releases.html)
2. Add the jar to the classpath of your application container running **CloverETL Server**, e.g. to directory `WEB-INF/lib`
3. Set value of the `security.job_parameters.encryptor.providerClassName` attribute to `org.bouncycastle.jce.provider.BouncyCastleProvider` in the `config.properties` file.
4. Set value of the `security.job_parameters.encryptor.algorithm` attribute to the desired algorithm (e.g. `PBEWITHSHA256AND256BITAES-CBC-BC`).

Example of configuration using Bouncy Castle:

```
security.job_parameters.encryptor.algorithm=PBEWITHSHA256AND256BITAES-CBC-BC
security.job_parameters.encryptor.providerClassName=org.bouncycastle.jce.provider.BouncyCastleProvider
```

---

## Chapter 20. Users and Groups

The CloverETL Server has a built-in security module that manages users and groups. User groups control access permissions to sandboxes and operations the users can perform on the Server, including authenticated calls to Server API functions. A single user can belong to multiple groups.

LDAP or Active Directory can be configured with the Server to authenticate users and optionally assign their effective groups (and permissions) from a global directory.

You can manage users and user groups in **Configuration/Users and Groups**. Please note that you need a “List users” (“List groups” respectively) permission for that.

---

## LDAP Authentication

[Configuration](#) (p. 153)

[Basic LDAP connection properties](#) (p. 153)

[Configuration of user and group lookup](#) (p. 153)

Since **CloverETL 3.2**, it is possible to configure **CloverETL Server** to use an LDAP server for users authentication. This way, the credentials of users registered in LDAP may be used for authentication to any **CloverETL Server** interface (API or web console).

However, the authorization (access levels to sandboxes content and privileges for operations) is still handled by the **CloverETL** security module. Each user, even when logged-in using LDAP authentication, must have their own "user" record (with related groups) in the **CloverETL** security module. So there must be a user with the same username and domain set to "LDAP". Such a record has to be created by a Server administrator before the user can log in.

To configure LDAP authentication, use the Setup GUI (p. 80).

### Use LDAP for user authentication only

If this **Authentication Policy** is selected, **CloverETL Server** uses LDAP directory to verify only the user's password:

1. The user specifies the LDAP credentials in the login form to the Server web console.
2. **CloverETL Server** looks up the user's record and checks whether it has the "LDAP" domain set.
3. The Server attempts to connect to the LDAP server using the user's credentials. If it succeeds, the user is logged in.

### Use LDAP for user authentication and user synchronization

In this mode, **CloverETL Server** verifies user's credentials and synchronizes additional information (group, name and email) with those stored in LDAP.

In case the Server is configured for user authentication and group synchronization, the procedure is as follows:

1. The user specifies the LDAP credentials in the login form to the Server web console.
2. **CloverETL Server** looks up the user's record and checks whether it has the "LDAP" domain set.
3. **CloverETL Server** connects to the LDAP server and checks whether the user exists (it uses specified search to lookup in LDAP).
4. If the user exists in LDAP, **CloverETL Server** performs authentication.
5. If the authentication is successful, **CloverETL Server** searches LDAP for user's groups.
6. **CloverETL** user is assigned to the **CloverETL** groups according to his current assignation to the LDAP groups.
7. User is logged-in.



### Note

Switching domains:

- If a user was **created as LDAP** and then switched to clover domain, you have to **set a password** for him in the **Change password** tab.

- If a user was **created as clover** and then switched to LDAP domain, they have a password in clover domain, but it is overridden by the LDAP password. After switching back to clover domain, the **original password is re-used**. It can be reset in the **Change password** tab, if needed.

## Configuration

By default CloverETL Server allows only its own internal mechanism for authentication. To enable authentication with LDAP, set the configuration property "security.authentication.allowed\_domains" properly. It is a list of user domains that are used for authentication.

Currently there are 2 authentication mechanism implemented: "LDAP" and "clover" ("clover" is an identifier of CloverETL internal authentication and may be changed by security.default\_domain property, but only for white-labeling purposes). To enable LDAP authentication, set value to "LDAP" (only LDAP) or "clover,LDAP". Users from both domain may login. It's recommended to allow both mechanisms together, until the LDAP is properly configured. So the admin user can still login to web GUI although the LDAP connection isn't properly configured.

You can use **Setup** to configure LDAP authentication. See [LDAP](#) (p. 80) in Chapter 12, [Setup](#) (p. 74).

## Basic LDAP connection properties

```
# Implementation of context factory.
security.ldap.ctx_factory=com.sun.jndi.ldap.LdapCtxFactory
# URL of LDAP server.
security.ldap.url=ldap://hostname:port
# User DN pattern that will be used to create LDAP user DN from login name.
security.ldap.user_dn_pattern=uid=${username},dc=company,dc=com
```

Depending on the LDAP server configuration the property security.ldap.user\_dn\_pattern can be pattern for user's actual distinguished name in the LDAP directory, or just the login name - in such case just set the property to \${username}.

## Configuration of user and group lookup

In order to be able to synchronize the Clover groups with those defined in LDAP directory, the security.ldap.user\_dn\_pattern has to be left unspecified. There are additional properties required so that the server is able to search the LDAP directory.

```
# User DN of a user that has sufficient privileges to search LDAP for users and groups.
security.ldap.userDN=cn=Manager,dc=company,dc=com
# The password for user mentioned above.
security.ldap.password=
```

There are optional settings affecting how the LDAP directory is searched.

```
# Timeout for queries searching the directory.
security.ldap.timeout=5000
# Maximal number of records that the query can return.
security.ldap.records_limit=2000
# How LDAP referrals are processed, possible values are: 'follow', 'ignore' and 'throw'.
# The default depends on the context provider.
security.ldap.referral=
```

Specified values work for this specific LDAP tree:

- dc=company,dc=com
  - ou=groups

- cn=admins (objectClass=groupOfNames,member=(uid=smith,dc=company,dc=com),member=(uid=jones,dc=company,dc=com))
- cn=developers (objectClass=groupOfNames,member=(uid=smith,dc=company,dc=com))
- cn=consultants (objectClass=groupOfNames,member=(uid=jones,dc=company,dc=com))
- ou=people
  - uid=smith (fn=John,sn=Smith,mail=smith@company.com)
  - uid=jones (fn=Bob,sn=Jones,mail=jones@company.com)

Following properties are necessary for lookup for the LDAP user by his username. (step [4] in the login process above)

```
# Base specifies the node of LDAP tree where the search starts.
security.ldap.user_search.base=dc=company,dc=eu
# Filter expression for searching the user by his username.
# Note, that this search query must return just one record.
# Placeholder ${username} will be replaced by username specified by the logging user.
security.ldap.user_search.filter=(uid=${username})
# Scope specifies type of search in "base". There are three possible values: SUBTREE | ONELEVEL | OBJECT
# http://download.oracle.com/javase/8/docs/api/javax/naming/directory/SearchControls.html
security.ldap.user_search.scope=SUBTREE
```

Following properties are names of attributes from the search defined above. They are used for getting basic info about the LDAP user in case the user record has to be created/updated by Clover security module: (step [6] in the login process above)

```
security.ldap.user_search.attribute.firstname=fn
security.ldap.user_search.attribute.lastname=sn
security.ldap.user_search.attribute.email=mail
# This property is related to the following step "searching for groups".
# Groups may be obtained from specified user's attribute, or found by filter (see next paragraph).
# Leave this property empty if the user doesn't have such attribute.
security.ldap.user_search.attribute.groups=memberOf
```

In the following step, clover tries to find groups which the user is assigned to. (step [4] in the login process above). There are two ways how to get list of groups which the user is assigned to. The user-groups relation is specified on the "user" side. The user record has some attribute with list of groups. It's "memberOf" attribute usually. Or the relation is specified on the "group" side. The group record has an attribute with list of assigned users. It's "member" attribute usually.

In case the relation is specified on users side, please specify property:

```
security.ldap.user_search.attribute.groups=memberOf
```

Leave it empty otherwise.

In case the relation is specified on the groups side, set properties for searching:

```
security.ldap.groups_search.base=dc=company,dc=com
# Placeholder ${userDN} will be replaced by user DN found by the search above.
# If the filter is empty, searching will be skipped.
security.ldap.groups_search.filter=(&(objectClass=groupOfNames)(member=${userDN}))
security.ldap.groups_search.scope=SUBTREE
```

Otherwise, please leave property `security.ldap.groups_search.filter` empty, so the search will be skipped.

Clover user record will be assigned to the clover groups according to the LDAP groups found by the search (or the attribute). (Groups synchronization is performed during each login)

```
# Value of the following attribute will be used for lookup for the Clover group by its code.
```

```
# So the user will be assigned to the Clover group with the same "code".  
security.ldap.groups_search.attribute.group_code=cn
```

## Users

The **Users** section of the **Configuration** menu allows you to:

[Edit Users Record](#) (p. 157)

[Change Users Password](#) (p. 157)

[Assign Users to Groups](#) (p. 157)

[Disable / Enable Users](#) (p. 158)

After default installation on an empty database, the admin user is created automatically.

Table 20.1. Admin user

User name	Description
clover	Clover user has admin permissions, thus default password <code>clover</code> should be changed after installation.

The screenshot shows the CloverETL Server Console web GUI. The 'Configuration' menu is active, and the 'Users' sub-menu is selected. The main content area displays a search bar for users, a 'New User' button, and a table of existing users. The 'Administrator' user is selected, and an 'Overview' panel is shown on the right, displaying user details such as Domain, Username, First name, Last name, and E-mail.

Full Name	Username	Login Details	Actions
Administrator	clover	Last login: 2017-12-06 11:09:47	...
Test User 1	user1	Last login: Never	...

Figure 20.1. Web GUI - section "Users" under "Configuration"

Table 20.2. User attributes

Attribute	Description
Domain	Domain which is the origin of the user. Currently, there are only two possible values: "clover" or "ldap".
Username	A common user identifier. Must be unique, cannot contain spaces or special characters, just letters and numbers.
Password	Case sensitive password. If the user loses his password, the new one must be set. The password is stored in an encrypted form for security reasons, so it cannot be retrieved from a database and must be changed by the user who has proper permission for such operation.
First name	The user's first name.
Last name	The user's last name.
E-mail	Email address which may be used by CloverETL administrator or by CloverETL Server for automatic notifications. See <a href="#">Send an Email</a> (p. 204) for details.

## Edit Users Record

A user with a "Create user" or "Edit user" permission can use this form to set basic user parameters.

Figure 20.2. Web GUI - edit user

## Change Users Password

If user loses his password, the new one must be set. So a user with the "Change passwords" permission can use this form to do it.

Figure 20.3. Web GUI - change password

## Assign Users to Groups

Assignment to groups gives the user proper permissions. Only logged user with the "Groups assignment" permission can access this form and specify groups which the user is assigned in. See [Groups](#) (p. 159) for details about permissions.

Figure 20.4. Web GUI - groups assignment

## Disable / Enable Users

---

Since a user record has various relations to the logs and history records, it can't be deleted. So it is disabled instead. This means that the record doesn't display in the list and the user can't login.

However, a disabled user may be enabled again. **Note** that the disabled user is removed from their groups, so groups should be assigned properly after re-enabling.

## Groups

Group is an abstract set of users, which gives assigned users some permissions. So it is not necessary to specify permissions for each single user.

There are independent levels of permissions implemented in CloverETL Server

- *permissions to Read/Write/eXecute in sandboxes* - The sandbox owner can specify different permissions for different groups. For details, see [Sandbox Content Security and Permissions](#) (p. 177).
- *permissions to perform some operation* - user with operation permission "Permission assignment" may assign specific permission to existing groups.
- *permissions to launch specific service* - For details, see Chapter 37, [Launch Services](#) (p. 282).

Table 20.3. Default groups created during installation

Group name	Description
admins	This group has operation permission "all" assigned, which means, that it has unlimited permission. Default user "clover" is assigned to this group, which makes him administrator.
all users	Every single CloverETL user is assigned to this group by default. It is possible to remove user from this group, but it is not a recommended approach. This group is useful for some permissions to sandbox or some operation, which you would like to make accessible for all users without exceptions.

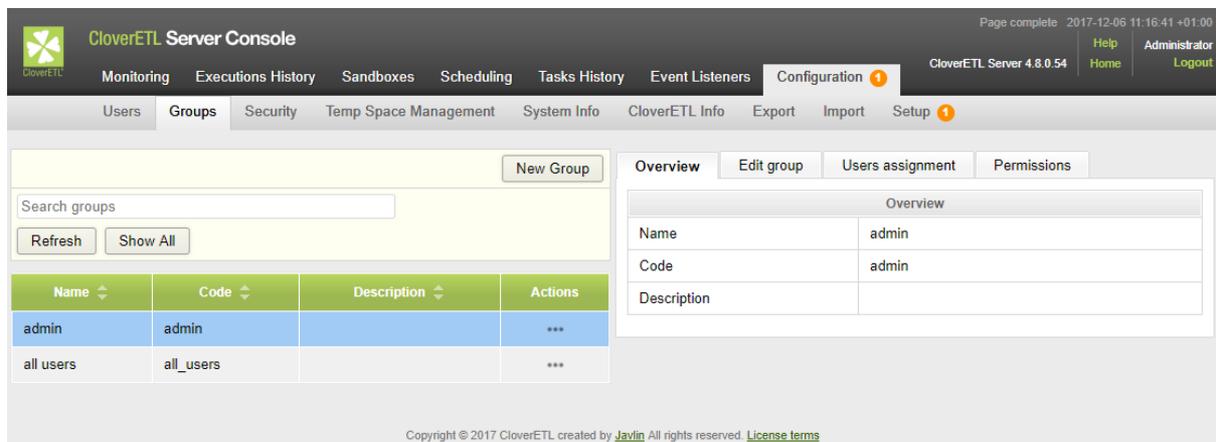


Figure 20.5. Web GUI - section "Groups"

## Users Assignment

Relation between users and groups is N:M. Thus in the same way, how groups are assignable to users, users are assignable to groups.

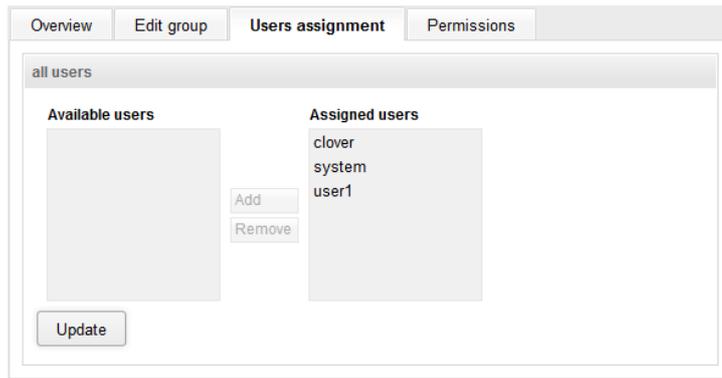


Figure 20.6. Web GUI - users assignment

## Groups permissions

Groups permissions are structured as a tree, where permissions are inherited from the root to leaves. Thus if some permission (tree node) is enabled (blue dot), all permissions in sub tree are automatically enabled (white dot). Permissions with red cross are disabled.

Thus for "admin" group just "all" permission is assigned, every single permission in sub tree is assigned automatically.

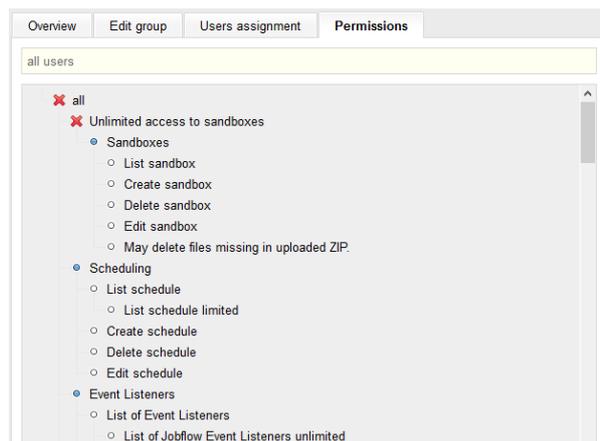


Figure 20.7. Tree of permissions

With no of the following privileges, user can: login to the server console, create server project (in Designer) from its own sandbox, create a file in its own existing sandbox, and run graphs.

- **all**

A user with this permission has all available permissions. Admin group has all permissions by default.

- **Unlimited access to sandboxes**

This permission allows user to perform operations on all sandboxes, even if the sandbox accessibility is not specified explicitly.

**Unlimited access to sandboxes** permission does not include the [suspend sandbox permission](#) (p. 167).

- **Sandboxes**

This permission allows user work with sandboxes. This permission contains all the permissions below. The user can perform operations only on sandboxes owned by himself or on sandboxes with explicitly added access to him.

See Chapter 21, [Sandboxes - Server Side Job Files](#) (p. 174).

- **List sandbox**

In server web interface, this permission allows user to list her sandboxes and list sandboxes with read permission granted to the user's group.

In server web interface, this permission is necessary to create, edit, or delete sandboxes.

Within a sandbox with write access granted, user can edit or remove files and create or delete directories even without this permission.

- **Create sandbox**

This permission allows user to create a new sandbox.

If the sandbox is to be created in web interface, the user is required to have the [list sandbox permission](#) (p. 161).

- **Delete sandbox**

This permission allows user to delete a sandbox.

If the sandbox is to be deleted in web interface, the user is required to have the [list sandbox permission](#) (p. 161).

- **Edit sandbox**

This permission allows user to edit a sandbox.

If the sandbox is to be modified in web interface, the user is required to have the [list sandbox permission](#) (p. 161).

- **May delete files missing in uploaded ZIP**

In **Sandbox** → **Upload ZIP**, this permission allows user to use a checkbox to delete files missing in the ZIP to be uploaded. If the user does not have this permission, the checkbox to delete mission files in ZIP is not displayed.

If the sandbox is to be uploaded from a ZIP file in server web interface, it is required to have the [list sandbox permission](#) (p. 161).

- **Scheduling**

This permission allows user to manage schedules.

See Chapter 28, [Scheduling](#) (p. 224).

- **List schedule**

This permission allows user to list all schedules.

- **List schedule limited**

This permission allows user to list the enabled schedules.

- **Create schedule**

This permission allows user to create a new schedule.

The user needs the [list schedule limited permission](#) (p. 161) to access the scheduling section to create a new schedule.

- **Delete schedule**

This permission allows user to delete the schedule.

User needs [list schedule limited permission](#)(p. 161) or [list schedule permission](#)(p. 161) to access the scheduling section to delete the schedule.

- **Edit schedule**

This permission allows user to edit the schedule.

User needs [list schedule limited permission](#)(p. 161) or [list schedule permission](#)(p. 161) to access the scheduling section to edit the schedule.

- **Event listeners**

This permission allows user to manage event listeners.

See Chapter 30, [Listeners](#) (p. 236).

- **List of Event Listeners**

This permission allows user to list all event listeners.

- **List of Jobflow Event Listeners unlimited**

This permission allows user to list jobflow event listeners.

See [Jobflow Event Listeners](#) (p. 244)

- **List of Jobflow Event Listeners limited**

This permission allows user to list jobflow event listeners of sandboxes the user can read from.

- **List of Graph Event Listeners unlimited**

This permission allows user to list all graph event listeners.

See [Graph Event Listeners](#) (p. 238).

- **List of Graph Event Listeners limited**

This permission allows user to list graph event listeners from sandboxes the user can read from.

- **List of File Event Listeners unlimited**

This permission allows user to list all file event listeners.

See [File Event Listeners \(remote and local\)](#) (p. 253).

- **List of File Event Listeners limited**

This permission allows user to list all file event listeners.

- **List of JMS Event Listeners unlimited**

This permission allows user to list all JMS listeners.

---

See [JMS Message Listeners](#) (p. 246).

- **List of JMS Event Listeners limited**

This permission allows user to list all JMS listeners.

- **List of Universal Event Listeners unlimited**

This permission allows user to list all universal event listeners.

See [Universal Event Listeners](#) (p. 251).

- **List of Universal Event Listeners limited**

This permission allows user to list all universal event listeners.

See [Universal Event Listeners](#) (p. 251).

- **Create Event Listener**

This permission allows user to create event listeners.

If the event listener is to be created in server web interface, the user needs to have permission to list the event listeners of the particular type.

- **Create Jobflow Event Listener**

This permission allows user to create a new Jobflow Event listener.

If the jobflow event listener is to be created in server web interface, the user needs to have the [list of jobflow event listeners limited permission](#) (p. 162).

See [Jobflow Event Listeners](#) (p. 244).

- **Create Graph Event Listener**

This permission allows user to create a graph event listener.

If the graph event listener is to be created in server web interface, the user needs to have the [list of graph event listeners limited permission](#) (p. 162).

See [Graph Event Listeners](#) (p. 238).

- **Create File Event Listener**

This permission allows user to create a graph event listener.

If the file event listener is to be created in server web interface, the user needs to have the [list of file event listeners limited permission](#) (p. 162).

See [File Event Listeners \(remote and local\)](#) (p. 253).

- **Create JMS Listener**

This permission allows user to create a JMS event listener.

If the JMS event listener is to be created in server web interface, the user needs to have the [list of JMS event listeners limited permission](#) (p. 163).

See [JMS Message Listeners](#) (p. 246).

- **Create Universal Event Listener**

This permission allows user to create a universal event listener.

If the universal event listener is to be created in server web interface, the user needs to have the [list of universal event listeners limited permission](#) (p. 163).

See [Universal Event Listeners](#) (p. 251).

- **Edit Event Listener**

This permission allow user to edit an event listener.

If the event listener is to be created in server web interface, the user needs to have permission to list event listener of the particular type.

- **Edit Jobflow Event Listener**

This permission allows user to edit a jobflow event listener.

If the jobflow event listener is to be edited in server web interface, the user needs to have the [list of jobflow event listeners limited permission](#) (p. 162).

See [Jobflow Event Listeners](#) (p. 244).

- **Edit Graph Event Listener**

This permission allows user to edit a graph event listener.

If the graph event listener is to be edited in server web interface, the user needs to have the [list of graph event listeners limited permission](#) (p. 162).

See [Graph Event Listeners](#) (p. 238).

- **Edit File Event Listener**

This permission allows user to edit a file event listener.

If the file event listener is to be edited in server web interface, the user needs to have the [list of file event listeners limited permission](#) (p. 162).

See [File Event Listeners \(remote and local\)](#) (p. 253).

- **Edit JMS Event Listener**

This permission allows user to edit a JMS event listener.

If the JMS event listener is to be edited in server web interface, the user needs to have the [list of JMS event listeners limited permission](#) (p. 163).

- **Edit Universal Event Listener**

This permission allows user to edit a universal event listener.

If the universal event listener is to be edited in server web interface, user needs to have permission [list of universal event listeners limited permission](#) (p. 163).

See [Universal Event Listeners](#) (p. 251).

- **Delete Event Listener**

This permission allows user to delete event listeners.

- **Delete Jobflow Event Listener**

This permission allows user to delete a jobflow event listener.

User needs to have the [delete\\_graph\\_event\\_listener\\_permission](#) (p. 165) to delete a jobflow event listener.

If the jobflow event listener is to be deleted in server web interface, the user needs to have the [list\\_of\\_jobflow\\_event\\_listeners\\_limited\\_permission](#) (p. 162)

- **Delete Graph Event Listener**

This permission allows user to delete a graph event listener.

If the graph event listener is to be deleted in server web interface, the user needs to have the [list\\_of\\_graph\\_event\\_listeners\\_limited\\_permission](#) (p. 162).

See [Graph Event Listeners](#) (p. 238).

- **Delete File Event Listener**

This permission allows user to delete a file event listener.

The user needs to have the [delete\\_graph\\_event\\_listener\\_permission](#) (p. 165) to delete a file event listener.

If the file event listener is to be deleted in server web interface, the user needs to have the [list\\_of\\_file\\_event\\_listeners\\_limited\\_permission](#) (p. 162).

See [File Event Listeners \(remote and local\)](#) (p. 253).

- **Delete JMS Event Listener**

This permission allows user to delete a JMS Event Listener.

The user needs to have the [delete\\_graph\\_event\\_listener\\_permission](#) (p. 165) to delete a JMS event listener.

If the graph event listener is to be deleted in server web interface, the user needs to have the [list\\_of\\_JMS\\_event\\_listeners\\_limited\\_permission](#) (p. 163).

- **Delete Universal Event Listener**

This permission allows user to delete a universal event listener.

The user needs to have the [delete\\_graph\\_event\\_listener\\_permission](#) (p. 165) to delete universal event listener.

If the universal event listener is to be deleted in server web interface, the user needs to have the [list\\_of\\_universal\\_event\\_listeners\\_limited\\_permission](#) (p. 163).

See [Universal Event Listeners](#) (p. 251).

- **Manual task Execution**

This permission allows user to manually execute a task (send an email, execute a script, etc.) with an immediate effect.

See Chapter 27, [Manual Task Execution](#) (p. 223).

- **Unlimited access to execution history**

This permission allows user to perform the same operations as [unlimited access to execution history list permission](#) (p. 166).

- **Unlimited access to execution history list**

This permission allows user to view execution history of all jobs.

- **Limited access to execution history list**

This permission allows user to view execution history of jobs from sandboxes the user can read from. In Designer, this permission is required to be able to view **Execution log** in Designer's console and execution history in **Execution** tab.

- **Launch Services**

This permission allows user to list, create, edit, and delete launch services.

See Chapter 37, [Launch Services](#) (p. 282).

- **List Launch Services unlimited**

This permission allows user to list all launch services.

- **List Launch Services Limited**

This permission allows user to list launch services from sandboxes the user can read from.

- **Create Launch service**

This permission allows user to create a new launch service.

User has to have the [create graph event listener permission](#)(p. 163) to bind the launch service with a graph.

If the launch service is to be created in server web interface, the user has to have the [list launch services limited permission](#) (p. 166)(or the [list launch services unlimited permission](#)(p. 166) to access the section with launch services.

- **Delete Launch Service**

This permission allows user to delete a launch service.

User has to have [delete graph event listener permission](#) (p. 165) to delete a launch service.

If the launch service is to be deleted in server web interface, the user needs to have the [list launch services limited permission](#) (p. 166) to access the section with launch services.

- **Edit Launch Service**

This permission allows user to edit a launch services.

User has to have [edit graph event listener](#) (p. 164) to edit the launch service.

If the launch service is to be edited in server web interface, the user needs to have the [list launch services limited permission](#) (p. 166) to choose the launch service in the server interface.

- **Tasks history**

This permission allows user to access **Tasks history** section.

See Chapter 26, [Tasks](#) (p. 203).

- **Monitoring**

**Monitoring** permission grants user all its subpermissions.

- **Monitoring section**

This permission allows user to access the monitoring section.

See Chapter 17, [Monitoring](#) (p. 131).

- **Suspend**

This permission allows user to suspend the server, a cluster node, or a sandbox.

The user needs to have the [monitoring section permission](#) (p. 167) to access the Monitoring section.

- **Suspend server**

This permission allows user to suspend or resume the server.

The user needs to have the [monitoring section permission](#) (p. 167) to access the monitoring section.

- **Suspend cluster node**

This permission allows user to suspend or resume a cluster node.

The user needs to have the [monitoring section permission](#) (p. 167) to access the monitoring section.

- **Suspend sandbox**

This permission allows user to suspend a sandbox. The user needs to have [list sandbox permission](#) (p. 161) to view the sandboxes to suspend them.

See also Chapter 21, [Sandboxes - Server Side Job Files](#) (p. 174).

- **Reset caches**

Deprecated.

- **Running jobs unlimited**

If the graph is to be run from server web interface, the user needs to have the [list sandbox permission](#) (p. 161) to list the graphs.

- **Running jobs limited**

If the graph is to be run from server web interface, the user needs to have the [list sandbox permission](#) (p. 161) to list the graphs.

- **Configuration**

This permission allows user to access the configuration section.

- **Users**

This permission allow user to access the **Users** section and configure user accounts.

- **List user**

This permission allows user to list users and access to the **Users** administration section (**Configuration** → **Users**)

- **Change passwords**

This permission allows user to change his password and to change password of another user.

To see list of users, the user needs the [list user permission](#) (p. 167).

- **Edit user**

This permission allows user to change group assignment.

To see the list of users, the user needs to have the [list user permission](#) (p. 167).

- **Edit own profile and password**

This permission allows user to change his profile (first name, last name, email, and password).

The user can access her profile in main web console view under username, in upper right corner of the page.

- **Delete user**

This permission allows user to disable a user.

The user needs to have the [list user permission](#) (p. 167) to list available users.

- **Create user**

This permission allows user to create a new user.

If the user is to be created in server web interface, the creating user needs to have the [list user permission](#) (p. 167) to list users to access this option.

- **Groups assignment**

This permission allows user to assign users to groups.

The user needs to have the [edit user permission](#) (p. 168) to successfully finish the assignment of users to groups.

If the user is to be created in server web interface, the creating user needs to have the [list user permission](#) (p. 167) to list users to access this option.

- **Groups**

This permission allows user to manage groups: user can list groups, create groups, delete groups, edit the group, assign users to the group, and change permissions of the group.

- **List groups**

This permission allows user to list groups. This permission is necessary for use of other options from the **Groups** group.

- **Create group**

This permission allows user to create a new user group.

If the user group is to be created in server web interface, the user needs to have the [list groups permission](#) (p. 168) to view a list of groups and to access this option.

- **Delete group**

This permission allows user to delete a user group.

Only empty groups can be deleted. You need to have the [list groups permission](#) (p. 168) to view list of groups and to access this option.

- **Edit group**

This permission allow user to edit user groups.

This permission does not include **User assignment** and **Permission assignment**.

If the user group is to be edited from server web interface, the user needs to have the [list groups permission](#) (p. 168).

- **Users assignment**

This permission allows user to assign users to groups.

The user needs [Edit group permission](#) (p. 169) to commit the changes in the assignment.

If the assignment is to be edited in server web interface, the user needs to have the [list groups permission](#) (p. 168) to list the groups.

- **Permission assignment**

This permission allows user to configure group **Permissions**.

The user needs have the [Edit group permission](#) (p. 169) to commit the changes.

If the permissions are to be edited in server web interface, the user needs to have the [list groups permission](#) (p. 168) to list the groups.

- **Secure parameters administration**

- **Secure params**

This permission allows user to change the value of a secure parameter.

The user can use secure parameters in graphs even without this permission.

- **CloverETL/System info sections**

This permission allows user to view **System Info** and **CloverETL Info** sections.

- **CloverETL Server properties**

This permission allows user to view **Server Properties** tab and **Data Profiler properties** tab in **CloverETL Info** section.

The user needs to have the [CloverETL/System info sections permission](#) (p. 169) to access **CloverETL Info** section.

- **Reload license**

This permission allows user to reload and view the server license.

The user needs to have the [CloverETL/System info sections permission](#) (p. 169) to access the **Configuration** section.

- **Upload license**

This permission allows user to update the server license.

The user needs to have the [CloverETL/System info sections permission](#) (p. 169) to access the **Configuration** section.

See [Activation](#) (p. 51).

- **Server Configuration Management**

This permission allows user to import and export the server configuration.

See Chapter 22, [Server Configuration Migration](#) (p. 188).

- **Export Server Configuration**

This permission allows user to export the server configuration.

See [Server Configuration Export](#) (p. 189).

- **Import Server Configuration**

This permission allows user to import the server configuration.

See [Server Configuration Import](#) (p. 190).

- **Temp Space Management**

This permission allows user to access **Temp Space Management** section.

See Chapter 18, [Temp Space Management](#) (p. 141).

- **Server Setup**

This permission allows user to access the server setup.

See Chapter 12, [Setup](#) (p. 74).

- **Heap Memory Dump**

This permission allows user to create a **Thread dump** and a **Heap Memory Dump**.

See Chapter 23, [Diagnostics](#) (p. 194).

- **Groovy Code API**

This permission allows user to run Groovy scripts.

- **Open Profiler Reporting Console**

This permission allows user to login to the **Profiler reporting console**.

The permission is necessary to view the results of Clover Profiling Jobs in Designer.

Even without this permission, a user can create and run `.cpj` jobs from Designer.

## User Lockout

**CloverETL** can lock a user access after a set number of unsuccessful login attempts as a way of protecting against brute force attacks on users' credentials.

The lockout occurs only in **CloverETL**. For example, it will not affect LDAP in the case of LDAP user authentication. By default, the feature is disabled.

Information regarding user lockout is stored in the **USER\_ACTION** server log. Notifications can be sent via email; however, it is necessary to set up a connection to an SMTP server in the E-mail (p. 80) tab of the **Setup GUI**.

The feature has several parameters which can be set by modifying the following lines in the configuration file; either directly or in the **Configuration File** tab of the Setup GUI (p. 74):

### **security.lockout.login.attempts**

Limits the number of login attempts of the user. The next failed login attempt will lock the user's access.

When setting the value, keep in mind that **CloverETL Designer** with several server projects can attempt to log in multiple times.

The recommended value is 50. Change the value to 0 to disable the feature.

### **security.lockout.reset.period**

Represents the period (in seconds) during which failed login attempts are counted. If no such attempt occurs during this period, the counter of failed login attempts is reset to 0. This way, the user does not have to worry about accidentally locking himself out of the system after a certain number of failed login attempts over an extended period of time.

The default value is 60 (1 minute). Change the value to 0 to set the period to infinity.

### **security.lockout.unlock.period**

Represents the period (in seconds) after which a successful login attempt will unlock the previously locked user. After this period, the user is able to login using his credentials again without the need to have his account unlocked by the administrator.

The parameter protects the system against denial of service (DoS) attacks and should be set to a reasonable value so you are not locked out of the system for too long in case the administrator's account is affected by the attack.

The default value is 300 (5 minutes). Change the value to 0 to set the period to infinity.

### **security.lockout.notification.email**

The parameter represents a comma separated list of emails of persons who should be notified when a user lockout occurs.

Note that the locked out user receives the notification email automatically (if the server's SMTP is configured and they have provided their email address). This parameter should therefore be set, for example, to an administrators' mail group so they are aware of the situation.

The recommended, default values are set in such a way as to efficiently protect the system against brute force attacks, prevent complete lockout of the administrator access and not limit users in standard usage of **CloverETL Server**.

The properties can be set in the following section of the properties file:

```
## Uncomment lines bellow to enable user lockout after number of failed logins
## Number of failed login attempts after which a next failed login attempt will lock the user
## 0 means feature is switched off
## default suggested value is 50
#security.lockout.login.attempts=50
## Periods are specified in seconds
## Period of time during which the failed login attempts are counted
## Default is 60s (1 min)
#security.lockout.reset.period=60
## Period of time after which a successful login attempt will unlock previously locked user
## Default is 300s (5 min)
#security.lockout.unlock.period=300
## Comma separated list of emails which will be notified when user is locked out.
#security.lockout.notification.email=
```

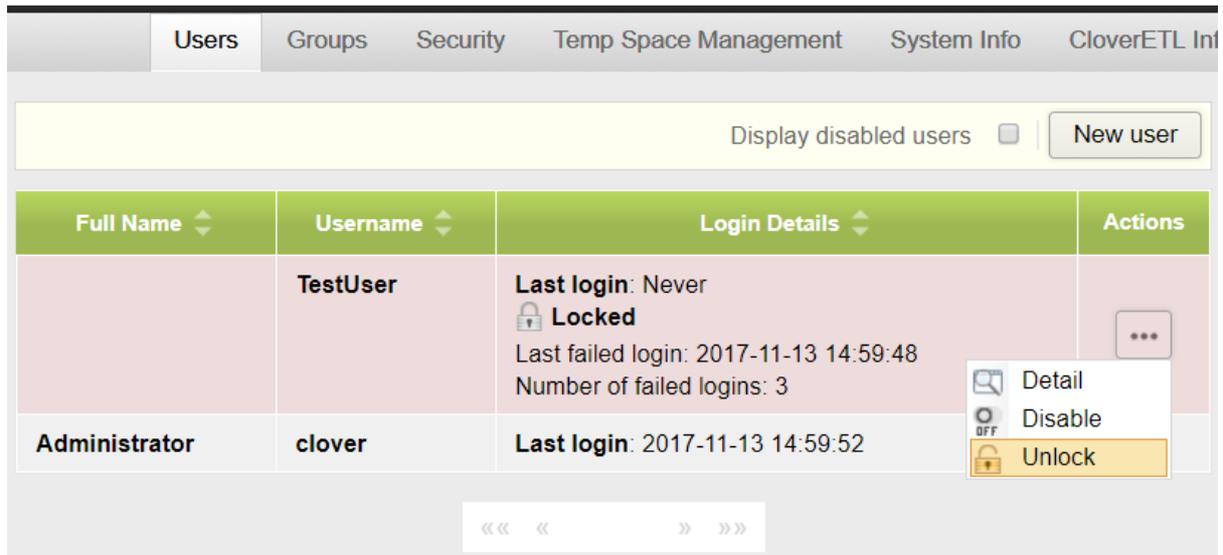
### Unlocking User

Once the user's access is locked, you can see the status in the **Users** tab of the **Configuration** section.

Full Name	Username	Login Details	Actions
	TestUser	Last login: Never Locked Last failed login: 2017-11-13 14:59:48 Number of failed logins: 3	...
Administrator	clover	Last login: 2017-11-13 14:59:52	...

Figure 20.8. Locked User

To unlock the user, click on the (...) button in the respective row of the **Action** column and choose **Unlock**.



The screenshot shows the 'Users' management interface in CloverETL. At the top, there are navigation tabs: 'Users', 'Groups', 'Security', 'Temp Space Management', 'System Info', and 'CloverETL Int'. Below the tabs is a search bar with the text 'Display disabled users' and a checkbox, and a 'New user' button. The main content is a table with the following columns: 'Full Name', 'Username', 'Login Details', and 'Actions'. The table contains two rows: one for 'TestUser' (locked) and one for 'Administrator' (active). The 'TestUser' row is highlighted in red, and its 'Actions' column shows a dropdown menu with options: 'Detail', 'Disable', and 'Unlock'. The 'Unlock' button is highlighted in orange. At the bottom of the table, there are navigation arrows: '<<< << >> >>>'.

Full Name	Username	Login Details	Actions
	TestUser	Last login: Never 🔒 Locked Last failed login: 2017-11-13 14:59:48 Number of failed logins: 3	⋮ Detail Disable Unlock
Administrator	clover	Last login: 2017-11-13 14:59:52	

Figure 20.9. Unlocking User

Once you click the **Unlock** button, you will be prompted to confirm the action.

Click **Yes** to unlock the user.

## Chapter 21. Sandboxes - Server Side Job Files

A sandbox is a place where you store all your project’s transformation graph files, jobflows, data, and other resources. It’s a server side analogy to a Designer project. The Server adds additional features to sandboxes, like user permissions management and global per-sandbox configuration options.

The Server and the Designer are integrated so that you are able to connect to a Server sandbox using a “Server Project” in your Designer workspace. Such a project works like a remote file system – all data is stored on the Server and accessed remotely. Nonetheless, you can do everything with Server Projects the same way as with local projects – copy and paste files, create, edit, and debug graphs, etcetera. See the **CloverETL Designer manual** for details on configuring a connection to the Server.

Technically, a sandbox is a dedicated directory on the Server host file system and its contents are managed by the Server. Advanced types of sandboxes, like “partitioned sandbox” have multiple locations to allow distributed parallel processing (more about that in Chapter 39, [Clustering Features](#) (p. 312)). A sandbox cannot contain another sandbox within – it’s a single root path for a project.

It’s recommended to put all sandboxes in a folder outside the CloverETL Server installation (by default the sandboxes would be stored in the `${user.data.home}/CloverETL/sandboxes`, where the “user.data.home” is automatically detected user home directory). However, each sandbox can be located on the file system independently of the others if needed. The containing folder and all its contents must have read/write permission for the user under which the CloverETL Server/application server is running.

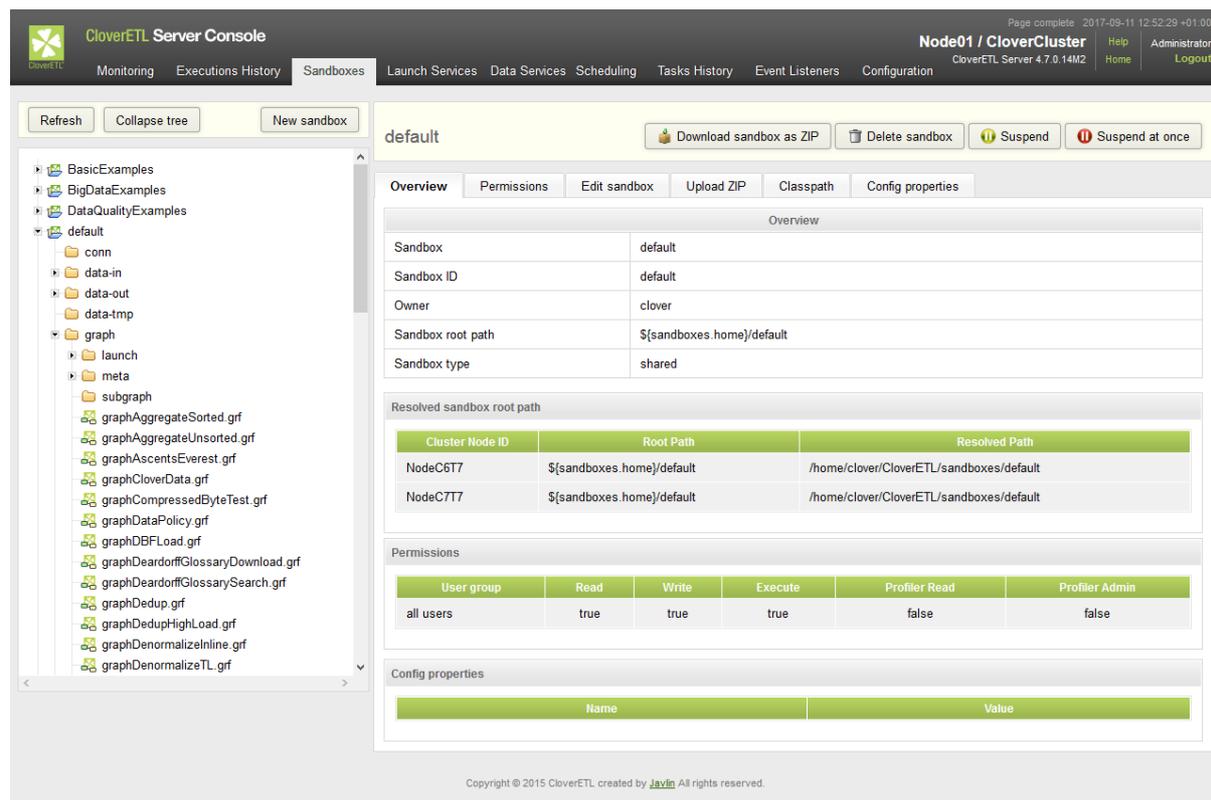


Figure 21.1. Sandboxes Section in CloverETL Server Web GUI

Each sandbox in non-cluster environment is defined by following attributes:

Table 21.1. Sandbox attributes

Sandbox	Sandbox name used just for display. It is specified by user in during sandbox creation and it can be modified later.
Sandbox ID	A unique "name" of the sandbox. It is used in server APIs to identify sandbox. It must meet common rules for identifiers. It is specified by user in during sandbox creation and it can be modified later. <i>Note: modifying is not recommended, because it may be already used by some APIs clients.</i>
Sandbox root path	Absolute server side file system path to sandbox root. It is specified by user during sandbox creation and it can be modified later. Instead of the absolute path, it's recommended to use <code>\${sandboxes.home}</code> placeholder, which may be configurable in the CloverETL Server configuration. So e.g. for the sandbox with ID "dataReports" the specified value of the "root path" would be <code>"\${sandboxes.home}/dataReports"</code> . Default value of "sandboxes.home" config property is <code>"\${user.data.home}/CloverETL/sandboxes"</code> where the "user.data.home" is configuration property specifying home directory of the user running JVM process - it's OS dependent). Thus on the unix-like OS, the fully resolved sandbox root path may be: <code>"/home/clover/CloverETL/sandboxes/dataReports"</code> . See Chapter 39, <a href="#">Clustering Features</a> (p. 312) for details about sandboxes root path in cluster environment.
Owner	It is set automatically during sandbox creation. It may be modified later.
Sandbox type	Type of the sandbox. It can be: shared, partitioned or local.

---

## Referencing Files from the ETL Graph or Jobflow

In some components you can specify file URL attribute as a reference to some resource on the file system. Also external metadata, lookup or DB connection definition is specified as reference to some file on the filesystem. With CloverETL Server there are more ways how to specify this relation.

- Relative path

All relative paths in your graphs are considered as relative paths to the root of the same sandbox which contains job file (ETL graph or Jobflow).

- sandbox:// URLs

Sandbox URL allows user to reference the resource from different sandboxes with standalone CloverETL Server or the cluster. In cluster environment, CloverETL Server transparently manages remote streaming if the resource is accessible only on some specific cluster node.

See [Using a Sandbox Resource as a Component Data Source](#) (p. 317) for details about the sandbox URLs.

## Sandbox Content Security and Permissions

Each sandbox has its owner which is set during sandbox creation. This user has unlimited privileges to this sandbox as well as administrators. Another users may have access according to sandbox settings.

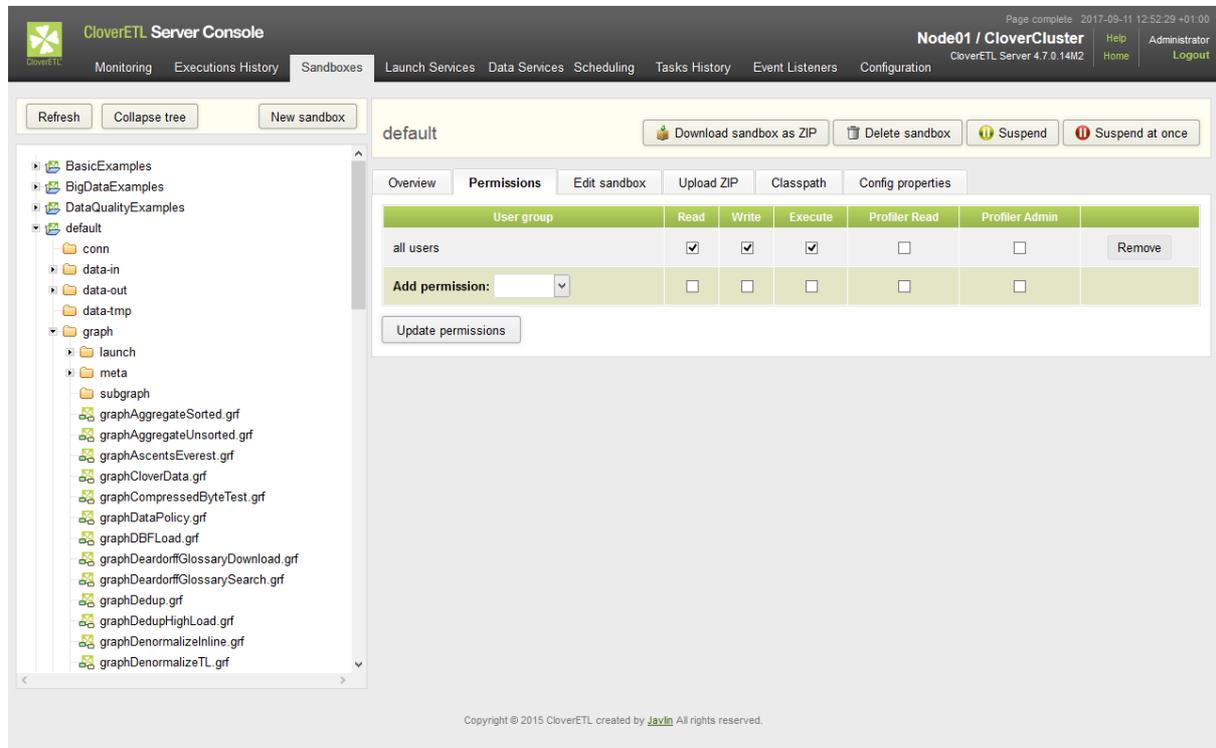


Figure 21.2. Sandbox Permissions in CloverETL Server Web GUI

Permissions to a specific sandbox are modifiable in **Permissions** tab in sandbox detail. In this tab, selected user groups may be allowed to perform particular operations.

