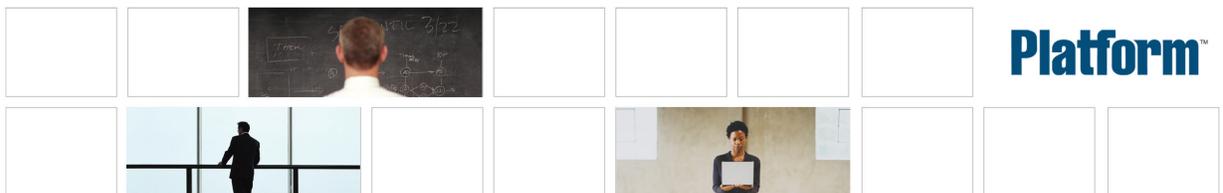

Upgrading your Platform Symphony Version 3.x Cluster

Platform Symphony
Version 5.0
November 2009



Copyright

© 1994-2009 Platform Computing Corporation

All rights reserved.

Although the information in this document has been carefully reviewed, Platform Computing Corporation ("Platform") does not warrant it to be free of errors or omissions. Platform reserves the right to make corrections, updates, revisions or changes to the information in this document.

UNLESS OTHERWISE EXPRESSLY STATED BY PLATFORM, THE PROGRAM DESCRIBED IN THIS DOCUMENT IS PROVIDED "AS IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. IN NO EVENT WILL PLATFORM COMPUTING BE LIABLE TO ANYONE FOR SPECIAL, COLLATERAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, INCLUDING WITHOUT LIMITATION ANY LOST PROFITS, DATA, OR SAVINGS, ARISING OUT OF THE USE OF OR INABILITY TO USE THIS PROGRAM.

We'd like to hear from you

You can help us make this document better by telling us what you think of the content, organization, and usefulness of the information. If you find an error, or just want to make a suggestion for improving this document, please address your comments to doc@platform.com.

Your comments should pertain only to Platform documentation. For product support, contact support@platform.com.

Document redistribution and translation

This document is protected by copyright and you may not redistribute or translate it into another language, in part or in whole.

Internal redistribution

You may only redistribute this document internally within your organization (for example, on an intranet) provided that you continue to check the Platform Web site for updates and update your version of the documentation. You may not make it available to your organization over the Internet.

Trademarks

®LSF is a registered trademark of Platform Computing Corporation in the United States and in other jurisdictions.

™ACCELERATING INTELLIGENCE, PLATFORM COMPUTING, PLATFORM SYMPHONY, PLATFORM JOBSCHEDULER, PLATFORM ENTERPRISE GRID ORCHESTRATOR, PLATFORM EGO, and the PLATFORM and PLATFORM LSF logos are trademarks of Platform Computing Corporation in the United States and in other jurisdictions.

®UNIX is a registered trademark of The Open Group in the United States and in other jurisdictions.

Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

Microsoft is either a registered trademark or a trademark of Microsoft Corporation in the United States and/or other countries.

®Windows is a registered trademark of Microsoft Corporation in the United States and other countries.

Intel®, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Other products or services mentioned in this document are identified by the trademarks or service marks of their respective owners.

Third-party license agreements

<http://www.platform.com/Company/third.part.license.htm>

Third-party copyright notices

<http://www.platform.com/Company/Third.Party.Copyright.htm>

Contents

Overview of the upgrade process	5
Obtain files from Platform Computing	6
Update the database schema	7
Update an Oracle database	7
Update a SQL Server database	8
Set up the new cluster in parallel with the original cluster	9
Plan to avoid conflicts with the original cluster (Linux)	9
Plan to avoid conflicts with the original cluster (Windows)	9
Install and configure Symphony 5.0 on all the hosts	10
Configure the cluster ports and database	12
Migrate your cluster configuration	13
Migrate users.xml	13
Migrate the Windows user passwords	13
Migrate any additional configuration changes	13
Start the new cluster	14
Test the new cluster	14
Migrate consumer and resource group definitions	15
Restart the new cluster	16
Migrate your application to the new cluster	17
Use Symphony 3.1 after installing Symphony 5.0	18
Prepare the environment scripts	18
Set the Symphony 3.1 environment and access Symphony 3.1	18
Remove the original Symphony Cluster	19

Overview of the upgrade process

Use this document if your cluster runs Symphony version 3.1 or 3.2. There is no direct upgrade of Symphony version 3.x hosts to Symphony version 5.0. Direct upgrade to version 5.0 is supported from Symphony version 4.1 only. If your cluster runs Symphony 4.0, upgrade to version 4.1 first, then upgrade to 5.0.

