
ISF Adaptive Cluster for Symphony User Guide

Platform ISF Adaptive Cluster for Symphony
Version 1.1
December 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 JOB SCHEDULER, PLATFORM ISF, 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

ISF Adaptive Cluster for Symphony	5
Scope	5
About ISF Adaptive Cluster	5
Install ISF Adaptive Cluster for Symphony on Windows	8
Install ISF Adaptive Cluster for Symphony on Linux	10
Verify and configure ISF Adaptive Cluster for Symphony	12
1. Check the ISF Adaptive Cluster installation	12
2. Configure ISF Adaptive Cluster for Symphony	12
Configuring IS_LICENSE_ACCT_PATH	13
Monitoring performance metrics	14
View the CPU and memory utilization of the cluster	14
Uninstall ISF Adaptive Cluster on Windows	15
Uninstall ISF Adaptive Cluster on Linux	16
Troubleshooting	17
Known issues	17
FAQ	17

ISF Adaptive Cluster for Symphony

Platform ISF Adaptive Cluster for Symphony maximizes utilization of computing resources by harvesting VMs for Symphony workload when the demand exceeds Symphony's dedicated cluster resources. The VMs, which are shared with virtual desktops and other non-Symphony applications, can be made available to Symphony during periods of low VM utilization by these applications.

Scope

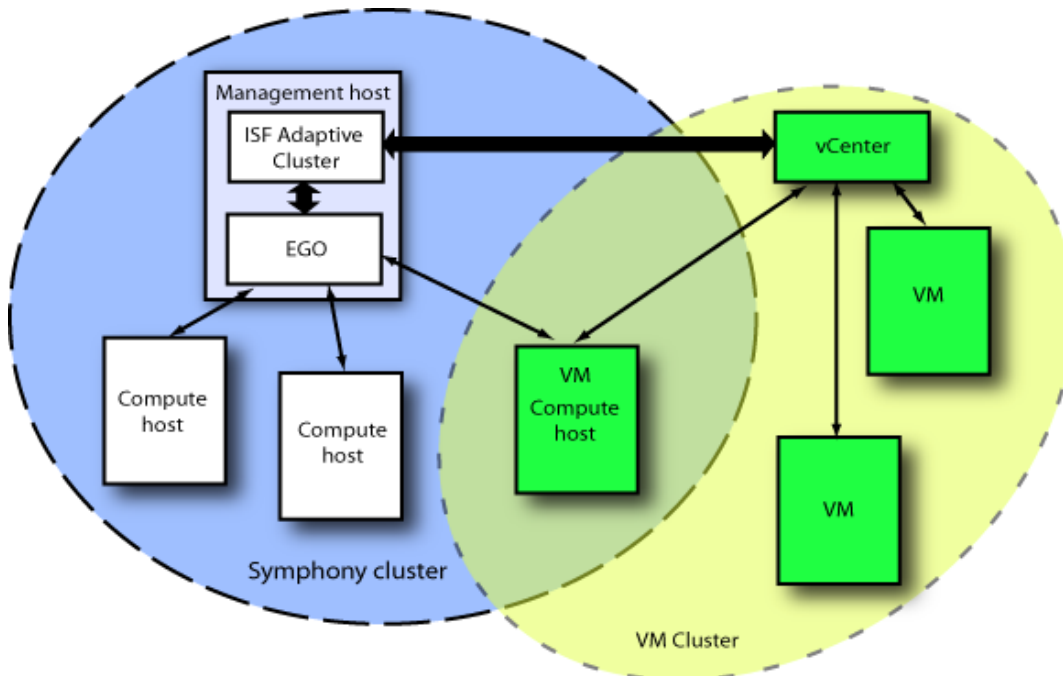
Applicability	Details
Operating system (management hosts)	<ul style="list-style-type: none"> 64-bit Windows 64-bit Linux2.6 and glibc 2.3 or above
Symphony	Version 5.0
VMWare	ESX 4.0, 3.5.x with vCenter 2.5.x, 4.0

About ISF Adaptive Cluster

The ISF Adaptive Cluster uses configurable CPU and memory utilization thresholds to determine if VMs can be made available for Symphony workload. Once VMs become available, the ISF Adaptive Cluster can manage their personalities according to the application's requirements.