There are the following types of operations:

Table 21.2. Sandbox permissions

Read	Users can see this sandbox in their sandboxes list.
Write	Users can modify files in the sandbox through CS APIs.
Execute	Users can execute jobs in this sandbox. <i>Note: jobs executed by "graph event listener" and similar features is actually executed by the same user as job which is source of event. See details in "graph event listener". Job executed by schedule trigger is actually executed by the schedule owner. See details in Chapter 28, <a href="#">Scheduling</a> (p. 224). If the job needs any files from the sandbox (e.g. metadata), user also must have read permission, otherwise the execution fails.</i>
Profiler Read	User can view results of profiler jobs executed from the sandbox.
Profiler Admin	User can administer results of profiler jobs executed from the sandbox.

Please note that, these permissions modify access to the content of specific sandboxes. In additions, it's possible to configure permissions to perform operations with sandbox configuration. e.g. create sandbox, edit sandbox, delete sandbox, etc. Please see Chapter 20, [Users and Groups](#) (p. 151) for details.

## Sandbox Content

Sandbox should contain jobflows, graphs, metadata, external connection and all related files. Files, especially graph or jobflow files, are identified by relative path from sandbox root. Thus you need two values to identify specific job file: sandbox and path in sandbox. Path to the Jobflow or ETL graph is often referred as "Job file".

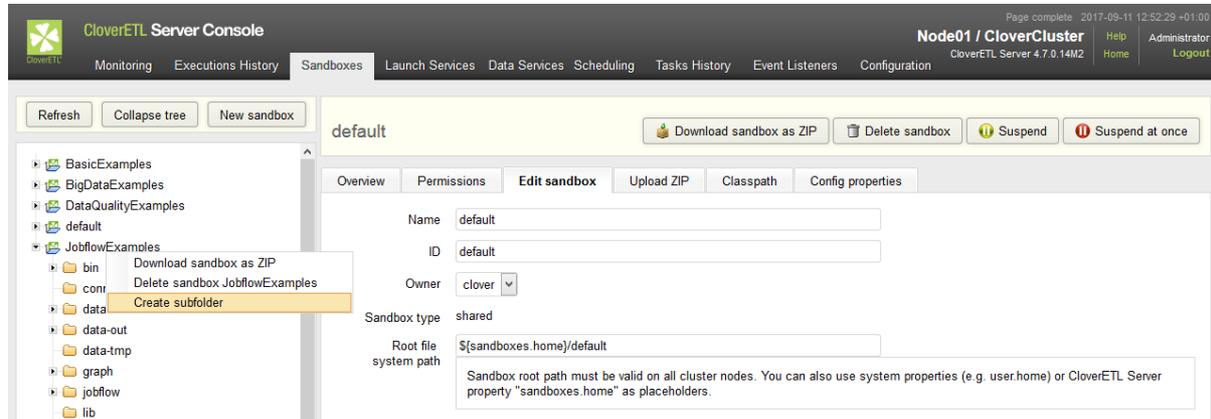


Figure 21.3. Web GUI - section "Sandboxes" - context menu on sandbox

Although web GUI section **sandboxes** isn't file-manager, it offers some useful features for sandbox management.

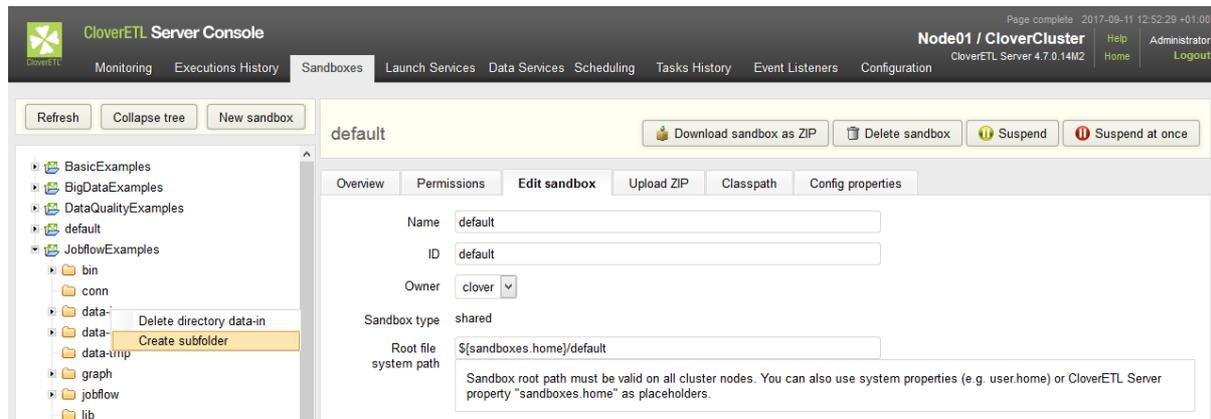


Figure 21.4. Web GUI - section "Sandboxes" - context menu on folder

## Download sandbox as ZIP

Select a sandbox in left panel, then web GUI displays button "Download sandbox as ZIP" in the tool bar on the right side.

Created ZIP contains all readable sandbox files in the same hierarchy as on file system. You can use this ZIP file for upload files to the same sandbox, or another sandbox on different server instance.

## Chapter 21. Sandboxes - Server Side Job Files

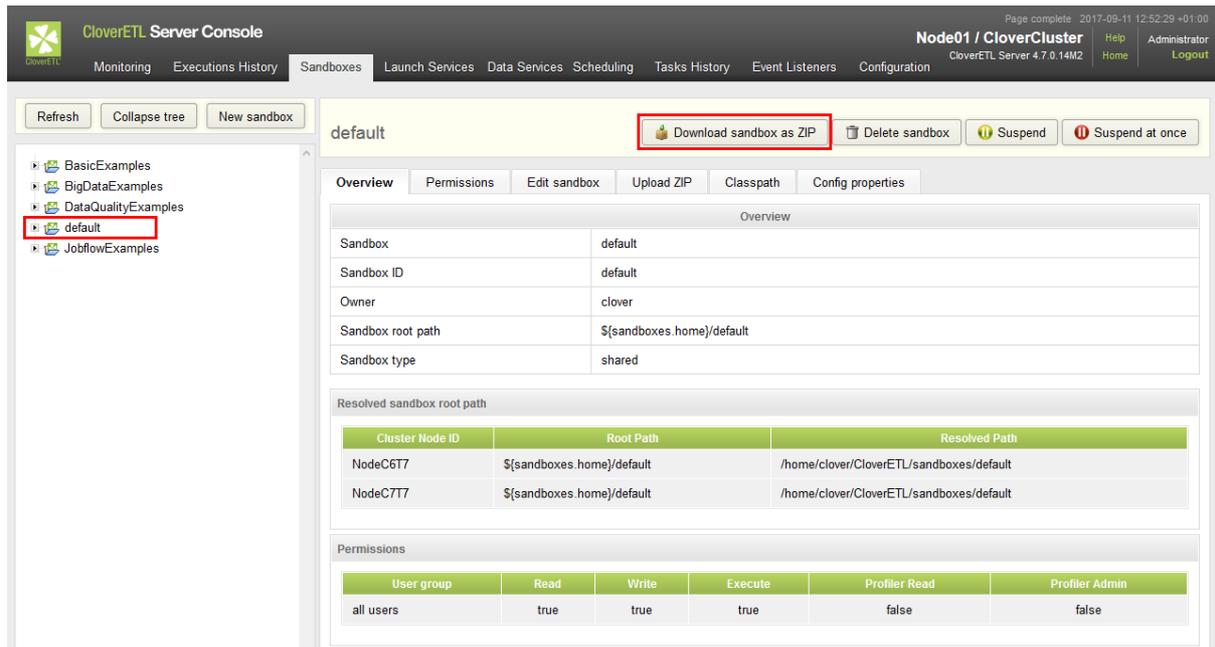


Figure 21.5. Web GUI - download sandbox as ZIP

### Upload ZIP to sandbox

Select a sandbox in left panel. You must have write permission to the selected sandbox. Then select tab "Upload ZIP" in the right panel. Upload of a ZIP is parametrized by couple of switches, which are described below. Open a common file chooser dialog by button "+ Upload ZIP". When you choose a ZIP file, it is immediately uploaded to the server and result message is displayed. Each row of the result message contains description of one single file upload. Depending on selected options, file may be skipped, updated, created or deleted.

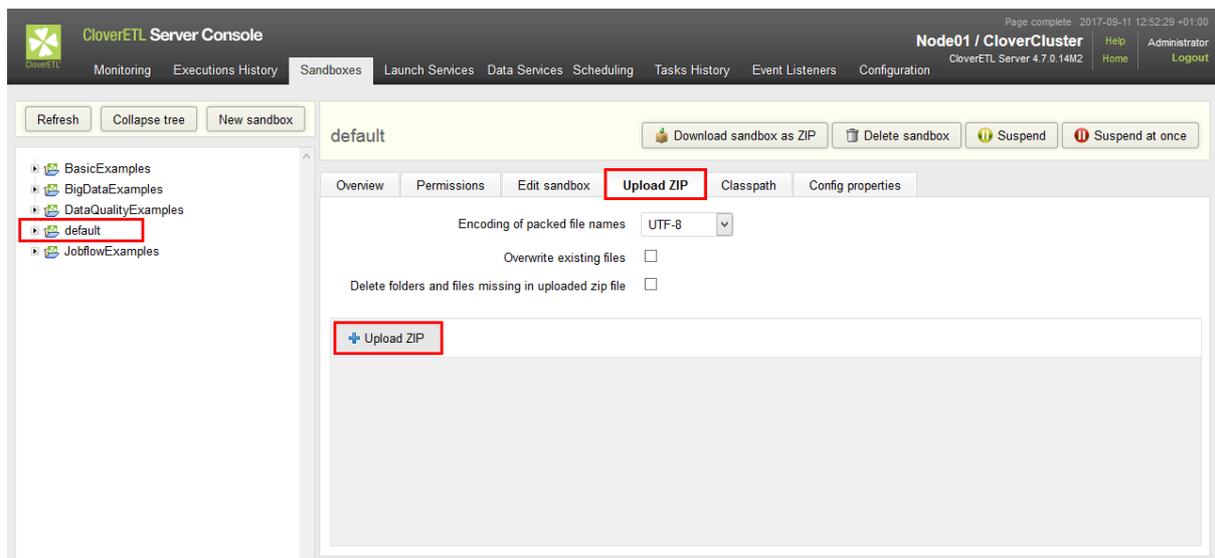


Figure 21.6. Web GUI - upload ZIP to sandbox

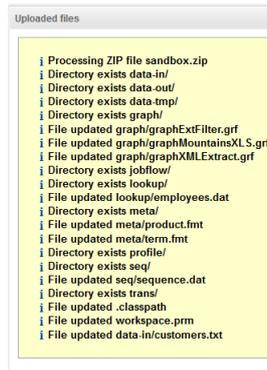


Figure 21.7. Web GUI - upload ZIP results

Table 21.3. ZIP upload parameters

Label	Description
Encoding of packed file names	File names which contain special characters (non ASCII) are encoded. By this select box, you choose right encoding, so filenames are decoded properly.
Overwrite existing files	If this switch is checked, existing file is overwritten by a new one, if both of them are stored in the same path in the sandbox and both of them have the same name.
Replace content	If this option is enabled, all files which are missing in uploaded ZIP file, but they exist in destination sandbox, will be deleted. This option might cause loose of data, so user must have special permission "May delete files, which are missing in uploaded ZIP" to enable it.

## Download file in ZIP

Select a file in the left panel, then web GUI displays button "Download file as ZIP" in the tool bar on the right side.

Created ZIP contains just selected file. This feature is useful for large files (i.e. input or output file) which cannot be displayed directly in web GUI. So user can download it.

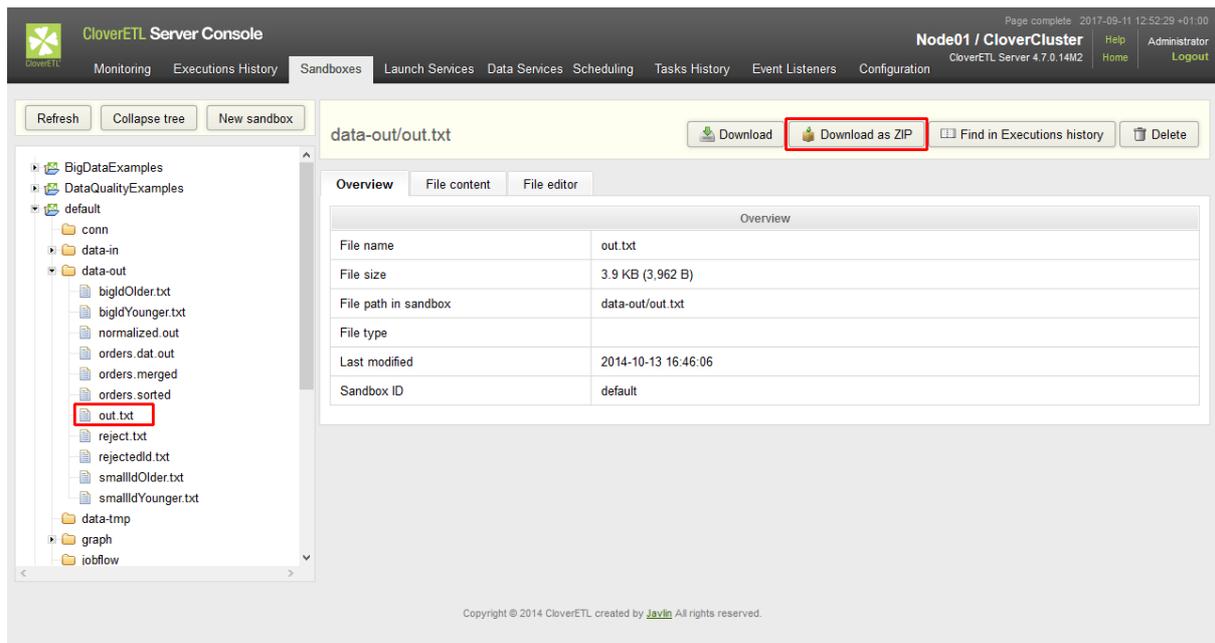


Figure 21.8. Web GUI - download file as ZIP

## Download file HTTP API

It is possible to download/view sandbox file accessing "download servlet" by simple HTTP GET request:

```
http://[host]:[port]/[Clover Context]/downloadFile?[Parameters]
```

Server requires BASIC HTTP Authentication. Thus with Linux command line HTTP client "wget" it would look like this:

```
wget --user=clover --password=clover  
http://localhost:8080/clover/downloadFile?sandbox=default&file=data-out/data.dat
```

Please note, that ampersand character is escaped by back-slash. Otherwise it would be interpreted as command-line system operator, which forks processes.

URL Parameters

- sandbox - Sandbox code. Mandatory parameter.
- file - Path to the file relative from sandbox root. Mandatory parameter.
- zip - If set to "true", file is returned as ZIP and response content type is "application/x-zip-compressed". By default it is false, so response is content of the file.

## Delete Sandbox

You can delete a sandbox by selecting the sandbox and clicking the **Delete sandbox** button on the top of the right pane,

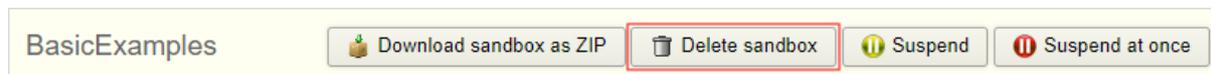


Figure 21.9. Delete Sandbox - Button

or by right-clicking the sandbox in the tree pane on the left and selecting the **Delete sandbox** option.

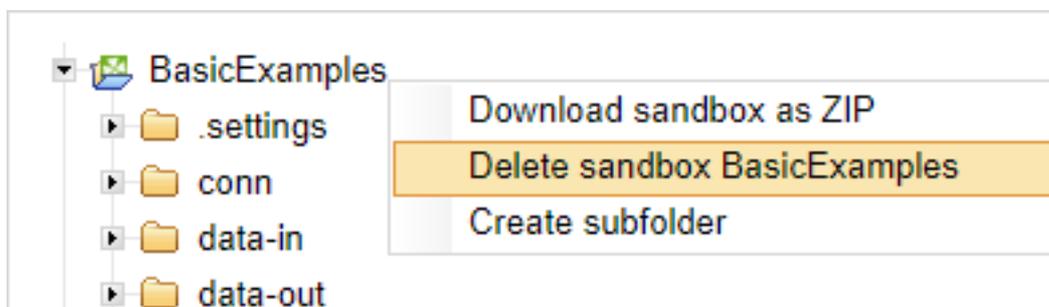
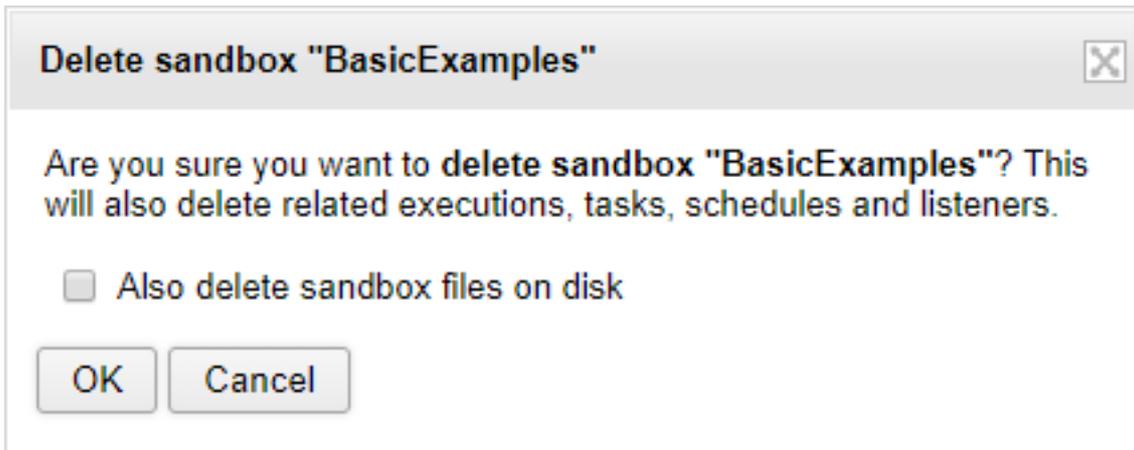


Figure 21.10. Delete Sandbox - Right-Click

After that, a confirmation dialog opens where you can choose to delete sandbox files on disk, as well.



*Figure 21.11. Delete Sandbox - Confirmation Dialog*

## Job Config Properties

Each ETL graph or Jobflow may have set of config properties, which are applied during the execution. Properties are editable in web GUI section "sandboxes". Select job file and go to tab "Config properties".

The same config properties are editable even for each sandbox. Values specified for sandbox are applied for each job in the sandbox, but with lower priority then config properties specified for the job.

If neither sandbox nor job have config properties specified, defaults from main server configuration are applied. Global config properties related to Job config properties have prefix "executor.". E.g. server property "executor.classpath" is default for Job config property "classpath". (See Part III, "[Configuration](#)" (p. 67) for details)

In addition, it is possible to specify additional job parameters, which can be used as placeholders in job XML. Please keep in mind, that these placeholders are resolved during loading and parsing of XML file, thus such job couldn't be pooled.

If you use a relative path, the path is relative to `${SANDBOX_ROOT}`.

In path definition, you can use system properties - e.g. `${java.io.tmpdir}` - and some of server config properties: `${sandboxes.home}`, `${sandboxes.home.partitioned}` and `${sandboxes.home.local}`.

Table 21.4. Job config parameters

Property name	Default value	Description
classloader_caching	false	Clover creates new classloaders whenever is necessary to load a class in runtime. For example, Reformat component with a Java transformation has to create a new classloader to load the class. It is worth noting that classloaders for JDBC drivers are not re-created. Classloader cache is used to avoid PermGen out of memory errors (some JDBC drivers automatically register itself to DriverManager, which can cause the classloader cannot be released by garbage collector). This behavior can be inconvenient for example if you want to share POJO between components. For example, a Reformat component creates an object (from a jar file on runtime classpath) and stores it into a dictionary. Another Reformat component get the object from the dictionary and tries to cast the object to expected class. ClassCastException is thrown due different classloaders used in the Reformat components. Using this flag you can force CloverServer to re-use classloader when possible.
classpath		List of paths or jar files which contain external classes used in the job file (transformations, generators, JMS processors). All specified resources will be added to runtime classpath of the transformation job. All Clover Engine libraries and libraries on application-server's classpath are automatically on the classpath. Separator is specified by Engine property "DEFAULT_PATH_SEPARATOR_REGEX". Directory path must always end with a slash character "/", otherwise ClassLoader doesn't recognize it's a directory. Server always automatically adds "trans" subdirectory of job's sandbox, so It doesn't have to be added explicitly.
compile_classpath		List of paths or jar files which contain external classes used in the job file (transformations, generators, JMS processors) and related libraries for their compilation. Please note, that libraries on application-server's classpath aren't included automatically. Separator is specified by Engine property "DEFAULT_PATH_SEPARATOR_REGEX". The directory path must always end with a slash character "/", otherwise ClassLoader doesn't recognize it's a directory. Server always automatically adds "SANDBOX_ROOT/trans/" directory and all JARs in "SANDBOX_ROOT/lib/" directory, so they don't have to be added explicitly.
debug_mode		If true, edges with debug enabled will store data into files in a debug directory.

Property name	Default value	Description
	false	Without explicit setting, running of a graph from Designer with server integration would set the debug_mode to true. On the other hand, running of a graph from the server console sets the debug_mode to false.
delete_obsolete_temp_files	false	If true, system will remove temporary files produced during previous finished runs of the respective job.  This property is useful together with enabled debug mode ensuring that obsolete debug files from previous runs of a job are removed from temp space. This property is set to "true" by default when executing job using designer-server integration.
enqueue_executions	false	Boolean value. If it is true, executions above max_running_concurrently are enqueued, if it is false executions above max_running_concurrently fail.
jobflow_token_tracking	true	If false, token tracking in jobflow executions will be disabled.
locale	DEFAULT_LOCALE engine property	Can be used to override the DEFAULT_LOCALE engine property.
log_level	INFO	Log4j log level for this graph executions. (ALL   TRACE   DEBUG   INFO   WARN   ERROR   FATAL) For lower levels (ALL, TRACE or DEBUG), also root logger level must be set to lower level. Root logger log level is INFO by default, thus transformation run log does not contain more detail messages then INFO event if job config parameter "log_level" is set properly. See Chapter 16, <a href="#">Logging</a> (p. 127) for details about log4j configuration.
max_graph_instance_age	0	A time interval in milliseconds which specifies how long may a transformation instance last in server's cache. 0 means that the transformation is initialized and released for each execution. The transformation cannot be stored in the pool and reused in some cases (a transformation uses placeholders using dynamically specified parameters)
max_running_concurrently	unlimited	Max number of concurrently running instances of this transformation. In cluster environment, the limit is per node.
password		This property is deprecated. Password for decoding of encoded DB connection passwords.
skip_check_config	default value is taken from engine property	Switch which specifies whether check config must be performed before transformation execution.
time_zone	DEFAULT_TIME_ZONE engine property	Can be used to override the DEFAULT_TIME_ZONE engine property.
tracking_interval	2000	Interval in milliseconds for sampling nodes status in running transformation.
use_jmx	true	If true, job executor registers jmx mBean of running transformation.
use_local_context_url	false	If true, the context URL of a running job will be a local "file:" URL. Otherwise, a "sandbox:" URL will be used.
verbose_mode	true	If true, more descriptive logs of job runs are generated.
worker_execution	true	Set to false to enforce execution in Server Core. Can be set per file or per sandbox.

## Chapter 21. Sandboxes - Server Side Job Files

The screenshot displays the CloverETL Server Console interface. The top navigation bar includes 'Monitoring', 'Executions History', 'Sandboxes', 'Launch Services', 'Data Services', 'Scheduling', 'Tasks History', 'Event Listeners', and 'Configuration'. The current page is titled 'graph/graphAggregateSorted.grf' and features buttons for 'Run graph', 'Download', 'Download as ZIP', 'Show in Executions history', and 'Delete'. The left sidebar shows a tree view of job files, with 'graph/graphAggregateSorted.grf' selected. The main content area is titled 'Config properties' and contains a table of job configuration properties. The table has columns for 'Name' and 'Value'. The properties listed are:

Name	Value	
tracking_interval	500 Interval in ms for sampling nodes status in running job.	Delete
verbose_mode	true Boolean value which specifies whether to put more detail messages to job run log.	Delete
log_level	debug Log4j log level for this job executions. (ALL   TRACE   DEBUG   INFO   WARN   ERROR   FATAL)	Delete
worker_execution	true Boolean value. Default value is true - jobs are executed in separate workers. You can force the job execution inside CloverETL Server using this config property.	Delete

Below the table is an 'Update' button and a section for 'Properties inherited from sandbox' with a table structure for Name and Value.

Figure 21.12. Job config properties

---

## WebDAV Access to Sandboxes

Since 3.1

WebDAV API allows you to access and manage sandbox content using a standard WebDAV specification.

Specifically, it allows for:

- Browsing a directory structure
- Editing files
- Removing files/folders
- Renaming files/folders
- Creating files/folders
- Copying files
- Moving files

The WebDAV interface is accessible from the URL: "http://[host]:[port]/clover/webdav".

Note: Although common browsers will open this URL, most of them are not rich WebDAV clients. Thus, you will only see a list of items, but you cannot browse the directory structure.

---

## WebDAV Clients

There are many WebDAV clients for various operating systems, some OS support WebDAV natively.

### Linux like OS

Great WebDAV client working on Linux systems is Konqueror. Please use different protocol in the URL: `webdav://[host]:[port]/clover/webdav`

Another WebDAV client is Nautilus. Use different protocol in the URL `dav://[host]:[port]/clover/webdav`.

### MS windows

Last distributions of MS Windows (Win XP and later) have native support for WebDAV. Unfortunately, it is more or less unreliable, so it is recommended to use some free or commercial WebDAV client.

- The best WebDAV client we've tested is BitKinex: <http://www.bitkinex.com/webdavclient>
- Another option is to use Total Commander (<http://www.ghisler.com/index.htm>) with WebDAV plugin: <http://www.ghisler.com/plugins.htm#filesys>

### Mac OS

Mac OS supports WebDAV natively and in this case it should be without any problems. You can use "finder" application, select "Connect to the server ..." menu item and use URL with HTTP protocol: "http://[host]:[port]/clover/webdav".

---

## WebDAV Authentication/Authorization

CloverETL Server WebDAV API uses the HTTP Basic Authentication by default. However it may be reconfigured to use HTTP Digest Authentication.

Digest Authentication may be useful, since some WebDAV clients can't work with HTTP Basic Authentication, only with Digest Authentication.

HTTP Digest Authentication is feature added to the version 3.1. If you upgraded your older CloverETL Server distribution, users created before the upgrade cannot use the HTTP Digest Authentication until they reset their passwords. So when they reset their passwords (or the admin does it for them), they can use Digest Authentication as well as new users.

The HTTP Digest Authentication is configured with `security.digest_authentication.*` configuration properties. To enable it, set `security.digest_authentication.features_list` to contain features that are listed in the `security.digest_authentication.features_list`. As items in `security.basic_authentication.features_list` have higher priority, you should empty it to allow HTTP Digest Authentication to be used.

See Chapter 14, [List of Configuration Properties](#) (p. 106) for details.

For details on authentication methods see <https://tools.ietf.org/html/rfc7617> and <https://tools.ietf.org/html/rfc2617>.

---

## Chapter 22. Server Configuration Migration

CloverETL Server provides means to migrate its configuration (e.g. event listeners, schedules etc.) or parts of the configuration between separate instances of the server. A typical use case is deployment from test environment to production - this involves not only deployment of CloverETL graphs, but also copying parts of configuration such as file event listeners etc.

Configuration migration is performed in 2 steps - export of the configuration from the source server, followed by import of the configuration at the destination server. After exporting, the configuration is stored as an XML file. The file can be modified manually before import, for example to migrate only parts of the configuration. Additionally, the configuration file can be stored in a versioning system (such as Subversion or Git) for versioning of the CloverETL Server configuration.

It is recommended to perform import of configuration on a suspended CloverETL Server and to plan for maintenance. Additionally, it is recommended to backup the CloverETL Server configuration database before the import.

The following items are parts of the *Server Configuration* and can be migrated between servers:

- Users & Groups (p. 151)
- Sandboxes (p. 174)
- Job Parameters (p. 199)
- Schedules (p. 224)
- [Graph Event Listeners](#)(p. 238) [Jobflow Event Listeners](#)(p. 244) [JMS Message Listeners](#)(p. 246) [File Event Listeners \(remote and local\)](#) (p. 253)
- Launch Services (p. 282)
- Temp Spaces (p. 141)

### Permissions for Configuration Migration

Whether a user is entitled to perform configuration migration is determined by having *Server Configuration Management* permission; this permission has two sub-permissions: *Export Server Configuration* and *Import Server Configuration* (see [Groups permissions](#) (p. 160) for further information on permissions). These permissions are of higher priority than permissions related to a particular migrated item type - so even if the user does not have a permission e.g. to list server's schedules, with *Export Server Configuration* he will be allowed to export all of defined schedules. The same is true for adding and changing items with the *Import Server Configuration* permission.

See [Server Configuration permission](#) (p. 170).

## Server Configuration Export

Export of a server configuration is performed from the Server Console - the screen for export can be found in section **Configuration > Export**. You can choose which items will be exported (see Figure 22.1 (p. 189)). After clicking on the **Export Configuration** an XML file will be offered for download. The name of the XML file reflects time when the configuration was exported.

In case user manually edits the exported XML file, you should ensure that the file has a valid content. This can be done by validation against XSD schema. The schema for a configuration XML document can be found at `http://[host]:[port]/[contextPath]/schemas/clover-server-config.xsd`.

The XML file contains selected items of the CloverETL server instance. The file can be modified before the import to another server instance - for example to import schedules only.

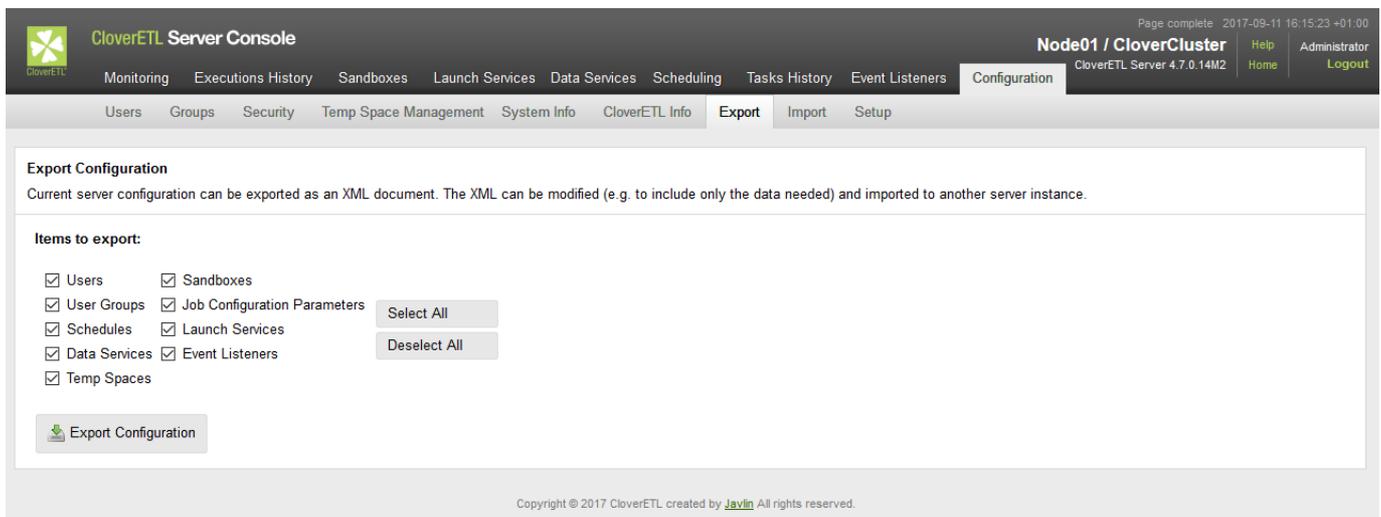


Figure 22.1. Server Configuration Export screen

## Server Configuration Import

Import of CloverETL Server configuration merges the configuration exported from another server into the running server instance where the import was initiated. The configuration to be imported is loaded from an XML file created by export, see [Server Configuration Export](#) (p. 189). Import of server configuration is performed from the Server Console - the screen for import can be found in **Configuration > Import** section.

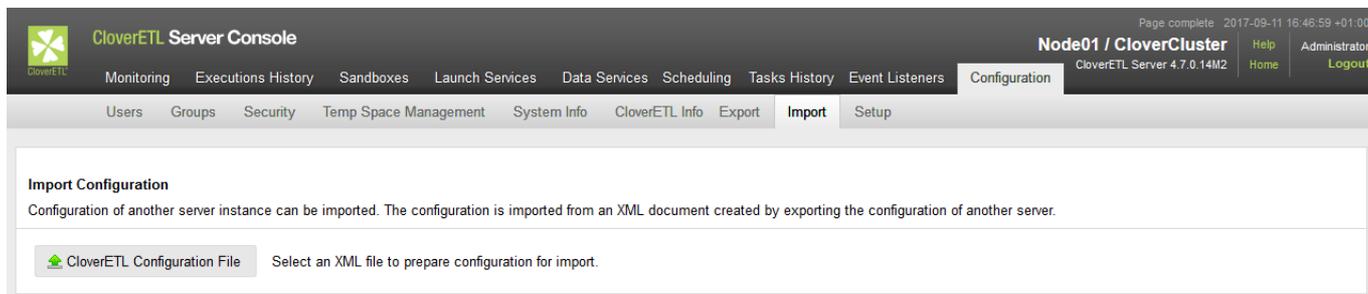


Figure 22.2. Server Configuration Import screen

The XML configuration file defines configuration items to be imported into the destination server. The items are matched against current configuration of the destination server. Depending on result of the matching, the items from the XML configuration are either added to the destination server or will update existing item with properties defined in the XML file. Matching of items is based on a key that depends on the item type:

- users - user code
- user groups - group code
- sandboxes - sandbox code
- job parameters - triplet (job parameter name, sandbox code, job file)
- event listeners - event listener name
- schedule - schedule description
- launch service - triplet (service name, server user, service user group)
- data service - owner's username, sandbox, path to job file, public availability (y/n), saving runrecords (y/n), state (enabled/disabled)
- temp spaces - pair (temp space node ID, temp space path)

## Configuration Import Process

### Uploading Configuration

The first step in the configuration import is to upload the XML file to the CloverETL server. After clicking on **CloverETL Configuration File** button a window is opened where user can select an XML file with the configuration to import. The file is uploaded automatically after the dialog is closed. Once upload is finished the name of the uploaded file is shown in the toolbar along with **CloverETL Configuration File** button. In case reading of configuration from XML has finished without error, additional controls are displayed in the toolbar:

- **Preview Import** button to perform "dry run" of the configuration import
- **Commit Import** button to perform actual import of configuration to server's database
- **Import Options** section to further configure import process:
  - **New only** option specifies that only new items will be imported leaving existing items on server untouched
- **Import Items** section to select what item types will be imported to the server

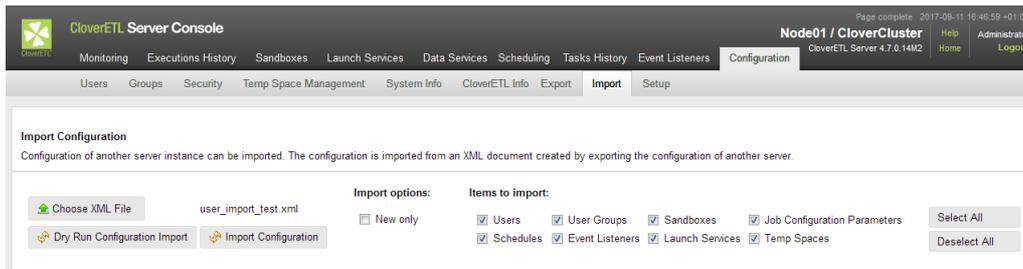


Figure 22.3. Server Configuration uploaded



### Note

When transferring configuration from one server instance to another, it is important that these instances are of compatible, preferably the same, version. The user is notified when the source and target differ at at least minor version number (e.g. 3.4.1 and 3.5.0). It is also recommended not to transfer configuration between clustered and non-clustered server instances.

## Verifying Configuration

Once the configuration is uploaded, the system executes "dry run" of the configuration import automatically. The *dry run* performs configuration import, but no data is actually written do the server's database. The outcome of the operation is **Import Log** with messages related either to the configuration as a whole or to particular imported items (see Figure 22.4 (p. 192)). There is also another view of **Imported Items** to list all added/updated items grouped into tables according to their types. Each item has an icon indicating result of the item import operation:

- + - the item has been added as a new one
- \$ - the item has been updated
- • - the item has been updated, but none of its properties has changed
- - - the item has been removed

For the updated items, the state of the item before update is shown in the lower of the rows with less contrast text color, the new values of item's properties are shown in upper of the rows with regular text color. Actual changes are highlighted by background color change on respective property and also on both rows. The **Imported Items** view can be filtered using radio buttons above it:

- **Changes only** button will display only items that have been either added or actually changed by update
- **All updates** button will display all of imported items, event those identical to already present ones

### Example 22.1. Example of simple configuration defining one new server user.

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<cloverServerConfiguration xmlns="http://cloveretl.com/server/data" timeZone="Europe/Berlin">
  <usersList>
    <user disabled="false">
      <username>johnsmith</username>
      <firstName>John</firstName>
      <lastName>Smith</lastName>/
      <email>smithj@tnet.com</email>
      <domain>clover</domain>
      <password>SITx8e1pjoo4em17As3LNw==</password>
      <passwordA1Digest>70151bae7488da4159ac5ccec97d0995</passwordA1Digest>
      <userGroups>
        <groupCode>all_users</groupCode>
        <groupCode>job_managers</groupCode>
      </userGroups>
    </user>
  </usersList>
</cloverServerConfiguration>
```

## Chapter 22. Server Configuration Migration

```
</user>
</usersList>
</cloverServerConfiguration>
```

**Import Configuration**  
Configuration of another server instance can be imported. The configuration is imported from an XML document created by exporting the configuration of another server.

CloverETL Configuration File: clover\_configuration\_2017-09-11\_17-04-13+0100.xml  
Export timestamp: Sep 11, 2017 5:04:13 PM  
Exported by user: clover  
Node ID: Node01  
Timezone: Europe/London  
CloverETL Version: 4.7.0.14M2

**Import options:**  
 New only

**Items to import:**  
 Users  
 Schedules  
 Sandboxes  
 Launch Services  
 Data Services  
 User Groups  
 Event Listeners  
 Job Configuration Parameters  
 Temp Spaces

Buttons: Preview Import, Commit Import, Import Options <<, Cancel, Select All, Deselect All

**Import Log**

**Summary**  
Preview of configuration import revealed one or more errors in the configuration. Fix them in the XML document, re-upload it, and execute import preview in order to verify its validity.

Hide Messages <<  Show imported items with info-only messages

**Message item:** User [id=1, username=johnsmith]  
**Messages:**  
User 'johnsmith' refers to group 'all\_users' but the user is not present among group's members. Either change the imported configuration so that group 'all\_users' references the user as well, or do not list the users of the group 'all\_users' at all.

Preview of server configuration import completed. 0 items would have been added, 0 items updated.  
Import preview has detected an error that would prevent configuration import.

**Imported Items**  
Filter view:  Changes only  All updates  
Text View Tabular View  
No changes.

Figure 22.4. Outcome of the import preview for configuration from Example 22.1 (p. 191)

The **Summary** in the **Import Log** says whether the dry run was successful. Should there be any problems with items imported, the item is displayed along with the cause of the error (see Figure 22.4 (p. 192)) .

**Import Configuration**  
 Configuration of another server instance can be imported. The configuration is imported from an XML document created by exporting the configuration of another server.

CloverETL Configuration File: user\_import\_test.xml  
 Export timestamp: <unknown> Exported by user: <unknown>  
 Timezone: Europe/Berlin CloverETL Version: <unknown>

Preview Import Import Options << Commit Import

**Import options:**  New only

**Items to import:**  Users  Schedules  Sandboxes  Launch Services  User Groups  Event Listeners  Job Configuration Parameters  Temp Spaces

Select All Deselect All

**Import Log**

**Summary**  
 Preview of configuration completed without errors. Now you can proceed with actual configuration import by clicking on 'Commit Import'.

Hide Messages <<  Show imported items with info-only messages

Configuration does not specify server version.  
 Preview of server configuration import completed. 1 items would have been added, 1 items updated.

**Imported Items**  
 Filter view:  Changes only  All updates

Text View Tabular View

**Users & Groups**

Users									
	User Name	First Name	Last Name	E-mail	Domain	Disabled	Password	Password A1 Digest	Groups
+	johnsmith	John	Smith	smith@tnet.com	clover	<input checked="" type="checkbox"/>	SITx8e1pjoo4em17As3LNw==	70151bae7488da4159ac5ccec97d0995	all_users

User Groups					
	Code	Name	Description	Users	Permissions
\$	all_users	all users		clover_user1, johnsmith	[sandbox_all, user_edit_me, scheduling_all, event_listener_all, run_record_all, launch_config_all, task_log_all, monitoring_all]
		all users		clover_user1	[sandbox_all, user_edit_me, scheduling_all, event_listener_all, run_record_all, launch_config_all, task_log_all, monitoring_all]

Figure 22.5. Outcome of import preview for configuration after fixing by removal of broken group reference.

User is expected follow the advice displayed in the **Import Log** and edit the XML until import preview has finished without errors.

### Committing Import

Once the import preview has finished without problems, one can proceed with actual configuration import. This is performed by clicking on the **Commit Import** button. After confirmation, **Import Log** will display outcome of the operation.

It is possible that some items will not be initialized properly after they have been imported (e.g. their initialization requires presence of a cluster node that went down in the meantime). User is notified about these problems in **Import Log** with link to the problematic item. One should check such items in appropriate section of the CloverETL Server console and change their settings to fix the issue or remove them.

---

## Chapter 23. Diagnostics

This chapter introduces you into diagnostic tools that help you localize performance and memory issues.

CloverETL Server allows you to create a thread dump or a heap dump. The thread and heap dumps are useful for investigation of performance and memory issues.

In server GUI, go to **Configuration** → **System Info** → **Diagnostics**.

### Heap Dump

A **heap Dump** is a content of the JVM process memory stored in a binary file.

#### Server Core

To download a **Heap Dump** of the Server Core, click the **Download** button under the **Heap Dump** section. The download of **Heap Dump** may take some time.

The **Dump live objects only** checkbox allows you to avoid dumping of objects awaiting garbage collection.

#### Worker

The heap dump of Worker can be created with jcmd command: `jcmd <pidOfWorker> GC.heap_dump <filename>` You should specify the file name with full path to avoid searching for the file as jcmd does not always create it in the working directory.

See [details on jcmd](#).

### Heap Dump Analysis

You can use **jvisualvm** or **jhat** to view and analyze the heap dump.



#### Important

A **heap dump** does not work on **WebSphere**. On WebSphere, you can create heap dumps using the administration console. For the instructions, see [IBM Knowledge Center](#).

### Thread Dump

A **thread dump** is a list of existing JVM threads with their callstacks and held locking objects (if supported). It can be viewed in a text editor.

#### Server Core

To download the thread dump, click the **Download** button under the **Thread Dump** section.

#### Worker

The thread dump of Worker can be created with jcmd command: `jcmd <pidOfWorker> Thread.print`

See [details on jcmd](#).

### Required Permissions

To create a **thread dump** or **heap dump**, the following permissions are required:

- **Configuration** → **Heap Memory Dump**

See also [Heap Memory Dump permission](#) (p. 170).

## Enabling GC Logging

Some memory and performance issues can be investigated with help of garbage collection logs.

### Core

To enable the logging of the Server Core, add `-verbose:gc -XX:+PrintGCTimeStamps -Xloggc:server_core_gc_log.txt` to `JAVA_OPTS` in `$CATALINA_HOME/bin/setenv.sh`. Restart of the Server is required to reflect the configuration change.

### Worker

To enable the garbage collection logging in Worker, add the flags `-verbose:gc -XX:+PrintGCTimeStamps -Xloggc:worker_gc_log.txt` to **Configuration > Setup > Worker**. For the GC logging of Worker, use a different file name than for the Server. The restart of Worker is required.

### More Details

Another useful garbage collector flags are `-XX:+PrintGCDetails -XX:+PrintGCDateStamps -XX:+PrintGCCause -XX:+UseGCLogFileRotation -XX:NumberOfGCLogFiles=10 -XX:GCLogFileSize=5M`

For details on JVM flags, see <http://www.oracle.com/technetwork/java/javase/tech/vmoptions-jsp-140102.html>.

---

# Chapter 24. Troubleshooting Worker

## Worker Logs

When investigating issues with Worker itself or jobs running in Worker, there are several logs with useful information:

Logs found in the **Monitoring** → **Server Logs** section of the Server Console.

- **COMMON** log - the main Server log contains also information related to Worker. This log contains the full command line used to start Worker, this allows you to check the command line arguments. Additionally, standard output of the Worker process is redirected to this log - this is useful especially if the Worker process crashes during startup.

The COMMON log file is located in `${java.io.tmpdir}/cloverlogs/all.log`

- **WORKER** log - the main Worker log provides information about Worker startup, initialization, executed jobs, runtime activities, etc. The initialization details contain information about Worker's JNDI resources, etc.

The WORKER log file is located in `${java.io.tmpdir}/cloverlogs/worker_[nodeID].log`

You can also open this log via the **Go to logs** action in the **Worker** section of the **Monitoring** page.

## Worker Command Line

The full command line that was used to start the Worker process can be found in:

- Monitoring section, use the **Show command line** action on Worker. For more details, see [Showing Worker's Command Line Arguments](#) (p. 140).
- the COMMON log of the Server (found in the **Monitoring** → **Server Logs** page). See section above for more details.

Investigate the command line options in case Worker does not correctly start or if the configuration of the running Worker is not correct.

## Worker Does Not Start

If Worker does not start, check the following:

- Server's COMMON log and the WORKER log, see above (p. 196). Look for errors during Worker startup and initialization.
- Worker's command line arguments, see above (p. 196) Look for invalid command line arguments. Additionally, check the custom JVM arguments set on Worker, in the Worker (p. 78) tab of Setup (p. 74) or via the `worker.jvmOptions` (p. 115) configuration property.

## Additional Diagnostic Tools

Below are additional useful diagnostic tools:

- Generate heap dump on Out of Memory errors - add `-XX:+HeapDumpOnOutOfMemoryError` to Worker's JVM arguments (p. 115). A dump file `java_pid.hprof` will be generated when an Out of Memory error occurs. The heap dump will be located in the working directory of the Worker's process (same as the Server Core's working directory). You can override the directory location with the `-XX:HeapDumpPath=/disk2/dumps` option. **Important:** the generated file can be large, its size is equal to the heap size.

The generated heap dump file can be investigated with tools like `jvisualvm` or `jhat`.

- To investigate usage of direct memory, add `-XX:NativeMemoryTracking=summary` to Worker's JVM arguments (p. 115). The details on native memory usage can be displayed with `jcmm <pid> VM.native_memory summary`. For more information, see [Native Memory Tracking tool](#).
- Enable detailed garbage collector logging. Add the following options to the Worker's JVM arguments (p. 115):  
`-Xloggc:/home/clover/worker-gc.log -XX:+PrintGCDetails -XX:+PrintGCDateStamps -XX:+PrintTenuringDistribution -XX:+PrintGCCause -XX:+UseGCLogFileRotation -XX:NumberOfGCLogFiles=10 -XX:GCLogFileSize=5M`

In the above recommended options, the `-Xloggc` options sets the path for the detailed garbage collector log. You can analyse the log file with various tools, eg. <http://gceasy.io/>.

- Enable remote JMX monitoring. Add the following options to the Worker's JVM arguments (p. 115):  
`-Dcom.sun.management.jmxremote=true -`  
`Dcom.sun.management.jmxremote.port=8687 -`  
`Dcom.sun.management.jmxremote.authenticate=false -`  
`Dcom.sun.management.jmxremote.ssl=false -`  
`Djava.rmi.server.hostname=example.com`

With the above options, you enable remote connection to JMX monitoring, which provides a wide range of information about the running JVM, JNDI resources etc. Change the value of `java.rmi.server.hostname` to the hostname of your Server.

## Restarting Worker

In case Worker gets into an unrecoverable state (e.g. out of heap memory, etc.) and you fix the source issue, you can restart it from the Monitoring section (For more details, see [Restarting the Worker](#) (p. 140).):

- restart immediately, which will abort jobs currently running in Worker;
- restart after running jobs finish, in case the currently running jobs are crucial.

## Issues with Classloading

To debug issues with classloading, add `-verbose:class` to **JVM arguments** of Worker. Loaded and unloaded classes will be printed to the output. The output can be seen in the common log.

---

# Part V. Using Graphs

---

## Chapter 25. Graph/Jobflow Parameters

The CloverETL Server passes a set of parameters to each graph or jobflow execution.

Keep in mind that `${paramName}` placeholders (parameters) are resolved only during the initialization (loading of XML definition file), so if you need the parameters to be resolved for each job execution, you cannot set the job to be pooled. However, current parameter values are always accessible by an inline Java code:

```
String runId = getGraph().>getGraphProperties().getProperty("RUN_ID");
```

Properties may be added or replaced:

```
getGraph().getGraphProperties().setProperty("new_property", value);
```

This is a set of parameters which are always set by CloverETL Server (for more information, see [Job Config Properties](#) (p. 183)):

*Table 25.1. Defaults for graph execution configuration*

key	description
SANDBOX_CODE	An identifier of a sandbox which contains the executed graph.
JOB_FILE	A path to the file (graph, subgraph, jobflow). The path relative to the sandbox root path.
SANDBOX_ROOT	An absolute path to sandbox root.
RUN_ID	An ID of the graph execution. In a standalone mode or cluster mode, it is always unique. It may be lower than 0 value, if the run record isn't persistent. For more information, see Chapter 37, <a href="#">Launch Services</a> (p. 282).
PARENT_RUN_ID	A run ID of the graph execution which is a parent to the current one. Useful when the execution is a subgraph, child-job of some jobflow or worker for distributed transformation in a cluster. When the execution doesn't have a parent, the PARENT_RUN_ID is the same as RUN_ID.
ROOT_RUN_ID	A run ID of the graph execution which is a root execution to the current one (the one which doesn't have a parent). Useful when the execution is a subgraph, child-job of some jobflow or worker for distributed transformation in a cluster. When the execution doesn't have a parent, the ROOT_RUN_ID is the same as RUN_ID.
CLOVER_USERNAME	The username of the user who launched the graph or jobflow.
NODE_ID	An ID of the node running the graph or jobflow.

## Parameters by Execution Type

Additional parameters are passed to the graph depending on how the graph is executed:

[Executed from Web GUI](#) (p. 200)

[Executed by Launch Service Invocation](#) (p. 200)

[Executed by HTTP API Run Graph Operation Invocation](#) (p. 200)

[Executed by RunGraph Component](#) (p. 200)

[Executed by WS API Method executeGraph Invocation](#) (p. 200)

[Executed by Task Graph Execution by Scheduler](#) (p. 200)

[Executed from JMS Listener](#) (p. 200)

[Executed by Task Start a graph by Graph/Jobflow Event Listener](#) (p. 201)

[Executed by Task Graph Execution by File Event Listener](#) (p. 201)

### Executed from Web GUI

Graphs executed from the web GUI have no additional parameters.

### Executed by Launch Service Invocation

Service parameters which have the **Pass to graph** attribute enabled are passed to the graph not only as "dictionary" input data, but also as a graph parameter.

### Executed by HTTP API Run Graph Operation Invocation

Any URL parameter with a `param_` prefix is passed to an executed graph but without the `param_` prefix, i.e. `param_FILE_NAME` specified in a URL is passed to the graph as a property named `FILE_NAME`.

### Executed by RunGraph Component

Since 3.0, only parameters specified by the **paramsToPass** attribute are passed from the parent graph to the executed graph. However common properties (`RUN_ID`, `PROJECT_DIR`, etc.) are overwritten with new values.

### Executed by WS API Method executeGraph Invocation

Parameters with values may be passed to the graph with a request for execution.

### Executed by Task Graph Execution by Scheduler

Table 25.2. passed parameters

key	description
EVENT_SCHEDULE_EVENT_TYPE	The type of a schedule: SCHEDULE_PERIODIC   SCHEDULE_ONETIME
EVENT_SCHEDULE_LAST_EVENT	Date/time of a previous event
EVENT_SCHEDULE_DESCRIPTION	A schedule description, which is displayed in the web GUI
EVENT_USERNAME	The owner of the event. For schedule it is the user who created the schedule.
EVENT_SCHEDULE_ID	An ID of the schedule which triggered the graph

### Executed from JMS Listener

There are many graph parameters and dictionary entries passed, depending on the type of incoming message. See details in [JMS Message Listeners](#) (p. 246).

### Executed by Task Start a Graph by Graph/Jobflow Event Listener

Since 3.0, only *specified* properties from a "source" job are passed to the executed job, by default. This behavior can be changed by the `graph.pass_event_params_to_graph_in_old_style` Server config property so that *all* parameters from a "source" job are passed to the executed job. This switch is implemented for backwards compatibility. With the default behavior, in the editor of graph event listener, you can specify a list of parameters to pass. For more information, see [Start a Graph](#) (p. 211).

The following parameters with current values are always passed to the target job

Table 25.3. *passed parameters*

key	description
EVENT_RUN_SANDBOX	A sandbox with the graph which is the source of the event
EVENT_JOB_EVENT_TYPE	GRAPH_STARTED   GRAPH_FINISHED   GRAPH_ERROR   GRAPH_ABORTED   GRAPH_TIMEOUT   GRAPH_STATUS_UNKNOWN, analogically JOBFLOW_* for jobflow event listeners.
EVENT_RUN_JOB_FILE	A jobFile of the job which is the source of the event
EVENT_RUN_ID	An ID of the graph execution which is the source of the event.
EVENT_TIMEOUT	A number of milliseconds which specifies an interval of timeout. Useful only for "timeout" graph event.
EVENT_RUN_RESULT	A result (or current status) of the execution which is the source of the event.
EVENT_USERNAME	The owner of the event. For graph events it is the user who created the graph event listener

### Executed by Task Graph Execution by File Event Listener

Table 25.4. *passed parameters*

key	description
EVENT_FILE_PATH	A path to the file which is the source of the event. Does not contain a file name. Does not end with a file separator. Is passed only for the local file event listener.
EVENT_FILE_NAME	A filename of the file which is the source of the event. Is passed only when the "grouping" mode is disabled. Otherwise there are more than one file event.
EVENT_FILE_URLS	Contains string, which may be used "as is" in the <b>file URL</b> attribute of various <b>CloverETL</b> components. It may contain a URL to one or more (if grouping is enabled) files. It may contain local path(s) or remote URL(s) where credentials are replaced by placeholders (due to security reasons).
EVENT_FILE_AUTH_USERNAME	A username/ID to the remote location.
EVENT_FILE_AUTH_USERNAME_URL_ENCODED	The same as EVENT_FILE_AUTH_USERNAME, but the value is also URL encoded, so it may be used in the URL.
EVENT_FILE_AUTH_PASSWORD	Password/key to the remote location. It's encrypted by the master password. It is passed only when the file listener uses user+password authentication.
EVENT_FILE_AUTH_PASSWORD_URL_ENCODED	

key	description
	The same as EVENT_FILE_AUTH_PASSWORD, but the value is also URL encoded, so it may be used in the URL (EVENT_FILE_URLS parameter).
EVENT_FILE_EVENT_TYPE	SIZE   CHANGE_TIME   APPEARANCE   DISAPPEARANCE
EVENT_FILE_PATTERN	A pattern specified in a file event listener
EVENT_FILE_LISTENER_ID	
EVENT_USERNAME	Owner of the event. For file events, it is the user who created the file event listener.

---

## Adding Another Graph Parameters

### Additional "Graph Config Parameters"

It is possible to add so-called additional parameters in the web GUI ( Chapter 21, [Sandboxes - Server Side Job Files](#) (p. 174)) for the selected graph or for all graphs in the selected sandbox. See details in [Job Config Properties](#) (p. 183).

### Task "execute\_graph" Parameters

The **execute graph** task may be triggered by a schedule, graph event listener, or file event listener. The task editor allows you to specify key=value pairs which are passed to the executed graph.

---

## Chapter 26. Tasks

A task is a graph, jobflow, Groovy script, etc. that can be started manually, started on scheduled time, or triggered by some event. A task specifies WHAT to do.

There are several tasks implemented for a schedule and graph event listener as follows:

- [Start a Graph](#) (p. 211)
- [Start a Jobflow](#) (p. 213)
- [Abort job](#) (p. 215)
- [Execute Shell Command](#) (p. 208)
- [Send an Email](#) (p. 204)
- [Execute Groovy Code](#) (p. 221)
- [Archive Records](#) (p. 216)

### Tasks in Cluster Environment

In the Cluster environment, you can specify a node where the task runs. The task can run on **Any node** or on **one of selected nodes**. If there is no node ID specified, the task may be processed on any cluster node; so in most cases, it will be processed on the same node where the event was triggered. If there are some nodeIDs specified, task will be processed on the first node in the list which is connected in cluster and ready.

Tasks are used in

Chapter 28, [Scheduling](#) (p. 224)

Chapter 30, [Listeners](#) (p. 236)

Chapter 27, [Manual Task Execution](#) (p. 223)

## Send an Email

The **send e-mail** task is useful for notifications about a result of graph execution. For example, you can create a listener with this task type to be notified about each failure in the specified sandbox or a failure of the particular graph.

This task is very useful, but for now only as a response for graph events. This feature is very powerful for monitoring. (for description of this task type, see [Graph Event Listeners](#) (p. 238)).

*Note: It seems useless to send emails periodically, but it may send current server status or daily summary. These features will be implemented in further versions.*

Table 26.1. Attributes of "Send e-mail" task

Task type	"Send an email"
To	The recipient's email address. It is possible to specify more addresses separated by a comma. It is also possible to use placeholders. For more information, see <a href="#">Placeholders</a> (p. 205).
Cc	Cc stands for 'carbon copy'. A copy of the email will be delivered to these addresses. It is possible to specify more addresses separated by a comma. It is also possible to use placeholders. For more information, see <a href="#">Placeholders</a> (p. 205).
Bcc	Bcc: stands for 'Blind carbon copy'. It is similar as Cc, but the others recipients aren't aware, that these recipients received a copy of the email.
Reply-to (Sender)	Email address of sender. It must be a valid address according to the SMTP server. It is also possible to use placeholders. For more information, see <a href="#">Placeholders</a> (p. 205).
Subject	An email subject. It is also possible to use placeholders. For more information, see <a href="#">Placeholders</a> (p. 205).
HTML	A body of the email in HTML. The email is created as multipart, so the HTML body should have a precedence. A plain text body is only for email clients which do not display HTML. It is also possible to use placeholders. For more information, see <a href="#">Placeholders</a> (p. 205).
Text	A body of the email in plain text. The email is created as multipart, so the HTML body should have a precedence. A plain text body is only for email clients which do not display HTML. It is also possible to use placeholders. For more information, see <a href="#">Placeholders</a> (p. 205).
Log file as attachment	If this switch is checked, the email will have an attachment with a packed log file of the related graph execution.

The screenshot shows the 'Create event listener' configuration interface. The 'Enabled' checkbox is checked. The 'Name' field is 'EventListener1', 'Owner' is 'clover', and 'Sandbox' is 'default'. The 'Job file' is 'graph/graphDataPolicy.grf' and the 'Choose event type' is 'GRAPH\_FINISHED'. The 'Choose task type' is 'Send an email'. The 'E-mail template' is empty. The 'To' field contains the placeholder '\$user.email'. The 'Cc' and 'Bcc' fields are empty. The 'Reply-to' field is 'clover.server@'. The 'Subject' is 'CloverETL Server notification - Graph run \${run.id} of \${run.graphId} finished'. The 'Text' field contains a list of placeholders: runId, Sandbox, Graph, Result, Started, Finished, Error node, Error message, and Error exception. The 'HTML' field contains an HTML template with placeholders for the same variables. A 'Log file as attachment' checkbox is at the bottom left.

Figure 26.1. Web GUI - send email

Note: Do not forget to configure a connection to an SMTP server (see Part III, “[Configuration](#)” (p. 67)).

## Placeholders

Placeholder may be used in some fields of tasks. They are especially useful for email tasks, where you can generate the content of email according to context variables.

Note: In most cases, you can avoid this by using email templates (See E-mail task for details)

These fields are preprocessed by Apache Velocity templating engine. See the Velocity project URL for syntax description <http://velocity.apache.org/>.

There are several context variables, which you can use in placeholders and even for creating loops and conditions.

- *event*
- *now*
- *user*
- *run*
- *sandbox*

Some of them may be empty depending on the type of the event. For example, if a task is processed because of a graph event, then *run* and *sandbox* variables contain related data, otherwise they are empty.

Table 26.2 Placeholders useful in email templates

Variable name	Contains
now	Current date-time
user	<p>The user, who caused this event. It may be an owner of a schedule, or someone who executed a graph. It contains sub-properties which are accessible using dot notation (i.e. <code>\${user.email}</code>) email:</p> <ul style="list-style-type: none"> <li>• user.email</li> <li>• user.username</li> <li>• user.firstName</li> <li>• user.lastName</li> <li>• user.groups (list of values)</li> </ul>
run	<p>A data structure describing one single graph execution. It contains sub-properties which are accessible using dot notation (i.e. <code>\${run.jobFile}</code>)</p> <ul style="list-style-type: none"> <li>• job.jobFile</li> <li>• job.status</li> <li>• job.startTime</li> <li>• job.stopTime</li> <li>• job.errNode</li> <li>• job.errMessage</li> <li>• job.errException</li> <li>• job.logLocation</li> </ul>
tracking	<p>A data structure describing a status of components in a graph execution. It contains sub-properties which are accessible using the Velocity syntax for loops and conditions.</p> <pre> #if (\${tracking}) &lt;table border="1" cellpadding="2" cellspacing="0"&gt; #foreach (\$phase in \$tracking.trackingPhases) &lt;tr&gt;&lt;td&gt;phase: \${phase.phaseNum}&lt;/td&gt;   &lt;td&gt;\${phase.executionTime} ms&lt;/td&gt;   &lt;td&gt;&lt;/td&gt;&lt;td&gt;&lt;/td&gt;&lt;/tr&gt; #foreach (\$node in \$phase.trackingNodes)   &lt;tr&gt;&lt;td&gt;\${node.nodeName}&lt;/td&gt;     &lt;td&gt;\${node.result}&lt;/td&gt;     &lt;td&gt;&lt;/td&gt;&lt;td&gt;&lt;/td&gt;&lt;/tr&gt; #foreach (\$port in \$node.trackingPorts)   &lt;tr&gt;&lt;td&gt;&lt;/td&gt;&lt;td&gt;&lt;/td&gt;     &lt;td&gt;\${port.type}:\${port.index}&lt;/td&gt;     &lt;td&gt;\${port.totalBytes} B&lt;/td&gt;     &lt;td&gt;\${port.totalRows} rows&lt;/td&gt;&lt;/tr&gt; #end #end #end &lt;/table&gt; #end } </pre>
sandbox	<p>A data structure describing a sandbox containing an executed graph. It contains sub-properties which are accessible using dot notation (i.e. <code>\${sandbox.name}</code>)</p> <ul style="list-style-type: none"> <li>• sandbox.name</li> <li>• sandbox.code</li> <li>• sandbox.rootPath</li> </ul>
schedule	<p>A data structure describing a schedule which triggered this task. It contains sub-properties which are accessible using dot notation (i.e. <code>\${schedule.description}</code>)</p> <ul style="list-style-type: none"> <li>• schedule.description</li> <li>• schedule.startTime</li> <li>• schedule.endTime</li> <li>• schedule.lastEvent</li> <li>• schedule.nextEvent</li> <li>• schedule.fireMisfired</li> </ul>

## Execute Shell Command

**Execute Shell Command** executes a system command or shell script.

This task is used in Chapter 28, [Scheduling](#) (p. 224) Chapter 30, [Listeners](#) (p. 236) and Chapter 27, [Manual Task Execution](#) (p. 223).

*Table 26.3. Attributes of "Execute shell command" task*

Task type	"Execute shell command"
Start on	Node IDs to process the task.  This attribute is accessible only in a cluster environment. If there are nodes specified, the task will be processed on the first node which is online and ready.
Shell script	Command line for execution of external process.
Working directory	A working directory for the process.  If not set, the working directory of the application server process is used.
Timeout	Timeout in milliseconds. After a period of time specified by this number, the external process is terminated and all results are logged.

The screenshot shows the 'New schedule' configuration window. The 'Triggered Task' section is highlighted with a red border. It contains the following fields:

- Task:** A dropdown menu set to 'Execute shell command'.
- Start on:** Radio buttons for 'Any node' (selected) and 'One of selected nodes'.
- Shell script:** A text area containing the command `du -hs . >> size.log`.
- Working directory:** A text field containing `/home/clover`.
- Timeout:** A text field containing `10000` with the unit `ms`.
- Available variables:** A button with a right-pointing arrow.

At the bottom of the window are 'Create' and 'Cancel' buttons.

*Figure 26.2. Web GUI - shell command*

## Execute Shell Command Parameters

Some parameters are available only in particular context: scheduling, event listeners, or manual task execution.

Table 26.4. Parameters of "Execute shell command" task

event	An event that has triggered the task
now	Current date-time
task	The triggered task
user	<p>The object representing the user who executed the graph/jobflow. It contains sub-properties that are accessible using dot notation (i.e. <code>\${user.email}</code>)</p> <ul style="list-style-type: none"> <li>• <code>user.email</code></li> <li>• <code>user.username</code></li> <li>• <code>user.firstName</code></li> <li>• <code>user.lastName</code></li> <li>• <code>user.groups</code> (list of values)</li> </ul>

Table 26.5. Parameters of "Execute shell command" task - available in scheduling

schedule	<p>The object representing the schedule that triggered this task. It contains sub-properties that are accessible using dot notation (i.e. <code>\${schedule.description}</code>)</p> <ul style="list-style-type: none"> <li>• <code>schedule.description</code></li> <li>• <code>schedule.startTime</code></li> <li>• <code>schedule.endTime</code></li> <li>• <code>schedule.lastEvent</code></li> <li>• <code>schedule.nextEvent</code></li> <li>• <code>schedule.fireMisfired</code></li> </ul>
EVENT_USERNAME	The name of the user who caused the event.
EVENT_USER_ID	The numeric ID of the user who caused the event.
EVENT_SCHEDULE_DESCRIPTION	A description of the schedule.
EVENT_SCHEDULE_EVENT_TYPE	The type of the schedule - <code>SCHEDULE_ONETIME</code> or <code>SCHEDULE_PERIODIC</code> .
EVENT_SCHEDULE_ID	The numeric ID of the schedule
EVENT_SCHEDULE_LAST_EVENT	Date-time of the latest schedule triggering (in <code>java.util.Date.toString()</code> format).

Table 26.6. Parameters of "Execute shell command" task - available in listeners

run	<p>The object representing a single graph/jobflow execution. It contains sub-properties that are accessible using dot notation (i.e. <code>\${run.jobFile}</code>).</p> <ul style="list-style-type: none"> <li>• <code>run.jobFile</code></li> <li>• <code>run.status</code></li> <li>• <code>run.startTime</code></li> <li>• <code>run.stopTime</code></li> <li>• <code>run.errNode</code></li> <li>• <code>run.errMessage</code></li> <li>• <code>run.errException</code></li> </ul>
-----	--

sandbox	The object representing a sandbox containing the executed graph/jobflow. It contains sub-properties that are accessible using dot notation (i.e. <code> sandbox.name </code> ) <ul style="list-style-type: none"> <li>• <code> sandbox.name </code></li> <li>• <code> sandbox.code </code></li> <li>• <code> sandbox.rootPath </code></li> </ul>
tracking	An object representing a status of components in a graph execution. It contains sub-properties that are accessible using Velocity syntax for loops and conditions.
EVENT_USERNAME	The name of the user who caused the event.
EVENT_USER_ID	A numeric ID of the user who caused the event.
EVENT_RUN_SANDBOX	A code of the sandbox containing the graph/jobflow.
EVENT_RUN_JOB_FILE	A sandbox-relative path to the graph/jobflow file.
EVENT_RUN_RESULT	The current status of the graph/jobflow execution <ul style="list-style-type: none"> <li>• <code> N_A </code></li> <li>• <code> READY </code></li> <li>• <code> RUNNING </code></li> <li>• <code> WAITING </code></li> <li>• <code> FINISHED_OK </code></li> <li>• <code> ERROR </code></li> <li>• <code> ABORTED </code></li> <li>• <code> TIMEOUT </code></li> <li>• <code> UNKNOWN </code></li> </ul>
EVENT_RUN_ID	A numeric ID of the run record representing graph/jobflow execution
EVENT_TIMEOUT	A specified timeout (in milliseconds) for the <code> TIMEOUT </code> event to occur.
EVENT_JOB_EVENT_TYPE	Graph event that triggered the task <ul style="list-style-type: none"> <li>• <code> GRAPH_STARTED </code></li> <li>• <code> GRAPH_PHASE_FINISHED </code></li> <li>• <code> GRAPH_FINISHED </code></li> <li>• <code> GRAPH_ERROR </code></li> <li>• <code> GRAPH_ABORTED </code></li> <li>• <code> GRAPH_TIMEOUT </code></li> <li>• <code> GRAPH_STATUS_UNKNOWN </code></li> </ul>

Table 26.7. Parameters of "Execute shell command" task - available in manual task execution

parameters	Task parameters - container for String-String key-value pairs passed to this task.
------------	--

## Start a Graph

**Start a Graph** starts a specified graph from a specified sandbox.

Table 26.8. Attributes of "Graph execution" task

Task type	"Start a graph"
Start on	Node(s) to process the task.
Sandbox	This select box contains sandboxes which are readable for the logged user. Select the sandbox which contains the graph to execute.
Graph	The graph to be executed.  This select box is filled with all graphs files accessible in the selected sandbox. Type a graph name or path to filter available items.
Save run record	Saves run record to database.  If the task runs too often (once in several seconds), you can increase the database performance by disabling this attribute.
Parameters	A list of parameters passed to the graph.  Event parameters like <code>EVENT_RUN_RESULT</code> , <code>EVENT_RUN_ID</code> , etc. are passed to the executed job without limitations. The <code>EVENT_RUN_RESULT</code> and <code>EVENT_RUN_ID</code> parameters are used in context of event listeners. They are not used in context of scheduling.

The screenshot shows a 'New schedule' dialog box. The 'Triggered Task' section is highlighted with a red border. It contains the following fields:

- Task:** Start a graph (dropdown)
- Start on:** Any node (radio selected), One of selected nodes (radio)
- Sandbox:** DataQualityExamples (dropdown)
- Graph:** graph/ValidateData.grf (dropdown)
- Save run record:**  Save run record ⓘ
- Parameters:** MAX\_ADDRESS\_LENGTH 25 (with a red minus icon and a green plus icon)
- Below the parameters: Enter parameter name (dropdown) Value (input field) with a green plus icon and a trash icon.

At the bottom of the dialog are 'Create' and 'Cancel' buttons.

Figure 26.3. Web GUI - Graph execution task

Please note that the behaviour of this task type is almost the same as [Start a Jobflow](#) (p. 213).

## Parameters

You can start a graph with parameters.

To start a graph with a parameter, choose an existing parameter from the list, set its value, and click the **plus sign** button at the end of line.

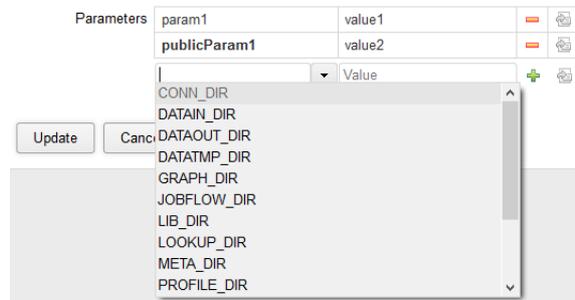


Figure 26.4. Web GUI - Graph execution task

If the graph is started by an event listener, it receives additional parameters from the triggering job.

### Parameters passed to graph by Event Listeners

Table 26.9. Additional parameters available in Event Listeners

EVENT_USERNAME	The name of the user who caused the event
EVENT_USER_ID	A numeric ID of the user who caused the event.
EVENT_RUN_SANDBOX	A code of the sandbox containing the graph/jobflow
EVENT_RUN_JOB_FILE	A sandbox-relative path to the graph/jobflow file.
EVENT_RUN_RESULT	The current status of the graph/jobflow execution (N_A, READY, RUNNING, WAITING, FINISHED_OK, ERROR, ABORTED, TIMEOUT or UNKNOWN).
EVENT_RUN_ID	A numeric ID of the run record representing graph/jobflow execution.
EVENT_JOB_EVENT_TYPE	A graph/jobflow event type that triggered the task.
EVENT_TIMEOUT	A specified timeout (in milliseconds) for the TIMEOUT event to occur

## Start a Jobflow

**Start a jobflow** starts a specified jobflow from a specified sandbox.

Table 26.10. Attributes of "Jobflow execution" task

Task type	"Start a jobflow"
Start on	Node(s) to process the task  This attribute is accessible only in the cluster environment. If there are nodes specified, the task will be processed on the first node which is online and ready.
Sandbox	This select box contains sandboxes which are readable for the logged user. Select sandbox which contains jobflow to execute.
Jobflow	This select box is filled with all jobflow files accessible in the selected sandbox. Type the jobflow name or path to filter available items.
Save run record	Saves run record to database. If the task runs too often (once in several seconds), you can increase the database performance by disabling this attribute.
Parameters	Key-value pairs which are passed to the executed job as parameters.  Event parameters like <code>EVENT_RUN_RESULT</code> , <code>EVENT_RUN_ID</code> , etc. are passed to the executed job without limitations. The <code>EVENT_RUN_RESULT</code> and <code>EVENT_RUN_ID</code> parameters are used in context of event listeners. They are not used in context of scheduling.

The screenshot shows a web form titled "New schedule" for configuring a jobflow execution task. The form is divided into several sections:

- General Settings:**
  - Enabled
  - Description: Jobflow Schedule Example
  - Owner: clover
  - Type:  onetime  periodic
  - Start date/time: [empty field] Europe/Prague
  - Fire misfired event as soon as possible
- Triggered Task:**
  - Task: Start a jobflow
  - Start on:  Any node  One of selected nodes
  - Sandbox: JobflowExamples
  - Jobflow: jobflow/01-Automation-BasicFileProcessing.jbf
  - Save run record
  - Parameters: Enter parameter name | Value

Buttons for "Create" and "Cancel" are located at the bottom left of the form.

Figure 26.5. Web GUI - Jobflow execution task

Please note that the behaviour of this task type is almost the same as [Start a Graph](#) (p. 211).

If the jobflow start is triggered by an event, the same set of parameters as in a graph event listener is passed to the jobflow. [Parameters](#) (p. 212).

---

## Start a Profiler Job

**Start a profiler job** starts a specified profiler job from a specified directory.

You can pass parameters to the profiler job in the same way as in case of starting a graph (p. 211) or jobflow (p. 213).

In case of triggering the profiler job by an event listener, the same set of additional parameters, as in case of execution of graph, is passed to the profiler job. See [Parameters passed to graph by Event Listeners](#) (p. 212).

## Abort job

This task kills/aborts a specified job (graph or jobflow), if it is currently running.

Table 26.11. Attributes of "Abort job" task

Task type	"Abort job"
Start on	Node(s) to process the task.  This attribute is accessible only in a cluster environment. If there are nodes specified, the task will be processed on the first node which is online and ready.
Kill source of event	If this switch is on, the task will kill the job which is the source of the event, which activated this task. Attributes <b>sandbox</b> and <b>job</b> are ignored.  This checkbox is useful only if <b>Abort job</b> is activated by some event.
Sandbox	Select a sandbox which contains the job to kill.  This attribute works only when the <b>Kill source of event</b> switch is off.
Job	This select box is filled with all jobs accessible in the selected sandbox. All instances of the selected job that are currently running and will be killed. This attribute works only when <b>Kill source of event</b> switch is off.

The screenshot shows the 'Create event listener' configuration window. The 'Choose task type' dropdown menu is highlighted with a red rectangular box and is currently set to 'Abort job'. Other visible settings include: 'Enabled' (checked), 'Sandbox' (default), 'Job file' (graph/graphCloverData.gif), 'Choose event type' (GRAPH\_TIMEOUT), 'Job timeout interval' (7200 seconds, 120.0 minutes, 2.0 hours), 'Kill source of event' (checked), and 'Sandbox' (empty dropdown). The 'Job file' field below the highlighted area shows 'no jobs available - choose another sandbox'. 'Create' and 'Cancel' buttons are at the bottom.

Figure 26.6. Web GUI - "Abort job"

## Archive Records

This task can archive (or delete) obsolete records from the database.

Table 26.12. Attributes of "Archivator" task

Task type	"Archivator"
Start on	This attribute specifies a cluster node on which the task may process.  This attribute is accessible only in a cluster environment.  If it is empty, it may be any node; if there are nodes specified, the task will be processed on the first node which is online and ready.
Archivator type	There are two possible values: <code>archive</code> or <code>delete</code> .  <code>Delete</code> removes records without any possibility of recovery.  <code>Archive</code> option removes records from the database, but creates a ZIP package with CSV files containing the deleted data.
Older than	Time period (in minutes) - specifies which records are evaluated as obsolete. Records older than the specified interval are stored in archives.
Output path for archives	This attribute is useful only for the <code>archive</code> option.
Include executions history	
Include temp files	If checked, the archivator removes all graph temporary files older than the given timestamp defined in <b>Older than</b> attribute. The temporary files are files with graph debug data, dictionary files and files created by graph components.
Include tasks history	If checked, the archivator will include run records. Log files of graph runs are included as well.
Include profiler runs	If checked, the archivator will include profiler job results.
Include server instance history	

New schedule

Enabled

Description

Owner

Type  onetime  periodic

Periodicity  by interval  by timetable

Not active before date/time

Not active after date/time

Interval

Fire misfired event as soon as possible

Choose task type

Node IDs to process the task (empty for any node)

Older than (minutes)

Archivator type

Output path for archives

Include executions history

Run records with status

Sandbox

Job file Select sandbox first

Include temp files

Temp files with record status

Include tasks history

Task type

Task result mask

Include profiler runs

Include server instance run history

Figure 26.7. Web GUI - archive records

## Send a JMS Message

This type of task is useful for notifications about result of a graph execution. For example, you can create a graph event listener with this task type to be notified about each failure in a specific sandbox or failure of a particular graph.

JMS messaging requires JMS API (jms.jar) and third-party libraries. All these libraries must be available on the application server classpath. Some application servers contain these libraries by default, some do not, thus the libraries must be added explicitly.

Table 26.13. Attributes of JMS message task

Task type	"JMS message"
Initial context	Choose between the <b>default</b> and <b>custom</b> initial context.
Initial context class name	A full class name of <code>javax.naming.InitialContext</code> implementation. Each JMS provider has its own implementation. For example, in case of Apache MQ, it is <code>org.apache.activemq.jndi.ActiveMQInitialContextFactory</code> . If it is empty, the Server uses the default initial context.
Broker URL	
Connection factory JNDI name	The JNDI name of a connection factory. It depends on a JMS provider.
Destination JNDI name	The JNDI name of a message queue/topic on the server
Username	A username for connection to a JMS message broker
Password	A password for connection to a JMS message broker
URL	A URL of a JMS message broker
JMS pattern	This select box is available only when the user is creating a new record. It contains all predefined JMS message patterns. If the user chooses any of them, the text field below is automatically filled with a value from the pattern.
Text	The body of a JMS message. It is also possible to use placeholders. See <a href="#">Placeholders</a> (p. 205) of <i>send email task for details</i> .

Figure 26.8. Web GUI - Task JMS message editor

Table 26.14. Parameters of "Send a JMS Message"

event	The event that triggered the task.
now	Current date-time
task	The triggered task.
user	The object representing the owner of the schedule. It contains sub-properties that are accessible using dot notation (i.e. \${user.email}) email, username, firstName, lastName, groups (list of values).
schedule	The object representing the schedule that triggered this task. It contains sub-properties that are accessible using dot notation (i.e. \${schedule.description}) description, startTime, endTime, lastEvent, nextEvent, fireMisfired.
EVENT_USERNAME	The username of the user who caused the event
EVENT_USER_ID	A numeric ID of the user who caused the event.
EVENT_SCHEDULE_DESCRIPTION	A description of the schedule

---

EVENT_SCHEDULE	EVENT_TYPE The type of the schedule - SCHEDULE_ONETIME or SCHEDULE_PERIODIC.
EVENT_SCHEDULE	ID A numeric ID of the schedule.
EVENT_SCHEDULE	LAST_EVENT Date-time of the latest schedule triggering (in java.util.Date.toString() format).

## Execute Groovy Code

This type of task allows to execute a code written in the Groovy script language. The script can be defined in place or using a path to external `.groovy` file. It is possible to use some variables.

The basic attribute of this task is a source code of written in Groovy.

If the source codes are provided from both a file and through the input form, only the code from the input form will be executed.

In a cluster environment, there is also one additional attribute **Node IDs to process the task**. If it is empty, it may be any node; if there are nodes specified, the task will be processed on the first node which is online and ready.

CloverETL Server contains Groovy version 2.0.0

*Table 26.15. List of variables available in Groovy code*

variable	class	description	availability
event	com.cloveretl.server.events.AbstractServerEvent		every time
task	com.cloveretl.server.persistent.Task		every time
now	java.util.Date	current time	every time
parameters	java.util.Properties	Properties of a task	every time
user	com.cloveretl.server.persistent.User	Same as event.getUser()	every time
run	com.cloveretl.server.persistent.RunRecord		When the event is an instance of GraphServerEvent
tracking	com.cloveretl.server.worker.common.persistent.TrackingGraphingGraph()	same as TrackingGraphingGraph()	When the event is an instance of GraphServerEvent
sandbox	com.cloveretl.server.persistent.Sandbox	same as run.getSandbox()	When the event is an instance of GraphServerEvent
schedule	com.cloveretl.server.persistent.Schedule	same as ((ScheduleServerEvent)event).getSchedule()	When the event is an instance of ScheduleServerEvent
servletContext	javax.servlet.ServletContext		every time
cloverConfiguration	com.cloveretl.server.spring.CloverConfiguration	Configuration values for <b>CloverETL Server</b>	every time
serverFacade	com.cloveretl.server.facade.api.ServerFacade	The reference to the facade interface. Useful for calling CloverETL Server core.  WAR file contains JavaDoc of facade API and it is accessible on URL: <a href="http://host:port/clover/javadoc/index.html">http://host:port/clover/javadoc/index.html</a>	every time
sessionToken	String	A valid session token of the user who owns the event. It is useful for authorisation to the facade interface.	every time

Variables `run`, `tracking` and `sandbox` are available only if the event is an instance of `GraphServerEvent` class. A variable `schedule` is only available for `ScheduleServerEvent` as an event variable class.

## Example of use Groovy script

This example shows a script which writes a text file describing the finished graph. It shows use of the 'run' variable.

```
import com.cloveretl.server.persistent.RunRecord;
String dir = "/tmp/";
RunRecord rr = (RunRecord)run;

String fileName = "report"+rr.getId()+"_finished.txt";