Caution:

Do not try to upgrade your cluster with the RPM option `-U`. This option is destructive because it removes other clusters on the host.

To move from Symphony 3.1 or 3.2 to Symphony 5.0, we strongly recommend that the 5.0 cluster have different management hosts from the original cluster. For compute hosts, we support a parallel installation. There are manual steps to migrate configuration files.

For Linux, when you configure automatic system startup (`egosetrc.sh`) for Symphony 5.0, it disables the feature for the original version.

If the original cluster was Symphony 3.1, the older EGO command prompt is disabled when you install Symphony 5.0.

Benefits:

- It can almost eliminate production downtime required for upgrade
- It is easy to roll back to the previous version

To prevent conflicts, the new Symphony 5.0 cluster has the following differences from the original cluster:

- Different installation directory
- Different ports
- Different syntax to grant root privileges to `egoadmin` (on Linux)
- Different cluster ID (on Windows)

Note:

(for parallel clusters) To maintain consistent performance while running two versions of the software at the same time, management hosts will need enough extra memory to run an additional management console and an additional loader controller.

Note:

(for parallel clusters) For Linux, to maintain automatic startup functionality for the original version of Symphony, link the original `rc` to a different alias before you install Symphony 5.0.

Note:

(for parallel clusters) For Windows, to use the Platform Management Console with the original cluster, input the URL in a browser manually.

Note:

(for parallel clusters) For Windows, to use the Symphony 5.0 CLI, launch it from the Symphony program group shortcut. To use Symphony 3.1 and 3.2, follow the procedure in this document.

Obtain files from Platform Computing

1. Determine which files you need to install the new version of Symphony on all of your existing management hosts and compute hosts. Contact Platform Computing to obtain all the appropriate packages.
 - Installation packages: obtain the package that matches the operating system of your hosts. Windows users require just one package, Linux users require separate packages for management and compute hosts.
 - Database schema and update scripts—obtain the appropriate database schema and update scripts for your supported external database for reports data collection. There are separate scripts to update EGO and Symphony components.
 - If there are any patches applicable to your installation, obtain the appropriate files for each operating system and install the patches after you install the full packages.

Depending on your needs, you may also choose to update Symphony DE, Symphony client, or other Platform software.

Update the database schema

You will use the same database for data from the original and new clusters. The updates you make to the schema for Symphony 5.0 and EGO 1.2.3 will not conflict with data logging from the original cluster.

Update an Oracle database

The database is properly configured and running.

- If your original cluster is Symphony 3.1, you have applied the required Symphony 3.1 patch (Solution 88352) from Platform.
- You have a user name, password, and URL to access the database server. You can create triggers, sequences, tables, and stored procedures.
- You installed the latest JDBC driver (`ojdbc14.jar` or newer) for the database. This driver is available from the following URL:

http://www.oracle.com/technology/software/tech/java/sqlj_jdbc/index.html

- You are able to run `sqlplus`.

To update a Platform product, get the corresponding database update scripts from Platform and update the database schema before you update the cluster. The database upgrade must be performed in two steps. First, upgrade from 3.x to 4.1, then upgrade from 4.1 to 5.0.

For a normal Platform Symphony cluster that consists of Symphony and EGO components, get the following packages:

- 3.1 to 4.1
 - `upgrade_ego_data_1_2_1_to_1_2_3.sql`
 - `upgrade_sym_data_3_1_to_4_1.sql`
- 3.2 to 4.1
 - `upgrade_sym_data_3_2_to_4_1.sql`
- 4.1 to 5.0
 - `upgrade_sym_data_4_1_to_5_0.sql`

If you have Symphony 3.1 with LSF in the cluster, get the following packages:

- `upgrade_ego_data_1_2_1_to_1_2_3_mh.sql`
- `upgrade_sym_data_3_1_to_4_1_mh.sql`
- `upgrade_sym_data_4_1_to_5_0.sql`