Functional flow

The ISF Adaptive Cluster starts as an EGO service immediately after the Symphony EGO master is up and running. On startup, the ISF Adaptive Cluster connects to both the Symphony master and to each configured VMware vCenter. After the connections have been established, the ISF Adaptive Cluster identifies all the Symphony virtual machines in the virtual clusters. These logical compute hosts are located through the EGO API in the Symphony cluster. Initially, these hosts are closed and therefore not available to Symphony.



The ISF Adaptive Cluster periodically monitors the VMware clusters to gather cluster-level utilization metrics (CPU and memory utilization) and compare them to preconfigured threshold values. If there is unsatisfied Symphony workload demand, three possible actions can be taken for each VMware cluster during each scheduling cycle:

1. If all of the metrics are between corresponding Low and High thresholds, no action is taken, i.e., no additional VMs are made available to Symphony and no VMs are removed.
2. If either cluster metric is above its High threshold value (either CPU or memory utilization), the ISF Adaptive Cluster goes through all the physical hosts in the cluster retrieving host-level values for the utilization metrics. If a value is above the High threshold for the specific host, one Symphony compute host residing in a VM at this physical host is removed by closing it. This sequence is repeated, i.e., one VM is removed at each physical host where the utilization is above the High threshold, up to the preconfigured maximum number of VMs that can be removed during a scheduling cycle. As a result, EGO reclaims these hosts and the Symphony workload on these hosts is queued.
3. If both cluster metrics are below the corresponding Low threshold values, the ISF Adaptive Cluster retrieves the unsatisfied demand from the Symphony applications' consumers. The ISF Adaptive Cluster cycles through all the physical hosts in the cluster arranging the Symphony workload VMs in a host list in decreasing order of the physical host's utilization values. The ISF Adaptive Cluster then opens one VM, which resides on the least loaded physical host, to make it available to Symphony. This sequence continues, i.e., one VM is made available at each physical host where the CPU and memory utilization values are below the Low thresholds, up to the preconfigured maximum number of VM hosts, but not more than the initial unsatisfied demand requires.

Regardless of the comparison results, each time the Symphony application's demand for compute resources is satisfied, the ISF Adaptive Cluster removes all VM hosts from Symphony that currently are not allocated to the applications' consumers. This cleanup operation prevents violation of utilization thresholds for non-Symphony virtual VMs as Symphony will not allocate and schedule any suddenly demanded workload to such hosts at any moment until the next scheduling cycle.

Personality management

Personality management involves maximizing the availability of VMs that match an application's requirements in an attempt to satisfy a consumer with the highest demand. For example, applications are often deployed in mixed clusters made up of resources running different OS platforms. When an application has demand, its workload needs to be routed to resources with the appropriate OS. In this case, the ISF Adaptive Cluster will provision the VMs that meet the OS requirements of the application. This enables cluster resources to be more versatile and better equipped to handle diverse application requirements.

Scheduling policy

The ISF Adaptive Cluster scheduling policy attempts to balance the unsatisfied demands of consumers. The objective of the scheduling policy is to make VMs available that satisfy the consumer with the most demand.

For example, assume there are two consumers, A and B, configured to use resource groups resGroupA and resGroupB, respectively. Consumer A has a demand for 20 slots and consumer B has a demand for 10 slots. The maximum number of VMs to make available or remove for each scheduling cycle is configured to be 4. If the ISF Adaptive Cluster decides it can make VMs available, it will open 4 VMs from resource group resGroupA because the demand for consumer A is more than the demand for consumer B. In each scheduling cycle, it will continue to make VMs available until either there are no VMs left to open or the demand for consumer A is lower than consumer B. If the ISF Adaptive Cluster decides it will remove VMs, it will remove 4 VMs from resource group resGroupB; during each scheduling cycle, it will continue to do so until either there are no VMs left to remove or the demand for consumer A is lower than consumer B.