FileWriter fw = new FileWriter(new File(dir+fileName));
fw.write("Run ID      :"+rr.getId()+"\n");
fw.write("Graph ID     :"+rr.getGraphId()+"\n");
fw.write("Sandbox       :"+rr.getSandbox().getName()+"\n");
fw.write("\n");
fw.write("Start time    :"+rr.getStartTime()+"\n");
fw.write("Stop time     :"+rr.getStopTime()+"\n");
fw.write("Duration      :"+rr.getDurationString()+"\n");
fw.write("Status       :"+rr.getStatus()+"\n");
fw.close();
```

---

## Chapter 27. Manual Task Execution

Since 3.1

A manual task execution allows you to invoke a task directly with an immediate effect, without defining and triggering an event.

There are a number of task types that are usually associated with a triggering event, such as a file listener or a graph/jobflow listener. You can execute any of these tasks manually.

Additionally, you can specify task parameters to simulate a source event that would normally trigger the task. The following is a figure displaying how a 'file event' could be simulated. The parameters for various event sources are listed in the Chapter 25, [Graph/Jobflow Parameters](#) (p. 199).

Manual Task Execution		
<i>i</i> Manual task execution allows you to invoke a task directly with an immediate effect, without defining a listener.		
Task parameters	EVENT_FILE_PATH	C:\Users\clover\data
	EVENT_FILE_NAME	file_data
	EVENT_FILE_EVENT_TYPE	size
Task parameters are passed to a task and can be used as parameters for example in a graph or in an e-mail.		
<b>Triggered Task</b>		
Task	Start a graph	
Start on	<input checked="" type="radio"/> Any node <input type="radio"/> One of selected nodes	
Sandbox	default	
Graph	graph/graphPartition.grf	
	<input checked="" type="checkbox"/> Save run record	
Parameters	Enter parameter name	Value
<input type="button" value="Execute task"/> <input type="button" value="Close"/>		

Figure 27.1. Web GUI - "Manual task execution" form

### Using Manual Task Execution

In the Server GUI, switch to the **Event Listeners** tab. In the **New Listener** drop-down menu, select the **Manual Task Execution** option.

Choose the task type you would like to use. See documentation on chosen tasks:

[Send an Email](#) (p. 204)

[Execute Shell Command](#) (p. 208)

[Start a Graph](#) (p. 211)

[Start a Jobflow](#) (p. 213)

[Start a Profiler Job](#) (p. 214)

[Abort job](#) (p. 215)

[Archive Records](#) (p. 216)

[Send a JMS Message](#) (p. 218)

[Execute Groovy Code](#) (p. 221)

To access the **Manual Task Execution** form, you need [Manual task execution permission](#) (p. 165).

## Chapter 28. Scheduling

The scheduling module allows you to create a time schedule for operations you need to trigger in a repetitive or timely manner.

Similar to cron from Unix systems, each schedule represents a separate time schedule definition and a task to perform.

In a Cluster, you can explicitly specify which node should execute the scheduled task using the **Node ID** parameter. However, if not set, the node will be selected automatically from all available nodes (but always just one).

The screenshot displays the CloverETL Server Console Scheduling interface. The top navigation bar includes 'Monitoring', 'Executions History', 'Sandboxes', 'Launch Services', 'Data Services', 'Scheduling', 'Tasks History', 'Event Listeners', and 'Configuration'. The 'Scheduling' tab is active. The main area is divided into two sections: a list of existing schedules and a 'New schedule' form.

**Existing Schedules Table:**

Description	Type	Task type	Previous event	Next event	Enabled	Actions
Daily accounting reports	timetable	execute_graph		2018-02-06 01:00:00	ON	Detail Run now Delete
Daily finance reports	timetable	execute_graph	2018-02-05 12:36:00	2018-02-06 01:03:00	ON	Detail Run now Delete
Daily marketing reports	timetable	execute_graph	2018-02-05 12:00:00		OFF	Detail Run now Delete
Daily sales reports	timetable	execute_graph		2018-02-06 02:26:00	ON	Detail Run now Delete
Delete old debug files	timetable	archivator	2018-02-05 12:30:00	2018-02-05 13:10:00	ON	Detail Run now Delete
Delete old execution artifacts	timetable	archivator	2018-02-05 12:00:00	2018-02-05 13:00:00	ON	Detail Run now Delete
Monthly accounting records	timetable	execute_graph			OFF	Detail Run now Delete
Monthly finance reports	timetable	execute_graph		2018-03-02 02:01:00	ON	Detail Run now Delete
Monthly sales reports	timetable	execute_graph		2018-03-02 03:19:12	ON	Detail Run now Delete
Monthly marketing reports	timetable	execute_graph	2018-02-05 12:00:00	2018-03-02 03:02:00	ON	Detail Run now Delete

**New Schedule Form:**

- Enabled
- Description: This year sales report
- Owner: sales
- Type:  onetime  periodic
- Start date/time: Europe/Prague
- Fire misfired event as soon as possible
- Triggered Task: Task: Execute shell command
- Start on:  Any node  One of selected nodes
- Shell script: Type the shell script to be executed
- Working directory: [empty]
- Timeout: 10000 ms
- Available variables: [dropdown]
- Buttons: Create, Cancel

Figure 28.1. Web GUI - section "Scheduling" - create new

Tasks you can schedule are described in Chapter 26, [Tasks](#) (p. 203).

[Send an Email](#) (p. 204)

[Execute Shell Command](#) (p. 208)

[Start a Graph](#) (p. 211)

[Start a Jobflow](#) (p. 213)

[Start a Profiler Job](#) (p. 214)

[Abort job](#) (p. 215)

[Archive Records](#) (p. 216)

[Send a JMS Message](#) (p. 218)

[Execute Groovy Code](#) (p. 221)

## Timetable Setting

This section describes how to specify when schedule should be triggered. Note that exact trigger times are not guaranteed. There may be couple of seconds delay. Schedule itself can be specified in different ways.

- [Onetime Schedule](#) (p. 225)
- [Periodical schedule by Interval](#) (p. 227)
- [Periodical schedule by timetable \(Cron Expression\)](#) (p. 228)

## Onetime Schedule

This schedule is triggered just once.

Table 28.1. Onetime schedule attributes

Type	"onetime"
Start date/time	Date and time, specified with minutes precision.
Fire misfired event as soon as possible	If checked and trigger time is missed because of any reason (e.g. server restart), it will be triggered immediately when it is possible. Otherwise it is ignored and it will be triggered at the next scheduled time.

The screenshot shows a web form titled "New schedule". It has several sections:

- General Settings:**
  - Enabled
  - Description: Onetime Schedule Example
  - Owner: clover
  - Type:  ontime  periodic
  - Start date/time: 2017-02-14 12:00:00 (with a calendar icon and "Europe/Prague" time zone)
  - Fire misfired event as soon as possible
- Triggered Task:**
  - Task: Execute shell command
  - Start on:  Any node  One of selected nodes
  - Shell script: Type the shell script to be executed
  - Working directory: (empty field)
  - Timeout: 10000 ms
  - Available variables: (dropdown menu)

At the bottom are "Create" and "Cancel" buttons. A red rectangular box highlights the "Type", "Start date/time", and "Fire misfired event as soon as possible" fields.

Figure 28.2. Web GUI - onetime schedule form

**New schedule**

Enabled

Description

Owner

Type  onetime  periodic

Start date/time  Europe/Prague

<< < February, 2017 > >> x ssible

	Sun	Mon	Tue	Wed	Thu	Fri	Sat
5	29	30	31	1	2	3	4
6	5	6	7	8	9	10	11
7	12	13	14	15	16	17	18
8	19	20	21	22	23	24	25
9	26	27	28	1	2	3	4
10	5	6	7	8	9	10	11

Triggered Task

Task

Start on

Shell script  executed

Working directory

Timeout  ms

Available variables »

Figure 28.3. Web GUI - schedule form - calendar

## Periodical schedule by Interval

This type of schedule is the simplest periodical type. Trigger times are specified by these attributes:

*Table 28.2. Periodical schedule attributes*

Type	"periodic"
Periodicity	"interval"
Not active before date/time	Date and time, specified with minutes precision.
Not active after date/time	Date and time, specified with minutes precision.
Interval (minutes)	Specifies interval between two trigger times. The next task is triggered even if the previous task is still running.
Fire misfired event as soon as possible	If checked and trigger time is missed because of any reason (e.g. server restart), it will be triggered immediately, when it is possible. Otherwise it is ignored and it will be triggered at the next scheduled time.

The screenshot shows a web form titled "New schedule". The form is divided into several sections:

- Enabled:** A checked checkbox.
- Description:** A text input field containing "Periodical Schedule Example".
- Owner:** A dropdown menu showing "clover".
- Type:** Radio buttons for "onetime" and "periodic", with "periodic" selected.
- Periodicity:** Radio buttons for "by interval" and "by timetable", with "by interval" selected.
- Not active before date/time:** A date/time input field showing "2017-02-14 12:00:00" and a region selector for "Europe/Prague".
- Not active after date/time:** A date/time input field showing "2017-02-21 12:00:00" and a region selector for "Europe/Prague".
- Interval:** A text input field showing "60" and a dropdown menu for "minutes".
- Fire misfired event as soon as possible:** A checked checkbox.
- Triggered Task:**
  - Task:** A dropdown menu showing "Execute shell command".
  - Start on:** Radio buttons for "Any node" and "One of selected nodes", with "Any node" selected.
  - Shell script:** A large text area with the placeholder text "Type the shell script to be executed".
  - Working directory:** An empty text input field.
  - Timeout:** A text input field showing "10000" and "ms".
  - Available variables:** A button with a right-pointing arrow.
- Buttons:** "Create" and "Cancel" buttons at the bottom.

A red rectangular box highlights the "Type", "Periodicity", "Not active before date/time", "Not active after date/time", "Interval", and "Fire misfired event as soon as possible" fields.

*Figure 28.4. Web GUI - periodical schedule form*

## Periodical schedule by timetable (Cron Expression)

Timetable is specified by powerful (though a bit tricky) cron expression.

Table 28.3. Cron periodical schedule attributes

Type	"periodic"
Periodicity	"by timetable"
Not active before date/time	Date and time, specified with minutes precision.
Not active after date/time	Date and time, specified with minutes precision.
Cron expression	Cron is powerful tool, which uses its own format for scheduling. This format is well known among UNIX administrators. i.e. "0 0/2 4-23 * * ?" means "every 2 minutes between 4:00am and 11:59pm".
Fire misfired event as soon as possible	If checked and trigger time is missed because of any reason (e.g. server restart), it will be triggered immediately when it is possible. Otherwise it is ignored and it will be triggered at the next scheduled time.

Figure 28.5. Cron periodical schedule form

The **Edit** button helps you to set up a cron expression even without knowledge of the exact cron expression format.

Seconds Minutes Hours Days Months Days of Week

Schedule period: Every 10 minutes

Apply Cancel

Figure 28.6. Editing the cron expression - minutes

Seconds Minutes Hours Days Months Days of Week

Schedule period: Every selected day of week

Monday	Tuesday	Wednesday
Thursday	Friday	Saturday
Sunday		

Apply Cancel

Figure 28.7. Editing the cron expression - days of week



### Note

Server cron expression for *Days of Week* differs from \*nix cron expression. Days in cron expression in Server start from 1 which corresponds to Sunday. \*nix cron expression uses 0 or 7 for Sunday.

## Allocations of Scheduled Task on Nodes

In a cluster environment, you can set the node on which the scheduled task will be launched.

The task can be scheduled on an arbitrary node or on one or more specified nodes.

### Allocation on Any Arbitrary Node

The screenshot shows a 'New schedule' dialog box with the following fields and options:

- Enabled
- Description: Graph Schedule Example
- Owner: clover
- Type:  onetime  periodic
- Start date/time: 2017-02-14 12:00:00 (Location: Europe/Prague)
- Fire misfired event as soon as possible
- Triggered Task: Start a graph
- Start on:  Any node  One of selected nodes

Figure 28.8. Schedule allocation - Any node

### Allocation on One or More Chosen Nodes

The screenshot shows the 'New schedule' dialog box with the 'Start on' option set to 'One of selected nodes'. The 'Available nodes' list contains 'alpha' and the 'Selected nodes' list contains 'tray'. There are 'Add' and 'Remove' buttons between the lists. At the bottom, there is a field for 'Other nodes (currently unavailable)' with the placeholder text 'List of nodes, separated by commas'.

Figure 28.9. Schedule allocation - One ore more specific nodes

## Scheduling the Tasks - Examples

---

### Start a graph at specific time

---

This example shows scheduling the start of a graph at specific time.

1. In server GUI, in **Scheduling** section, click on the **New schedule** button.
2. Enter the description that allows you to successfully identify the task among other existing tasks, e.g. **My one-time processing example**.
3. Choose the **onetime Task** to start the graph just once.
4. Enter start date and time. Use the calendar  to enter the required date and time in a correct format.
5. In the **Task** field, choose **Start a graph**.
6. Select the **sandbox** and **graph** within the sandbox.

### Start a Jobflow once an hour

---

This example shows scheduling of a periodic task.

Create a new schedule that runs **UserStats.jbf** jobflow from the **reports** sandbox once an hour.

1. Enter a **description**, e.g. **Hourly user stats**.
2. Choose the **periodic type** and **periodicity by interval**.
3. The task will be started once an hour within some time period. Enter the beginning of this period (**not active before**) and end of the period (**not active after**).
4. Enter the **interval** between two jobflow starts.
5. Select the **sandbox** and **graph** within the sandbox.

### Complex Scheduling

---

This example shows a complex scheduling using a cron expression.

Start a graph **WeekendGraph.grf** every Saturday and Sunday at 22:15.

1. Enter the description of the schedule, e.g. **The Weekend Graph**.
2. Select the **periodic Type** and **periodicity by timetable**.
3. Click the **Edit** button on the **Cron expression** line.

On the **Minutes** tab, select **Every selected minute** and select **15**.

On the **Hours** tab, select **Every selected hour** and select **22**.

On the **Days of Week** tab, select **Every selected day of week** and choose **Saturday** and **Sunday**.

4. In the **Task** field, choose **Start a graph**.
5. Select the **sandbox** and **graph**.

# Chapter 29. Viewing Job Runs - Executions History

[Filtering and ordering](#) (p. 232)

[Tracking](#) (p. 234)

[Log File](#) (p. 235)

**Executions History** shows the history of all jobs that the Server has executed – transformation graphs, jobflows, and Data Profiler jobs. You can use it to find out why a job failed, see the parameters that were used for a specific run, and much more.

The table shows basic information about the job: Run ID, Node, Job file, Executed by, Status, and time of execution. After clicking on a row in the list, you can see additional details of the respective job, such as associated log files, parameter values, tracking and more.



## Note

Some jobs might not appear in the Executions History list. These are jobs that have disabled persistency for increased performance (for example, some Launch Services disable storing the run information in order to increase service responsiveness).

## Filtering and ordering

Use the Filter panel to filter the view. By default, only parent tasks are shown (Show executions children) – e.g. master nodes in a cluster and their workers are hidden by default.

Use the up and down arrows in the table header to sort the list. By default, the latest job is listed first.

The screenshot shows the CloverETL Server Console interface. At the top, there's a navigation bar with 'Monitoring' and 'Executions History' selected. The main content area features a 'Filter' panel with various search criteria like Run ID, Sandbox, From date, Status, and Error message. Below the filter is a table of job executions. The table has columns for Run ID, Node, Job, Executed by, Status, and Started. The table contains five rows of data, each with a status icon (play, checkmark, or exclamation mark) and a 'Kill' button. At the bottom of the table, there's a pagination control showing '1 2' and a 'Refresh' button.

Run ID	Node	Job	Executed by	Status	Started
221	Node01	BasicExamples graph/GEOCoding.grf	clover	▶	2017-09-13 09:45:40
220	Node01	BigDataExamples graph/GenerateReport.grf	clover	✓	2017-09-13 09:43:22
219	Node01	BigDataExamples graph/CheckParameters-MapReduce.grf	clover	!	2017-09-13 09:43:07
218	Node01	BigDataExamples graph/BigDataExample_HDFS.grf	clover	!	2017-09-13 09:41:47
214	Node01	JobflowExamples jobflow/01-Automation-BasicFileProcessing.jbf	clover	✓	2017-09-13 09:40:03

Figure 29.1. Executions History - executions table

When some job execution is selected in the table, the detail info is shown on the right side.

Table 29.1. Persistent run record attributes

Attribute	Description
Run ID	A unique number identifying the run of the job. Server APIs usually return this number as a simple response to the execution request. It is useful as a parameter of subsequent calls for specification of the job execution.
Execution type	A type of a job as recognized by the Server. STANDALONE for ETL graph, JOBFLOW for Jobflow, PROFILER_JOB for profiler, MASTER for the main record of partitioned execution in a cluster, PARTITIONED_WORKER for the worker record of partitioned execution in a cluster
Parent run ID	A run ID of the parent job. Typically the jobflow which executed this job, or master execution which encapsulates this worker execution.
Root run ID	A run ID of the root parent job. Job execution which wasn't executed by another parent job.
Execution group	Jobflow components may group sub-jobs using this attribute. See the description of Jobflow components for details.
Nested jobs	Indication that this job execution has or has not any child execution.
Node	In cluster mode, it shows the ID of the cluster node which this execution was running on.
Executor	If it runs on worker, it contains the text "worker".
Executed by	The user who executed the job. Either directly using some API/GUI or indirectly using the scheduling or event listeners.
Sandbox	The sandbox containing a job file. For jobs which are sent together with an execution request, so the job file doesn't exist on the Server site, it is set to the "default" sandbox.
Job file	A path to a job file, relative to the sandbox root. For jobs which are sent together with an execution request, so the job file doesn't exist on the Server site, it is set to generated string.
Job version	The revision of the job file. A string generated by <b>CloverETL Designer</b> and stored in the job file.
Status	Status of the job execution. READY - waiting for execution start, RUNNING - processing the job, FINISHED OK - the job finished without any error, ABORTED - the job was aborted directly using some API/GUI or by the parent jobflow, ERROR - the job failed, N/A (not available) - the server process died suddenly, so it couldn't properly abort the jobs. After restart, the jobs with unknown status are set as N/A
Started	Server date-time (and timezone) of the execution start.
Finished	Server date-time (and timezone) of the execution finish.
Duration	Execution duration
Error in component ID	If the job failed due the error in a component, this field contains the ID of the component.
Error in component type	If the job failed due the error in a component, this field contains type of the component.
Error message	If the job failed, this field contains the error description.
Exception	If the job failed, this field contains error stack trace.
Input parameters	A list of input parameters passed to the job. A job file can't be cached, since the parameters are applied during loading from the job file. The job file isn't cached, by default.
Input dictionary	A list of dictionary elements passed to the job. A dictionary is used independently of job file caching.
Output dictionary	A list of dictionary elements at the moment the job ends.

For jobs which have some children executions, e.g. partitioned or jobflows also an executions hierarchy tree is shown.

## Chapter 29. Viewing Job Runs - Executions History

The screenshot shows the CloverETL Server Console interface. The top navigation bar includes 'Monitoring', 'Executions History', 'Sandboxes', 'Launch Services', 'Data Services', 'Scheduling', 'Tasks History', 'Event Listeners', and 'Configuration'. The right side shows 'Node01 / CloverCluster' and 'CloverETL Server 4.7.0.14M2'.

The main area is divided into two panels. The left panel is a table of job executions:

Job ID	Node	Job	Executed by	Status	Started
5	Node01	BasicExamples graph/ExecutingExternal.grf	clover	✓	2017-09-13 09:58:02
4	Node01	BasicExamples graph/EmailValidation.grf	clover	✓	2017-09-13 09:54:51
1	Node01	BasicExamples graph/GEOCoding.grf	clover	✓	2017-09-13 09:45:40
0	Node01	BigDataExamples graph/GenerateReport.grf	clover	!	2017-09-13 09:43:22
9	Node01	BigDataExamples graph/CheckParameters-MapReduce.grf	clover	!	2017-09-13 09:43:07

The right panel shows the 'Run Record detail' for 'graph/ExecutingExternal.grf (Run ID: 225)'. It includes fields for Run ID (225), Job type (ETL\_GRAPH), Parent run ID, Root run ID, Execution group, Execution label, Nested jobs, Node (Node01), Job file (graph/ExecutingExternal.grf), Job version, Status (FINISHED\_OK), Started (2017-09-13 09:58:02), Finished (2017-09-13 09:58:02), Duration (698 ms), Failed component ID, Failed component type, Error message, and Exception.

Figure 29.2. Executions History - overall perspective



### Tip

Since the detail panel, and especially job logs, may be wide, it may be useful to hide a table on the left, so the detail panel spreads. Click on the minimize icon on the top of the list panel to hide the panel. Then to show list panel again, click to the "Expand panel" icon on the left.

The screenshot shows the 'Executions hierarchy' for 'Jobflow01-Automation-BasicFileProcessing.jbl (Run ID: 214)'. The hierarchy is as follows:

- 214 Node01 Jobflow01-Automation-BasicFileProcessing.jbl (Children)
  - 215 Node01 Jobflow01-Automation-BasicFileProcessing.jbl (Process Each File)
    - 216 Node01 Jobflow01-Automation-BasicFileProcessing.jbl (Process Each File)
      - 217 Node01 Jobflow01-Automation-BasicFileProcessing.jbl (Process Each File)

The 'Log file' tab is selected, showing a detailed log of the job execution, including timestamps, node IDs, and process IDs. The log shows the job starting at 2017-09-13 09:40:09 and finishing at 2017-09-13 09:40:10. The log includes information about the job configuration, the job parameters, and the job status.

Figure 29.3. Executions Hierarchy with docked list of jobs

Executions hierarchy may be rather complex, so it's possible to filter the content of the tree by a fulltext filter. However when the filter is used, the selected executions aren't hierarchically structured.

## Tracking

The **Tracking** tab, contains details about the selected job:

Table 29.2. Tracking table information

Attribute	Description
Component ID	The ID of the component.
Component name	The name of the component.
Status	Status of data processing in the respective component. <ul style="list-style-type: none"> <li>FINISHED_OK: data was successfully processed</li> <li>ABORTED: data processing has been aborted</li> <li>N_A: status unknown</li> <li>ERROR: an error occurred while data was processed by the component</li> </ul>
CPU	CPU usage of the component.
Port	Component's ports (both input and output) that were used for data transfer.
Records	The number of records transferred through the port of the component.
kB	Amount of data transferred in kB.
Records/s	The number of records processed per second
KB/s	Data transfer speed in KB.
Records/s peak	The peak value of <b>Records/s</b> .
KB/s peak	The peak value of <b>KB/s</b> .

Component ID	Component name	Status	CPU	Port	Records	kB	Records/s	KB/s	Records/s peak	KB/s peak
Phase 0 ( 2 secs), Memory utilization 199MB										
LOAD_TRANSACTIONS	Load Transactions	FINISHED_OK	0.00%							
				OUTPUT:0	40000	781.25	13404	261.81	0	0.00
SORT_TRANSACTIONS	Sort Transactions	FINISHED_OK	0.00%							
				INPUT:0	40000	781.25	13404	261.81	0	0.00
				OUTPUT:0	40000	781.25	13404	261.81	0	0.00
LOAD_ONLINE_EXCHANGE_RATES1	Load Online Exchange Rates	FINISHED_OK	0.00%							
				OUTPUT:0	32	0.47	10	0.16	0	0.00
FORMAT_FIELDS	Format Fields	FINISHED_OK	0.00%							

Figure 29.4. Executions History - Tracking

## Log File

In the **Log file** tab, you can see the log of the job run with detailed information. A log with a green background indicates a successfully run job, while a red background indicates an error.

You can download the log as a plain text file by clicking the **Download log** button or as a zip archive by clicking the **Download log (zipped)** button.

```

2018-03-26 12:59:12,010 INFO 811943 [qtp1357563986-22] *** CloverETL, (c) 2002-2018 Javlin a.s. ***
2018-03-26 12:59:12,010 INFO 811943 [qtp1357563986-22] Running with CloverETL version 4.9.0.H2 build#003 compiled 23/03/2018 22:33:08
2018-03-26 12:59:12,011 INFO 811943 [qtp1357563986-22] Running on 2 CPU(s), OS Linux, architecture amd64
2018-03-26 12:59:12,011 INFO 811943 [qtp1357563986-22] License No. : Javlin-Internal-License
2018-03-26 12:59:12,167 INFO 811943 [qtp1357563986-22] Server uptime 12,504,565 and fingerprint:
jYEU0NuUjMBV2cuV2YpxULsFmbyVgduLUlUG2Fm5AKNuafGyZyJMay
2018-03-26 12:59:12,168 INFO 811943 [qtp1357563986-22] Executing job: BasicExamples/graph/CreditCardFraudDetection.grf with
permissions of user "clover"
2018-03-26 12:59:12,168 INFO 811943 [qtp1357563986-22] Triggered via Web GUI
2018-03-26 12:59:12,270 INFO 811943 [JobStarter_811943] Job is executed on worker: [worker@virt-alpha:10504]
2018-03-26 12:59:12,317 INFO 811943 [threadPoolFactoryBean-3] Running on Java(TM) SE Runtime Environment, 1.8.0_45, Oracle
Corporation (20k), max available memory for JVM is 473,024 KB
2018-03-26 12:59:12,471 INFO 811943 [threadPoolFactoryBean-3] Graph Initialization (CreditCards)
2018-03-26 12:59:12,471 INFO 811943 [threadPoolFactoryBean-3] Initializing phase 0
2018-03-26 12:59:12,708 INFO 811943 [threadPoolFactoryBean-3] Phase 0 initialized successfully.
2018-03-26 12:59:12,711 INFO 811943 [threadPoolFactoryBean-3] Executing job #811943 -
sandbox://BasicExamples/graph/CreditCardFraudDetection.grf
2018-03-26 12:59:12,713 INFO 811943 [MatchDog_811943] Job execution type: ETL_GRAPH
2018-03-26 12:59:12,714 INFO 811943 [MatchDog_811943] Job parameters:
CLOVER_USERNAME=clover

```

Figure 29.5. Executions History - Tracking

# Chapter 30. Listeners

Listeners can be seen as 'hooks'. They wait for a specific event and take a user-defined action if the event occurs.

Created listeners are shown in the **Event Listener** list (see Figure below). The list has the following columns:

Table 30.1. Event Listeners Table Description

Column name	Description
Listener type	Shows the graphic representation of each listener type.
Name	Show the name of the listener.
Event source description	Describes the event the listener is waiting for.
Task type	Shows the graphic representation of each task type.
Task description	Describes the task that will be activated if the set event occurs.
Monitoring details	Shows detailed monitoring information. By default, the column is collapsed (☰). Clicking the column header expands it, revealing details on the date of the last run and the number of successes and failures in the last 24 hours.
Enabled	Indicates whether the listener is enabled (✔) or disabled (✘). Clicking the icon enables/disables the listener.
Actions	By clicking the ⋮ button, you can see details of the selected listener or delete it.

In case of listeners' events and tasks working with graphs, jobflows and profiler jobs, the 📊 button appears in their respective description fields. This button allows you to see the graph in the Sandboxes section of the Server Console GUI.

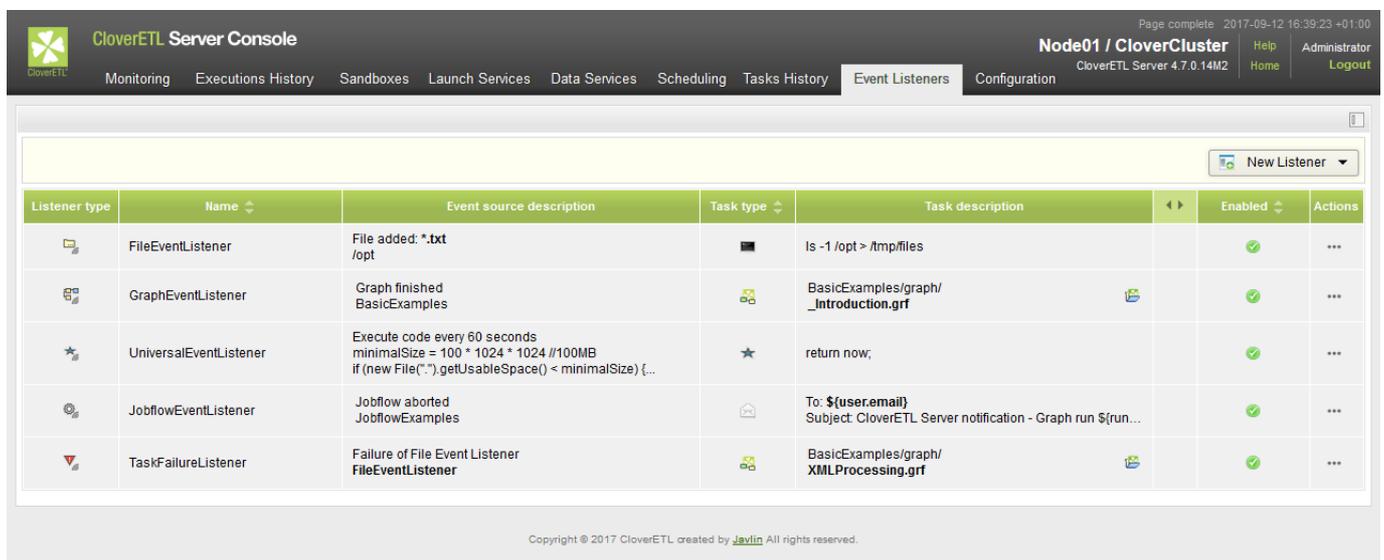


Figure 30.1. Listeners

The event is specific to the particular listener

- 📊 [Graph Event Listeners](#) (p. 238)
- 🔄 [Jobflow Event Listeners](#) (p. 244)
- 📧 [JMS Message Listeners](#) (p. 246)
- ★ [Universal Event Listeners](#) (p. 251)
- 📁 [File Event Listeners \(remote and local\)](#) (p. 253)
- 🚫 [Task Failure Listeners](#) (p. 261)

The available actions taken by the listeners are common for all listeners. The actions, that can be taken are:

-  Send an email (p. 204)
-  Execute a shell command (p. 208)
-  Start a graph (p. 211)
-  Start a jobflow (p. 213)
-  Start a profiler job (p. 214)
-  Abort a job (p. 215)
-  Archivator (p. 216)
-  Send a JMS Message (p. 218)
-  Execute a Groovy code (p. 221)

## Graph Event Listeners

Graph Event Listeners allow you to define a task that the Server will execute as a reaction to the success, failure or other event of a specific job (a transformation graph).

Each listener is bound to a specific graph and is evaluated no matter whether the graph was executed manually, scheduled, or via an API call, etc.

You can use listeners to chain multiple jobs (creating a success listener that starts the next job in a row). However, we recommend using Jobflows to automate complex processes because of its better development and monitoring capabilities.

Graph Event Listeners are similar to Jobflow Event Listeners ([Jobflow Event Listeners](#) (p. 244)) – for CloverETL Server both are simply 'jobs'.

In the Cluster, the event and the associated task are executed on the same node the job was executed on, by default. If the graph is distributed, the task will be executed on the master worker node. However, you can override where the task will be executed by explicitly specifying a Node IDs in the task definition.

## Graph Events

---

Each event carries properties of a graph, which is the source of the event. If there is an event listener specified, the task may use these properties. For example the next graphs in a chain may use "EVENT\_FILE\_NAME" placeholder which was activated by the first graph in the chain. Graph properties, which are set specifically for each graph run (e.g. RUN\_ID), are overridden by the last graph.

### Types of graph events

#### Graph started

The **Graph started** event is created, when an ETL graph execution successfully started.

#### Graph phase finished

The **Graph phase finished** event is created, everytime a graph phase is finished and all its nodes are finished with status FINISHED\_OK.

#### Graph finished

The **Graph finished** event is created, when all phases and nodes of a graph are finished with FINISHED\_OK status.

#### Graph error

The **Graph error** event is created, when a graph cannot be executed for some reason, or when any node of graph fails.

#### Graph aborted

The **Graph aborted** event is created, when a graph is explicitly aborted.

#### Graph timeout

The **Graph timeout** event is created, when a graph runs longer than a specified interval. Thus you should specify the **Job timeout interval** for each listener of a graph timeout event. You can specify the interval in seconds, minutes or hours.

Graph to Check

Sandbox default

Graph graph/graphHTTPConnector.grf

Event type Graph timeout

Job timeout interval 1800 seconds  
30.0 minutes  
0.5 hours

Figure 30.2. Web GUI - graph timeout event

## Graph unknown status

The **Graph unknown status** event is created when the Server, during the startup, detects run records with undefined status in the executions history. Undefined status means, that the Server has been killed during the graph run. The Server automatically changes the state of the graph to *Not Available* and sends a *graph unknown status* event.

Please note, that this works just for executions, which have a persistent record in the executions history. It is possible to execute a transformation without a persistent record in the executions history, typically for better performance of fast running transformations (e.g. using Launch Services).

## Listener

User may create a listener for a specific event type and graph (or all graphs in sandbox). The listener is actually a connection between a graph event and a task, where the graph event specifies *when* and the task specifies *what* to do.

Event handling consists of the following course of actions:

- the event is created
- listeners for this event are notified
- each listener performs the related task

## Tasks

Task types are described in Chapter 26, [Tasks](#) (p. 203).

In the Cluster environment, all tasks have an additional attribute **Node IDs** to process the task. If there is no node ID specified, the task may be processed on any cluster node. In most cases, it will be processed on the same node where the event was triggered. If there are some nodeIDs specified, the task will be processed on the first node in the list which is connected in cluster and ready.

[Send an Email](#) (p. 204)

[Execute Shell Command](#) (p. 208)

[Start a Graph](#) (p. 211)

[Start a Jobflow](#) (p. 213)

[Start a Profiler Job](#) (p. 214)

[Abort job](#) (p. 215)

[Archive Records](#) (p. 216)

[Send a JMS Message](#) (p. 218)

[Execute Groovy Code](#) (p. 221)

## Use Cases

Possible use cases are:

- [Execute graphs in chain](#) (p. 240)
- [Email notification about graph failure](#) (p. 241)
- [Email notification about graph success](#) (p. 242)
- [Backup of data processed by graph](#) (p. 243)

## Execute graphs in chain

For example, we have to execute graph B, only if another graph A finished without any error. So there is a relation between these graphs. We can achieve this behavior by creating a graph event listener. We create a listener for graph finished OK event of graph A and choose an `execute_graph` task type with graph B specified for execution. If we create another listener for graph B with the `execute_graph` task with graph C specified, it will work as a chain of graphs.

**New Graph Event Listener**

Graph event listeners allow you to define a task that will be executed as a reaction to the success or failure of executing a specific graph.

Name:

Owner:

Enabled

**Graph to Check**

Sandbox:

Graph:

Event type:

**Triggered Task**

Task:

Start on:  Any node  One of selected nodes

Sandbox:

Graph:

Save run record

Pass parameters from the checked graph

Parameters:

Parameters passed from the listener to the task

*Figure 30.3. The event source graph isn't specified, thus the listener works for all graphs in the specified sandbox*

## Email notification about graph failure

**New Graph Event Listener**

**i** Graph event listeners allow you to define a task that will be executed as a reaction to the success or failure of executing a specific graph.

Name

Owner

Enabled

**Graph to Check**

Sandbox

Graph

Event type

**Triggered Task**

Task

Start on  Any node  One of selected nodes

E-mail template

To

[Show Cc, Bcc and Reply-to](#)

Subject

HTML

[Show plain text](#)

Log file as attachment (if it's available)

Available variables »

Figure 30.4. Web GUI - email notification about graph failure

## Email notification about graph success

 **New Graph Event Listener**

**i** Graph event listeners allow you to define a task that will be executed as a reaction to the success or failure of executing a specific graph.

Name:

Owner:

Enabled

**Graph to Check**

Sandbox:

Graph:

Event type:

**Triggered Task**

Task:

Start on:  Any node  One of selected nodes

E-mail template:

To:

[Show Cc, Bcc and Reply-to](#)

Subject:

HTML: 

```
<h1>Graph run ${run.id} of ${sandbox.code} / ${run.jobFile}
finished</h1>

<p>runId: ${run.id}</p>
<p>User: ${run.user.username}</p>
<p>Result: ${run.status}</p>
<p>Started: ${run.startTime}</p>
<p>Finished: ${run.stopTime}</p>

#if( ${run.errNode} )
<p>Error node: ${run.errNode}</p>
#end
```

[Show plain text](#)

Log file as attachment (if it's available)

Available variables »

Figure 30.5. Web GUI - email notification about graph success

## Backup of data processed by graph

**New Graph Event Listener**

*i* Graph event listeners allow you to define a task that will be executed as a reaction to the success or failure of executing a specific graph.

Name: Backup Processed Data

Owner: clover

Enabled

**Graph to Check**

Sandbox: default

Graph: graph/graphDataPolicy.grf

Event type: Graph finished

**Triggered Task**

Task: Execute shell command

Start on:  Any node  One of selected nodes

Shell script: /opt/scripts/backup\_data.sh

Working directory: /var/data

Timeout: 10000 ms

Available variables »

Create Cancel

Figure 30.6. Web GUI - backup of data processed by graph

## Jobflow Event Listeners

Jobflow Event Listeners allow you to define a task that the Server will execute as a reaction to the success or failure of executing a specific job (a jobflow).

Each listener is bound to a specific jobflow and is evaluated every time the jobflow is executed (no matter whether manually, through another jobflow, via a schedule, API call, etc.).

Jobflow Event Listeners work very similarly to [Graph Event Listeners](#) (p. 238) in many ways, since ETL Graphs and Jobflows are both 'jobs' from the point of view of the CloverETL Server.

In the Cluster, the event and the associated task are executed on the same node the job was executed on. If the jobflow is distributed, the task will be executed on the master worker node. However, you can override the default setting by explicitly specifying a Node ID in the task definition.

## Jobflow Events

---

Each event carries properties of the event source job. If there is an event listener specified, a task may use these properties. For example, the next job in the chain may use "EVENT\_FILE\_NAME" placeholder which activated the first job in the chain. Job properties, which are set specifically for each run (e.g. RUN\_ID), are overridden by the last job.

### Types of jobflow events

#### Jobflow started

A **Jobflow started** event is created, when jobflow execution successfully started.

#### Jobflow phase finished

The **Jobflow phase finished** event is created everytime a jobflow phase is finished and all its nodes are finished with the FINISHED\_OK status.

#### Jobflow finished

The **Jobflow finished** event is created, when all phases and nodes of a jobflow are finished with the FINISHED\_OK status.

#### Jobflow error

The **Jobflow error** event is created, when a jobflow cannot be executed for some reason, or when any node of the jobflow fails.

#### Jobflow aborted

The **Jobflow aborted** event is created, when a jobflow is explicitly aborted.

#### Jobflow timeout

The **Jobflow timeout** event is created when a jobflow runs longer then the specified interval. Thus you have to specify the **Job timeout interval** for each listener of the jobflow timeout event. You can specify this interval in seconds, minutes or hours.

The screenshot shows a web form titled "Jobflow to Check". It contains the following fields:

- Sandbox:** A dropdown menu with "JobflowExamples" selected.
- Jobflow:** A text input field with the placeholder "Type to search or leave empty to match any jobflow".
- Event type:** A dropdown menu with "Jobflow timeout" selected.
- Job timeout interval:** A table with three rows:

1800	seconds
30.0	minutes
0.5	hours

Figure 30.7. Web GUI - jobflow timeout event

## Jobflow unknown status

The **Jobflow unknown status** event is created, when the Server, during the startup, detects run records with undefined status in the executions history. Undefined status means, that the Server has been killed during the jobflow run. The server automatically changes the state of the jobflow to *Not Available* and sends the *jobflow status unknown* event.

Please note, that this works just for executions, which have a persistent record in the executions history. It is possible to execute a transformation without a persistent record in executions history, typically for better performance of fast running transformations (e.g. using Launch Services).

## Listener

The user may create a listener for the specified event type and jobflow (or all jobflows in sandbox). The listener is actually a connection between the jobflow event and task, where the jobflow event specifies *when* and the task specifies *what* to do.

Event handling consist of the following course of actions:

- event is created
- listeners for this event are notified
- each listener performs the related task

## Tasks

A task specifies an operation which should be performed as a reaction to a triggered event.

Task types are described in Chapter 26, [Tasks](#) (p. 203).

*Note:* You can use a task of any type for a jobflow event listener. The description of task types is divided into two sections just to show the most evident use cases.

[Send an Email](#) (p. 204)

[Execute Shell Command](#) (p. 208)

[Start a Graph](#) (p. 211)

[Start a Jobflow](#) (p. 213)

[Start a Profiler Job](#) (p. 214)

[Abort job](#) (p. 215)

[Archive Records](#) (p. 216)

[Send a JMS Message](#) (p. 218)

[Execute Groovy Code](#) (p. 221)

## JMS Message Listeners

JMS Message Listeners allow you to listen for incoming JMS messages. You specify the source of the messages (JMS Topic or JMS Queue) and a task that will be executed for each incoming message.

JMS messaging requires a JMS API (jms.jar) and specific third-party libraries. Every one of these libraries must be available on a classpath of an application server. Some application servers contain these libraries by default; however, some do not. In such a case, libraries must be added explicitly before starting the **CloverETL Server**.



### JMS Message Listeners on Worker

JMS Message Listeners can be used with Worker as well. In such a case, make sure that the required .jar file and libraries are available on the Worker's classpath as well. For more information, see [Adding Libraries to the Worker's Classpath](#) (p. 64).

JMS is a complex topic that goes beyond the scope of this document. For more detailed information about JMS, refer to the Oracle website: <https://docs.oracle.com/javaee/7/tutorial/jms-concepts.htm#BNCDQ>

Note that the JMS implementation is dependent on the application server that the **CloverETL Server** is running in.

In Cluster, you can either explicitly specify which node will listen to JMS or not. If unspecified, all nodes will register as listeners. In the case of JMS Topic, all nodes will get the message and trigger the task (multiple instances) or, in the case of JMS Queue, a random node will consume the message and run the task (just one instance).

Table 30.2. Attributes of JMS message task

Attribute	Description
Initialize by	<p>This attribute is useful only in a cluster environment. It is a node ID where the listener should be initialized. If it is not set, the listener is initialized on all nodes in the cluster.</p> <p>In the Cluster environment, each JMS event listener has a <b>Node IDs</b> attribute which may be used to specify the cluster node which will consume messages from the queue/topic. There are the following possibilities:</p> <ul style="list-style-type: none"> <li>• <b>No failover:</b> Just one node ID specified - Only the specified node may consume messages, however the node status must be "ready". When the node isn't ready, messages aren't consumed by any cluster node.</li> <li>• <b>Failover with node concurrency:</b> No node ID specified (empty input) - All cluster nodes with status "ready" consume messages concurrently.</li> <li>• <b>Failover with node reservation:</b> More node IDs specified (separated by a comma) - Just one of the specified nodes consumes messages at a time. If the node fails for any reason (or its status isn't "ready"), any other "ready" node from the list continues with consuming messages.</li> </ul> <p>In a standalone environment, the <b>Node IDs</b> attribute is ignored.</p>
<b>JNDI Access</b>	
Initial context	Default or custom
Initial context factory class	<p>A full class name of the <code>javax.naming.InitialContext</code> implementation. Each JMS provider has its own implementation. For example, Apache MQ has <code>org.apache.activemq.jndi.ActiveMQInitialContextFactory</code>. If it is empty, the Server uses a default initial context.</p> <p>The specified class must be on the web-app classpath or application-server classpath. It is usually included in one library with a JMS API implementation for each specific JMS broker provider.</p>

Attribute	Description
Broker URL	A URL of a JMS message broker
<b>Listen To</b>	
Connection factory	A JNDI name of a connection factory. It depends on a JMS provider.
Username	A username for a connection to a JMS message broker
Password	A password for a connection to JMS message broker
Queue/Topic	A JNDI name of a message queue/topic on the Server
Durable subscriber	<p>If false, the message consumer is connected to the broker as 'non-durable', so it receives only messages which are sent while the connection is active. Other messages are lost.</p> <p>If the attribute is true, the consumer is subscribed as 'durable' so it receives even messages which are sent while the connection is inactive. The broker stores such messages until they can be delivered or until the expiration is reached.</p> <p>This switch is useful <i>only for Topics</i> destinations, because Queue destinations always store messages until they can be delivered or the expiration is reached.</p> <p>Please note that consumer is inactive e.g. during server restart and during short moment when the user updates the "JMS message listener" and it must be re-initialized. So during these intervals, the message in the Topic may get lost if the consumer does not have the durable subscription.</p> <p>If the subscription is durable, client must have <b>ClientId</b> specified. This attribute can be set in different ways in dependence on JMS provider. E.g. for ActiveMQ, it is set as a URL parameter <code>tcp://localhost:1244?jms.clientID=TestClientID</code>.</p>
Subscriber name	Available only when <b>Durable subscriber</b> is <code>true</code> . By default, a durable subscriber name is generated automatically in the <code>subscr_[clusterNodeId]_[listenerId]</code> format; therefore, a subscriber has a different name on each cluster node. Using this attribute, you can specify a custom subscriber name that will be identical on all cluster nodes.
Message selector	<p>This query string can be used as a specification of conditions for filtering incoming messages. Syntax is well described on <a href="#">Java EE API</a> web site. It has different behavior depending on the type of consumer (queue/topic):</p> <p>Queue: Messages that are filtered out remain in the queue.</p> <p>Topic: Messages filtered out by a Topic subscriber's message selector will never be delivered to the subscriber. From the subscriber's perspective, they do not exist.</p>
<b>Message Processing</b>	
Number of consumers	E.g. 1
Groovy code	A Groovy code may be used for additional message processing and/or for refusing a message. Both features are described below.

## Optional Groovy code

Groovy code may be used for additional message processing or for refusing a message.

- **Additional message processing** Groovy code may modify/add/remove values stored in the containers "properties" and "data".
- **Refuse/acknowledge the message** If the Groovy code returns `Boolean.FALSE`, the message is refused. Otherwise, the message is acknowledged. A refused message may be redelivered, however the JMS broker

should configure a limit for redelivering messages. If the Groovy code throws an exception, it's considered a coding error and the JMS message is NOT refused because of it. So, if the message refusal is to be directed by some exception, it must be handled in Groovy.

Table 30.3. Variables accessible in groovy code

type	key	description
javax.jms.Message	msg	instance of a JMS message
java.util.Properties	properties	See below for details. It contains values (String or converted to String) read from a message and it is passed to the task which may then use them. For example, the <b>execute graph</b> task passes these parameters to the executed graph.
java.util.Map<String, Object>	data	See below for details. Contains values (Object, Stream, etc.) read or proxied from the message instance and it is passed to the task which may then use them. For example, the <b>execute graph</b> task passes it to the executed graph as dictionary entries.
javax.servlet.ServletContext	servletContext	An instance of ServletContext.
com.cloveretl.server.api.ServerFacade	serverFacade	An instance of serverFacade usable for calling CloverETL Server core features.
java.lang.String	sessionToken	Session Token needed for calling serverFacade methods

## Message data available for further processing

A JMS message is processed and the data it contains is stored into two data structures: Properties and Data.

Table 30.4. Properties Elements

key	description
JMS_PROP_[property key]	For each message property, one entry is created where "key" is made of the JMS_PROP_ prefix and property key.
JMS_MAP_[map entry key]	If the message is an instance of MapMessage, for each map entry, one entry is created where "key" is made of the JMS_MAP_ prefix and map entry key. Values are converted to String.
JMS_TEXT	If the message is an instance of TextMessage, this property contains content of the message.
JMS_MSG_CLASS	A class name of a message implementation
JMS_MSG_CORRELATIONID	Correlation ID is either a provider-specific message ID or an application-specific String value
JMS_MSG_DESTINATION	If the JMSDestination header field contains the destination to which the message is being sent.
JMS_MSG_MESSAGEID	MessageID is a String value that should function as a unique key for identifying messages in a historical repository. The exact scope of uniqueness is provider-defined. It should at least cover all messages for a specific installation of a provider, where an installation is some connected set of message routers.
JMS_MSG_REPLYTO	A destination to which a reply to this message should be sent.
JMS_MSG_TYPE	A message type identifier supplied by the client when the message was sent.
JMS_MSG_DELIVERYMODE	DeliveryMode value specified for this message.
JMS_MSG_EXPIRATION	Time the message expires, which is the sum of the time-to-live value specified by the client and the GMT at the time of the send.
JMS_MSG_PRIORITY	The JMS API defines ten levels of priority value (0 = lowest, 9 = highest). In addition, clients should consider priorities 0-4 as gradations of normal priority and priorities 5-9 as gradations of expedited priority.
JMS_MSG_REDELIVERED	True if this message is being redelivered.
JMS_MSG_TIMESTAMP	Time a message was handed off to a provider to be sent. It is not the time the message was actually transmitted, because the actual send may occur later due to transactions or other client-side queueing of messages.

Note that all values in the "Properties" structure are stored as a String type – however they are numbers or text.

For backwards compatibility, all listed properties can also be accessed using lower-case keys; however, it is a deprecated approach.

Table 30.5. "Data" elements

key	description
JMS_DATA_MSG	An instance of javax.jms.Message.
JMS_DATA_STREAM	An instance of java.io.InputStream. Accessible only for TextMessage, BytesMessage, StreamMessage, ObjectMessage (only if a payload object is an instance of String). Strings are encoded in UTF-8.
JMS_DATA_TEXT	An instance of String. Only for TextMessage and ObjectMessage, where a payload object is an instance of String.
JMS_DATA_OBJECT	An instance of java.lang.Object - message payload. Only for ObjectMessage.

The **Data** container is passed to a task that can use it, depending on its implementation. For example, the task **execute graph** passes it to the executed graph as dictionary entries.

In a Cluster environment, you can explicitly specify node IDs, which can execute the task. However, if the data payload is not serializable and the receiving and executing node differ, an error will be thrown as the Cluster cannot pass the data to the executing node.

Inside a graph or a jobflow, data passed as dictionary entries can be used in some component attributes. For example, the **File URL** attribute would look like: `"dict:JMS_DATA_STREAM:discrete"` for reading the data directly from the incoming JMS message using a proxy stream.



### Note

If the graph is **executed on Worker**, the dictionary entries must be **serialized**; otherwise, they cannot be passed to the graph.

For backwards compatibility, all listed dictionary entries can also be accessed using lower-case keys; however, it is a deprecated approach.

## Universal Event Listeners

Since 2.10

Universal Event Listeners allow you to write a piece of Groovy code that controls when an event is triggered, subsequently executing a predefined task. The Groovy code is periodically executed and when it returns `TRUE`, the task is executed.

Table 30.6. Attributes of Universal message task

Attribute	Description
Node IDs to handle the event	<p>In a Cluster environment, each universal event listener has a <b>Node IDs</b> attribute which may be used to specify which cluster node performs the Groovy code. There are following possibilities:</p> <ul style="list-style-type: none"> <li>• <b>No failover:</b> Just one node ID specified - Only the specified node performs the Groovy code, however node status must be "ready". When the node isn't ready, the code isn't performed at all.</li> <li>• <b>Failover with node concurrency:</b> No node ID specified (empty input) - All cluster nodes with the status "ready" concurrently perform the Groovy code. So the code is executed on each node in the specified interval.</li> <li>• <b>Failover with node reservation:</b> More node IDs specified (separated by a comma) - Just one of the specified nodes performs the Groovy code. If the node fails for any reason (or its status isn't "ready"), any other "ready" node from the list continues with periodical Groovy code processing.</li> </ul> <p>In a standalone environment, the <b>Node IDs</b> attribute is ignored.</p>
Interval of check in seconds	Periodicity of Groovy code execution.
Groovy code	Groovy code that evaluates either to <code>TRUE</code> (execute the task) or <code>FALSE</code> (no action). See below for more details.

### Groovy code

A piece of Groovy is repeatedly executed and evaluated; based on the result, the event is either triggered and the task executed or no action is taken.

For example, you can continually check for essential data sources before starting a graph. Or, you can do complex checks of a running graph and, for example, decide to kill it if necessary. You can even call the CloverETL Server core functions using the ServerFacade interface, see Javadoc: <http://host:port/clover/javadoc/index.html>

### Evaluation Criteria

If the Groovy code returns `Boolean.TRUE`, the event is triggered and the associated task is executed. Otherwise, nothing happens.

If the Groovy code throws an exception, it is considered a coding error and the event is NOT triggered. Thus, exceptions should be properly handled in the Groovy code.

Table 30.7. Variables accessible in Groovy code

type	key	description
java.util.Properties	properties	An empty container which may be filled with String-String key-value pairs in your Groovy code. It is passed to the task which may use them somehow. For example, the task <b>execute graph</b> passes these parameters to the executed graph.
java.util.Map<String, Object>	data	An empty container which may be filled with String-Object key-value pairs in your Groovy code. It is passed to the task which may use them somehow according to its implementation - e.g. the task <b>execute graph</b> passes it to the executed graph as dictionary entries. Note that it is not serializable, thus if the task is relying on it, it can be processed properly only on the same cluster node.
javax.servlet.ServletContext	context	An instance of ServletContext in which CloverETL Server is running.
com.cloveretl.server.api.ServerFacade	serverFacade	An instance of serverFacade usable for calling CloverETL Server core features.
java.lang.String	sessionToken	A sessionToken needed for calling methods on the serverFacade.

## File Event Listeners (remote and local)

Local file-system changes: Since 1.3

Remote file-system changes: Since 4.2

File Event Listeners allow you to monitor changes on a specific local file system path or remote URL – for example, new files appearing in a folder – and react to such an event with a predefined task.

You can either specify an exact file name or use a wildcard or regexp, then set a checking interval in seconds and define a task to process the event.

There is a global minimum check interval that you can change if necessary in the configuration (the `clover.event.fileCheckMinInterval` property). See Chapter 14, [List of Configuration Properties](#) (p. 106).

The screenshot shows a web-based configuration form for creating a new file event listener. The form is organized into several sections:

- Name:** FileEventListener\_1
- Owner:** clover
- Initialize by:**  Any node  One of selected nodes
- Enabled:**
- Check Location:**
  - File system: Local file system
  - Path: /opt/unprocessedLogs
  - Resolved paths: virt-alpha: /opt/unprocessedLogs, virt-gray: /opt/unprocessedLogs
  - Buttons: Validate Accessibility, Paste URL
  - Send e-mail on check failure
- File to Check:**
  - Type of check: File added
  - Filename match type: Wildcards filename match
  - Filename pattern: \*.bt
  - Check every: 60 seconds
  - Trigger task when file has not changed for 3 checks (180 seconds)
  - Ignore empty files
  - Trigger one task for multiple changed files
- Triggered Task:**
  - Task: Start a graph
  - Node IDs to process the task (empty for any node):
  - Sandbox: default
  - Graph: graph/launch/productSearch.grf
  - Save run record
  - Parameters: Name, Value (with a plus icon to add more)
  - Parameters passed from the listener to the task: »

At the bottom of the form are 'Create' and 'Cancel' buttons.

Figure 30.8. Web GUI - creating a File Event listener

Table 30.8. Parameters passed from the listener to the task

Parameter	Description
EVENT_USERNAME	The name of the user who caused the event.
EVENT_USER_ID	A numeric ID of the user who caused the event.
EVENT_FILE_NAME	The name of the file (without the path or URL) that triggered the event. Present only when <b>Trigger one task for multiple changed files</b> is disabled.
EVENT_FILE_PATH	A resolved (without placeholders) path to the observed directory on the local filesystem. Valid only for a local file listener.
EVENT_FILE_PATTERN	A filename pattern.
EVENT_FILE_EVENT_TYPE	The type of the file event. Possible values: SIZE, CHANGE_TIME, APPEARANCE, DISAPPEARANCE.
EVENT_FILE_LISTENER_ID	An ID of the listener which triggered the event.
EVENT_FILE_URLS	Full URLs to access the files, e.g. in the File URL attribute of components. If <b>Trigger one task for multiple changed files</b> is enabled and there are multiple URLs, they are separated by a separator specified by CloverETL Engine property <code>DEFAULT_PATH_SEPARATOR_REGEX</code> .
EVENT_FILE_AUTH_USERNAME	A username/ID to the remote location.
EVENT_FILE_AUTH_USERNAME_URL_ENCODED	The same as <code>EVENT_FILE_AUTH_USERNAME</code> , but the value is also URL encoded, so it may be used in a URL.
EVENT_FILE_AUTH_PASSWORD	a password/key to the remote location. It is encrypted by the master password.
EVENT_FILE_AUTH_PASSWORD_URL_ENCODED	The same as <code>EVENT_FILE_AUTH_PASSWORD</code> , but the value is also URL encoded, so it may be used in a URL.

## Cluster environment

In a Cluster environment, each file event listener has a **Node IDs** attribute which may be used to specify which cluster node will perform the checks on its local file system. There are following possibilities:

- **No failover:** Just one node ID specified - Only the specified node observes the local/remote filesystem; however, the node status must be "ready". When the node isn't ready, the file system isn't checked at all.

To create a file event listener with no failover, select **One of selected nodes** in **Initialize by** and select one node from the table below.

- **Failover with node concurrency:** No node ID specified (empty input) - All cluster nodes with the status "ready" concurrently check the local/remote filesystem according to file event listener attributes settings. In this mode, when the listener is configured to observe the local filesystem, each cluster node observes its own local file system. So it's useful only when the observed path is properly shared among the cluster nodes. It may behave unpredictably otherwise. On the other hand, when the listener is configured to observe the remote filesystem, listeners running on different cluster nodes may connect to the same remote resource. The nodes use a locking mechanism when accessing the local or remote filesystem, so no conflict between listeners running concurrently on different nodes can occur.

To create file event listener with node concurrency, select **Any node** in **Initialize by**.

- **Failover with node reservation:** More node IDs specified (separated by comma) - Just one of the specified nodes checks its filesystem. If the node fails for any reason (or its status isn't "ready"), any other "ready" node

from the list continues with checking. Please note, that when file event listener is re-initialized on another cluster node, it compares the last directory content detected by the failed node with the current directory content.

To create a file event listener with node reservation, select **One of selected nodes** in **Initialize by** and select more nodes.

In a standalone environment, the **Node IDs** attribute is ignored.

## Supported filesystems and protocols

### Local filesystem

The user may specify a path to the directory which the listener shall observe. The listener doesn't read the directory content recursively. The directory must exist.

If the listener can run concurrently on more nodes, the directory must be shared among all these nodes and the directory must exist on all these nodes. In a cluster environment, the directory must exist on each cluster node where the listener may run.

It is recommended to use placeholders to unify the configuration on all nodes. The recommended placeholders are: CloverETL Server config property `${sandboxes.home}` and JVM system property `${java.io.tmpdir}`. It is possible to use any JVM system property or Environment variable.

Figure 30.9. File available on local file system

### Remote filesystem

The user may specify a URL to the directory which the listener shall observe. The supported protocols are: FTP, S3, SFTP and SMB. Different protocols may use different authentication methods: none, username+password and keystore. The listener doesn't read the directory content recursively. The directory must exist.

Figure 30.10. File available on remote location

Currently the subset of the protocols allowed by file-operations is supported:

- **FTP - File Transfer Protocol** (no authentication or username+password authentication) URL example:

```
ftp://host:23/observed/path/
```

- **SFTP (SSH/FTP) - SSH File Transfer Protocol** (username+private key authentication) URL example:

```
sftp://host:23/observed/path/
```

It is recommended to use placeholders to unify the path configuration on all nodes. The recommended placeholders are: CloverETL Server config property `${sandboxes.home}`, JVM system property `${user.home}`. It is possible to use any JVM system property or Environment variable.

- **S3 - Amazon S3 Storage** (AWS Access Key ID + Secret Access Key authentication) URL example:

```
s3://s3.amazonaws.com/bucketname/path/
```

Please specify the AWS Access Key ID as a username and Secret Access Key as a password.

- **Microsoft SMB/CIFS Protocol** (username+password authentication) URL example:

```
smb://host/path/
```

- **Microsoft SMBv2/v3 Protocol** (username+password authentication) URL example:

```
smb2://host/path/
```

## Observed file

---

The local observed file is specified by a directory path and file name pattern.

The remote observed file is specified by a URL, credentials and file name pattern.

The user may specify just one exact file name or file name pattern for observing more matching files in specified directory. If there are more changed files matching the pattern, separated event is triggered for each of these files.

There are three ways how to specify file name pattern of observed file(s):

- [Exact match](#) (p. 256)
- [Wildcards](#) (p. 256)
- [Regular expression](#) (p. 257)

### Exact match

You specify the exact name of the observed file.

### Wildcards

You can use wildcards common in most operating systems (\*, ?, etc.)

- \* - Matches zero or more instances of any character
- ? - Matches one instance of any character
- [ . . . ] - Matches any of characters enclosed by the brackets
- \ - Escape character

Examples

- \*.csv - Matches all CSV files
- input\_\*.csv - Matches i.e. input\_001.csv, input\_9.csv
- input\_???.csv - Matches i.e. input\_001.csv, but does not match input\_9.csv

## Regular expression

Examples

- (.\*)\.(jpg|jpeg|png|gif)\$ - Matches image files

## Notes

- It is strongly recommended to use absolute paths with placeholders. It is possible to use a relative path, but the working directory depends on an application server.
- Use forward slashes as file separators, even on MS Windows. Backslashes might be evaluated as escape sequences.

## File Events

---

For each listener you have to specify event type, which you are interested in.

Please note that since CloverETL 4.2, the grouping mode may be enabled for the file listener, so all file changes detected by a single check produce just one 'grouped' file event. Otherwise each single file produces its own event.

There are four types of file events:

- [File added](#) (p. 257)
- [File removed](#) (p. 257)
- [File size changed](#) (p. 257)
- [File timestamp changed](#) (p. 257)

### File added

Event of this type occurs, when the observed file is created or copied from another location between two checks. Please keep in mind that event of this type occurs immediately when a new file is detected, regardless if it is complete or not. Thus task which may need a complete file is executed when the file is still incomplete. Recommended approach is to save the file to a different location and when it is complete, rename it or move it to an observed location where CloverETL Server may detect it. File moving/renaming should be an atomic operation.

An event of this type does not occur when the file has been updated (change of timestamp or size) between two checks. Appearance means that the file didn't exist during the previous check and it exists now, during the current check.

### File removed

Event of this type occurs, when observed file is deleted or moved to another location between two checks.

### File size changed

Event of this type occurs when the size of the observed file has changed between two checks. Event of this type is never produced when the file is created or removed. The file must exist during both checks.

### File timestamp changed

Event of this type occurs, when timestamp of the observed file has changed between two checks. Event of this type is never produced when the file is created or removed. The file must exist during both checks.

## Check Interval, Task and Use Cases

---

- The user may specify the minimum time interval between two checks. Use the **Check every** field to specify the interval in seconds.
- Each listener defines a task which will be processed as a reaction to the file event. All task types and their attributes are described in the Scheduling (p. 224) and Graph Event Listeners (p. 238) sections.
  - Graph Execution when a file with input data is accessible.
  - Graph Execution when a file with input data is updated.
  - Graph Execution when a file with generated data is removed and must be recreated.

## How to use source of event during task processing

A file(s) which caused the event (considered as a source of the event) may be used during task processing. CloverETL graph/jobflow components with the **File URL** attribute (e.g. reader or writer components) may directly use an event source by parameter placeholder: `${EVENT_FILE_URLS}`. For another parameters, see [Executed by Task Graph Execution by File Event Listener](#) (p. 201).

Note that previous versions used lower-case placeholders. Since version 3.3, placeholders are upper-case, however lower-case still work for backward compatibility.

For **graph execution** task this works only if the graph is not pooled. Thus `keep in pool interval` must be set to 0 (default value).

## Delayed triggering for incomplete files

It is possible to delay task execution for incomplete files. This is useful in cases when the condition to execute the listener's task has been met, but there is still some work that needs to be done on the file, e.g. the whole file needs to be uploaded.

### Ignore empty files

If the process creating the file creates an empty file, then switch to different task for several minutes or even hours and finally writes the content, tick **Ignore empty files** checkbox. The task will be triggered only if a non-empty file appears.

### Trigger task when file has not changed for n checks

If the file size slowly rises until the file is complete, tick the checkbox **Trigger task when file has not changed**. Then specify the number of additional file size checks that are to be performed on the file. The listener's task will not be triggered until the checks are performed and the file's size stays the same between these checks.

### Combination

If you use **Ignore empty files** and **Trigger task when file has not changed for n checks** together, the first one filters out empty files and the latter one the files that are being changed. The task will be triggered only on files that are not empty and that have not changed for the specified number of checks.

## Howtos

---

[Create a file event listener listening to changes on local file system](#) (p. 259)

[Observe file from one cluster node](#) (p. 259)

[Quickly setup failure notification](#) (p. 259)

[Quickly enable or disable file event listener](#) (p. 260)

## Create a file event listener listening to changes on local file system

This howto shows a way to create a new listener checking appearance of a file (`new_invoices.txt`) on a local file system (`/mnt/sdb2/`). The appearance will trigger a graph (`graph/checkInvoices.grf`) from the **InvoicesProcessing** sandbox.

In **Event Listeners** → **File Event Listeners**, click **New Listener**.

Enter the **Name** of the listener, e.g. **Invoices**.

Enter the **Path** to the directory where files will appear: `/mnt/sdb2`. You can check that Clover can access this directory (the directory exists and permissions are set up properly) with the **Validate Accesibility** button.

If the observed directory becomes inaccessible, CloverETL Server can send you an email. To do so, tick **Send email on check failure** and enter recipient(s).

The event should be triggered on file appearance - set **Type of check** to **File added**.

Enter the file name `new_invoices.txt` to **Filename pattern**.

If the file is created empty, but the content is written after some time, tick **Ignore empty files**. Doing so, the task will be executed after the file contains some data.

If it takes a long time to copy the whole file to the observed position, the Clover Server can perform several check to ensure that the file to process is not to be changed. Tick **Trigger task when file has not changed for** and enter the number of checks. If you tick **Ignore empty files**, this checks will be performed after the file is not empty.

Choose **Sandbox** with the graph (**InvoicesProcessing**) and the **graph** (`graph/checkInvoices.grf`).

To save the changes, click on the **Create** button.

## Observe file from one cluster node

Create the listener in the same way as on the Server.

Switch **Initialize by** to **One of selected nodes**.

Add the particular node(s) from **Available nodes** to **Selected nodes**.

## Quickly setup failure notification

To create a notification for when the file event listener fails, click on the **Create notification** button. Pressing the button opens up a popup dialog where email addresses can be entered separated by commas.

The entered email addresses are remembered and pre-filled the next time the button is pressed. If the popup is closed with invalid email addresses entered, the field is cleared.

When creating the notification, a Task Failure Listener is created with an email task listening to the selected File Event Listener. The first entered email address will be used as the Reply-to(Sender) address. The subject and body of the email is as predefined by the Task Failure template. The trigger limit is set to 5.

## Editing failure notification

If there is a Task Failure Listener listening to given File Event Listener then instead of the **Create Notification** button a **Notification Detail** button is displayed. This button redirects to the Task Failure Listener page and shows

the details of the Task Failure Listener listening to the File Event Listener. If more than one Task Failure Listeners are listening to the File Event Listener, then the details of the first one is shown.

### Quickly enable or disable file event listener

In **Event Listeners** → **File Event Listeners**, there is a table with event listeners. In this table, click the icon in the **Enabled** column.

### Pasting URL

The whole URL including user name and password can be pasted at once. Click **Paste URL** and paste the string.

If the name or password in URL contain special characters, e.g. +, the special characters should be encoded:  
`ftp://anonymous:test%2B@example.com/dir/file.txt`

Note: use encoding accepted by the `java.net.URLDecoder.decode()` function.

## Task Failure Listeners

[Task Choice](#) (p. 261)

[Task Failed E-mail Template](#) (p. 262)

Since 4.4

Task Failure Listeners allow you to detect and react to failures in your server when a task you set up in a listener fails to execute, e.g. a File Listener is set up to detect changes on an FTP server, but it fails to connect to the FTP server.

Task Failure Listeners do not detect failures of the task itself, e.g. a File Listener is set up to detect changes on an FTP server and send an email if the change is detected. If the File Listener fails to send the email for some reason, the Task Failure Listener won't detect it.

The same tasks to be executed are available as with all the other listeners, the difference is that when creating a new Task Failure Listener the pre-selected task is **Sending an email** if the email service is configured in Configuration.

**New Task Failure Listener**

*Task failure listeners allow you to detect and react to failures of a listener's check, e.g. when a file listener is set up to detect changes on an FTP server, but it fails to connect to the FTP server.*

Name:

Owner:

Enabled

**Listener to Check**

Listener type:

Listener:   
Listens to: File added: \*.txt  
 Action: Start a graph

Trigger limit:

**Triggered Task**

Task:

Sandbox:

Graph:

Save run record ⓘ

Parameters:   + ⓘ

»

Figure 30.11. Web GUI - creating a Task Failure listener

## Task Choice

There are three options to choose from: Listening to any task failure, listening to failures from a group (such as File Event Listeners) or listening to a failure of a chosen listener.

Selecting an option from the **Listen to type** menu restricts the **Listen to** combobox to event sources of the chosen category. If there are no event sources in the category, the Failure Listener can still be created, it will react to failures of tasks that will be created in that category.

When selecting an event source, you can type into the text field to filter the dropdown menu. After selecting an event source, some information is presented about the listener.

When sending an email, you can use the 'task failed' template by selecting it from the **E-mail template** dropdown menu.

## Task Failed E-mail Template

The default email template may be modified using placeholders described in [Placeholders](#) (p. 205) and parameters in Table 30.9, “[Parameters usable in task failure email template](#)” (p. 262). Furthermore, some additional parameters can be used if the failed listener is a File Event Listener, see Table 30.10, “[File Event Listener specific parameters usable in task failure email template](#)” (p. 262).

Table 30.9. Parameters usable in task failure email template

Parameter	Description
TASK_LISTENER_ID	The ID of the failed listener.
TASK_LISTENER_NAME	The name of the failed listener.
TASK_LISTENER_TYPE	The type of the failed listener.
TASK_LISTENER_TYPE_TEXT	The full name of the failed listener's type.
TASK_LISTENER_OWNER_USERNAME	The username of the failed listener.

Table 30.10. File Event Listener specific parameters usable in task failure email template

Parameter	Description
FILE_EVENT_LISTENER_FILE_PATH	The path of the observed directory.
FILE_EVENT_LISTENER_FILE_REMOTE_URI	The URI of the observed directory.
FILE_EVENT_LISTENER_FILE_NAME_PATTERN	The file name pattern files the listener observes.
FILE_EVENT_LISTENER_FILE_CHECK_TYPE	The type of check the listener performs.
FILE_EVENT_LISTENER_FILE_MATCH_TYPE	The list match type.
FILE_EVENT_LISTENER_FILE_NODES	The nodes of nodes the listener can be initialized by.

---

## Chapter 31. Recommendations for Transformations Developers

---

### Add external libraries to app-server classpath

Connections (JDBC/JMS) may require third-party libraries. We strongly recommend adding these libraries to the app-server classpath.

CloverETL allows you to specify these libraries directly in a graph definition so that CloverETL can load these libraries dynamically. However, external libraries may cause memory leak, resulting in `java.lang.OutOfMemoryError: PermGen space`, in this case.

In addition, app-servers should have the JMS API on their classpath – and the third-party libraries often bundle this API as well. So it may result in classloading conflicts if these libraries are not loaded by the same classloader.

---

### Another graphs executed by RunGraph component may be executed only in the same JVM instance

In the server environment, all graphs are executed in the same VM instance. The attribute **same instance** of the **RunGraph** component cannot be set to false.

---

## Chapter 32. Extensibility - CloverETL Engine Plugins

Since 3.1.2

The CloverETL Server can use external engine plugins loaded from a specified source. The source is specified by `engine.plugins.additional.src` config property.

See details about the possibilities with CloverETL configuration in **Part III, “[Configuration](#)” (p. 67)**

This property must be the absolute path to the directory or zip file with additional CloverETL engine plugins. Both the directory and zip file must contain a subdirectory for each plugin. These plugins are not a substitute for plugins packed in a WAR file. Changes in the directory or the ZIP file apply only when the Server is restarted.

Each plugin has its own class-loader that uses a parent-first strategy by default. The parent of plugins' classloaders is web-app classloader (content of [WAR]/WEB-INF/lib). If the plugin uses any third-party libraries, there may be some conflict with libraries on the parent-classloaders classpath. These are common exceptions/errors suggesting that there is something wrong with classloading:

- `java.lang.ClassCastException`
- `java.lang.ClassNotFoundException`
- `java.lang.NoClassDefFoundError`
- `java.lang.LinkageError`

There are several ways you can get rid of such conflicts:

- Remove your conflicting third-party libraries and use libraries on parent classloaders (web-app or app-server classloaders)
- Use a different class-loading strategy for your plugin.
  - In the plugin descriptor `plugin.xml`, set attribute `greedyClassLoader="true"` in the element `plugin`
  - It means that the plugin classloader will use a self-first strategy
- Set an inverse class-loading strategy for selected Java packages.
  - In the plugin descriptor `plugin.xml`, set attribute `excludedPackages` in the element `plugin`.
  - It is a comma-separated list of package prefixes – for example: `excludedPackages="some.java.package,some.another.package"`
  - In the previous example, all classes from `some.java.package`, `some.another.package` and all their sub-packages would be loaded with the inverse loading strategy, then the rest of classes on the plugins classpath.

The suggestions above may be combined. Finding the best solution for these conflicts may depend on the libraries on app-server classpath.

For more convenient debugging, it is useful to set a TRACE log level for related class-loaders.

```
<logger name="org.jetel.util.classloader.GreedyURLClassLoader">
  <level value="trace"/>
</logger>
<logger name="org.jetel.plugin.PluginClassLoader">
  <level value="trace"/>
</logger>
```

See Chapter 16, [Logging](#) (p. 127) for details about overriding a server log4j configuration.

---

## Chapter 33. Troubleshooting

---

### Graph hangs and is un-killable

A graph can sometimes hang and be un-killable if some network connection in it hangs. Set a shorter `tcp-keepalive` so that the connection times out earlier. The default value on Linux is 2 hours (7,200 seconds). You can set it to 10 minutes (600 seconds).

See [Using TCP keepalive under Linux](#).

The file descriptor can be closed manually using `gdb`. See [How to close file descriptor via Linux shell command](#).

---

## SSL/TLS Issues

### Java 7 and HTTPS

There may be issues when using Java 7 and HTTPS-based communication. There are servers that do not support TLS 1.0 (Transport Layer Security) protocol anymore and the Java 7 does not support newer versions of the TLS by default. To activate newer versions of the TLS protocol, the user has to specify the system property `https.protocols` and set the value to `TLSv1`, `TLSv1.1`, `TLSv1.2`. See [Production Server](#) (p. 16) to find out how to set the system property on a particular application server.

### SSL-related Failures on WebLogic 12

Certain graphs using SSL-encrypted connections may fail on WebLogic 12 due to damaged library distributed with this application server. The issue can be identified by a SHA-1 digest error in the graph execution stacktrace:

```
...
Caused by: java.io.IOException: Could not convert socket to TLS
    at com.sun.mail.pop3.Protocol.stls(Protocol.java:659)
    at com.sun.mail.pop3.POP3Store.getPort(POP3Store.java:269)
    at com.sun.mail.pop3.POP3Store.protocolConnect(POP3Store.java:207)
Caused by: javax.net.ssl.SSLException: java.lang.SecurityException:
    SHA1 digest error for org/bouncycastle/jce/provider/JCECPublicKey.class
...
```

To fix the issue, replace the library `[MW_HOME]/oracle_common/modules/bcprov-jdk16-1.45.jar` with the one downloaded directly from Bouncy Castle home page. Restart the application server to load the new library.

### Graph run in Worker is Slow

It may be caused by slow data storage. Use `vmstat`, e.g. `vmstat 1 30`. If you see high values under `io/bi` or `io/bo` columns, it might be that case. Another tool to confirm or disconfirm slow data storage as possible cause is `iotop`.

---

# Part VI. API

---

---

## Chapter 34. Simple HTTP API

The Simple HTTP API is a basic Server automation tool that lets you control the Server from external applications using simple HTTP calls.

Most of operations is accessible using the HTTP GET method and return plain text. Thus, both `request` and `response` can be conveniently sent and parsed using very simple tools (`wget`, `grep`, etc.).

If global security is on (on by default), the Basic HTTP authentication is used. Authenticated operations will require valid user credentials with corresponding permissions.

Note that the ETL graph-related operations `graph_run`, `graph_status` and `graph_kill` also work for jobflows and Data Profiler jobs.

The generic pattern for a request URL:

```
http://[domain]:[port]/[context]/[servlet]/[operation]?[param1]=[value1]&[param2]=[value2]...
```

**example:** `http://localhost:8080/clover/simpleHttpApi/help`



### Note

For backward compatibility, you can also use `http://localhost:8080/clover/request_processor/help`.

For a `wget` client, you can use the following command line:

```
wget --user=$USER --password=$PASS -O ./$OUTPUT_FILE $REQUEST_URL
```

- [Operation help](#) (p. 268)
- [Operation graph\\_run](#) (p. 268)
- [Operation graph\\_status](#) (p. 269)
- [Operation graph\\_kill](#) (p. 269)
- [Operation server\\_jobs](#) (p. 270)
- [Operation sandbox\\_list](#) (p. 270)
- [Operation sandbox\\_content](#) (p. 270)
- [Operation executions\\_history](#) (p. 271)
- [Operation suspend](#) (p. 272)
- [Operation resume](#) (p. 272)
- [Operation sandbox\\_create](#) (p. 273)
- [Operation sandbox\\_add\\_location](#) (p. 273)
- [Operation sandbox\\_remove\\_location](#) (p. 274)
- [Operation download\\_sandbox\\_zip](#) (p. 274)
- [Operation upload\\_sandbox\\_zip](#) (p. 274)
- [Operation cluster\\_status](#) (p. 275)
- [Operation export\\_server\\_config](#) (p. 275)
- [Operation import\\_server\\_config](#) (p. 276)

The HTTP API is enabled by default. You can disable it with the configuration property `http.api.enabled`. In the Server GUI, switch to **Configuration** → **Setup** and add the following line

```
http.api.enabled=false
```

to the properties file.

## Operation help

### parameters

no

### returns

a list of possible operations and parameters with its descriptions

### example

```
http://localhost:8080/clover/simpleHttpApi/help
```

## Operation graph\_run

Call this operation to start an execution of the specified job. The operation is called `graph_run` for backward compatibility, however it may execute ETL graph, jobflow or profiler job.

### parameters

Table 34.1. Parameters of `graph_run`

parameter name	mandatory	default	description
graphID	yes	-	A file path to the job file, relative to the sandbox root.
sandbox	yes	-	Text ID of sandbox.
additional job parameters	no		Any URL parameter with the <code>param_</code> prefix is passed to the executed job and may be used in transformation XML as a placeholder, but without the <code>param_</code> prefix. e.g. <code>param_FILE_NAME</code> specified in URL may be used in the XML as <code>\${FILE_NAME}</code> . These parameters are resolved only during loading of XML, so it cannot be pooled.
additional config parameters	no		URL parameters prefixed with <code>config_</code> can set some of the execution parameters. For graphs, the following parameters are supported: <ul style="list-style-type: none"> <li><code>config_skipCheckConfig</code> - when set to <code>false</code>, graph configuration will be checked before the execution.</li> <li><code>config_logLevel</code> - log level of the executed graph, one of OFF, FATAL, ERROR, WARN, INFO, DEBUG, TRACE, ALL.</li> <li><code>config_clearObsoleteTempFiles</code> - when set to <code>true</code>, temp files of previous runs of this graph will be deleted before the execution.</li> <li><code>config_debugMode</code> - when set to <code>true</code>, debug mode for a given graph will be enabled. For more information, see <a href="#">Job Config Properties</a> (p. 183).</li> </ul>
nodeID	no	-	In cluster mode, it is the ID of a node which should execute the job. However it is not final. If the graph is distributed or the node is disconnected, the graph may be executed on another node.
verbose	no	MESSAGE	MESSAGE   FULL - how verbose should possible error message be.

**returns**

run ID: incremental number, which identifies each execution request

**example**

```
http://localhost:8080/clover/simpleHttpApi/graph_run?graphID=graph/graphDBExecute.grf&sandbox=mva
```

## Operation graph\_status

Call this operation to obtain a status of a specified job execution. The operation is called `graph_status` for backward compatibility; however, it may return status of an ETL graph or jobflow.

**parameters**

Table 34.2. Parameters of `graph_status`

parameter name	mandatory	default	description
runID	yes	-	Id of each graph execution
returnType	no	STATUS	STATUS   STATUS_TEXT   DESCRIPTION   DESCRIPTION_XML
waitForStatus	no	-	Status code which we want to wait for. If it is specified, this operation will wait until the graph is in the required status.
waitTimeout	no	0	If <code>waitForStatus</code> is specified, it will wait only for the specified amount of milliseconds. Default 0 means forever, but it depends on an application server configuration. When the specified timeout expires and graph run still isn't in a required status, the server returns code 408 (Request Timeout). 408 code may be also returned by an application server if its HTTP request timeout expires before.
verbose	no	MESSAGE	MESSAGE   FULL - how verbose should possible error message be.

**returns**

Status of a specified graph. It may be a number code, text code or a complex description in dependence on the optional parameter `returnType`. Description is returned as a plain text with a pipe as a separator, or as XML. A schema describing XML format of the XML response is accessible on CloverETL Server URL: `http://[host]:[port]/clover/schemas/executions.xsd` Depending on the `waitForStatus` parameter, it may return a result immediately or wait for a specified status.

**example**

```
http://localhost:8080/clover/simpleHttpApi/graph_status ->
-> ?runID=123456&returnType=DESCRIPTION&waitForStatus=FINISHED&waitTimeout=60000
```

## Operation graph\_kill

Call this operation to abort/kill a job execution. The operation is called `graph_kill` for backward compatibility, however it may abort/kill ETL graph, jobflow or profiler job.

**parameters**

Table 34.3. Parameters of `graph_kill`

parameter name	mandatory	default	description
runID	yes	-	The ID of each graph execution
returnType	no	STATUS	STATUS   STATUS_TEXT   DESCRIPTION
verbose	no	MESSAGE	MESSAGE   FULL - how verbose should possible error message be.

**returns**

The status of the specified graph after an attempt to kill it. It may be a number code, text code or a complex description in dependence on optional parameter.

**example**

```
http://localhost:8080/clover/simpleHttpApi/graph_kill?runID=123456&returnType=DESCRIPTION
```

## Operation `server_jobs`

**parameters**

no

**returns**

a list of runIDs of currently running jobs.

**example**

```
http://localhost:8080/clover/simpleHttpApi/server_jobs
```

## Operation `sandbox_list`

**parameters**

no

**returns**

List of all sandbox text IDs. In the next versions, it will return only accessible ones.

**example**

```
http://localhost:8080/clover/simpleHttpApi/sandbox_list
```

## Operation `sandbox_content`

**parameters**Table 34.4. Parameters of `sandbox_content`

parameter name	mandatory	default	description
sandbox	yes	-	text ID of sandbox
verbose	no	MESSAGE	MESSAGE   FULL - how verbose should possible error message be.

**returns**

A list of all elements in the specified sandbox. Each element may be specified as a file path relative to the sandbox root.

**example**

```
http://localhost:8080/clover/simpleHttpApi/sandbox_content?sandbox=mva
```

**Operation executions\_history****parameters**

Table 34.5. Parameters of executions\_history

parameter name	mandatory	default	description
sandbox	yes	-	The text ID of a sandbox.
from	no		Lower datetime limit of start of execution. The operation will return only records after (and equal to) this datetime. Format: "yyyy-MM-dd HH:mm" (must be URL encoded).
to	no		The upper datetime limit of start of execution. The operation will return only records before (and equal to) this datetime. Format: "yyyy-MM-dd HH:mm" (must be URL encoded).
stopFrom	no		The lower datetime limit of stop of execution. The operation will return only records after (and equal to) this datetime. Format: "yyyy-MM-dd HH:mm" (must be URL encoded).
stopTo	no		The upper datetime limit of stop of execution. The operation will return only records before (and equal to) this datetime. Format: "yyyy-MM-dd HH:mm" (must be URL encoded).
status	no		Current execution status. The operation will return only records with specified STATUS. The values are RUNNING   ABORTED   FINISHED_OK   ERROR
sandbox	no		Sandbox code. The operation will return only records for graphs from a specified sandbox.
graphId	no		The text Id, which is unique in a specified sandbox. The file path is relative to the sandbox root.
orderBy	no		An attribute for list ordering. Possible values: id   graphId   status   startTime   stopTime. By default, there is no ordering.
orderDescend	no	true	A switch which specifies ascending or descending ordering. If true (default), ordering is descending.
returnType	no	IDs	Possible values are: IDs   DESCRIPTION   DESCRIPTION_XML
index	no	0	an index of the first returned records in a whole record set. (starting from
records	no	infinite	The maximum amount of returned records.
verbose	no	MESSAGE	MESSAGE   FULL - how verbose should possible error message be.

**returns**

List of executions according to filter criteria.

For returnType==IDs returns a simple list of runIDs (with new line delimiter).

For `returnType==DESCRIPTION` returns complex response which describes current status of selected executions, their phases, nodes and ports.

```
execution|[runID]|[status]|[username]|[sandbox]|[graphID]|[startedDatetime]|[finishedDatetime]|[clusterNode]|[grap
phase|[index]|[execTimeInMilis]
node|[nodeID]|[status]|[totalCpuTime]|[totalUserTime]|[cpuUsage]|[peakCpuUsage]|[userUsage]|[peakUserUsage]
port|[portType]|[index]|[avgBytes]|[avgRows]|[peakBytes]|[peakRows]|[totalBytes]|[totalRows]
```

### example of request

```
http://localhost:8080/clover/simpleHttpApi/executions_history ->
-> ?from=&to=2008-09-16+16%3A40&status=&sandbox=def&graphID=&index=&records=&returnType=DESCRIPTION
```

### example of DESCRIPTION (plain text) response

```
execution|13108|FINISHED_OK|clover|def|test.grf|2008-09-16 11:11:19|2008-09-16 11:11:58|nodeA|2.4
phase|0|38733
node|DATA_GENERATOR1|FINISHED_OK|0|0|0.0|0.0|0.0|0.0
port|Output|0|0|0|0|130|10
node|TRASH0|FINISHED_OK|0|0|0.0|0.0|0.0|0.0
port|Input|0|0|0|5|0|130|10
node|SPEED_LIMITER0|FINISHED_OK|0|0|0.0|0.0|0.0|0.0
port|Input|0|0|0|0|0|130|10
port|Output|0|0|0|5|0|130|10
execution|13107|ABORTED|clover|def|test.grf|2008-09-16 11:11:19|2008-09-16 11:11:30
phase|0|11133
node|DATA_GENERATOR1|FINISHED_OK|0|0|0.0|0.0|0.0|0.0
port|Output|0|0|0|0|0|130|10
node|TRASH0|RUNNING|0|0|0.0|0.0|0.0|0.0
port|Input|0|5|0|5|0|52|4
node|SPEED_LIMITER0|RUNNING|0|0|0.0|0.0|0.0|0.0
port|Input|0|0|0|0|0|130|10
port|Output|0|5|0|5|0|52|4
```

For `returnType==DESCRIPTION_XML` returns a complex data structure describing one or more selected executions in XML format. A schema describing XML format of the XML response is accessible on CloverETL Server URL: `http://[host]:[port]/clover/schemas/executions.xsd`

## Operation suspend

Suspends the Server or sandbox (if specified). No graphs may be executed on suspended Server/sandbox.

### parameters

Table 34.6. Parameters of suspend

parameter name	mandatory	default	description
sandbox	no	-	The text ID of a sandbox to suspend. If not specified, it suspends the whole Server.
atonce	no		If this param is set to true, running graphs from suspended Server (or just from sandbox) are aborted. Otherwise it can run until it is finished in standard way.

### returns

Result message

## Operation resume

### parameters

Table 34.7. Parameters of resume

parameter name	mandatory	default	description
sandbox	no	-	The text Id of a sandbox to resume. If not specified, the Server will be resumed.
verbose	no	MESSAGE	MESSAGE   FULL - how verbose should the possible error message be.

**returns**

Result message

## Operation sandbox\_create

This operation creates a specified sandbox. If it is a sandbox of "partitioned" or "local" type, it also creates locations by "sandbox\_add\_location" operation.

**parameters**

Table 34.8. Parameters of sandbox create

parameter name	mandatory	default	description
sandbox	yes	-	The text ID of a sandbox to be created.
path	no	-	A path to the sandbox root if the Server is running in a standalone mode.
type	no	shared	Sandbox type: shared   partitioned   local. For a standalone Server may be left empty, since the default "shared" is used.
createDirs	no	true	Switch whether to create a directory structure of the sandbox (only for a standalone Server or "shared" sandboxes in a cluster environment).
verbose	no	MESSAGE	MESSAGE   FULL - how verbose should possible error message be.

**returns**

Result message

## Operation sandbox\_add\_location

This operation adds a location to the specified sandbox. Can be only used with partitioned or local sandboxes.

**parameters**

Table 34.9. Parameters of sandbox add location

parameter name	mandatory	default	description
sandbox	yes	-	A sandbox which we want to add a location to.
nodeId	yes	-	A location attribute - a node which has direct access to the location.
path	yes	-	A location attribute - a path to the location root on the specified node.
location	no	-	A location attribute - a location storage ID. If not specified, a new one will be generated.
verbose	no	MESSAGE	MESSAGE   FULL - how verbose should possible error message be.

**returns**

Result message

---

**Operation sandbox\_remove\_location**

This operation removes a location from the specified sandbox. Only sandboxes of the partitioned or local type can have locations associated.

**parameters***Table 34.10. Parameters of sandbox add location*

parameter name	mandatory	default	description
sandbox	yes	-	Removes a specified location from its sandbox.
location	yes	-	A location storage ID. If the specified location isn't attached to the specified sandbox, the sandbox won't be changed.
verbose	no	MESSAGE	MESSAGE   FULL - how verbose should possible error message be.

**returns**

Result message

---

**Operation download\_sandbox\_zip**

This operation downloads the content of a specified sandbox as a ZIP archive.

**parameters***Table 34.11. Parameters*

parameter name	mandatory	default	description
sandbox	yes	-	A code of the sandbox to be downloaded.

**returns**

a content of a specified sandbox as a ZIP archive

**example**

```
wget --http-user=username --http-password=password http://localhost:8080/clover/simpleHttpApi/download_sandbox_zi
```

---

**Operation upload\_sandbox\_zip**

This operation uploads the content of a ZIP archive into a specified sandbox.

**parameters**

Table 34.12. Parameters

parameter name	mandatory	default	description
sandbox	yes	-	A code of the sandbox the ZIP file will be expanded to.
zipFile	yes	-	The ZIP archive file.
overwriteExisting	no	false	If <code>true</code> , the files already present in the sandbox will be overwritten.
deleteMissing	no	false	If <code>true</code> , the files not present in the ZIP file will be deleted from the sandbox.
fileNameEncoding	no	UTF-8	The encoding that was used to store file names in the ZIP archive.

**returns**

Result message

**an example of request (with using curl CLI tool (<http://curl.haxx.se/>))**

```
curl -u username:password -F "overwriteExisting=true"
-F "zipFile=@/tmp/my-sandbox.zip"
http://localhost:8080/clover/simpleHttpApi/upload_sandbox_zip
```

---

## Operation cluster\_status

This operation displays cluster's nodes list.

**parameters**

no

**returns**

Cluster's nodes list.

---

## Operation export\_server\_config

This operation exports a current server configuration in XML format.

**parameters**

Table 34.13. Parameters of server configuration export

parameter name	mandatory	default	description
include	no	all	<p>Selection of items that will be included in the exported XML file; the parameter may be specified multiple times. Possible values are:</p> <ul style="list-style-type: none"> <li>• all - include items of all types</li> <li>• users - include a list of users</li> <li>• userGroups - include a list of user groups</li> <li>• sandboxes - include a list of sandboxes</li> <li>• jobConfigs - include a list of job configuration parameters</li> <li>• schedules - include a list of schedules</li> <li>• eventListeners - include a list of event listeners</li> <li>• launchServices - include a list of launch services</li> <li>• tempSpaces - include a list of temporary spaces</li> </ul>

**returns**

Current server configuration as an XML file.

**example**

```
wget --http-user=username --http-password=password http://localhost:8080/clover/simpleHttpApi/export_server_config
```

---

## Operation import\_server\_config

This operation imports server configuration.

**parameters**

Table 34.14. Parameters of server configuration import

parameter name	mandatory	default	description
xmlFile	yes	-	An XML file with server's configuration.
dryRun	no	true	If <code>true</code> , a dry run is performed with no actual changes written.
verbose	no	MESSAGE	MESSAGE   FULL - how verbose should the response be: MESSAGE for a simple message, FULL for a full XML report.
newOnly	no	false	If <code>true</code> only new items will be imported to the Server; the items already present on the Server will be left untouched.
include	no	all	Selection of items that will be imported from the XML; the parameter may be specified multiple times. Possible values are: <ul style="list-style-type: none"> <li>• <code>all</code> - import items of all types</li> <li>• <code>users</code> - import users</li> <li>• <code>userGroups</code> - import user groups</li> <li>• <code>sandboxes</code> - import sandboxes</li> <li>• <code>jobConfigs</code> - import job configuration parameters</li> <li>• <code>schedules</code> - import schedules</li> <li>• <code>eventListeners</code> - import listeners</li> <li>• <code>launchServices</code> - import launch services</li> <li>• <code>tempSpaces</code> - import temporary spaces</li> </ul>

**returns**

Result message or XML report

**an example of request (with using curl CLI tool (<http://curl.haxx.se/>))**

```
curl -u username:password -F "dryRun=true" -F "verbose=FULL"
-F "xmlFile=@/tmp/clover_configuration_2013-07-10_14-03-23+0200.xml"
http://localhost:8080/clover/simpleHttpApi/import_server_config
```

---

## Chapter 35. JMX mBean

The CloverETL Server JMX mBean is an API that can be used for monitoring the internal status of the Server.

MBean is registered with the name:

```
com.cloveretl.server.api.jmx:name=cloverServerJmxMBean
```

---

### JMX Configuration

Application's JMX MBeans aren't accessible outside of JVM by default. It needs some changes in an application server configuration to make JMX Beans accessible.

This section describes how to configure a JMX Connector for development and testing. Thus authentication may be disabled. For production deployment, authentication should be enabled. For more information, see for example [Password Authentication](#)

Configurations and possible problems:

- [How to configure JMX on Apache Tomcat](#) (p. 278)
- [How to Configure JMX on Glassfish](#) (p. 279)
- [How to Configure JMX on WebSphere](#) (p. 279)
- [How to Configure JMX on Worker](#) (p. 280)
- [Possible Problems](#) (p. 280)

---

### How to configure JMX on Apache Tomcat

Tomcat's JVM must be executed with these parameters:

1. `-Dcom.sun.management.jmxremote=true`
2. `-Dcom.sun.management.jmxremote.port=8686`
3. `-Dcom.sun.management.jmxremote.ssl=false`
4. `-Dcom.sun.management.jmxremote.authenticate=false`
5. `-Djava.rmi.server.hostname=your.server.domain` (necessary only for remote JMX connections)

On UNIX like OS set environment variable CATALINA\_OPTS i.e. like this:

```
export CATALINA_OPTS="-Dcom.sun.management.jmxremote=true
-Dcom.sun.management.jmxremote.port=8686
-Dcom.sun.management.jmxremote.ssl=false
-Dcom.sun.management.jmxremote.authenticate=false
-Djava.rmi.server.hostname=your.server.domain.com"
```

File TOMCAT\_HOME/bin/setenv.sh (if it does not exist, you may create it) or TOMCAT\_HOME/bin/catalina.sh

On Windows each parameter must be set separately:

```
set CATALINA_OPTS=-Dcom.sun.management.jmxremote=true
set CATALINA_OPTS=%CATALINA_OPTS% -Dcom.sun.management.jmxremote.port=8686
```

```
set CATALINA_OPTS=%CATALINA_OPTS% -Dcom.sun.management.jmxremote.authenticate=false
set CATALINA_OPTS=%CATALINA_OPTS% -Dcom.sun.management.jmxremote.ssl=false
set CATALINA_OPTS=%CATALINA_OPTS% -Djava.rmi.server.hostname=your.server.domain
```

File TOMCAT\_HOME/bin/setenv.bat (if it does not exist, you may create it) or TOMCAT\_HOME/bin/catalina.bat

With these values, you can use the URL `service:jmx:rmi:///jndi/rmi://localhost:8686/jmxrmi` for connection to JMX server of JVM. No user/password is needed

## How to Configure JMX on Glassfish

Go to Glassfish admin console (by default accessible on `http://localhost:4848` with `admin/adminadmin` as user/password)

Go to the section "Configuration" > "Admin Service" > "system" and set attributes like this:

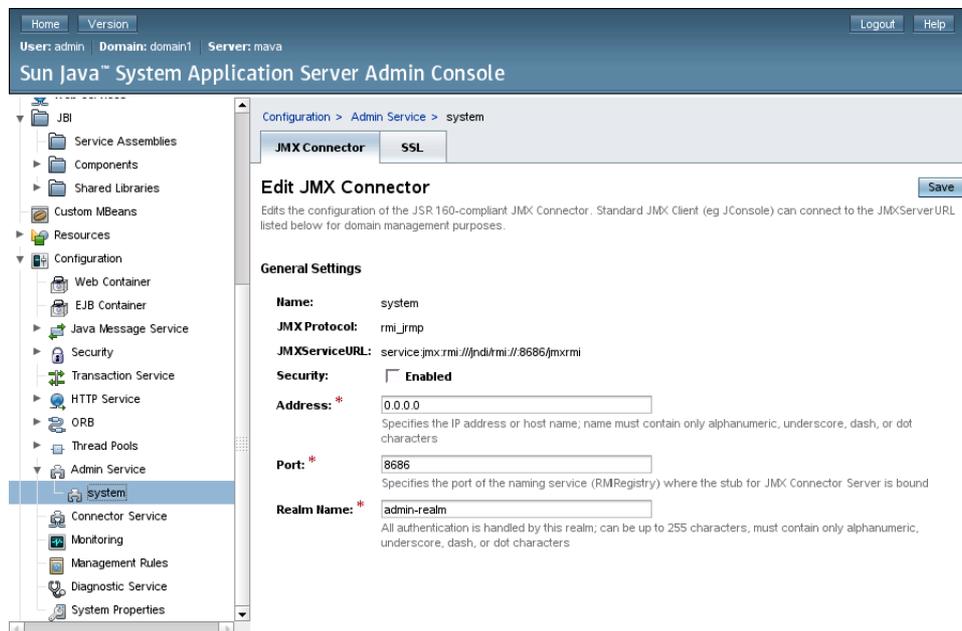


Figure 35.1. Glassfish JMX connector

With these values, you can use the URL

```
service:jmx:rmi:///jndi/rmi://localhost:8686/jmxrmi
```

for connection to JMX server of JVM.

Use `admin/adminadmin` as user/password. (`admin/adminadmin` are default glassfish values)

## How to Configure JMX on WebSphere

WebSphere does not require any special configuration, but the Clover MBean is registered with a name that depends on application server configuration:

```
com.cloveretl.server.api.jmx:cell=[cellName],name=cloverServerJmxMBean,node=[nodeName],
process=[instanceName]
```

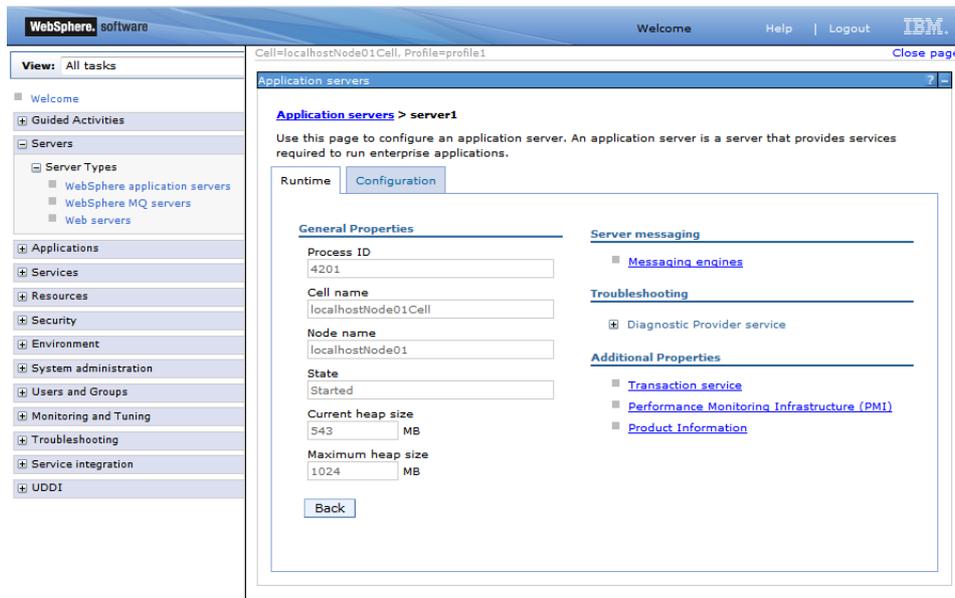


Figure 35.2. WebSphere configuration

The URL for connecting to JMX server is:

```
service:jmx:iiop://[host]:[port]/jndi/JMXConnector
```

where *host* is the host name you are connecting to and *port* is an RMI port number. If you have a default WebSphere installation, the JNDI port number will likely be 9100, depending on how many servers there are installed on one system and the specific one you want to connect to. To be sure, when starting WebSphere, check the logs for a line similar to this:

```
0000000a RMICConnectorC A   ADMC0026I: The RMI Connector is available at port 9100
```

You will also need to set on the classpath the following jar files from WebSphere home directory:

```
runtimes/com.ibm.ws.admin.client_8.5.0.jar
runtimes/com.ibm.ws.ejb.thinclient_8.5.0.jar
runtimes/com.ibm.ws.orb_8.5.0.jar
```

## How to Configure JMX on Worker

See Additional Diagnostic Tools (p. 196) for details on how to enable JMX on Worker.

## Possible Problems

- Default JMX mBean server uses RMI as a transport protocol. Sometimes RMI cannot connect remotely when one of peers uses Java version 1.6. As a solution, simply set these two system properties: `-Djava.rmi.server.hostname=[hostname or IP address] -Djava.net.preferIPv4Stack=true`

## Operations

For details about operations, see the JavaDoc of the MBean interface:

JMX API MBean JavaDoc is accessible in the running CloverETL Server instance on URL: `http://[host]:[port]/[contextPath]/javadoc-jmx/index.html`

---

## Chapter 36. SOAP WebService API

The CloverETL Server SOAP Web Service is an advanced API that provides an automation alternative to the Simple HTTP API. While most of the HTTP API operations are available in the SOAP interface too, the SOAP API provides additional operations for manipulating sandboxes, monitoring, etc.

The SOAP API service is accessible on URL:

```
http://[host]:[port]/clover/webservice
```

The SOAP API service descriptor is accessible on URL:

```
http://[host]:[port]/clover/webservice?wsdl
```

Protocol HTTP can be changed to secured HTTPS based on the web server configuration.

---

### SOAP WS Client

Exposed service is implemented with the most common binding style "document/literal", which is widely supported by libraries in various programming languages.

To create client for this API, only WSDL document (see the URL above) is needed together with some development tools according to your programming language and development environments.

JavaDoc of the WebService interface with all related classes is accessible in a running CloverETL Server instance on URL `http://[host]:[port]/[contextPath]/javadoc-ws/index.html`

If the web server has an HTTPS connector configured, the client must also meet the security requirements according to web server configuration, i.e. client trust + key stores configured properly.

---

### SOAP WS API Authentication/Authorization

Since exposed service is stateless, an authentication "sessionToken" has to be passed as a parameter to each operation. The client can obtain the authentication sessionToken by calling the `login` operation.

# Chapter 37. Launch Services

[Launch Services Overview](#) (p. 282)

[Deploying Graph in Launch Service](#) (p. 283)

[Using Dictionary in ETL Graph/Jobflow for a Launch Services](#) (p. 283)

[Configuring the Job in CloverETL Server Web GUI](#) (p. 283)

[Launch Services Authentication](#) (p. 287)

[Sending the Data to Launch Service](#) (p. 287)

[Results of the Graph Execution](#) (p. 287)

Launch Services allow you to publish a transformation graph or a jobflow as a Web Service. With Launch Services, CloverETL transformations can be exposed to provide a request-response based data interface (e.g. searches, complex lookups, etc.) for other application or directly to users.



## Important

Since 4.9.0, the **Launch Services** have been deprecated and will be removed in future versions. We recommend to use Data Services (p. 289) instead.

The Launch Services UI is hidden by default. If you still need it, you can make it visible using the [launchservices.visible](#) (p. 113) configuration property.

---

## Launch Services Overview

The architecture of a Launch Service is layered. It follows the basic design of multi-tiered applications utilizing a web browser.

Launch services let you build a user-friendly form that the user fills in and sends to the CloverETL Server for processing.

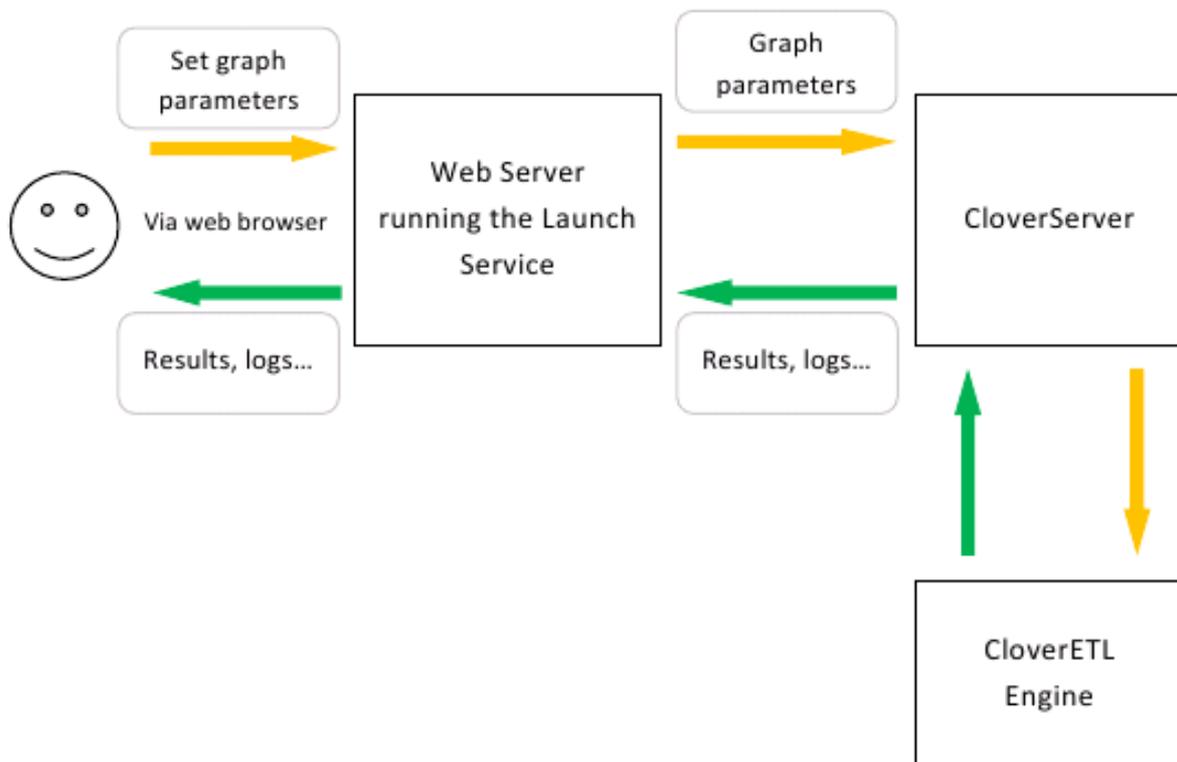


Figure 37.1. Launch Services and CloverETL Server as web application back-end

---

## Deploying Graph in Launch Service

To prepare a graph for publishing as a Launch Service, keep this in mind during the design process:

1. You can define a graph/jobflow listeners to create parameterized calls. Parameters are passed to the graph as Dictionary entries – design the graph so that it uses the Dictionary as input/output for parameters (e.g. file names, search terms, etc.).
2. The graph will need to be published in the Launch Services section, where you provide the configuration and binding for parameters to dictionary entries.

---

## Using Dictionary in ETL Graph/Jobflow for a Launch Services

A graph or a jobflow published as a service usually means that the caller sends request data (parameters or data) and the transformation processes it and returns back the results.

In a Launch Service definition, you can bind service's parameters to Dictionary entries. These need to be predefined in the transformation.

Dictionary is a key-value temporary data interface between the running transformation and the caller. Usually, although not restricted to, Dictionary is used to pass parameters in and out the executed transformation.

For more information about Dictionary, read the “Dictionary” section in the CloverETL Designer User's Guide.

---

## Passing Files to Launch Services

If a Launch service is designed to pass an input file to a graph or jobflow, the input dictionary entry has to be of type `readable.channel`.

---

## Configuring the Job in CloverETL Server Web GUI

Each Launch Service configuration is identified by its name, user, and group restriction. You can create several configurations with the same name, which is valid as long as they differ in their user or group restrictions.

User restrictions can then be used to launch different jobs for different users, even though they use the same launch configuration (i.e. name). For example, developers may want to use a debug version of the job, while the end customers will want to use a production job. The user restriction can also be used to prohibit certain users from executing the launch configuration completely.

Similarly, a group restriction can be used to differentiate jobs based on the user's group membership.

If multiple configurations match the current user/group and configuration name, the most specific one is picked. (The user name has a higher priority than the group name.)

---

## Adding New Launch Configuration

Use the **New launch configuration** button on the Launch Services tab to create a new Launch Service.

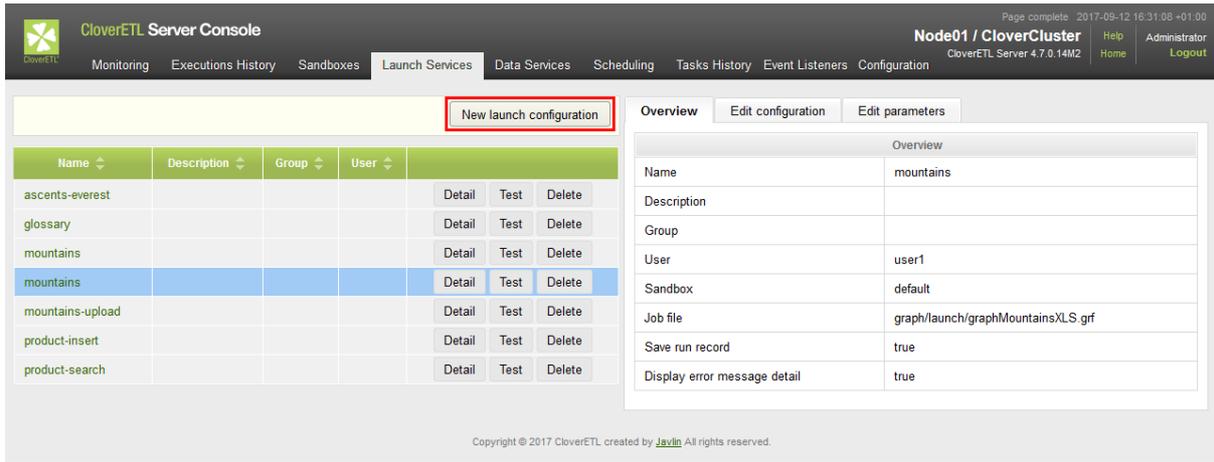


Figure 37.2. Launch Services section

The name is the identifier for the service and will be used in the service URL. Then, select a sandbox and either a transformation graph or a jobflow that you want to publish.

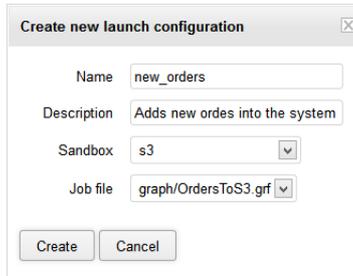


Figure 37.3. Creating a new launch configuration

Once you create the new Launch Service, you can set additional attributes like:

1. User and group access restrictions and additional configuration options (Edit Configuration)
2. Bind Launch Service parameters to Dictionary entries (Edit Parameters)

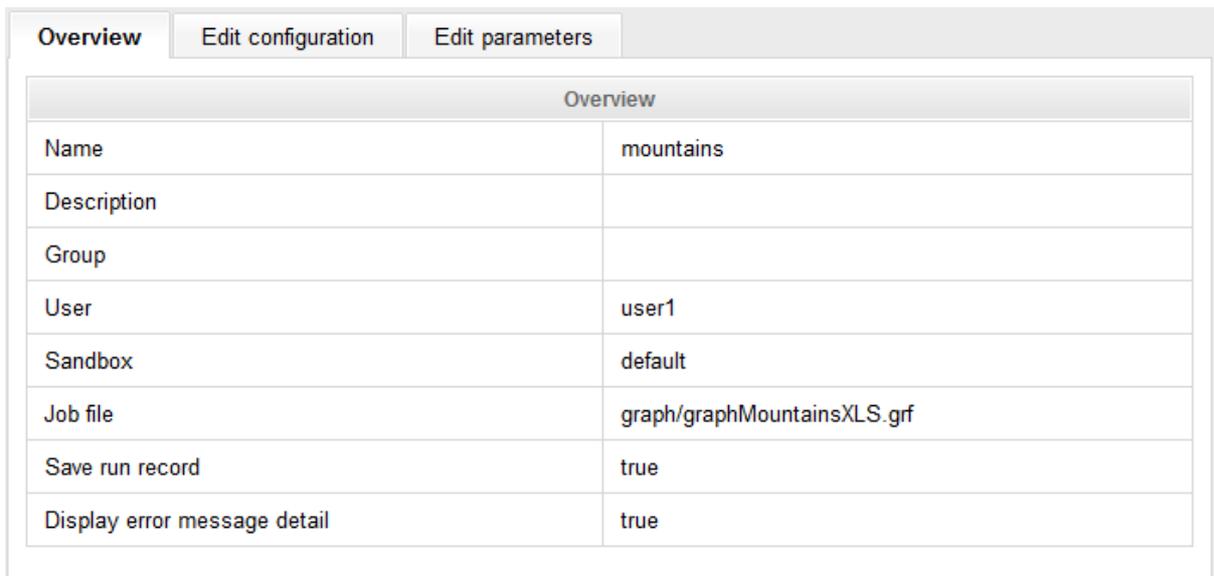


Figure 37.4. Overview tab

The Overview tab shows the basic details about the launch configuration. These can be modified in the Edit Configuration tab:

### Edit Configuration

The screenshot shows the 'Edit configuration' tab with the following fields and values:

- Name: mountains
- Description: (empty)
- Group: (dropdown menu)
- User: user1
- Sandbox: default
- Job file: graph/graphMountainsXLS.grf
- Save run record:
- Display error message detail:

An 'Update' button is located at the bottom left of the form.

Figure 37.5. Edit Configuration tab

Editing configurations:

- *Name* - The name (identifier) under which the configuration will be accessible from the web.
- *Description* - The description of the configuration.
- *Group* - Restricts the configuration to a specific group of users.
- *User* - Restricts the configuration to a specified user.
- *Sandbox* - The CloverETL Sandbox where the configuration will be launched.
- *Job file* - Selects the job to run.
- *Save run record* - If checked, the details about the launch configuration will be stored in the Execution History. Uncheck this if you need to increase performance – storing a run record decreases response times for high frequency calls.
- *Display error message detail* - Check this if you want to get a verbose message in case the launch fails.

### Edit Parameters

The **Edit parameters** tab can be used to configure parameter mappings for the launch configuration. The mappings are required for the Launch Service to be able to correctly assign parameters values based on the values sent in the launch request.

Overview Edit configuration **Edit parameters**

Create parameter

Dictionary entry name

HTTP request parameter name

HTTP request parameter required

Pass HTTP request body  ?

Trim parameter value

Empty parameter is null

Parameter value format string  ?

Parameter value locale  ?

Date parameter time zone  ?

Default parameter value

Pass value as graph parameter  ?

Figure 37.6. Creating new parameter

To add a new parameter binding, click on the **Add parameter** button. Every required graph/jobflow listener property defined by the job needs to be created here.

Overview Edit configuration **Edit parameters**

Dictionary entry name	HTTP request parameter name	HTTP request parameter required	Pass HTTP request body	Pass value as graph parameter	Default parameter value		
heightMin	heightMin	false	false	false		Delete	Detail

Figure 37.7. Edit Parameters tab

You can set the following fields for each property:

- *Dictionary entry name* - The name of the Dictionary entry defined in the graph/jobflow that you want to bind.

- *HTTP request parameter name* - The name of this property as it will be visible in the published service. This name can be different from **Name**.
- *HTTP request parameter required* - If checked, the parameter is mandatory and an error will be reported if it is omitted.
- *Pass HTTP request body* - If checked, the request body is set to dictionary entry as a readable channel.
- *Pass value as graph parameter* - If checked, the property value will be passed to the job also as a parameter (`{ParameterName}`, where `ParameterName` is equal to `Name`). This lets you use the parameter anywhere in the job definition (not just places that support Dictionary). However, parameters are evaluated during job initialization. Thus, such a job cannot be pooled which decreases performance for high frequency repetitive calls to the service. In this case, consider redesigning the transformation to use Dictionary instead, allowing for pooling.
- *Default parameter value* - The default value applied in case the parameter is omitted in the launch request.

---

## Launch Services Authentication

If you are using launch services, you have two ways how to be logged in: using form-based authentication of Server console or HTTP basic authentication of Launch services.

The form-based authentication of Server console enables the user to create or modify Launch services. If you are logged in this way, you act as an administrator of Launch services.

To insert data into the Launch service form, you should be logged in using HTTP basic authentication. Follow the link to the Launch service form and web browser will request your credentials. If you are logged in using HTTP basic authentication, you act as an user of Launch services forms.

---

## Sending the Data to Launch Service

A launch request can be sent via HTTP GET or POST methods. A launch request is simply a URL which contains the values of all parameters that should be passed to the job. The request URL is composed of several parts:

(You can use a Launch Services test page, accessible from the login screen, to test drive Launch Services.)

```
[Clover Context]/launch/[Configuration name]?[Parameters]
```

- `[Clover Context]` is the URL to the context in which the CloverETL Server is running. Usually this is the full URL to the CloverETL Server (for example, for CloverETL Demo Server this would be `http://server-demo.cloveretl.com:8080/clover`).
- `[Configuration name]` is the name of the launch configuration specified when the configuration was created. In our example, this would be set to “mountains” (case-sensitive).
- `[Parameters]` is the list of parameters the configuration requires as a query string. It is a URL-encoded [RFC 1738] list of `name=value` pairs separated by the “&” character.

Based on the above, the full URL of a launch request for our example with mountains may look like this: `http://server-demo.cloveretl.com:8080/clover/launch/NewMountains?heightMin=4000`. In the request above, the value of `heightMin` property is set to 4000.

---

## Results of the Graph Execution

After the job terminates, the results are sent back to the HTTP client as content of an HTTP response.

Output parameters are defined in the job’s Dictionary. Every Dictionary entry marked as “Output” is sent back as a part of the response.

Depending on the number of output parameters, the following output is sent to the HTTP client:

- *No output parameters* - Only a summary page is returned. The page contains: when the job was started, when it finished, the user name, and so on. The format of the summary page cannot be customized.
- *One output parameter* - In this case, the output is sent to the client as in the body of the HTTP response with its MIME content type defined by the property type in Dictionary.
- *Multiple output parameters* - In this case, each output parameter is sent to the HTTP client as a part of the multipart HTTP response. The content type of the response is either multipart/related or multipart/x-mixed-replace, depending on the HTTP client (the client detection is fully automatic). The multipart/related type is used for browsers based on Microsoft Internet Explorer and the multipart/x-mixed-replace is sent to browsers based on Gecko or Webkit.

Launch requests are recorded in the log files in the directory specified by the `launch.log.dir` property in the CloverETL Server configuration. For each launch configuration, one log file named `[Configuration name]#[Launch ID].log` is created. For each launch request, this file will contain only one line with following tab-delimited fields:

(If the property `launch.log.dir` is not specified, log files are created in the temp directory `[java.io.tmpdir]/cloverlog/launch` where `java.io.tmpdir` is system property)

- *Launch start time*
- *Launch end time*
- *Logged-in user name*
- *Run ID*
- *Execution status* FINISHED\_OK, ERROR or ABORTED
- *IP Address* of the client
- *User agent* of the HTTP client
- *Query string* passed to the Launch Service (full list of parameters of the current launch)

If the configuration is not valid, the same launch details are saved into the `_no_launch_config.log` file in the same directory. All unauthenticated requests are saved to the same file as well.

# Chapter 38. Data Services

[Overview](#) (p. 289)

[User Interface](#) (p. 290)

[Using Data Services](#) (p. 301)

## Overview

**Data Services** allow you to deploy a web service. The architecture of the data service is described in the documentation on Designer. This section describes the server-side functionality of the Data Services.

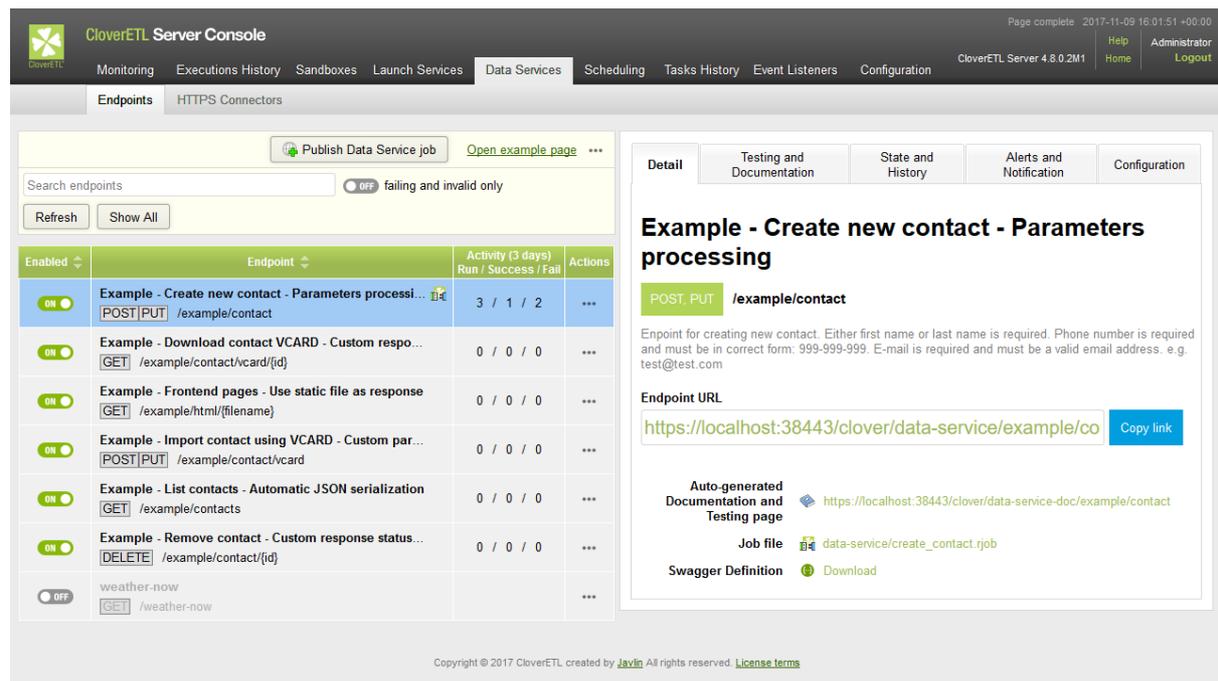


Figure 38.1. Data Services

The **Data Service** can be accessible via HTTP or HTTPS protocol.

By default, it accepts requests via HTTP protocol on the same port as the Server. This is suitable for Data Services that do not require authentication.

If you need a secure connection, you should configure Data Service to listen on HTTPS: create an HTTPS Connector and use it in one or more data service endpoints. This way, you can configure a service listening on HTTPS port without restarting the Server. You can create multiple HTTPS Connectors and use it, for example, per consumer service. One Data Service endpoint can use only one HTTPS Connector.

The advantage of HTTPS Connectors is that you do not need to restart the Server to allow HTTPS connection.

Data Service can send you a notification in case of failure. You can set the threshold (number of subsequent failures or percentage) and way of notification (in the Server's UI or via email).

To investigate failed requests, you can use history of the particular endpoint. Optionally, you can set the Data Service endpoint run to be recorded in **Executions History**.

## User Interface

Data Services user interface contains two main tabs: [Endpoints](#) (p. 291) and [HTTPS Connectors](#) (p. 298).

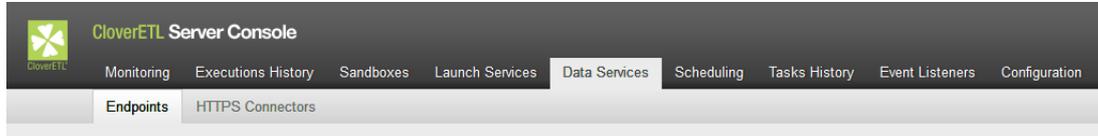


Figure 38.2. Data Service tabs

## Endpoints

**Endpoints** tab consists of useful [buttons](#) (p. 290) in the top, [list of data services](#) (p. 291) and tabs with configuration of the particular data service.

Enabled	Endpoint	Activity (3 days) Run / Success / Fail	Actions
ON	Example - Create new contact - Parameters processing [POST PUT] /example/contact	0 / 0 / 0	...
ON	Endpoint is failing. Example - Download contact VCARD - Custom response body serialization [GET] /example/contact/vcard/{id}	6 / 4 / 2	...
ON	Example - Frontend pages - Use static file as response [GET] /example/html/{filename}	0 / 0 / 0	...
ON	Example - Import contact using VCARD - Custom parsing of request body [POST PUT] /example/contact/vcard	0 / 0 / 0	...
ON	Example - List contacts - Automatic JSON serialization [GET] /example/contacts	0 / 0 / 0	...
OFF	Example - Remove contact - Custom response status code [DELETE] /example/contact/{id}		...

Figure 38.3. Endpoints

## Buttons

In the top of **Endpoints** tab, there are **Publish Data Service job** button, link to **Open example page** and a three-dot-button to open a menu.



Figure 38.4. Data Service Endpoints

The three-dot-button has the following menu.



Figure 38.5. Three-dot-button menu

*Open Endpoint Documentation Catalog* displays list of data services.

*Unpublish Data Services Examples* unpublishes the example Data Services.

*Export Data Services Configuration* exports the Data Services configuration.

*Import Data Services Configuration* imports the Data Services configuration.

## List of Data Services

The **Data Services** tab contains list of all data services on the Server.

Enabled	Endpoint	Activity (3 days) Run / Success / Fail	Actions
<input checked="" type="radio"/>	Example - Create new contact - Parameters processing [POST PUT] /example/contact	0 / 0 / 0	...
<input checked="" type="radio"/>	Endpoint is failing. Example - Download contact VCARD - Custom response body serialization [GET] /example/contact/vcard/{id}	6 / 4 / 2	...
<input checked="" type="radio"/>	Example - Frontend pages - Use static file as response [GET] /example/html/{filename}	0 / 0 / 0	...
<input checked="" type="radio"/>	Example - Import contact using VCARD - Custom parsing of request body [POST PUT] /example/contact/vcard	0 / 0 / 0	...
<input checked="" type="radio"/>	Example - List contacts - Automatic JSON serialization [GET] /example/contacts	0 / 0 / 0	...
<input type="radio"/>	Example - Remove contact - Custom response status code [DELETE] /example/contact/{id}		...

Figure 38.6. List of Data Services

The **Search endpoints** text field serves to enter a keyword to find the corresponding Data Services. The keyword is searched in a title of a Data Service, in its method and in the path that the Data Service uses.

The **failing and invalid only** button serves to display endpoints that are failing or incorrectly configured. This can be combined with the previous filter.

The **Refresh** button serves to apply the filter from **Search endpoints** text field.

The **Show All** button serves to clear the filter and display all endpoints.

- The button in the left column serves to *enable* or *disable* the service. Disabling the Data Service serves to temporarily disable the service, e.g. due to maintenance. The disabled Data Service returns the HTTP status code 503. Icon decorators in the first column indicate these states:

- The service is enabled.

- The service is disabled.

- In the second column, there is *Endpoint title*, *method* or *methods* and part of *endpoint URL*.

If the endpoint is failing, the row is highlighted with an exclamation mark icon and the red *Endpoint is failing* text works as a link to the **State and History** tab. Icon decorators indicate these endpoint states:

- The Data Service does not require authentication.

- The Data Service saves the job execution record in Execution History.

- ❗ - The Data Service is marked as failing.
- 🔑 - The Data Service is available on HTTPS.
- The third column contains query statistics from last three days. There is a number of all queries, successfully called queries and queries that failed.
- The last column contains a menu with actions to be performed on the data service: *Detail*, *Disable*, *Unpublish*, *Reset Endpoint State* and *Swagger file*.

## Detail

The **Detail** tab contains overview of the particular endpoint. To display the **Detail** tab, click the particular line in the list of endpoints or select the **Detail** option from the menu in the 4th column.

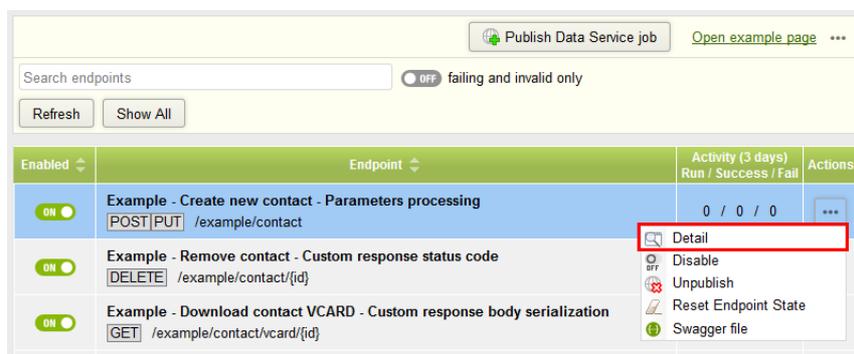


Figure 38.7. Selecting the particular Data Service

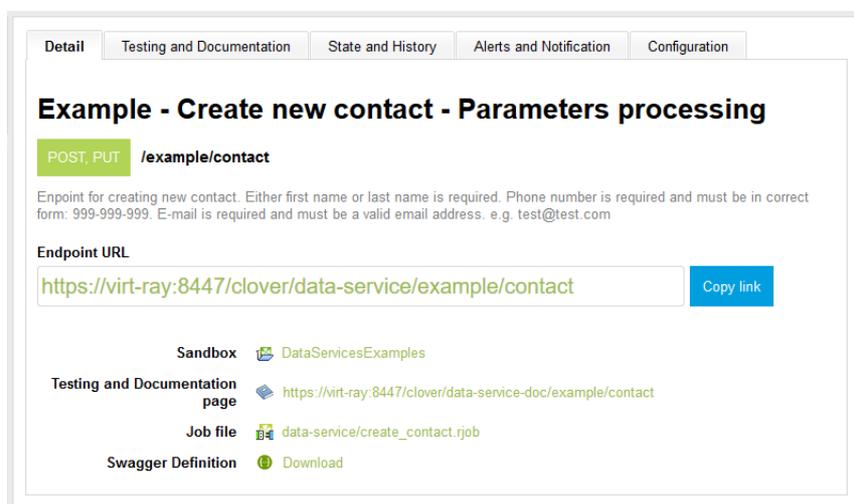


Figure 38.8. Data Service Detail tab

*Endpoint title* is in the top of the tab's pane. It is the endpoint title specified in Designer on the *Endpoint Configuration* tab.

The green rectangle contains a list of *Endpoint methods*.

Next to the green rectangle, there is a configurable part of *Endpoint URL*. It can be set from Designer.

The paragraphs below *Description* contain a user-defined description of the Data Service. It can be set in Designer.

*Endpoint URL* is URL of the endpoint. This URL serves the requests.

The **Copy link** button copies the link to the clipboard. This functionality is not supported in Internet Explorer 8.

*Sandbox* is the sandbox containing the data service .xjob file.

*Testing and Documentation* page links documentation of the endpoint. You can test the service there. This URL can be passed down to consumer of the service. The consumer can use the information from this URL to implement the client system.

*REST job file* is a file name and path relative to the sandbox.

The *Swagger file definition* allows you to download a [swagger file](#) with the definition of the Data Service.

## Testing and Documentation

The *Testing and Documentation* tab displays a user-defined documentation to the data service. The testing of the service is accessible under the **Try it out** button.

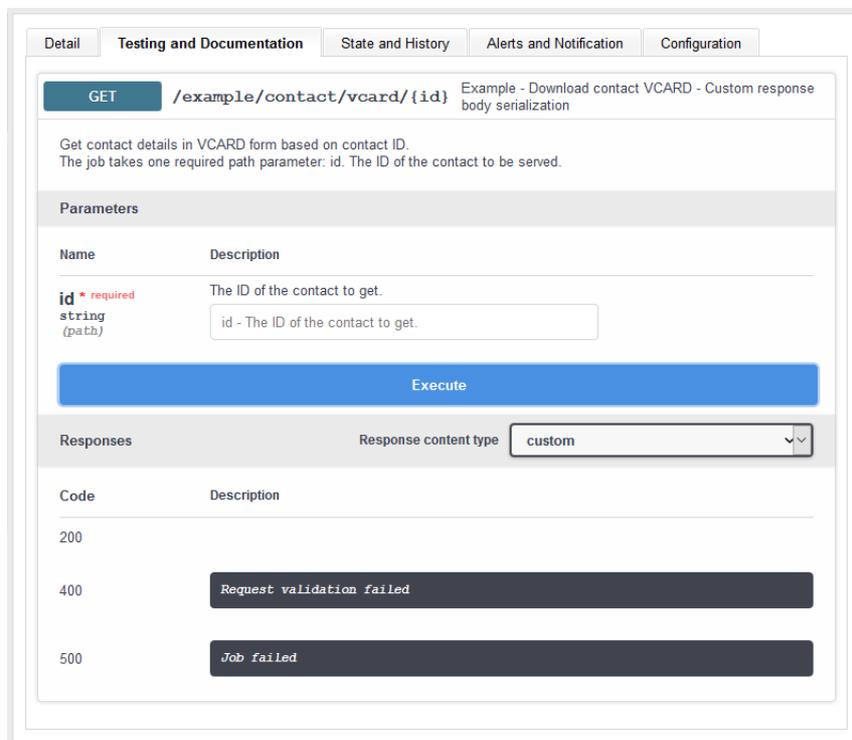


Figure 38.9. Data Service - Testing and Documentation tab

If you call the service from the Server UI, the Data Service will save its run record to the Executions History.

## State and History

The **State and History** shows invocation history of the particular Data Service. It contains a summary of the endpoint state in the top and a list of query details in the bottom. If the job is configured to save a record in execution history, the list also contains link to the **Execution History**.

You can filter records based on the time interval or you can list only the failures.

Here you can reset the state of the data service. For example, the data service endpoint was failing, you fixed it and you would like to be notified if it fails again.

HTTP Status	Time	Detail
✓ 201	2017-12-20 12:24:22	IP: 10.0.3.2 User: clover
✓ 201	2017-12-20 12:24:14	IP: 10.0.3.2 User: clover
✓ 201	2017-12-20 12:24:06	IP: 10.0.3.2 User: clover
✓ 201	2017-12-20 12:23:16	IP: 10.0.3.2 User: clover
! 400	2017-12-20 12:22:42	IP: 10.0.3.2 User: clover

Figure 38.10. Data Service - State and History

## Alerts and Notification

[No failure notification](#) (p. 294)

[Failure](#) (p. 294)

[Threshold specification](#) (p. 294)

[Any failure](#) (p. 294)

[Threshold](#) (p. 295)

[Failure Filtering](#) (p. 295)

[Failure Notification](#) (p. 296)

[E-mail Notification](#) (p. 296)

The **Alerts and Notification** tab serves to set the threshold meaning the failure of the Data Service endpoint and way to notify you about it. You can set when the endpoint is marked as failing or disable this notification completely.

### No failure notification

The **Never mark endpoint as failing** disables the failure notification in the Server UI. If you set this option and the request to the endpoint fails, there will be no red circle notification. Only the number of failures in activity column in the list of Data Service endpoints will be increased.

### Failure

In this context, any response state from 4xx and 5xx range is considered as a failure.

### Threshold specification

You can set threshold to

- any failure
- percentage of unsuccessful queries within interval
- fixed number of errors in a row

### Any failure

The **Any failure will mark endpoint as failing** option considers the endpoint as failing even if a single failure occurs.

This choice is suitable for infrequently called Data Service endpoints.

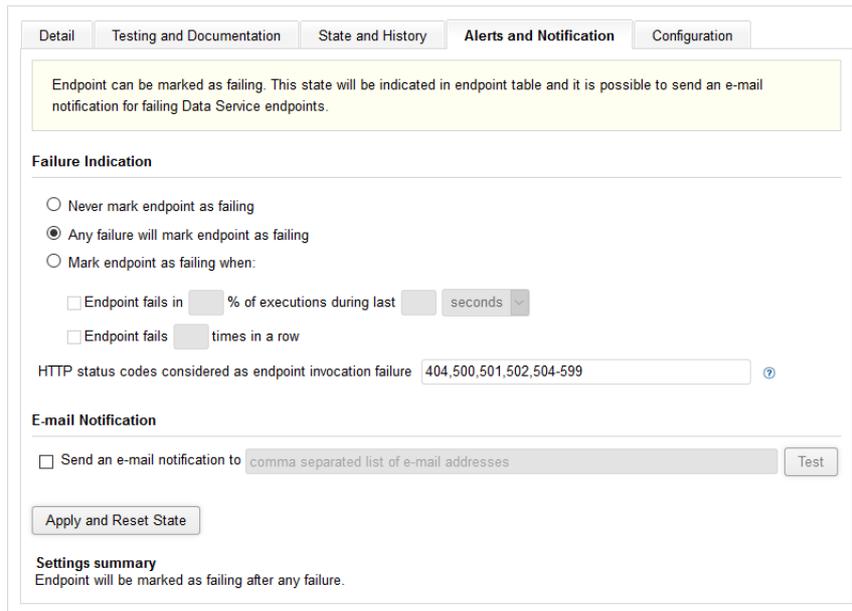


Figure 38.11. Data Service - Alerts and Notification

## Threshold

The **Mark endpoint as failing when** option sets the endpoint as failing when a threshold is reached. The threshold can be specified as a *percentage of jobs is failing* or as a *number of jobs in row is failing*.

This choice is suitable for frequently called Data Service endpoints.

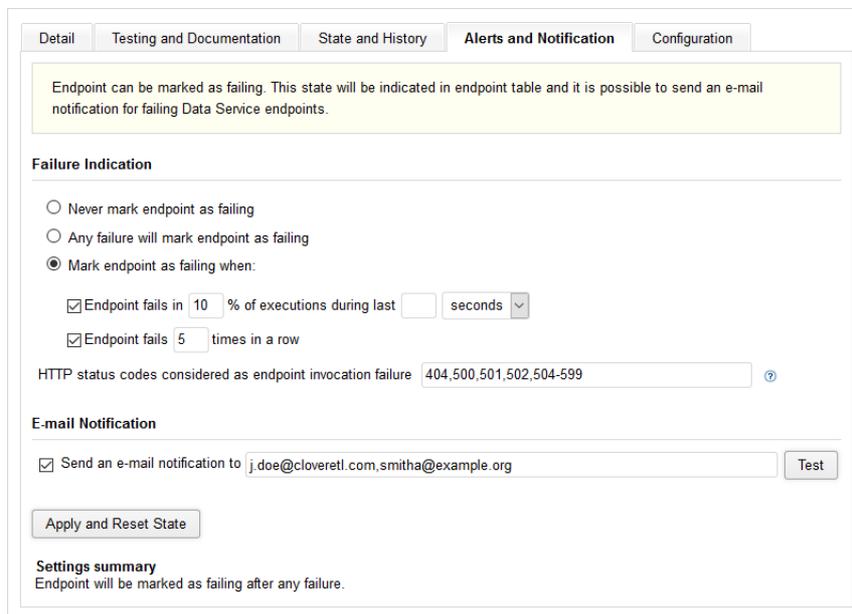


Figure 38.12. Data Service - Alerts and Notification

## Failure Filtering

You can also select HTTP status codes which should be considered by **CloverETL Server** as an endpoint invocation failure.

Select the HTTP status codes by entering individual codes or ranges of codes separated by commas. By default, codes 404,500,501,502,504-599 are considered as an endpoint invocation failure. Leaving the field blank means that **CloverETL Server** considers all HTTP status codes from the range 400-599 as invocation failure.



## Note

By default, **CloverETL Server** does **not** consider HTTP status code 503 as a failure, because the code is returned in the case of invocation of a manually disabled endpoint.

## Failure Notification

If the Data Service endpoint fails, it is shown in the list of Data Services.

Enabled	Endpoint	Activity (3 days) Run / Success / Fail	Actions
ON	Endpoint is failing. Get mountains GET /mountains	1 / 0 / 1	...

Figure 38.13. Data Service - Health status

The number of failing Data Service endpoints is shown in the main menu.

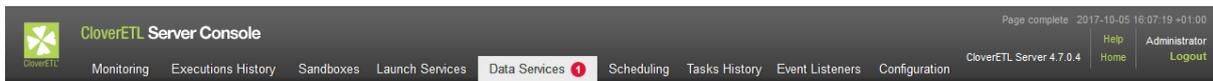


Figure 38.14. Some Data Service is failing

## E-mail Notification

You can also set an email notification. This email notification works additionally to the notification in the Server UI. It sends an email when the endpoint's state changes to *failing*. An email is also sent if the endpoint was failing and you manually reset the endpoint state.

**E-mail Notification**

Send an e-mail notification to

Figure 38.15. E-mail Notification

With the **Test** button, you can send a testing email to the addresses of the recipients.

To have the email notifications working, you should have a connection to an SMTP server set up.

## Configuration

The **Configuration** tab allows you to disable the endpoint authentication or to enable saving records in Executions History.

Detail   Testing and Documentation   State and History   Alerts and Notification   **Configuration**

**HTTP Authentication**

Authentication method: Do not require authentication

Run as: clover

**SSL/TLS Secured Transport (HTTPS)**

HTTPS connector: <none>

**Execution History**

Save job execution record in Execution History

Figure 38.16. Data Service

The Data Service can be configured to require credentials or not. If the Data Service does not require credentials, the user to run it should be set in its configuration with the **Run As** option.

You can configure the data service to **save job execution record in Execution History**. The saving job execution record has a performance impact. Use this option only for:

- infrequently called endpoints
- endpoints that are not in production environment
- endpoints to be debugged

## Catalog of Services

The *Catalog of services* is a list of data services allowing the user to view the documentation and test the service.

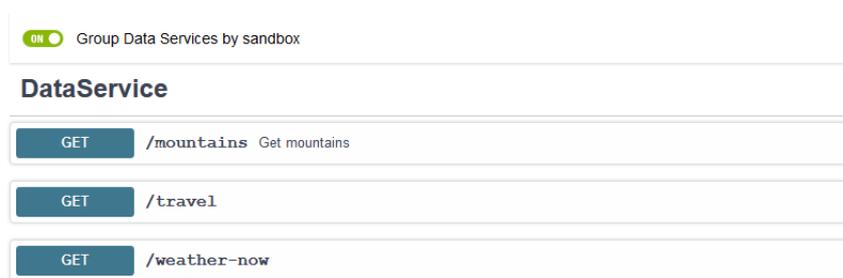


Figure 38.17. Glogal Catalogue of Services

The details can be accessed by clicking the header. The first click displays the details, the second one fold the details back.

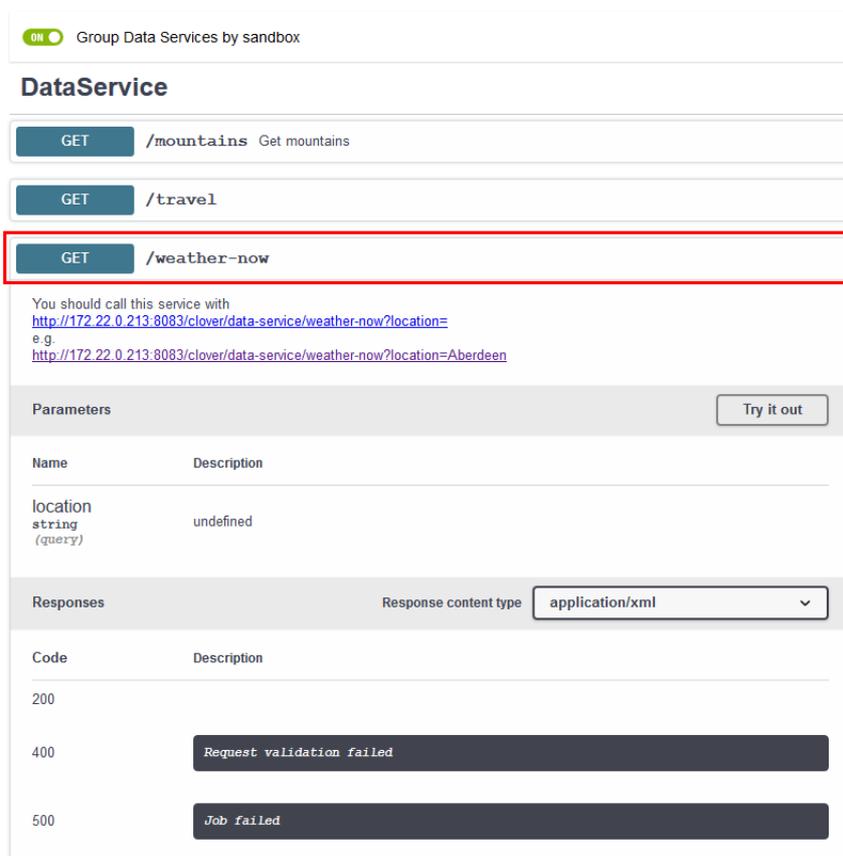


Figure 38.18. Glogal Catalogue of Services

In the *Catalog of Services*, the end points can be grouped by sandbox or ordered by URL.

## Built-in Data Service Examples

CloverETL Server contains built-in set of Data Service examples. The Data Service examples can be published from the **Data Services** tab.



Figure 38.19. Data Services - Publishing the examples

The published examples are displayed among the others in the list of Data Services.

Enabled	Endpoint	Activity (3 days) Run / Success / Fail	Actions
ON	Create contact [POST PUT] /example/contact	0 / 0 / 0	...
ON	Delete contact [DELETE] /example/contact/{id}	0 / 0 / 0	...
ON	Get contact VCARD [GET] /example/contact/vcard/{id}	0 / 0 / 0	...
ON	Get contacts [GET] /example/contacts	0 / 0 / 0	...
ON	Get static file [GET] /example/html/{filename}	0 / 0 / 0	...
ON	Import contact using VCARD [POST PUT] /example/contact/vcard	0 / 0 / 0	...

Figure 38.20. Data Services - Published the examples

## HTTPS Connectors

The **Data Service** can be accessible via HTTPS. The configuration of HTTPS is in **Data Services** → **HTTPS Connectors**.

Enabled	Name	Actions
ON	Customers Connector Port: 8448 Key store: /opt/CloverETL_Server/conf/certs/connector-customers.jks	...
ON	Reports Connector Port: 8447 Key store: /opt/CloverETL_Server/conf/certs/connector-reports.jks	...
ON	Sales Connector Port: 8449 Key store: /opt/CloverETL_Server/conf/certs/connector-sales.p12	...

**Reports Connector**

This HTTPS Connector is running on port 8447.

Port: 8447

Key store: /opt/CloverETL\_Server/conf/certs/connector-reports.jks

Figure 38.21. HTTPS Connectors

As a key store, we support the Java key store (.jks) and PKCS 12 key store (.p12 or .pfx) formats.

As a trust store, we support the Java key store (.jks) format.

On the left hand side, there is a list of available HTTPS Connectors. On the right, there are details of the connector selected from the list. The **New HTTPS Connector** button creates a new HTTPS Connector.

## List of HTTPS Connectors

The list of HTTPS Connectors shows available connectors. You can change the order by clicking on **Name** or **Enabled** in the header.

The button in the first column enables or disables the connector. Disabling the connector that is being used by Data Service makes the Data Service invalid.

The middle column shows the connector's name, port, path to the key store and path to the trust store.

The last column contains a menu with actions to display details on the connector and to delete it. The menu is accessible under the three-dot button.

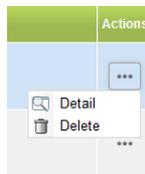


Figure 38.22. Tree-dot button menu

## New HTTPS Connector

The **New HTTPS Connector** tab serves to create a new HTTPS Connector that can be used by one or more Data Services. One Data Service can use only one HTTPS Connector.

 A screenshot of a web-based configuration form for a new HTTPS Connector. The form has two tabs: 'Detail' (selected) and 'Edit'. Below the tabs is a yellow informational box stating 'HTTPS Connector allows you to use Data Services via HTTPS protocol.' The form contains several sections:
 

- A checkbox labeled 'HTTPS Connector enabled' which is checked.
- A 'Name' text input field containing 'Reports Connector'.
- A 'Port' text input field containing '8447'.
- A section titled 'Key Store File' containing:
  - A checkbox 'Key store is located in sandbox' which is unchecked.
  - A 'Key store path' text input field containing '/opt/CloverETL\_Server/conf/certs/connector-reports.jks'.
  - A 'Key store password' text input field with masked characters.
  - A 'Key password' text input field with masked characters.
- A section titled 'Trust Store File' containing:
  - A checkbox 'Trust store is located in sandbox' which is unchecked.
  - A 'Trust store path' text input field.
  - A 'Password' text input field.

 At the bottom left of the form is an 'Apply' button.

Figure 38.23. HTTPS Connectors

The **HTTPS Connector enabled** checkbock enables or disables the HTTPS Connector to listen on the specified port. Stopping an HTTPS Connector that is being used by a Data Service makes the Data Service invalid.

**Name** is a name of the HTTPS Connector. The name should be unique. It is displayed in the list of HTTPS Connectors on the Endpoint's Configuration tab.

**Port** is a TCP port used by the HTTPS Connector. The port must not be occupied by another HTTPS Connector or any other program. If the Data Service is Deployed on CloverETL Cluster, it listens on this port on all cluster nodes. If you use a firewall, set it to allow incoming connections to this port. If you use SELinux, it must be configured to allow CloverETL Server to use this TCP port.

The **Key store is located in sandbox** checkbox switches between absolute paths to key store and paths relative to the Server sandbox. If you tick this option, **Sandbox** and **Key store** items are displayed. Otherwise, you will see **Key store path**. The recommended way is to store the key stores out of the sandbox.

**Sandbox** is the sandbox with the key store.

**Key store** is the key store within the sandbox.

**Key store path** is the absolute path to the Java key store. You can use environment variables, system properties of JVM and configuration parameters of the Server as a part of the path. Usually, you will use `${sandboxes.home}` here.

**Key store password** is the password to the Java key store.

**Key password** is the password to the key in the key store.

The **Trust store is located in sandbox** checkbox switches between absolute paths to trust store and paths relative to the Server sandbox. If enabled, you can enter **Sandbox** and **Trust store** options. Otherwise, you will see **Trust store path**.

**Sandbox** is the sandbox containing the trust store.

**Trust store** is the trust store within the sandbox.

**Trust store path** is the absolute path to the trust store. You can use environment variables, system properties of JVM and configuration parameters of the Server as a part of the path. Ususally, you will use `${sandboxes.home}` here.

**Password** is the password to the trust store.

To create a data service listening on HTTPS, you need a keystore with a server certificate. You can create one with the following command.