1. Extract the database update scripts to any directory.
2. In the command console, change to the `Oracle/Patch` subdirectory. For example, if you extracted the scripts to `sym50_temp`:

```
cd sym50_temp/DBschema/Oracle/Patch
```

3. Run the scripts in this directory to update the database schema.

```
sqlplus user_name/password@connect_string @update_script
```

where

Update the database schema

- *user_name* is the user name on the database server
 - *password* is the password for this user name on the database server
 - *connect_string* is the named SQLNet connection for this database
 - *update_script* is the name of the update script
4. Additionally, after installation, you must copy the jdbc driver to:

```
EGO_TOP/gui/1.2.3u2/lib
```

```
EGO_TOP/perf/1.2.3u2/lib
```

Update a SQL Server database

The database is properly configured and running.

- You have a user name, password, and URL to access the database server. You can create triggers, sequences, tables, and stored procedures.
- You installed the latest JDBC driver for the database. This driver is available from the following URL:

<http://msdn2.microsoft.com/en-us/data/aa937724.aspx>

To update a Platform product, get the corresponding database update scripts from Platform and update the database schema before you update the cluster. The database upgrade must be performed in two steps. First, upgrade from 3.2 to 4.1, then upgrade from 4.1 to 5.0.

- 3.2 to 4.1
 - `upgrade_sym_data_3_2_to_4_1.sql`
- 4.1 to 5.0
 - `upgrade_sym_data_4_1_to_5_0.sql`

1. Extract the database update scripts to any directory.
2. In the command console, change to the Sql Server/Patch subdirectory. For example, if you extracted the scripts to `sym50_temp`:

```
cd sym50_temp/DBschema/SqlServer/Patch
```

3. Run the scripts in this directory to update the database schema.

```
osql -U user_name -P password -d db_name -i @update_script
```

where

- *user_name* is the user name on the database server
 - *password* is the password for this user name on the database server
 - *db_name* is the reporting database name
 - *update_script* is the name of the update script
4. Additionally, after installation, you must copy the jdbc driver to:

```
EGO_TOP/gui/1.2.3u2/lib
```

```
EGO_TOP/perf/1.2.3u2/lib
```

Set up the new cluster in parallel with the original cluster

Plan to avoid conflicts with the original cluster (Linux)

Plan to avoid conflicts with the original cluster by performing extra steps when installing hosts in the new cluster.

Consider the following when installing Linux hosts.

- Before installation, specify a new connection base port that is different to the one you are using in your original cluster.

For example,

```
setenv BASEPORT 9100
```

- Before installation, specify advanced workload execution mode so the new cluster will have the same functionality as the original cluster.

For example,

```
setenv SIMPLIFIEDWEM N
```

- When granting root privileges to egoadmin, the `ego.sudoers` file already exists, so use the `-p` option.

```
egosetsudoers.sh -p
```

- Do not start hosts until you configure the ports.

Plan to avoid conflicts with the original cluster (Windows)

Plan to avoid conflicts with the original cluster by performing some extra steps when installing hosts in the new cluster.

Consider the following when installing Windows hosts.

- Install on different management hosts and follow the section of the installation documentation that describes using the alternate method of installation. Using the alternate method allows you to specify different configurations from the management hosts in the original cluster. To use the alternate method of installation, you need MSI packages for Symphony and EGO.

You need to specify the following parameters when you install. All management hosts in the new cluster must use the same value for each parameter:

- a) For the `INSTALLDIR` parameter, specify a path to the new installation directory.
- b) For the `BASEPORT` parameter, specify a new connection base port.
- c) You must also specify the `CLUSTERID` parameter, and specify a different cluster ID from the original cluster. You must use the same `CLUSTERID` on every host in the new cluster.

For example, use the cluster name as the cluster ID.

Note:

We do not recommend using dates or versions to identify the services because updates you do in future may change the version of the cluster without changing the cluster ID (setting cluster ID is only required if you have to install parallel clusters).

- Do not start hosts until you configure the ports.

If you expand the new cluster at a later date by adding more hosts, remember to use the alternate method of installation and specify the cluster ID if you are adding management hosts. You must use the same CLUSTERID on every host in the new cluster.

Install and configure Symphony 5.0 on all the hosts

These steps are described in detail in the installation documentation for a new cluster. However, do not follow all of the steps as if you are installing a cluster on new hosts. The following steps describe how the installation is different when you have an existing cluster.