Install ISF Adaptive Cluster for Symphony on Windows

1. Plan and prepare for installing ISF Adaptive Cluster for Symphony on Windows hosts

1. Requirements.
 - MSI version 2.0 or later is required.
 - Symphony cluster is already installed and working properly. Symphony master is a Windows host.
 - You must use the same Symphony installation directory as the existing Symphony cluster.
2. Obtain the "ISF Adaptive Cluster for Symphony" distribution file
ISFAdapt i veCl usterForSymphony_1. 1. msi .

2. Install "ISF Adaptive Cluster for Symphony" on all Windows management hosts

Use the same installation user account as Symphony to log onto all management hosts and run the installer.

Install parameters (also syntax for MSI command-line installation) Notes

Installation Directory	Cluster installation directory.
command-line syntax:	You must input the Symphony installation directory.
INSTALLDIR= <i>path</i>	
e.g.	
I NSTALLDI R=C: \Symphony	

3. Configure the master host

All EGO management hosts should use ISF Adaptive Cluster files in the shared directory, instead of using ISF Adaptive Cluster files on the local host. To update the host, run "i sconfi g mghost i s" on the master host, for example:

EGO_TOP\is\1.1\bin\isconfig mghost i s

where EGO_TOP represents the installation directory.

Exit the command console to reset your environment.

4. Update the cluster license

Follow these steps when you want to add or update a license for "ISF Adaptive Cluster for Symphony".

1. Acquire an "ISF Adaptive Cluster for Symphony" license key from Platform.
2. Append the license key to %EGO_CONFDIR%/license.dat.

5. Run the cluster configuration wizard

The cluster configuration wizard automatically determines the correct configuration for your cluster. Log on to the Platform Management Console as cluster administrator (Admin). The wizard will start automatically (you must enable pop-up windows in your browser). Follow all the steps.

Install ISF Adaptive Cluster for Symphony on Linux

1. Plan and prepare for installing ISF Adaptive Cluster for Symphony on Linux hosts

1. Requirements.
 - Symphony cluster is already installed and working properly. Symphony master is a Linux host.
 - You must use the same Symphony installation directory as the existing Symphony cluster.
2. Obtain the "ISF Adaptive Cluster for Symphony" Linux distribution file
ISFAdaptiveClusterForSymphony-Linux-1.1.0-buildnumber.rpm.

2. Install "ISF Adaptive Cluster for Symphony" on all Linux management hosts

Use the same installation user account as Symphony to log onto all management hosts.

Run the installer:

```
rpm -iv ISFAdaptiveClusterForSymphony-Linux-1.1.0-buildnumber.rpm --prefix EGO_TOP
```

where EGO_TOP represents the Symphony installation directory.

3. Configure the master host

All EGO management hosts should use ISF Adaptive Cluster files in the shared directory, instead of using ISF Adaptive Cluster files on the local host. To update the host, run "isconfig mghost is" on the master host, for example:

```
EGO_TOP/is/1.1/bin/isconfig mghost is
```

where EGO_TOP represents the installation directory.

Exit the command console to reset your environment.

4. Update the cluster license

Follow these steps when you want to add or update a license for "ISF Adaptive Cluster for Symphony".

1. Acquire an "ISF Adaptive Cluster for Symphony" license key from Platform.
2. Append the license key to %EGO_CONFDIR%/license.dat.

5. Run the cluster configuration wizard

The cluster configuration wizard automatically determines the correct configuration for your cluster. Log on to the Platform Management Console as cluster administrator (Admin). The wizard will start automatically (you must enable pop-up windows in your browser). Follow all the steps.