```
keytool -keystore service.jks -genkey -keyalg rsa -keysize 3072 -alias serverName
```

As a key store, we support the Java key store (`.jks`) and PKCS 12 key store (`.p12` or `.pfx`) formats.

As a trust store, we support the Java key store (`.jks`) format.

For security reasons, we recommend you to put the keystore outside the Server sandbox.

## Using Data Services

- [Deploying Data Service](#) (p. 301)
- [Publishing and Unpublishing Data Service from Sandbox](#) (p. 301)
- [Publishing Data Service Examples](#) (p. 302)
- [Changing Data Service to Anonymous](#) (p. 303)
- [Running Data Service on HTTPS](#) (p. 303)
- [Running Data Service on HTTPS on Cluster](#) (p. 304)
- [Monitoring Data Service](#) (p. 305)
- [Testing Data Service](#) (p. 305)
- [Performance Tuning](#) (p. 307)
- [Exporting Data Service Configuration](#) (p. 308)
- [Importing Data Service Configuration](#) (p. 308)
- [Avoiding Premature Marking of Data Service as Failing](#) (p. 308)
- [Looking up Particular Data Service](#) (p. 308)
- [Resetting State of Failing Data Service Endpoint](#) (p. 309)

## Deploying Data Service

To deploy **Data Service** from the Server, go to **Data Services** tab, click **Publish Data Service job** and choose a sandbox and .rjob file.

Figure 38.24. Publishing Data Service job

You can choose between Data Service that requires authentication and that does not. In the latter case, the Data Service will run under the specified account.

Figure 38.25. Publishing Data Service job that does not require authentication

## Publishing Multiple Jobs

To deploy multiple jobs, tick the **Publish another** checkbox. After deploying one job, the dialog for publishing Data Service is displayed again to let you enter the next one.

## Publishing and Unpublishing Data Service from Sandbox

You can deploy Data Service directly from a sandbox. To do so, you need read access to the sandbox and **List Data Services** and **Manage Data Services** privileges.

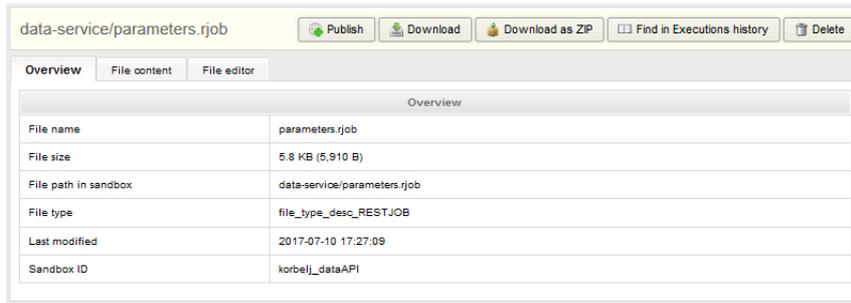


Figure 38.26. Publishing Data Service from Sandbox

You also undeploy Data Service from here. You need read access to the sandbox and **List Data Services** and **Manage Data Services** privileges.

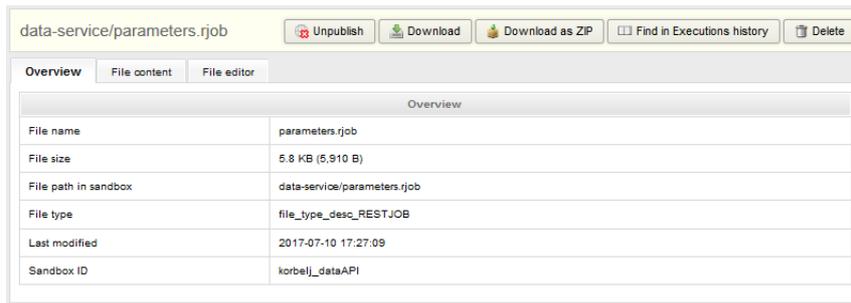


Figure 38.27. Unpublishing Data Service from Sandbox

## Publishing Data Service Examples

CloverETL Server contains a built-in set of Data Service examples. These examples are not deployed by default.

The Data Service examples can be deployed directly from the **Data Services** tab. If you do not have any Data Service deployed, click the **Publish Data service Examples** link.

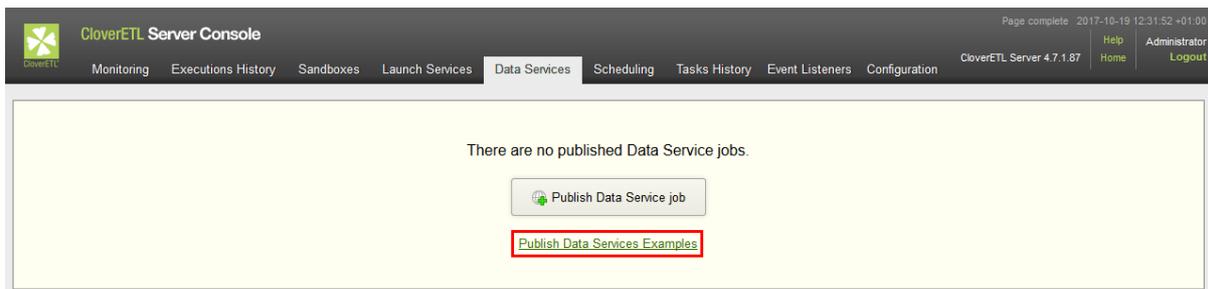


Figure 38.28. Publishing Data Service examples

If there is an existing Data Service, the button to publish examples is in the menu accessible under the three-dot button.

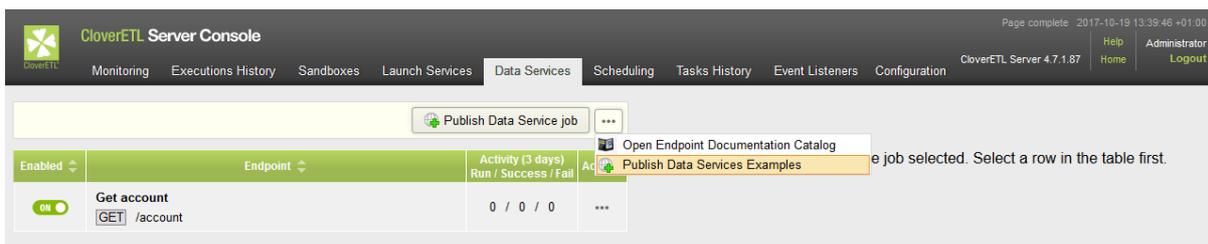


Figure 38.29. Publishing Data Service examples - II

See [Built-in Data Service Examples](#) (p. 298).

## Changing Data Service to Anonymous

By default, the Data Service requires a client to send the credentials. To create the Data Service that does not require authentication, switch to the **Configuration** tab in the Data Service **Detail** pane and change **Authentication method** to *Do not require authentication*. The Data service runs with privileges of an existing user; therefore, you should set the *Run as* field to the suitable user. This user should have permissions necessary to run the Data Service.

The screenshot shows the 'Configuration' tab of a Data Service configuration page. Under the 'HTTP Authentication' section, the 'Authentication method' dropdown is set to 'Do not require authentication' and the 'Run as' dropdown is set to 'clover'. Both dropdowns are highlighted with a red rectangle. Below this, there is an 'Execution History' section with a checkbox for 'Save job execution record in Execution History' which is currently unchecked. An 'Apply' button is located at the bottom left of the configuration area.

Figure 38.30. Configuring Anonymous Data Service

In the list of Data Services, the Data Service that does not require credentials is indicated by unlocked padlock icon.

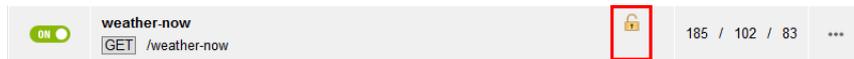


Figure 38.31. Data Service without authentication

## Running Data Service on HTTPS

By default, the Data Service runs on HTTP and you can configure it to run on HTTPS.

To run Data Service on HTTPS, create a new **HTTPS Connector**.

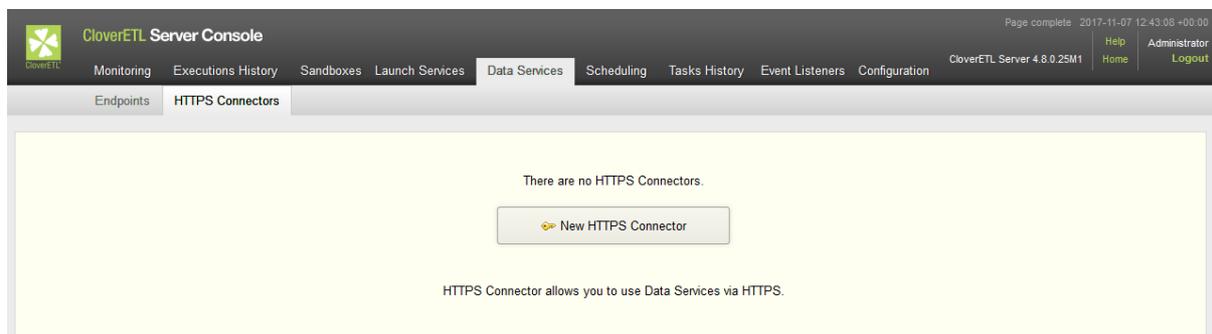


Figure 38.32. Creating a new Data Service Connector

Enter a name, port, keystore path, and keystore and key passwords.

Figure 38.33. Creating a new Data Service Connector II

In **Data Services** → **Endpoints**, select the Data Service to be running on HTTPS and switch to the **Configuration** tab.

Select the HTTPS connector from the combo box and click the **Apply** button. Now, the Data Service runs on HTTPS.

Figure 38.34. Using the HTTPS Connector in Data Service endpoint

You can have more independent HTTPS contexts running on one Server. There can be multiple Data Services running on the same HTTPS context.

## Running Data Service on HTTPS on Cluster

This case extends the case of [Running Data Service on HTTPS](#) (p. 303). Different cluster nodes have different domain names, but the Java key store has to have one certificate. There are two way to solve the problem with certificates.

- Use a wildcard certificate. The key store file should be placed on the shared file system.

- Use different certificates for each cluster node. The keystores with the certificates must be on the same path on all cluster nodes.

## Monitoring Data Service

To see the activity of Data Service, use the list of Data Services. There you can see the main overview of data services.

The state of a particular Data Service is on the *State and History* tab.

## Testing Data Service

To test the Data Service, select the Data Service in the list, switch to the **Testing and Documentation** tab and click the **Try it out** button.

The screenshot shows the CloverETL Server Console interface. The 'Data Services' tab is selected, displaying a table of services. The 'weather-now' service is highlighted in blue. The 'Testing and Documentation' tab is active, showing the service details for the GET /weather-now endpoint. The 'Try it out' button is highlighted with a red box.

Enabled	Endpoint	Activity (3 days) Run / Success / Fail	Actions
ON	Get mountains GET /mountains	0 / 0 / 0	...
ON	travel GET /travel	0 / 0 / 0	...
ON	weather-now GET /weather-now	20 / 18 / 2	...

Detail | **Testing and Documentation** | State and History | Alerts and Notification | Configuration

GET /weather-now

You should call this service with  
<http://172.22.0.213:8083/clover/data-service/weather-now?location=>  
 e.g. <http://172.22.0.213:8083/clover/data-service/weather-now?location=Aberdeen>

Parameters Try it out

Name	Description
location string (query)	undefined

Responses Response content type: application/xml

Code	Description
200	
400	Request validation failed
500	Job failed

Figure 38.35. Data Service - Testing and Documentation

Enter the parameter values and click **Execute**.

Detail **Testing and Documentation** State and History Alerts and Notification Configuration

GET /weather-now

You should call this service with  
<http://172.22.0.213:8083/clover/data-service/weather-now?location=>  
e.g.  
<http://172.22.0.213:8083/clover/data-service/weather-now?location=Aberdeen>

Parameters Cancel

Name	Description
location string (query)	undefined
	<input type="text" value="London"/>

**Execute**

Responses Response content type

Code	Description
200	
400	<b>Request validation failed</b>
500	<b>Job failed</b>

Figure 38.36. Data Service - Testing and Documentation - Parameters

The screenshot displays the 'Testing and Documentation' tab of a web service interface. At the top, there are navigation tabs: 'Detail', 'Testing and Documentation' (selected), 'State and History', 'Alerts and Notification', and 'Configuration'. Below these, the method 'GET' and path '/weather-now' are shown. A note indicates the service URL and provides an example with 'location=Aberdeen'. The 'Parameters' section shows a 'location' parameter of type 'string (query)' with the value 'York'. There are 'Execute' and 'Clear' buttons. The 'Responses' section shows the 'Response content type' set to 'application/xml'. Below this, the 'Curl' command is displayed: `curl -X GET "http://172.22.0.216:8083/clover/data-service/weather-now?location=London" -H "accept: application/xml"`. The 'Server response' section shows a '200' status code and the following XML body: 

```
<?xml version="1.0" encoding="UTF-8"?>
<root>
  <weather_cities>
    <City>York</City>
    <TemperatureC>12</TemperatureC>
    <PrecipitationsMm>0</PrecipitationsMm>
  </weather_cities>
</root>
```

 The 'Response headers' section lists: `server: Apache-Coyote/1.1`, `cache-control: no-cache, no-store, must-revalidate`, `pragma: no-cache`, `expires: 0`, `content-type: application/xml; charset=UTF-8`, `transfer-encoding: chunked`, and `date: Thu, 07 Sep 2017 11:43:50 GMT`. At the bottom, a 'Responses' table lists status codes and their descriptions: 200 (no description), 400 (Request validation failed), and 500 (Job failed).

Figure 38.37. Data Service - Testing and Documentation - Result

## Performance Tuning

To improve performance, do not save job execution records in Executions History (p. 232). To do so, do not tick *Save job execution records in Execution History* on Configuration tab.

## Exporting Data Service Configuration

You can export the Data Service configuration from **Data Services** → **Endpoints** tab. Click the three-dot-button and select **Export Data Services Configuration** from context menu.



Figure 38.38. Data Services - Export

The Data Services configuration will be exported.

You can also export Data Service configuration in **Configuration** → **Export** See [Server Configuration Export](#) (p. 189).

## Importing Data Service Configuration

You can import the Data Service configuration from **Endpoints**. Click the three-dot-button and select **Import Data Services Configuration** from the context menu.

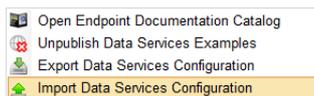


Figure 38.39. Data Services - Import

You can also import the Data Service configuration directly in **Configuration** → **Import** See [Server Configuration Import](#) (p. 190).

## Avoiding Premature Marking of Data Service as Failing

Data Service might prematurely switch to a failing state if the failure indication is set up to switch to a failing state after a given percentage of executions fails in a given time window. E.g. First execution fails.

To avoid this, you can set the minimum number of events necessary to be taken into account when calculating the change of Data Service state. It can be set with the `dataservice.failure.ratio.min.record.count` configuration property. The default value is 10 executions.

It can be set in **Configuration** → **Setup** → **Configuration File**. Add a line containing

```
dataservice.failure.ratio.min.record.count=10
```

to the configuration file.

You can set it to any reasonable positive integer. This configuration is valid for all Data Services available on the Server.

See also Chapter 14, [List of Configuration Properties](#) (p. 106).

## Looking up Particular Data Service

If you have multiple Data Services available, you can search for a specific Data Service:

If you know the endpoint name, you can look it up. Enter the text into the **Search endpoints** field and click the **Refresh** button. The Data Services will be filtered.

The entered text will be searched in the title of the Data Service, in the name of the request method, in the name of the `.x` job file and in the path that the Data Service uses.

If you would like to see invalid endpoints only, click the **failing and invalid only** icon. The both filters can be combined.

To switch off the filters, click the **Show All** button.

## **Resetting State of Failing Data Service Endpoint**

---

If the Data Service endpoint is in the failing state and the problem has been fixed, you can reset the endpoint state manually.

To reset the state, open the details of the endpoint, switch to the **Alerts and Notification** tab and click the **Apply and Reset State** button.

If the endpoint has an email address set, a notification email will be sent to this address.

---

## Custom HTTP Headers

Data Services accept a custom HTTP header `X-Clover-Save-Run-Record`. The possible values of the header are `TRUE` and `FALSE`. CloverETL Server accepts them case insensitively.

This header overrides the endpoint's configuration to save the run record or not.

**Testing and Documentation** page now automatically sends the header with value set to true. This means that all invocations from Testing and Documentation page are saved to the executions history.

Testing Data Service from Designer creates a record in Executions History regardless of it being published or not.

---

## Data Services on Cluster

Data Service jobs can run on Cluster in the same way as they run on **CloverETL Server**. Parallel run of one Data Service job on multiple cluster nodes is not supported.

---

## **Part VII. Cluster**

---

---

## Chapter 39. Clustering Features

There are two common Cluster features: high availability and scalability. Both are implemented by the CloverETL Server on different levels. This section should clarify the basics of CloverETL Clustering.

The CloverETL Server only works in the Cluster if your license allows it.

---

### High Availability

CloverETL Server does not recognize any differences between cluster nodes. Thus, there are no "master" or "slave" nodes meaning all nodes can be virtually equal. There is no single point of failure (SPOF) in the CloverETL cluster itself; however, SPOFs may be in the input data or some other external element.

Clustering offers high availability (HA) for all features accessible through HTTP, for event listeners and scheduling. Regarding the HTTP accessible features: it includes sandbox browsing, modification of services configuration (scheduling, launch services, listeners) and primarily job executions. Any cluster node may accept incoming HTTP requests and process them itself or delegate it to another node.

Since all nodes are typically equal, almost all requests may be processed by any cluster node:

- All job files, metadata files, etc. are located in shared sandboxes. Thus all nodes have access to them. A shared filesystem may be a SPOF, so it is recommended to use a replicated filesystem instead.
- The database is shared by all cluster nodes. Again, a shared DB might be a SPOF, however it may be clustered as well.

But there is still a possibility that a node itself cannot process a request. In such cases, it completely and transparently delegates the request to a node which can process the request.

These are the requests which are limited to one (or more) specific node(s):

- a request for the content of a partitioned or local sandbox. These sandboxes aren't shared among all cluster nodes. Note that this request may come to any cluster node which then delegates it transparently to a target node; however, this target node must be up and running.
- A job is configured to use a partitioned or local sandbox. These jobs need nodes which have a physical access to the required sandboxes.
- A job has allocation specified by specific cluster nodes. Concept of allocation is described in the following sections.

Thus an inaccessible cluster node may cause a failure of the request; So, if possible, it is better to avoid using specific cluster nodes or resources accessible only by specific cluster node.

CloverETL itself implements a load balancer for executing jobs. So a job which isn't configured for some specific node(s) may be executed anywhere in the cluster and the CloverETL load balancer decides, according to the request and current load, which node will process the job. All this is done transparently for the client side.

To achieve HA, it is recommended to use an independent HTTP load balancer. Independent HTTP load balancers allow transparent fail-overs for HTTP requests. They send requests to the nodes which are running.

---

### Scalability

There are two independent levels of scalability implemented. Scalability of transformation requests (and any HTTP requests) and data scalability (parallel data processing).

Both of these scalability levels are horizontal. Horizontal scalability means adding nodes to the cluster, whereas vertical scalability means adding resources to a single node. Vertical scalability is supported natively by the CloverETL engine and it is not described here.

## Transformation Requests

Basically, the more nodes we have in the cluster, the more transformation requests (or HTTP requests in general) we can process at one time. This type of scalability is the CloverETL Server's ability to support a growing number of clients. This feature is closely related to the use of an HTTP load balancer which is mentioned in the previous section.

## Parallel Data Processing

[ETL Graph Allocation](#) (p. 313)

[Component Allocation](#) (p. 314)

[Partitioning/Gathering Data](#) (p. 315)

[Node Allocation Limitations](#) (p. 315)

[Partitioned and Local Sandboxes](#) (p. 315)

[Using a Sandbox Resource as a Component Data Source](#) (p. 317)

This type of scalability is currently only available for ETL graphs. Jobflow and Profiler jobs can't run in parallel.

When a transformation is processed in parallel, the whole graph (or its parts) runs in parallel on multiple cluster nodes having each node process just a part of the data.

So the more nodes we have in a cluster, the more data can be processed in a specified time.

The data may be split (partitioned) before the graph execution or by the graph itself on the fly. The resulting data may be stored in partitions or gathered and stored as one group of data.

The curve of scalability may differ according to the type of transformation. It may be almost linear, which is almost always ideal, except when there is a single data source which cannot be read by multiple readers in parallel limiting the speed of further data transformation. In such cases, it is not beneficial to have parallel data processing since it would actually wait for input data.

## ETL Graph Allocation

Each ETL graph executed in a cluster environment is automatically subjected to a transformation analysis. The main goal of this analysis is to find so called ETL graph **allocation**. The graph allocation is a set of instructions for cluster environment defining how the transformation should be executed. For better understanding how the parallel data processing works, it is necessary to get detailed information about the graph analysis and resulted allocation.

First of all, the analysis needs to find allocation for each individual component. The component allocation is a set of cluster nodes where the component should be running. There are several ways how the component allocation can be specified, see following section of the documentation. But important information for now is, that a component can be requested to run in multiple instances - this is necessary for parallel data processing. Next step of analysis is to find optimal graph decomposition to ensure all component allocation will be satisfied and try to minimise number of remote edges between graph instances.

Resulted analysis says how many instances (workers) of the graph need to be executed, on which cluster nodes these instances will be running and which components will be present in the instances. In other words, one executed graph can be running in many instances, each instance can be processed on an arbitrary cluster node and moreover each instance contains only convenient components.

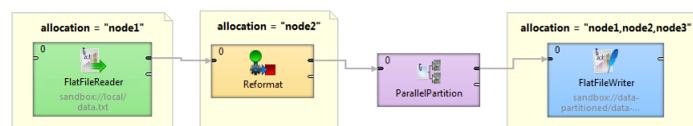


Figure 39.1. Component allocations example

This figure shows a sample graph with a few components with various component allocations. The first component - **FlatFileReader** - requests to be executed on node1, the following **Reformat** component should be running

on cluster node2, the **ParallelPartition** component is a special component which makes it possible to change cardinality of allocation of two interconnected components (detailed description of cluster partitioning and gathering follows this section). **FlatFileWriter**, the last component, requires to be executed on three cluster nodes node1, node2 and node3. Visualisation of the transformation analysis is shown in the following figure. Three workers (graphs) will be executed, each on a different cluster node (which is not necessary, even multiple workers can be associated with a single node). Worker on cluster node1 contains only **FlatFileReader** and first of three instances of the **FlatFileWriter** component. Both components are connected by remote edges with components which are running on node2. The worker running on node3 contains only **FlatFileWriter** fed by data remotely transferred from **ParallelPartitioner** running on node2.

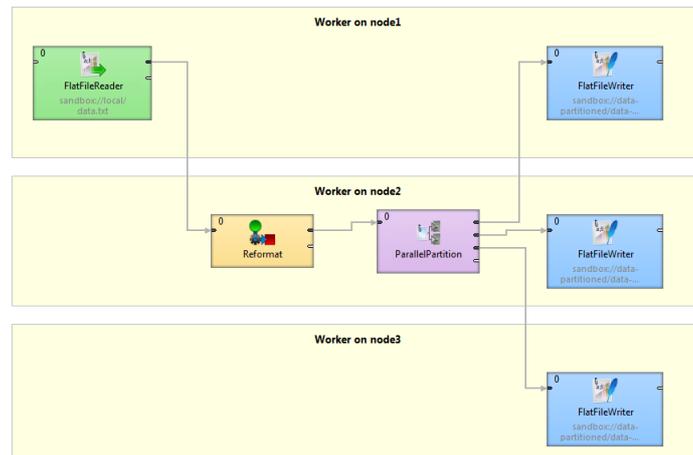


Figure 39.2. Graph decomposition based on component allocations

### Component Allocation

Allocation of a single component can be derived in several ways (list is ordered according to priority):

- **Explicit definition** - all components have a common attribute **Allocation**. CloverETL Designer allows user to use convenient dialog.

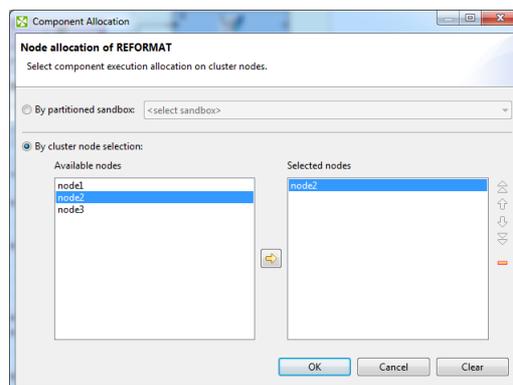


Figure 39.3. Component allocation dialog

Three different approaches are available for explicit allocation definition:

- **Allocation based on the number of workers** - the component will be executed in requested instances on some cluster nodes which are preferred by CloverETL Cluster. Server can use a build-in loadbalancing algorithm to ensure the fastest data processing.
- **Allocation based on reference on a partitioned sandbox** - component allocation corresponds with locations of given partitioned sandbox. Each partitioned sandbox has a list of locations, each bound to a specific cluster node. Thus allocation would be equivalent to list of locations. For more information, see "Partitioned sandbox" in [Partitioned and Local Sandboxes](#) (p. 315).

- **Allocation defined by a list of cluster node identifiers** (a single cluster node can be used more times)
- **Reference to a partitioned sandbox** FlatFileReader, FlatFileWriter and ParallelReader components derives their allocation from the `fileURL` attribute. In case the URL refers to a file in a partitioned sandbox, the component allocation is automatically derived from locations of the partitioned sandbox. So in case you manipulate with one of these components with a file in partitioned sandbox suitable allocation is used automatically.
- **Adoption from neighbour components** By default, allocation is inherited from neighbour components. Components on the left side have a higher priority. Cluster partitioners and cluster gathers are nature bounds for recursive allocation inheritance.

## Partitioning/Gathering Data

As mentioned before, data may be partitioned and gathered in multiple ways. It may be prepared before the graph is executed or it may be partitioned on the fly.

### Partitioning/gathering "on the fly"

There are six special components to consider: `ParallelPartition`, `ParallelLoadBalancingPartition`, `ParallelSimpleCopy`, `ParallelSimpleGather`, `ParallelMerge` and `ParallelRepartition`. All the components work similarly to their non-cluster variation. But their splitting or gathering nature is used to change data flow allocation, so they may be used to change distribution of the data among workers.

**ParallelPartition** and **ParallelLoadBalancingPartition** work similar to a common partitioner, they change the data allocation from 1 to N. Component preceding the `ParallelPartitioner` run on just one node, whereas the component behind the `ParallelPartitioner` run in parallel according to node allocation. **ParallelSimpleCopy** component can be used in similar locations. This component does not distribute the data records, but copies them to all output workers.

**ParallelGather** and **ParallelMerge** work in the opposite way. They change the data allocation from N to 1. The component preceding the gather/merge run in parallel while the Component behind the gather run on just one node.

### Partitioning/gathering data by external tools

Partitioning data on the fly may in some cases be an unnecessary bottleneck. Splitting data using low-level tools can be much better for scalability. The optimal case being that each running worker reads data from an independent data source. Thus there does not have to be a `ParallelPartitioner` component and the graph runs in parallel from the beginning.

Or the whole graph may run in parallel, however the results would be partitioned.

## Node Allocation Limitations

As described above, each component may have its own node allocation specified which may result in conflicts.

- *Node allocation of neighbouring components must have the same cardinality.* So it doesn't have to be the same allocation, but the cardinality must be the same. For example, there is an ETL graph with 2 components: **DataGenerator** and **Trash**. **DataGenerator** allocated on NodeA sending data to **Trash** allocated on NodeB works. **DataGenerator** allocated on NodeA sending data to **Trash** allocated on NodeA and NodeB fails.
- *Node allocation behind the **ParallelGather** and **ParallelMerge** must have cardinality 1.* So it may be of any allocation, but the cardinality must be just 1.
- *Node allocation of components in front of the **ParallelPartition**, **ParallelLoadBalancingPartition** and **ParallelSimpleCopy** must have cardinality 1.*

## Partitioned and Local Sandboxes

Partitioned and local sandboxes were mentioned in previous sections. These new sandbox types were introduced in version 3.0 and they are vital for parallel data processing.

Together with shared sandboxes, we have three sandbox types in total.

### Shared sandbox

This type of sandbox must be used for all data which is supposed to be accessible on all cluster nodes. This includes all graphs, jobflows, metadata, connections, classes and input/output data for graphs which should support HA, as described above. All shared sandboxes reside in the directory, which must be properly shared among all cluster nodes. You can use a suitable sharing/replicating tool according to the operating system and filesystem.

Figure 39.4. Dialog form for creating new shared sandbox

As you can see in the screenshot above, you can specify the root path on the filesystem and you can use placeholders or absolute path. Placeholders available are environment variables, system properties or CloverETL Server configuration property intended for this use: `sandboxes.home`. Default path is set as `[user.data.home]/CloverETL/sandboxes/[sandboxID]` where the `sandboxID` is an ID specified by the user. The `user.data.home` placeholder refers to the home directory of the user running the JVM process (`/home` subdirectory on Unix-like OS); it is determined as the first writable directory selected from the following values:

- `USERPROFILE` environment variable on Windows OS
- `user.home` system property (user home directory)
- `user.dir` system property (JVM process working directory)
- `java.io.tmpdir` system property (JVM process temporary directory)

Note that the path must be valid on all cluster nodes. Not just nodes currently connected to the cluster, but also on nodes that may be connected later. Thus when the placeholders are resolved on a node, the path must exist on the node and it must be readable/writable for the JVM process.

### Local sandbox

This sandbox type is intended for data, which is accessible only by certain cluster nodes. It may include massive input/output files. The purpose being, that any cluster node may access content of this type of sandbox, but only one has local (fast) access and this node must be up and running to provide data. The graph may use resources from multiple sandboxes which are physically stored on different nodes since cluster nodes are able to create network streams transparently as if the resources were a local file. For details, see [Using a Sandbox Resource as a Component Data Source](#) (p. 317).

Do not use a local sandbox for common project data (graphs, metadata, connections, lookups, properties files, etc.). It would cause odd behavior. Use shared sandboxes instead.

Node ID	Root path	
node3	`\${sandboxes.home.local}/uploadedData	add

Figure 39.5. Dialog form for creating a new local sandbox

The sandbox location path is pre-filled with the `sandboxes.home.local` placeholder which, by default, points to `[user.data.home]/CloverETL/sandboxes-local`. The placeholder can be configured as any other CloverETL configuration property.

### Partitioned sandbox

This type of sandbox is actually an abstract wrapper for a couple of physical locations existing typically on different cluster nodes. However, there may be multiple locations on the same node. A partitioned sandbox has two purposes which are both closely related to parallel data processing.

1. **node allocation** specification - locations of a partitioned sandbox define the workers which will run a graph or its parts. So each physical location will cause a single worker to run. This worker does not have to actually store any data to its location. It is just a way to tell the CloverETL Server: "execute this part of ETL graph in parallel on these nodes"
2. **storage for part of the data** during parallel data processing. Each physical location contains only part of the data. Typically, we have input data split in more input files, so we put each file into a different location and each worker processes its own file.

Node ID	Root path	
node3	\${sandboxes.home.partitioned}/data	delete
node4	\${sandboxes.home.partitioned}/data	add

Figure 39.6. Dialog form for creating new local sandbox

As you can see on the screenshot above, for a partitioned sandbox, you can specify one or more physical locations on different cluster nodes.

The sandbox location path is pre-filled with the `sandboxes.home.partitioned` placeholder which, by default, points to `[user.data.home]/CloverETL/sandboxes-partitioned`. The `sandboxes.home.partitioned` config property may be configured as any other CloverETL Server configuration property. Note that the directory must be readable/writable for the user running JVM process.

Do not use a partitioned sandbox for common project data (graphs, metadata, connections, lookups, properties files, etc.). It would cause odd behavior. Use shared sandboxes instead.

### Using a Sandbox Resource as a Component Data Source

A sandbox resource, whether it is a shared, local or partitioned sandbox (or ordinary sandbox on standalone server), is specified in the graph under the **fileURL** attributes as a so called sandbox URL like this:

```
sandbox://data/path/to/file/file.dat
```

where `data` is a code for the sandbox and `path/to/file/file.dat` is the path to the resource from the sandbox root. The URL is evaluated by CloverETL Server during job execution and a component (reader or writer) obtains the opened stream from the Server. This may be a stream to a local file or to some other remote resource. Thus, a job does not have to run on the node which has local access to the resource. There may be more sandbox resources used in the job and each of them may be on a different node.

The sandbox URL has a specific use for parallel data processing. When the sandbox URL with the resource in a *partitioned sandbox* is used, that part of the graph/phase runs in parallel, according to the node allocation specified by the list of partitioned sandbox locations. Thus, each worker has its own local sandbox resource.

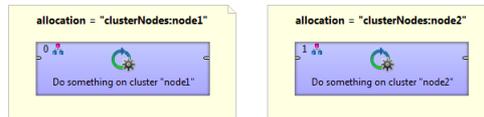
CloverETL Server evaluates the sandbox URL on each worker and provides an open stream to a local resource to the component.

The sandbox URL may be used on the standalone Server as well. It is an excellent choice when graph references some resources from different sandboxes. It may be metadata, lookup definition or input/output data. A referenced sandbox must be accessible for the user who executes the graph.

## Graph Allocation Examples

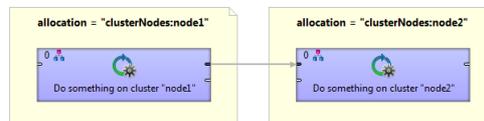
### Basic component allocation

This example shows a graph with two components, where allocation ensures that the first component will be executed on cluster node1 and the second component will be executed on cluster node2.



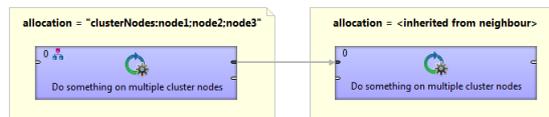
### Basic component allocation with remote data transfer

Two components connected with an edge can have a different allocation. The first is executed on node1 and the second is executed on node2. Cluster environment automatically ensures remote data records transfer.



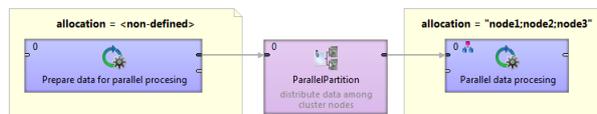
### Multiple execution

A graph with a multiple node allocation is executed in parallel. In this example, both components have same allocation, so three identical transformations will be executed on cluster node1, node2 and node3.



### Cluster data partitioning

A graph with two allocations. The first component has a single node allocation, which is not specified and is automatically derived to ensure a minimal number of remote edges. The **ParallelPartition** component distribute records for further data processing on the cluster node1, node2 and node3.



### Cluster data gathering

A graph with two allocations. Resulted data records of parallel data processing in the first component are collected in the **ParallelSimpleGather** component and passed to the cluster node4 for further single node processing.

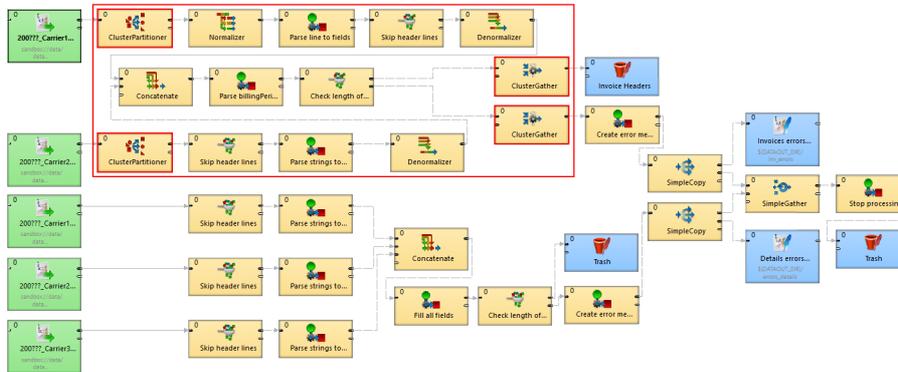


## Example of Distributed Execution

[Details of the Example Transformation Design](#) (p. 320)

[Scalability of the Example Transformation](#) (p. 323)

The following diagram shows a transformation graph used for parsing invoices generated by cell phone network providers.



The size of these input files may be up to a few gigabytes, so it is very beneficial to design the graph to work in the cluster environment.

## Details of the Example Transformation Design

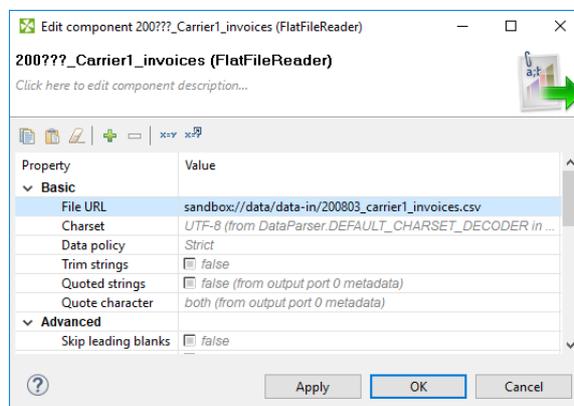
Note there are four cluster components in the graph and these components define a point of "node allocation" change, so the part of the graph demarcated by these components is highlighted by the red rectangle. The allocation of these components should be performed in parallel. This means that the components inside the rectangle should have convenient allocation. The rest of the graph runs on a single node.

### Specification of "node allocation"

There are 2 node allocations used in the graph:

- node allocation for components running in parallel (demarcated by the four cluster components)
- node allocation for the outer part of the graph which runs on a single node

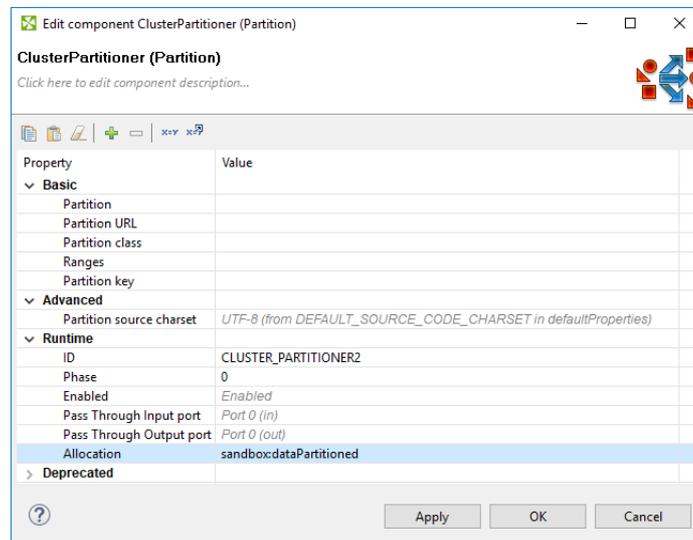
The single node is specified by the sandbox code used in the URLs of input data. The following dialog shows the File URL value: `sandbox://data/path-to-csv-file`, where `data` is an ID of the server sandbox containing the specified file. And it is the `data local` sandbox which defines the single node.



The part of the graph demarcated by the four cluster components may have specified its allocation by the file URL attribute as well, but this part does not work with files at all, so there is no file URL. Thus, we will use the

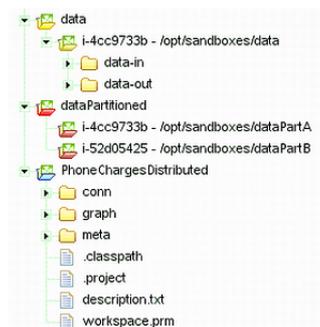
**node allocation** attribute. Since components may adopt the allocation from their neighbours, it is sufficient to set it only for one component.

Again, `dataPartitioned` in the following dialog is the sandbox ID.



Let's investigate our sandboxes. This project requires 3 sandboxes: **data**, **dataPartitioned** and **PhoneChargesDistributed**.

- data
  - contains input and output data
  - local sandbox (yellow folder), so it has only one physical location
  - accessible only on node `i-4cc9733b` in the specified path
- dataPartitioned
  - partitioned sandbox (red folder), so it has a list of physical locations on different nodes
  - does not contain any data and since the graph does not read or write to this sandbox, it is used only for the definition of "nodes allocation"
  - on the following figure, the allocation is configured for two cluster nodes
- PhoneChargesDistributed
  - common sandbox containing the graph file, metadata, and connections
  - shared sandbox (blue folder), so all cluster nodes have access to the same files



If the graph was executed with the sandbox configuration of the previous figure, the node allocation would be:

- components which run only on a single node, will run only on the `i-4cc9733b` node according to the "data" sandbox location.
- components with an allocation according to the **dataPartitioned** sandbox will run on nodes `i-4cc9733b` and `i-52d05425`.

## Scalability of the Example Transformation

The example transformation has been tested in an Amazon Cloud environment with the following conditions for all executions:

- the same master node
- the same input data: 1.2GB of input data, 27 million records
- three executions for each "node allocation"
- "node allocation" changed between every 2 executions
- all nodes has been of "c1.medium" type

We tested "node allocation" cardinality from 1 single node, all the way up to 8 nodes.

The following figure shows the functional dependence of run-time on the number of nodes in the cluster:

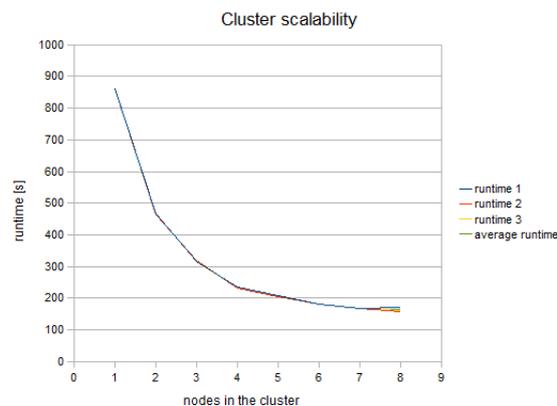


Figure 39.7. Cluster Scalability

The following figure shows the dependency of a speedup factor on the number of nodes in the cluster. The speedup factor is the ratio of the average runtime with one cluster node and the average runtime with x cluster nodes. Thus:

$$\text{speedupFactor} = \text{avgRuntime}(1 \text{ node}) / \text{avgRuntime}(x \text{ nodes})$$

We can see, that the results are favourable up to 4 nodes. Each additional node still improves the cluster performance; however, the effect of the improvement decreases. Nine or more nodes in the cluster may even have a negative effect because their benefit for performance may be lost in the overhead with the management of these nodes.

These results are specific for each transformation, there may be a transformation with a much better or possibly worse function curve.

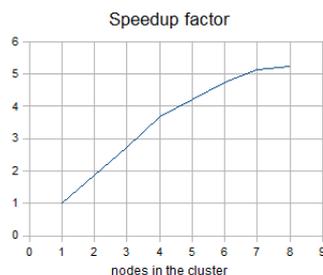


Figure 39.8. Speedup factor

Table of measured runtimes:

nodes	runtime 1 [s]	runtime 2 [s]	runtime 3 [s]	average runtime [s]	speedup factor
1	861	861	861	861	1
2	467	465	466	466	1.85
3	317	319	314	316.67	2.72
4	236	233	233	234	3.68
5	208	204	204	205.33	4.19
6	181	182	182	181.67	4.74
7	168	168	168	168	5.13
8	172	159	162	164.33	5.24

## Remote Edges

Data transfer between graphs running on different nodes is performed by a special type of edge - remote edge. The edge utilizes buffers for sending data in fixed-sized chunks. Each chunk has a unique number; therefore, in case of an I/O error, the last chunk sent can be re-requested.

You can set up values for various remote edge parameters via configuration properties. For list of properties, their meaning and default values, see [Optional Remote Edge Properties](#) (p. 329).

The following figure shows how nodes in a cluster communicate and transfer data - the client (graph running on Node 2) issues an HTTP request to Node 1 where a servlet accepts the request and checks the status of the source buffer. The source buffer is the buffer filled by the component writing to the left side of the remote edge. If the buffer is full, its content is transmitted to the Node 2, otherwise the servlet waits for configurable time interval for the buffer to become full. If the interval has elapsed without data being ready for download, the servlet finishes the request and Node 2 will re-issue the request at later time. Once the data chunk is downloaded, it is made available via the target buffer for the component reading from the right side of the remote edge. When the target buffer is emptied by the reading component, Node 2 issues new HTTP request to fetch the next data chunk.

This communication protocol and its implementation have consequences for the memory consumption of remote edges. A single remote edge will consume 3 x chunk size (1.5MB by default) of memory on the node that is the source side of the edge and 1 x chunk size (512KB by default) on the node that is the target of the edge. A smaller chunk size will save memory; however, more HTTP requests will be needed to transfer the data and the network latency will lower the throughput. Large data chunks will improve the edge throughput at the cost of higher memory consumption.

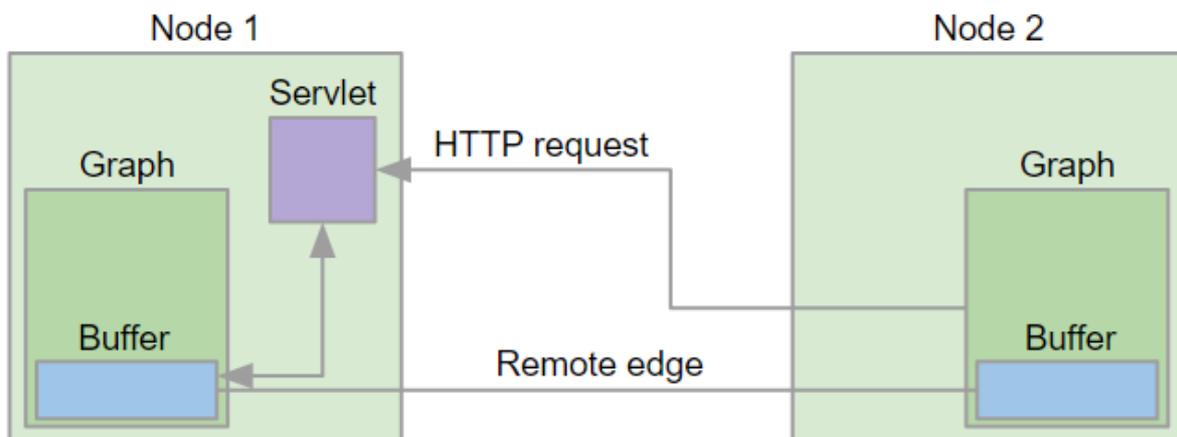


Figure 39.9. Remote Edge Implementation

---

## Chapter 40. Cluster Configuration

Cluster can work properly only if each node is properly configured. Clustering must be enabled, nodeID must be unique on each node, all nodes must have access to shared DB (direct connection or proxied by another cluster node) and shared sandboxes, and all properties for inter-node cooperation must be set according to network environment.

Properties and possible configuration are the following:

- [Mandatory Cluster Properties](#) (p. 326)
- [Optional Cluster Properties](#) (p. 327)
- [Example of 2 Node Cluster Configuration](#) (p. 331)
- [Jobs Load Balancing Properties](#) (p. 337)

---

## Mandatory Cluster Properties

Besides mandatory cluster properties, you need to set other necessary properties which are not specifically related to the cluster environment. Database connection must be also configured; however, besides direct connection, it is alternatively possible to configure proxying using another cluster node/nodes. For details, see the property [cluster.datasource.type](#) (p. 329).

### Mandatory properties - these properties must be properly set on each node of the cluster

<b>cluster.enabled</b>	Switch whether the Server should start in the standalone or cluster node mode. The property isn't set at all (empty value) by default, which means that the mode is chosen according to the loaded license. It is strongly recommended to set the property to <code>true</code> if the other cluster properties are configured, as well. Thus the cluster node will be initialized regardless of the license.  <b>Type:</b> boolean
<b>cluster.node.id</b>	Each cluster node must have unique ID.  <b>Type:</b> String  <b>Default:</b> node01
<b>cluster.jgroups.bind_address</b>	An IP address of the ethernet interface which is used for communication with another cluster nodes. Necessary for inter-node messaging.  <b>Type:</b> String, IP address  <b>Default:</b> 127.0.0.1
<b>cluster.jgroups.start_port</b>	Port where a jGroups server listens for inter-node messages.  <b>Type:</b> int, port  <b>Default:</b> 7800
<b>cluster.http.url</b>	A URL of the CloverETL cluster node. It must be an HTTP/HTTPS URL to the root of a web application. Typically it would be <code>http://[hostname]:[port]/clover</code> . Primarily, it is used for synchronous inter-node communication from other cluster nodes. It is recommended to use a fully qualified hostname or IP address, so it is accessible from a client browser or CloverETL Designer.  <b>Type:</b> String, URL  <b>Default:</b> <code>http://localhost:8080/clover</code>
<b>cluster.jgroups.tcping.initial_hosts</b>	List of IP addresses (with ports) where we expect running and listening nodes. It is related to other nodes "bind_address" and "start_port" properties. Necessary for inter-node messaging.  <b>Type:</b> String, in format: "IPAddress1[port1],IPAddress2[port2],..."  <b>Default:</b> 127.0.0.1[7800]

---

## Optional Cluster Properties

[Optional General Properties](#) (p. 327)

[Optional Remote Edge Properties](#) (p. 329)

### Optional General Properties

These properties are not vital for a cluster configuration - default values are sufficient.

**cluster.jgroups.external\_address** An IP address of the cluster node. Configure this only if the cluster nodes are on different sub-nets, so the IP address of the network interface isn't directly accessible from the other cluster nodes.

**Type:** String, IP address

**cluster.jgroups.external\_port** A port for asynchronous messaging. Configure this only if the cluster nodes are on different sub-nets, and the port opened on the IP address is different than the port opened on the node's network interface IP address.

**Type:** int, port

**cluster.jgroups.protocol.NAKACK.gc\_lag**

A number of delivered messages kept in the "sent messages buffer" of each jGroups view member. Messages are kept in a sender cache even though they were reported as delivered by existing view members, because there may be some other member temporarily not in the view. The higher the number, the higher the chance of reliable messages delivery in an unreliable network environment. However the messages consume memory: approximately 4kB for each message.

**Type:** int

**Default:** 10000

**cluster.jgroups.protocol.NAKACK.xmit\_table\_obsolete\_member\_timeout**

How long (in milliseconds) we keep obsolete member in the xmit-table. It is necessary for recognition of member temporarily unaccessible and removed from the view. With previous NAKACK implementation, the member removed from the view was also automatically removed from xmit-table, so it appeared as a new member when it re-joined the view. With current modified implementation the member is kept in the xmit-table for a configured interval longer, so when it re-joins the view, it is a known member and undelivered messages may be re-delivered to it. A member in the xmit-table isn't consuming memory.

**Type:** long

**Default:** 3600000

**cluster.jgroups.protocol.AUTH.value**

String used by a jgroups member to authenticate to the group. Must be the same on all cluster nodes. It is a protection against fake messages.

**Type:** String

**sandboxes.home.partitioned**

This property is intended to be used as a placeholder in the location path of partitioned sandboxes. So the sandbox path is specified with the placeholder

and it is resolved to the real path just before it is used. The default value uses configuration property `user.data.home` which points to the home directory of the user which runs the JVM process. Directory depends on the OS. On unix-like systems it is typically `/home/[username]`.

**Type:** String

**Default:** `${user.data.home}/CloverETL/sandboxes-partitioned`

**sandboxes.home.local**

This property is intended to be used as a placeholder in the location path of local sandboxes. So the sandbox path is specified with the placeholder and it is resolved to the real path just before it is used. The default value uses `user.data.home` configuration property which points to the home directory of the user which runs the JVM process. The directory location depends on the OS. On Unix-like systems it is typically `/home/[username]`.

**Type:** String

**Default:** `${user.data.home}/CloverETL/sandboxes-local`

**cluster.shared\_sandboxes\_path**

This property is deprecated. This property still works but is used only when shared sandbox doesn't have its own path specified. It is just for backward compatibility and it is not recommended for new deployments. Since 3.5, we recommend to specify the sandbox path explicitly and use the `sandboxes.home` property/placeholder.

**Type:** String

**cluster.node.sendinfo.interval**

A time interval in milliseconds; each node sends a heart-beat with information about itself to another nodes; this interval specifies how often the information is sent under common circumstances.

**Type:** int

**Default:** 2000

**cluster.node.sendinfo.cluster.node.sendinfo.min\_interval**

A time interval in milliseconds; a specified minimum interval between two heart-beats. A heart-beat may be send more often than specified by `cluster.node.sendinfo.interval`, e.g. when jobs start or finish. However the interval will never be shorter then this minimum.

**Type:** int

**Default:** 500

**cluster.node.sendinfo.history.interval**

A time interval in milliseconds, for which each node stores a heart-beat in the memory; It is used for rendering figures in the web GUI-monitoring section.

**Type:** int

**Default:** 240000 (4 minutes)

**cluster.node.remove.interval**

A time interval in milliseconds; if no node info comes in this interval, the node is considered as lost and it is removed from the cluster.

**Type:** int

**Default:** 50000

**cluster.max\_allowed\_time\_shift\_between\_nodes**

The maximum allowed time shift between nodes. All nodes must have system time synchronized. Otherwise the cluster may not work properly. So if this threshold is exceeded, node will be set as invalid.

**Type:** int

**Default:** 2000

**cluster.group.name**

Each cluster has its unique group name. If you need 2 clusters in the same network environment, each of them would have its own group name.

**Type:** String

**Default:** cloverCluster

**cluster.jgroups.protocol.AUTH.value**

An authentication string/password used for verification cluster nodes accessing the group. If this property is not specified, the cluster should be protected by firewall settings.

**Type:** String

**cluster.datasource.type**

Change this property to `remote` if the node doesn't have a direct connection to the CloverETL Server database, so it has to use some other cluster node as proxy to handle persistent operations. In such a case, the `cluster.datasource.delegate.nodeIds` property must be properly configured, as well. Properties `jdbc.*` will be ignored. Please note, that scheduler is active only on nodes with a direct connection.

**Type:** String

**Default:** local

**cluster.datasource.delegate.nodeIds** A list of cluster node IDs (separated by a comma ",") which this node may use as a proxy to handle persistent operations. At least one of the listed node IDs must be running, otherwise this node will fail. All listed node IDs must have a direct connection to CloverETL Server database properly configured. Property `cluster.datasource.delegate.nodeIds` is ignored by default. Property `cluster.datasource.type` must be set to `remote` to enable the feature.

**Type:** String

### Optional Remote Edge Properties

Below is a list of names and default values of properties used to configure remote edges in a clustered environment.

**cluster.edge.chunkSize**

Specifies the size of a chunk created by the right side of a remote edge (in bytes).

**default:** 524288

**cluster.edge.chunkWaitTimeout**

Specifies how long should the servlet wait for a next chunk to become available (in milliseconds).

**default:** 60000

<b>cluster.edge.connectTimeout</b>	Specifies a socket connection timeout when fetching a chunk (in milliseconds).  <b>default:</b> 30000
<b>cluster.edge.readTimeout</b>	Specifies a socket read timeout when fetching a chunk (in milliseconds).  <b>default:</b> 90000
<b>cluster.edge.handshakeTimeout</b>	Specifies how long should the client wait until a remote edge is registered by a data producing job (in milliseconds).  <b>default:</b> 120000
<b>cluster.edge.chunkReadRetries</b>	Specifies how many times should be a chunk fetch re-attempted before reporting an error to the consumer.  <b>default:</b> 2
<b>cluster.edge.disableChunkProtocol</b>	Disables the chunked data transfer protocol, switching to the old implementation.  <b>default:</b> false

## Example of 2 Node Cluster Configuration

[Basic 2-nodes Cluster Configuration](#) (p. 331)

[2-nodes Cluster with Proxied Access to Database](#) (p. 332)

[2-nodes Cluster with Load Balancer](#) (p. 333)

This section contains examples of CloverETL cluster nodes configuration. We assume that the user "clover" is running the JVM process and the license will be uploaded manually in the web GUI. In addition it is necessary to configure:

- sharing or replication of file system directory which the property "sandboxes.home" is pointing to. E.g. on Unix-like systems it would be typically /home/[username]/CloverETL/sandboxes.
- connection to the same database from both nodes

### Basic 2-nodes Cluster Configuration

This example describes a simple cluster: each node has a direct connection to a database.

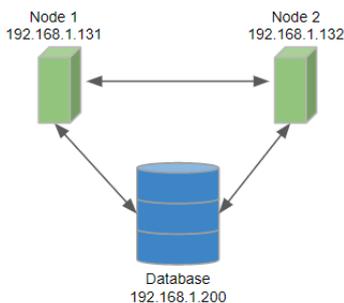


Figure 40.1. Configuration of 2-nodes cluster, each node has access to a database

Configuration of Node 1 on 192.168.1.131

```

jdbc.driverClassName=org.postgresql.Driver
jdbc.url=jdbc:postgresql://192.168.1.200/clover_db?charSet=UTF-8
jdbc.dialect=org.hibernate.dialect.PostgreSQLDialect
jdbc.username=clover
jdbc.password=clover

cluster.enabled=true
cluster.node.id=node01
cluster.http.url=http://192.168.1.131:8080/clover
cluster.jgroups.bind_address=192.168.1.131
cluster.jgroups.start_port=7800
cluster.jgroups.tcppping.initial_hosts=192.168.1.132[7800]

cluster.group.name=TheCloverCluster1

sandboxes.home=/home/clover/shared_sandboxes
  
```

Configuration of Node 2 on 192.168.1.132

```

jdbc.driverClassName=org.postgresql.Driver
jdbc.url=jdbc:postgresql://192.168.1.200/clover_db?charSet=UTF-8
jdbc.dialect=org.hibernate.dialect.PostgreSQLDialect
jdbc.username=clover
jdbc.password=clover
  
```

```

cluster.enabled=true
cluster.node.id=node02
cluster.http.url=http://192.168.1.132:8080/clover
cluster.jgroups.bind_address=192.168.1.132
cluster.jgroups.start_port=7800
cluster.jgroups.tcpping.initial_hosts=192.168.1.131[7800]

cluster.group.name=TheCloverCluster1

sandboxes.home=/home/clover/shared_sandboxes

```

The configuration is done in a **properties file**. The file can be placed either on a default (p. 71) or specified (p. 70) location.

## 2-nodes Cluster with Proxied Access to Database

This cluster configuration is similar to the previous one, but only one node has direct access to a database. The node2 has to use node1 as a proxy.

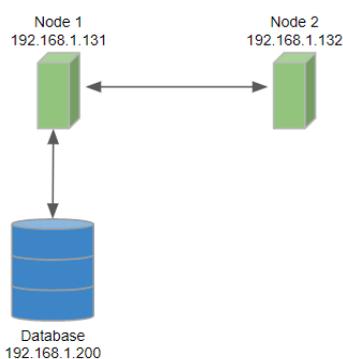


Figure 40.2. Configuration of 2-nodes cluster, one node without direct access to database

Configuration of Node 1 on 192.168.1.131

```

jdbc.driverClassName=org.postgresql.Driver
jdbc.url=jdbc:postgresql://192.168.1.200/clover_db?charSet=UTF-8
jdbc.dialect=org.hibernate.dialect.PostgreSQLDialect
jdbc.username=clover
jdbc.password=clover

cluster.enabled=true
cluster.node.id=node01
cluster.http.url=http://192.168.1.131:8080/clover
cluster.jgroups.bind_address=192.168.1.131
cluster.jgroups.start_port=7800

cluster.group.name=TheCloverCluster2

sandboxes.home=/home/clover/shared_sandboxes

```

Configuration of Node 2 on 192.168.1.132

```

cluster.datasource.type=remote (1)
cluster.datasource.delegate.nodeIds=node01

cluster.enabled=true
cluster.node.id=node02
cluster.http.url=http://192.168.1.132:8080/clover
cluster.jgroups.bind_address=192.168.1.132

```

```

cluster.jgroups.start_port=7800
cluster.jgroups.tcpping.initial_hosts=192.168.1.131[7800]

cluster.group.name=TheCloverCluster2

sandboxes.home=/home/clover/shared_sandboxes

```

❗ These two lines describe access to database via another node.

## 2-nodes Cluster with Load Balancer

If you use any external load balancer, the configuration of CloverETL Cluster will be same as in the first example.

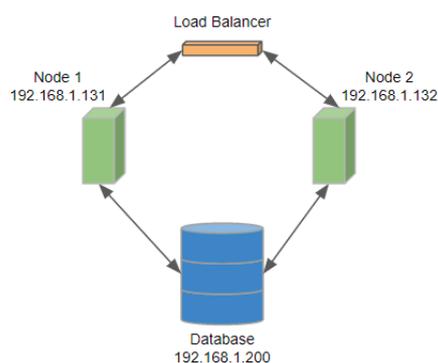


Figure 40.3. Configuration of 2-nodes cluster, one node without direct access to a database

The `cluster.http.url` and `cluster.jgroups.bind_address` are URLs of particular cluster nodes even if you use a load balancer.

Configuration of Node 1 on 192.168.1.131

```

jdbc.driverClassName=org.postgresql.Driver
jdbc.url=jdbc:postgresql://192.168.1.200/clover_db?charSet=UTF-8
jdbc.dialect=org.hibernate.dialect.PostgreSQLDialect
jdbc.username=clover
jdbc.password=clover

cluster.enabled=true
cluster.node.id=node01
cluster.http.url=http://192.168.1.131:8080/clover
cluster.jgroups.bind_address=192.168.1.131
cluster.jgroups.start_port=7800
cluster.jgroups.tcpping.initial_hosts=192.168.1.132[7800]

cluster.group.name=TheCloverCluster3

sandboxes.home=/home/clover/shared_sandboxes

```

Configuration of Node 2 on 192.168.1.132

```

jdbc.driverClassName=org.postgresql.Driver
jdbc.url=jdbc:postgresql://192.168.1.200/clover_db?charSet=UTF-8
jdbc.dialect=org.hibernate.dialect.PostgreSQLDialect
jdbc.username=clover
jdbc.password=clover

cluster.enabled=true
cluster.node.id=node02
cluster.http.url=http://192.168.1.132:8080/clover
cluster.jgroups.bind_address=192.168.1.132

```

```
cluster.jgroups.start_port=7800
cluster.jgroups.tcpping.initial_hosts=192.168.1.131[7800]

cluster.group.name=TheCloverCluster3

sandboxes.home=/home/clover/shared_sandboxes
```

## Example of 3 Node Cluster Configuration

### Basic 3-nodes Cluster Configuration

This example describes a cluster with three nodes where each node has a direct connection to a database.

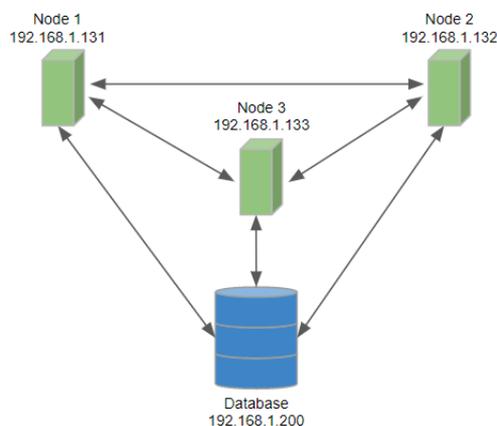


Figure 40.4. Configuration of 3-nodes cluster, each node has access to a database

Configuration of Node 1 on 192.168.1.131

```

jdbc.driverClassName=org.postgresql.Driver
jdbc.url=jdbc:postgresql://192.168.1.200/clover_db?charSet=UTF-8
jdbc.dialect=org.hibernate.dialect.PostgreSQLDialect
jdbc.username=clover
jdbc.password=clover

cluster.enabled=true
cluster.node.id=node01
cluster.http.url=http://192.168.1.131:8080/clover
cluster.jgroups.bind_address=192.168.1.131
cluster.jgroups.start_port=7800
cluster.jgroups.tcpping.initial_hosts=192.168.1.132[7800],192.168.1.133[7800]

cluster.group.name=TheCloverCluster4

sandboxes.home=/home/clover/shared_sandboxes
  
```

Configuration of Node 2 on 192.168.1.132

```

jdbc.driverClassName=org.postgresql.Driver
jdbc.url=jdbc:postgresql://192.168.1.200/clover_db?charSet=UTF-8
jdbc.dialect=org.hibernate.dialect.PostgreSQLDialect
jdbc.username=clover
jdbc.password=clover

cluster.enabled=true
cluster.node.id=node02
cluster.http.url=http://192.168.1.132:8080/clover
cluster.jgroups.bind_address=192.168.1.132
cluster.jgroups.start_port=7800
cluster.jgroups.tcpping.initial_hosts=192.168.1.131[7800],192.168.1.133[7800]

cluster.group.name=TheCloverCluster4

sandboxes.home=/home/clover/shared_sandboxes
  
```

Configuration of Node 3 on 192.168.1.133

```
jdbc.driverClassName=org.postgresql.Driver
jdbc.url=jdbc:postgresql://192.168.1.200/clover_db?charset=UTF-8
jdbc.dialect=org.hibernate.dialect.PostgreSQLDialect
jdbc.username=clover
jdbc.password=clover

cluster.enabled=true
cluster.node.id=node03
cluster.http.url=http://192.168.1.133:8080/clover
cluster.jgroups.bind_address=192.168.1.133
cluster.jgroups.start_port=7800
cluster.jgroups.tcpping.initial_hosts=192.168.1.131[7800],192.168.1.132[7800]

cluster.group.name=TheCloverCluster4

sandboxes.home=/home/clover/shared_sandboxes
```

## Jobs Load Balancing Properties

Multiplicators of load balancing criteria. A load balancer decides which cluster node executes the graph. It means, that any node may process a request for execution, but a graph may be executed on the same or on different node according to current load of the nodes and according to these multiplicators.

The higher the number, the higher the relevance for decision. All multiplicators must be greater than 0.

Each node of the cluster may have different load balancing properties. Any node may process incoming requests for transformation execution and each may apply criteria for loadbalancing in a different way according to its own configuration.

These properties aren't vital for cluster configuration - default values are sufficient

*Table 40.1. Load balancing properties*

property	type	default	description
cluster.lb.balance.runnfloatgraphs	float	3	Specify importance of running graphs for load balancing.
cluster.lb.balance.memfloat	float	0.5	Specify importance of used memory for load balancing.
cluster.lb.balance.cpufloat	float	1.5	Specify importance of number of CPUs for load balancing.
cluster.lb.balance.requestfloat	float	2	Specify importance of the fact that the node is the same which processes the request for execution. The same node which decides where to execute the graph. If you specify this multiplier great enough, the graph will be always executed on the same node which processes the request for execution.
cluster.lb.balance.nodefloat	float	1	Overall ratio bonus for configured node. Values greater then "1" increase probability of the node to be chosen by the loadbalancer. Value "1" means no bonus or penalty. "0" means that the node will never be chosen by the loadbalancer; however, it still may execute graphs, e.g. when there is no other node in the cluster or when the graph is designed to run on the node.

---

## Running More Clusters

If you run more clusters, each cluster has to have its own unique name. If the name is not unique, the cluster nodes of different clusters may consider foreign cluster nodes as part of the same cluster. The cluster name is configured using `cluster.group.name` option. See [Optional Cluster Properties](#) (p. 327).

---

## Cluster Reliability in Unreliable Network Environment

CloverETL Server instances must cooperate with each other to form a cluster together. If the connection between nodes doesn't work at all, or if it is not configured, cluster can't work properly. This chapter describes cluster nodes behavior in an environment where the connection between nodes is somehow unreliable.

### Nodes use three channels to exchange status info or data

#### 1. synchronous calls (via HTTP/HTTPS)

Typically NodeA requests some operation on NodeB, e.g. job execution. HTTP/HTTPS is also used for streaming data between workers of parallel execution

#### 2. asynchronous messaging (TCP connection on port 7800 by default)

Typically heart-beat or events, e.g. job started or finished.

#### 3. shared database – each node must be able to create DB connection

Shared configuration data, execution history, etc.

**Following scenarios are described below one by one, however they may occur together:**

- [NodeA Cannot Establish HTTP Connection to NodeB](#) (p. 339)
- [NodeA Cannot Establish TCP Connection \(Port 7800 by Default\) to NodeB](#) (p. 340)
- [NodeB is Killed or It Cannot Connect to the Database](#) (p. 340)
- [Auto-Resuming in Unreliable Network](#) (p. 341)
- [Long-Term Network Malfunction May Cause Jobs to Hang on](#) (p. 341)

## NodeA Cannot Establish HTTP Connection to NodeB

---

When HTTP request can't be established between nodes, jobs which are delegated between nodes or jobs running in parallel on more nodes will fail. The error is visible in the executions history. Each node periodically executes a check-task which checks the HTTP connection to other nodes. If the problem is detected, one of the nodes is suspended, since they can't cooperate with each other.

### Time-line describing the scenario:

- 0s network connection between NodeA and NodeB is down
- 0-40s a check-task running on NodeA can't establish HTTP connection to NodeB; check may last for 30s until it times-out; there is no re-try, if connection fails even just once, it is considered as unreliable, so the nodes can't cooperate.
- status of NodeA or NodeB (the one with shorter uptime) is changed to “suspended”

**The following configuration properties set the time intervals mentioned above:**

`cluster.node.check.checkPeriodicity` Periodicity of cluster node checks, in milliseconds.

**Default:** 20000

`cluster.sync.connection.readTimeout` An HTTP connection response timeout, in milliseconds.

**Default:** 30000

`cluster.sync.connection.connectTimeout` Establishing HTTP connection timeout, in milliseconds.

**Default:** 7000

## NodeA Cannot Establish TCP Connection (Port 7800 by Default) to NodeB

TCP connection is used for asynchronous messaging. When the NodeB can't send/receive asynchronous messages, the other nodes aren't notified about started/finished jobs, so a parent jobflow running on NodeA keeps waiting for the event from NodeB. A heart-beat is vital for meaningful load-balancing, the same check-task mentioned above also checks a heart-beat from all cluster nodes.

### Time-line describing the scenario:

- 0s network connection between NodeA and NodeB is down
- 60s NodeA uses the last available NodeB heart-beat
- 0-40s check-task running on NodeA detects missing heart-beat from NodeB
- status of NodeA or NodeB (the one with shorter uptime) is changed to `suspended`

### The following configuration properties set the time intervals mentioned above:

`cluster.node.check.checkInterval` Periodicity of cluster node checks, in milliseconds.

**Default:** 40000

`cluster.node.sendinfo.interval` Periodicity of heart-beat messages, in milliseconds.

**Default:** 2000

`cluster.node.sendinfo.min_heartbeat` An interval that may occasionally be sent more often than specified by `cluster.node.sendinfo.interval`. This property specifies the minimum interval in milliseconds.

**Default:** 500

`cluster.node.remove.interval` The maximum interval for missing a heart-beat, in milliseconds.

**Default:** 50000

## NodeB is Killed or It Cannot Connect to the Database

Access to a database is vital for running jobs, running scheduler and cooperation with other nodes. Touching a database is also used for detection of dead process. When the JVM process of NodeB is killed, it stops touching the database and the other nodes may detect it.

### Time-line describing the scenario:

- 0s-30s last touch on DB
- NodeB or its connection to the database is down
- 90s NodeA sees the last touch
- 0-40s check-task running on NodeA detects obsolete touch from NodeB
- status of NodeB is changed to `stopped`, jobs running on the NodeB are `solved`, which means that their status is changed to `UNKNOWN` and the event is dispatched among the cluster nodes. The job result is considered as `error`.

### The following configuration properties set the time intervals mentioned above:

`cluster.node.touch.interval` Periodicity of a database touch, in milliseconds.

**Default:** 20000

`cluster.node.touch.forced_interval` An interval when the other nodes accept the last touch, in milliseconds.

**Default:** 60000

`cluster.node.check.checkInterval` Periodicity of cluster node checks, in milliseconds.

**Default:** 40000

`cluster.node.touch.forced_at` A boolean value which can switch the resolving of running jobs mentioned above.

## Auto-Resuming in Unreliable Network

In version 4.4, auto-resuming of suspended nodes was introduced.

### Time-line describing the scenario:

- NodeB is suspended after connection loss
- **0s** NodeA successfully reestablishes connection to NodeB
- **120s** NodeA changes NodeB status to `forced_resume`
- NodeB attempts to resume itself if the maximum auto-resume count is not reached.
- If the connection is lost again, the cycle repeats; if the maximum auto-resume count is exceeded, the node will remain suspended until the counter is reset, to prevent suspend-resume cycles.
- **240m** auto-resume counter is reset

### The following configuration properties set the time intervals mentioned above:

`cluster.node.check.intervalBeforeNodeHasResume` Time before node has to resume, in milliseconds.

**Default:** 120000

`cluster.node.check.maxAutoResumeAttempts` How many times a node may try to auto-resume itself.

**Default:** 3

`cluster.node.check.intervalBeforeReasonableCounter` Time before the reasonable counter will be reset, in minutes.

**Default:** 240

## Long-Term Network Malfunction May Cause Jobs to Hang on

Jobflow or master execution executing child jobs on another cluster nodes must be notified about status changes of their child jobs. When the asynchronous messaging doesn't work, events from the child jobs aren't delivered, so parent jobs keep running. When the network works again, the child job events may be re-transmitted, so hung parent job may be finished. However, the network malfunction may be so long, that the event can't be re-transmitted.

### See following time-line to consider proper configuration:

- job A running on NodeA executes job B running on NodeB
- network between NodeA and NodeB is down from some reason
- job B finishes and sends the `finished` event, however it can't be delivered to NodeA – the event is stored in the `sent_events` buffer
- Since the network is down, a heart-beat can't be delivered as well and maybe HTTP connections can't be established, the cluster reacts as described in the sections above. Even though the nodes may be suspended, parent job A keeps waiting for the event from job B.
- **now, there are 3 possibilities:**

- a. Network finally starts working and since all undelivered events are in the `sent_events` buffer, they are re-transmitted and all of them are finally delivered. Parent job A is notified and proceeds. It may fail later, since some cluster nodes may be suspended.
- b. Network finally starts working, but the number of the events sent during the malfunction exceeded the `sent_events` buffer limit size. So some messages are lost and won't be re-transmitted. Thus the buffer size limit should be higher in the environment with unreliable network. Default buffer size limit is 10,000 events. It should be sufficient for thousands of simple job executions; basically, it depends on number of job phases. Each job execution produces at least 3 events (job started, phase finished, job finished). Please note that there are also other events fired occasionally (configuration changes, suspending, resuming, cache invalidation). Also messaging layer itself stores own messages to the buffer, but the number is negligible (tens of messages per hour). Heart-beat is not stored in the buffer.

There is also an inbound events buffer used as a temporary storage for events, so events may be delivered in correct order when some events can't be delivered at the moment. When the cluster node is inaccessible, the inbound buffer is released after timeout, which is set to 1 hour, by default.

- c. Node B is restarted, so all undelivered events in the buffer are lost.

**The following configuration properties set the time intervals mentioned above:**

`cluster.jgroups.protocol.NA.incoming.buffer.size` - Limit the size of the sent events buffer; Note that each stored message takes 2kB of heap memory.

**Default:** 10000

`cluster.jgroups.protocol.NA.incoming.buffer.timeout` - Amount of time before the incoming buffer is released if the cluster node is inaccessible.

---

## Chapter 41. Recommendations for Cluster Deployment

1. All nodes in the cluster should have a synchronized system date-time.
2. All nodes share sandboxes stored on a shared or replicated filesystem. The filesystem shared among all nodes is a single point of failure. Thus, the use of a replicated filesystem is strongly recommended.
3. All nodes share a DB, thus it must support transactions. I.e. The MySQL table engine, MyISAM, may cause unusual behavior because it is not transactional.
4. All nodes share a DB, which is a single point of failure. Use of a clustered DB is strongly recommended.
5. Configure the license by `license.file` property or upload it in the Web GUI, so it is stored in the database. Do not use `clover-license.war`.

---

## Chapter 42. Multiple CloverServer Instances on the same Host

Running multiple **CloverETL Server** instances on the same host is not recommended. If you do so, you should ensure that the instances do not interfere with each other.

- Each instance must run in a separate application server.
- Each instance must have its own port to listen. Application server can have some additional opened ports, therefore you might change the configuration of the application server as well.
- The instances must have a different `java.io.tmpdir` directory.
- Each instance needs a separate `sandboxes.home` directory.

See the documentation of particular application server for further information on running multiple instances of the application server on the same computer.

---

# Part VIII. Security

---

---

## Chapter 43. Security Recommendations for CloverETL Server

To improve security of CloverETL Server, you should:

- Change the default password for **clover** user. Without changing the password, everybody would be able to log in as **clover**. See [Change Users Password](#) (p. 157).
- Create a user different from **clover** and add it to the **admin** group. If there are more administrators, create a user account for each. See [Users](#) (p. 156).
- Set the **master password**. Without the master password, you cannot use secure parameters. See Chapter 19, [Secure Parameters](#) (p. 147).
- Run CloverETL Server with privileges of an ordinary user, e.g. create a system account `clover` used only for running CloverETL Server. Do not run CloverETL Server under the **root** account.
- Communication with system database may be unencrypted. Consider encrypting the connection to system database too.
- If database provides you with a root/admin account, do not use this account for CloverETL Server. Create a separate database user account, e.g. **clover**.
- Run **CloverETL Server** on HTTPS. If you communicate over HTTP, your data is sent unencrypted and eavesdroppers can easily see it.
- Disable the HTTP API if you do not need it. See Chapter 34, [Simple HTTP API](#) (p. 267).
- In Data Services, put keystores outside a sandbox and run the service on HTTPS. If you have a keystore in a sandbox, a user with write permissions could replace it with another key store. [HTTPS Connectors](#) (p. 298).
- Enable **user lockout** after repeated failed login attempts. If you use this feature in Cluster, make sure that all cluster nodes have the same lockout configuration. See [User Lockout](#) (p. 171)

---

## List of Figures

1.1. CloverETL Server User Interface .....	2
2.1. System Architecture .....	4
6.1. <code>setenv.sh</code> edited in Linux. ....	14
6.2. <code>setenv.bat</code> edited in Windows. ....	14
6.3. Apache Tomcat welcome page .....	19
6.4. Jetty welcome page .....	25
6.5. Adjusting Maximum heap size limit .....	29
6.6. Glassfish welcome page .....	33
6.7. JBoss AS 7 welcome page .....	37
6.8. JBoss EAP 7 welcome page .....	42
6.9. WebLogic welcome page .....	48
6.10. Login page of CloverETL Server without license .....	51
6.11. Add new license form .....	52
6.12. Update license form .....	53
6.13. CloverETL Server as the only running application on IBM WebSphere .....	58
10.1. CloverETL Server's System Database Configuration .....	69
12.1. Setup GUI with decorators .....	76
12.2. Example of the Server Configuration file .....	77
12.3. The License tab .....	77
12.4. Database connection configuration for JDBC (left) and JNDI (right) .....	78
12.5. The Worker tab .....	79
12.6. Sandbox path configuration with clustering disabled (left) and enabled (right) .....	79
12.7. Encryption configuration .....	80
12.8. E-mail configuration .....	80
12.9. LDAP configuration .....	81
12.10. Cluster configuration .....	82
14.1. MBean for a JNDI datasource in <code>jconsole</code> .....	117
17.1. Standalone server detail .....	131
17.2. Performance .....	132
17.3. CPU Load .....	133
17.4. Running jobs .....	133
17.5. System .....	133
17.6. Worker .....	133
17.7. Status History .....	134
17.8. Users' Accesses panel .....	134
17.9. Classloader cache .....	134
17.10. Status .....	135
17.11. Heartbeat .....	135
17.12. Threads .....	135
17.13. Quartz .....	136
17.14. Cluster overview .....	137
17.15. Node detail .....	138
17.16. Server Logs .....	139
18.1. Configured temp spaces overview - one default temp space on each cluster node .....	141
18.2. Newly added global temp space. ....	143
18.3. Temp spaces using environment variables and system properties .....	144
18.4. Disable operation reports action performed .....	145
18.5. Remove operation asks for confirmation in case there are data present in the temp space .....	146
19.1. Master password initialization .....	147
19.2. Graph parameters tab with initialized master password .....	148
20.1. Web GUI - section "Users" under "Configuration" .....	156
20.2. Web GUI - edit user .....	157
20.3. Web GUI - change password .....	157
20.4. Web GUI - groups assignment .....	157
20.5. Web GUI - section "Groups" .....	159

20.6. Web GUI - users assignment .....	160
20.7. Tree of permissions .....	160
20.8. Locked User .....	172
20.9. Unlocking User .....	173
21.1. Sandboxes Section in CloverETL Server Web GUI .....	174
21.2. Sandbox Permissions in CloverETL Server Web GUI .....	177
21.3. Web GUI - section "Sandboxes" - context menu on sandbox .....	178
21.4. Web GUI - section "Sandboxes" - context menu on folder .....	178
21.5. Web GUI - download sandbox as ZIP .....	179
21.6. Web GUI - upload ZIP to sandbox .....	179
21.7. Web GUI - upload ZIP results .....	180
21.8. Web GUI - download file as ZIP .....	180
21.9. Delete Sandbox - Button .....	181
21.10. Delete Sandbox - Right-Click .....	181
21.11. Delete Sandbox - Confirmation Dialog .....	182
21.12. Job config properties .....	185
22.1. Server Configuration Export screen .....	189
22.2. Server Configuration Import screen .....	190
22.3. Server Configuration uploaded .....	191
22.4. Outcome of the import preview for configuration from Example 22.1 .....	192
22.5. Outcome of import preview for configuration after fixing by removal of broken group reference. ....	193
26.1. Web GUI - send email .....	205
26.2. Web GUI - shell command .....	208
26.3. Web GUI - Graph execution task .....	211
26.4. Web GUI - Graph execution task .....	212
26.5. Web GUI - Jobflow execution task .....	213
26.6. Web GUI - "Abort job" .....	215
26.7. Web GUI - archive records .....	217
26.8. Web GUI - Task JMS message editor .....	219
27.1. Web GUI - "Manual task execution" form .....	223
28.1. Web GUI - section "Scheduling" - create new .....	224
28.2. Web GUI - onetime schedule form .....	225
28.3. Web GUI - schedule form - calendar .....	226
28.4. Web GUI - periodical schedule form .....	227
28.5. Cron periodical schedule form .....	228
28.6. Editing the cron expression - minutes .....	229
28.7. Editing the cron expression - days of week .....	229
28.8. Schedule allocation - Any node .....	230
28.9. Schedule allocation - One ore more specific nodes .....	230
29.1. Executions History - executions table .....	232
29.2. Executions History - overall perspective .....	234
29.3. Executions Hierarchy with docked list of jobs .....	234
29.4. Executions History - Tracking .....	235
29.5. Executions History - Tracking .....	235
30.1. Listeners .....	236
30.2. Web GUI - graph timeout event .....	239
30.3. The event source graph isn't specified, thus the listener works for all graphs in the specified sandbox ..	240
30.4. Web GUI - email notification about graph failure .....	241
30.5. Web GUI - email notification about graph success .....	242
30.6. Web GUI - backup of data processed by graph .....	243
30.7. Web GUI - jobflow timeout event .....	245
30.8. Web GUI - creating a File Event listener .....	253
30.9. File available on local file system .....	255
30.10. File available on remote location .....	255
30.11. Web GUI - creating a Task Failure listener .....	261
35.1. Glassfish JMX connector .....	279
35.2. WebSphere configuration .....	280
37.1. Launch Services and CloverETL Server as web application back-end .....	282

---

37.2. Launch Services section .....	284
37.3. Creating a new launch configuration .....	284
37.4. Overview tab .....	284
37.5. Edit Configuration tab .....	285
37.6. Creating new parameter .....	286
37.7. Edit Parameters tab .....	286
38.1. Data Services .....	289
38.2. Data Service tabs .....	290
38.3. Endpoints .....	290
38.4. Data Service Endpoints .....	290
38.5. Three-dot-button menu .....	290
38.6. List of Data Services .....	291
38.7. Selecting the particular Data Service .....	292
38.8. Data Service Detail tab .....	292
38.9. Data Service - Testing and Documentation tab .....	293
38.10. Data Service - State and History .....	294
38.11. Data Service - Alerts and Notitification .....	295
38.12. Data Service - Alerts and Notitification .....	295
38.13. Data Service - Health status .....	296
38.14. Some Data Service is failing .....	296
38.15. E-mail Notification .....	296
38.16. Data Service .....	296
38.17. Glogal Catalogue of Services .....	297
38.18. Glogal Catalog of Services .....	297
38.19. Data Services - Publishing the examples .....	298
38.20. Data Services - Published the examples .....	298
38.21. HTTPS Connectors .....	298
38.22. Tree-dot button menu .....	299
38.23. HTTPS Connectors .....	299
38.24. Publishing Data Service job .....	301
38.25. Publishing Data Service job that does not require authentication .....	301
38.26. Publishing Data Service from Sandbox .....	302
38.27. Unpublishing Data Service from Sandbox .....	302
38.28. Publishing Data Service examples .....	302
38.29. Publishing Data Service examples - II .....	302
38.30. Configuring Anonymous Data Service .....	303
38.31. Data Service without authentication .....	303
38.32. Creating a new Data Service Connector .....	303
38.33. Creating a new Data Service Connector II .....	304
38.34. Using the HTTPS Connector in Data Service endpoint .....	304
38.35. Data Service - Testing and Documentation .....	305
38.36. Data Service - Testing and Documentation - Parameters .....	306
38.37. Data Service - Testing and Documentation - Result .....	307
38.38. Data Services - Export .....	308
38.39. Data Services - Import .....	308
39.1. Component allocations example .....	313
39.2. Graph decomposition based on component allocations .....	314
39.3. Component allocation dialog .....	314
39.4. Dialog form for creating new shared sandbox .....	316
39.5. Dialog form for creating a new local sandbox .....	316
39.6. Dialog form for creating new local sandbox .....	317
39.7. Cluster Scalability .....	323
39.8. Speedup factor .....	323
39.9. Remote Edge Implementation .....	324
40.1. Configuration of 2-nodes cluster, each node has access to a database .....	331
40.2. Configuration of 2-nodes cluster, one node without direct access to database .....	332
40.3. Configuration of 2-nodes cluster, one node without direct access to a database .....	333
40.4. Configuration of 3-nodes cluster, each node has access to a database .....	335

---

---

## List of Tables

5.1. Hardware requirements of CloverETL Server .....	9
5.2. CloverETL Server Compatibility Matrix .....	10
7.1. JVM Memory Structure .....	60
7.2. ....	61
7.3. Firewall Exceptions .....	62
14.1. General configuration .....	106
14.2. Server - Worker configuration .....	114
14.3. Properties for JDBC JNDI Resources in Worker .....	118
14.4. Properties for JMS JNDI Resources in Worker .....	119
14.5. Properties for SSL communication in Worker .....	120
14.6. Defaults for job execution configuration - see <a href="#">Job Config Properties</a> for details .....	121
15.1. Parameters .....	125
17.1. Resource Utilization .....	132
19.1. Secure parameters configuration parameters .....	149
20.1. Admin user .....	156
20.2. User attributes .....	156
20.3. Default groups created during installation .....	159
21.1. Sandbox attributes .....	175
21.2. Sandbox permissions .....	177
21.3. ZIP upload parameters .....	180
21.4. Job config parameters .....	183
25.1. Defaults for graph execution configuration .....	199
25.2. passed parameters .....	200
25.3. passed parameters .....	201
25.4. passed parameters .....	201
26.1. Attributes of "Send e-mail" task .....	204
26.2. Placeholders useful in email templates .....	207
26.3. Attributes of "Execute shell command" task .....	208
26.4. Parameters of "Execute shell command" task .....	209
26.5. Parameters of "Execute shell command" task - available in scheduling .....	209
26.6. Parameters of "Execute shell command" task - available in listeners .....	209
26.7. Parameters of "Execute shell command" task - available in manual task execution .....	210
26.8. Attributes of "Graph execution" task .....	211
26.9. Additional parameters available in Event Listeners .....	212
26.10. Attributes of "Jobflow execution" task .....	213
26.11. Attributes of "Abort job" task .....	215
26.12. Attributes of "Archivator" task .....	216
26.13. Attributes of JMS message task .....	218
26.14. Parameters of "Send a JMS Message" .....	219
26.15. List of variables available in Groovy code .....	221
28.1. Onetime schedule attributes .....	225
28.2. Periodical schedule attributes .....	227
28.3. Cron periodical schedule attributes .....	228
29.1. Persistent run record attributes .....	233
29.2. Tracking table information .....	235
30.1. Event Listeners Table Description .....	236
30.2. Attributes of JMS message task .....	246
30.3. Variables accessible in groovy code .....	248
30.4. Properties Elements .....	249
30.5. "Data" elements .....	249
30.6. Attributes of Universal message task .....	251
30.7. Variables accessible in Groovy code .....	252
30.8. Parameters passed from the listener to the task .....	254
30.9. Parameters usable in task failure email template .....	262
30.10. File Event Listener specific parameters usable in task failure email template .....	262

34.1. Parameters of graph_run .....	268
34.2. Parameters of graph_status .....	269
34.3. Parameters of graph_kill .....	270
34.4. Parameters of sandbox_content .....	270
34.5. Parameters of executions_history .....	271
34.6. Parameters of suspend .....	272
34.7. Parameters of resume .....	273
34.8. Parameters of sandbox create .....	273
34.9. Parameters of sandbox add location .....	273
34.10. Parameters of sandbox add location .....	274
34.11. Parameters .....	274
34.12. Parameters .....	275
34.13. Parameters of server configuration export .....	276
34.14. Parameters of server configuration import .....	277
40.1. Load balancing properties .....	337

---

## List of Examples

22.1. Example of simple configuration defining one new server user. ....	191
--	-----