1. Plan and prepare for installation. You must customize the installation to avoid conflicts with the original cluster.
2. Install the package on the master host and configure but do not start or test. Consider the following:
 - Do not set up a new database, you will reuse the existing database to preserve the original data.
 - License the cluster as usual.
 - On Linux, you need to specify a different installation directory and different base port. You will probably install with the `--prefix` option.
 - On Linux, set up automatic startup as usual. This will break the feature in the original cluster.
 - On Linux, run `egosetsudoers.sh -p`.
 - On Linux, join the cluster as usual.
 - On Linux, define the host as a management host, as usual.
 - On Windows, specify a different installation directory and different base port, specify cluster ID, and choose advanced workload execution mode.
3. Install the package on the the remaining management hosts and configure but do not start or test. Consider the following:
 - On Linux, you probably need to specify a different installation directory and different base port. You will probably install with the `--prefix` option.
 - On Linux, set up automatic startup as usual. This will break the feature in the original cluster.
 - On Linux, run `egosetsudoers.sh -p`.
 - On Linux, join the cluster as usual.
 - On Linux, define the host as a management host, as usual.
 - On Windows, specify a different installation directory and different base port, specify cluster ID, and choose advanced workload execution mode.
4. Install the package on the compute hosts and configure but do not start or test. Consider the following:
 - On Linux, you probably need to specify a different installation directory and different base port. You will probably install with the `--prefix` option.
 - On Linux, set up automatic startup as usual. This will break the feature in the original cluster.
 - On Linux, run `egosetsudoers.sh -p`.

Set up the new cluster in parallel with the original cluster

- On Linux, join the cluster as usual.
- On Windows, specify a different installation directory and different base port, specify cluster ID, and choose advanced workload execution mode.

Enable master host failover later on, after you have the new cluster running.

Configure the cluster ports and database

Symphony 5.0 will install using some default ports. Change these before you start the new cluster.

1. **Web server ports.** Edit `$EGO_CONFDIR/. . . /gui /conf /server. xml` or `%EGO_CONFDIR%\ . . \ . \gui \conf \server. xml` in all management hosts, replace all occurrences of the default port numbers (8080, 8005, and 8009) with the actual ports you want to use.
2. **Data collection loader controller port.** Edit `$EGO_CONFDIR/. . . /perf /conf /pl c. xml` or `%EGO_CONFDIR%\ . . \ . \perf \conf \pl c. xml` and replace “<Port>4046</Port>” value with the actual port you want to use.
3. **Web service gateway port.** Edit `$EGO_CONFDIR/wsg. conf` or `%EGO_CONFDIR%\wsg.conf` and set `WSG_PORT` to the actual port you want to use.
4. **Database configuration.** Copy the `j dbc` driver to:

`EGO_TOP/gui /1. 2. 3u2/li b`

`EGO_TOP/perf /1. 2. 3u2/li b`

Migrate your cluster configuration

Migrate users.xml

1. Back up the file `$EGO_CONFDIR/users.xml` installed with the new cluster.
2. Copy `users.xml` from the original cluster to `SEGO_CONFDIR/users.xml` (in the new cluster).

Migrate the Windows user passwords

- Copy the hidden file `passwd.ego` from the original cluster to `SEGO_CONFDIR/passwd.ego` (in the new cluster).
- If you do not copy the file, you need to run `egosh ego execpasswd` to set a password for all consumer execution users on Windows in the new cluster.

Migrate any additional configuration changes

1. If you modified any parameters in the `ego.shared` file of the original cluster, back up `ego.shared` under `EGO_CONFDIR` in the new cluster, then manually recreate your changes in the new cluster. You might have added entries to the `HostType`, `HostModel`, or `Resource` sections.
2. If you modified any parameters in the `ego.cluster.cluster_name` file of the original cluster, back up `ego.cluster.cluster_name` under `EGO_CONFDIR` in the new cluster, then manually recreate your changes in the new cluster.
3. If you modified any parameters in the `ego.conf` file of the original cluster, back up `ego.conf` under `EGO_CONFDIR` in the new cluster, then manually recreate your changes in the new cluster.
4. If you modified any parameters in the service profiles of the original cluster (for example `sd.xml` or `rs.xml`), back up the files in the new cluster, then manually recreate your changes in the new cluster.

Start the new cluster

1. Start the new cluster.

On the master host, run:

```
egosh ego start all
```

The users and resources from your original cluster will be recognized by the new cluster, so you can use them when you configure consumers and resource groups.

Test the new cluster

These steps are described in detail in the Windows or UNIX installation documentation for a new cluster.

1. You may test the management and compute hosts after the cluster has started. Test the Platform Management Console and test workload. Make sure the cluster configuration is as desired.

Migrate consumer and resource group definitions

1. Back up files installed with the new cluster:

- `SEGO_CONFDIR/ResourceGroups.xml`
- `SEGO_CONFDIR/ConsumerTrees.xml`

2. Migrate the consumers you created in the original cluster to the new cluster. To do so, manually copy `<Consumer>` sections from the original `ConsumerTrees.xml` file to the new `ConsumerTrees.xml` file. For each consumer you created, you need to migrate three sections: `ConsumerHierarchy`, `DistributionTree for ManagementHosts`, and `DistributionTree` for the appropriate compute host resource group.

The built-in consumer has changed. The original cluster has `/SampleApplications/SOATesting` or `/SampleApplications/SOATesting32`. The new cluster has `/SymTesting/Symping50`.

3. Migrate the resource groups you created in the original cluster to the new cluster. To do so, manually copy `<ResourceGroup>` sections from the original `ResourceGroups.xml` file to the new `ResourceGroups.xml` file.

If your resource groups use static host lists, make sure those hosts belong to the new cluster.

Restart the new cluster

Restart the new cluster

1. Restart the new cluster.

On the master host, run:

```
egosh ego restart
```

2. Use the Platform Management Console to check your changes in the new cluster.

Migrate your application to the new cluster

1. Deploy the application service package to the new cluster.
You do not need to make changes to the service package for it to run on the new version.
2. If the original application is from Symphony 3.1, change SOAM_HOME and VERSION_NUM in the application profile to point to the new cluster.
3. Upgrade the application profile to be compatible with the new Symphony version. Importing the application profile upgrades it automatically.
 - a) Navigate to Symphony Workload > Configure Application and select the consumer for the application.
 - b) On the Applications tab, open the Global Actions menu and choose Add Application (Advanced).
 - c) On the Register a new application window, click Import.
 - d) On the Register/Update an application window, click Browse.
 - e) Browse to the application profile.
 - f) Click Import.
 - g) When prompted, click OK to convert the version.
 - h) On the Register a new application window, click Register.
4. Test the application in the new cluster.

Note:

3.x DE clients can still be used to deploy and register applications on the 5.0 grid.

Use Symphony 3.1 after installing Symphony 5.0

After you install Symphony 5.0, Symphony 3.1 Windows program shortcuts will not be functional. Take the following steps to use Symphony 3.1

1. Prepare the environment scripts (do this once).
2. Set the environment and access Symphony 3.1 (do this each time you want to use Symphony 3.1).

Prepare the environment scripts

1. Copy egoenv. bat from %EGO_TOP%\1. 2. 3u2\bin in the new 5.0 cluster.
2. Edit the copied file egoenv. bat:
 - modify EGOLIB, EGOBIN and EGOETC to point to the Symphony 3.1 cluster
 - modify EGO_CONFDIR and PATH to point to the Symphony 3.1 cluster
3. Copy symenv. bat from %EGO_TOP%\soam\5. 0\win32-vc7\bin (32-bit) or %EGO_TOP%\soam\5. 0\w2k3_x64-vc7-psdk\bin (64-bit)
4. Edit the copied file symenv. bat:
 - modify SOAM_HOME to point to the Symphony 3.1 cluster
 - modify SOAM_LIB, SOAMBIN and SOAMETC to point to the Symphony 3.1 cluster
 - modify PATH to point to the Symphony 3.1 cluster

Set the Symphony 3.1 environment and access Symphony 3.1

1. Open a Windows command console, and run:
 1. egoenv. bat
 2. symenv. bat
2. Run Symphony 3.1 and EGO CLI as usual.

Remove the original Symphony Cluster

1. Once the new cluster is working properly, complete the upgrade process by moving all workload to the new cluster and removing the original cluster.

Index

A

application profiles
 migrate to new version
 migration
 of application profile 17
upgrading 17

R

reports
 database
 Oracle 7, 8
 schema 7, 8