Verify and configure ISF Adaptive Cluster for Symphony

Test the cluster, then configure and test ISF Adaptive Cluster for Symphony.

1. Check the ISF Adaptive Cluster installation

After installation, there will be one new EGO service "IS" registered in EGO and started. Log on to the Platform Management Console as cluster administrator (Admin) and check "System Services". The "IS" service should be listed and its state should be "STARTED". Also, one new quick link "ISF Adaptive Cluster" will be listed in the dashboard and in the first tab of the PMC main console.

2. Configure ISF Adaptive Cluster for Symphony

1. Add new vCenter servers

To share resources between Symphony and VM Clusters, the IS service needs to get vCenter information and connect to it to get VM Cluster load information.

1. Log on to the Platform Management Console as cluster administrator (Admin).
2. Click on the ISF Adaptive Cluster quick link on the dashboard.
3. Click the vCenter Server tab.
4. On the Global Actions dropdown list, select Add vCenter Server.

The Add new vCenter Server page displays.

5. Input vCenter server, user name, and password information.
6. Click Add to add the vCenter server.

The Add New vCenter page closes and the original vCenter page gets refreshed with the new vCenter server information.

7. On the VM clusters page, check the VM cluster status. The VM clusters information should display after the vCenter service is added.

2. Configure ISF Adaptive Cluster parameters

After adding the vCenter server into the ISF Adaptive Cluster, you must configure the ISF Adaptive Cluster parameters.

1. Log on to the Platform Management Console as cluster administrator (Admin).
2. Click on the ISF Adaptive Cluster quick link on the dashboard.
3. Click the Configuration tab to configure global settings.

On the Configuration page, you can set which resource groups to monitor, configure the maximum number of VMs to open/close, and configure the scheduling interval. For thresholds, you should change the default to a proper value and, if necessary, add a new timeslot.

4. To configure individual VM clusters, follow these steps:
 - a) Click on the VM Clusters tab.
A list of VM clusters displays.
 - b) Click on the VM Cluster name that you want to configure.
A list of hosts displays.
 - c) Click on the Configuration tab.
On the VM cluster configuration page, you can customize threshold settings for this VM cluster.

3. Test resource sharing between Symphony and VM Clusters

Submit Symphony workload to demonstrate sharing between Symphony and VM clusters.

1. Check VMs host status when there is no workload:
 - a) Log on to the Platform Management Console as cluster administrator (Admin).
 - b) Click on the ISF Adaptive Cluster quick link on the dashboard.
The VM Clusters page displays.
 - c) Click on any VM Clusters to open the VMs page.
All VM hosts are "Closed" when no workload is running. Also, on the VM Clusters page, if no other application is running, the VM Cluster CPU ut and Mem ut are both below the configured thresholds.
2. Submit Symphony workload:
symping -u Admin -x Admin -r 10000 -m 100
3. Monitor the system to observe resource sharing:
 - Monitor workload
On the Symphony Workload page, select Application Name > SessionID > Task page. Monitor the hosts that are running the tasks. Initially, only Symphony-dedicated hosts will start and run Symphony workload. Subsequently, you can observe new VMs hosts becoming available and starting to run Symphony workload.
 - Monitor resource allocation
Select ISF Adaptive Cluster > VM Cluster > VMs. You can observe VMs, which are shared between Symphony and VM Clusters, are available for Symphony to run workload.

Configuring IS_LICENSE_ACCT_PATH

If you want to change the default location for the `is.<clustername>.license.acct` file, you need to define the `IS_LICENSE_ACCT_PATH` environment variable in the `IS.xml` service definition file. You can define the environment variable using the PMC; refer to the "Update a service" topic in the Symphony Cluster and Application Management Guide.

Monitoring performance metrics

The CPU and memory utilization metrics for each cluster can be viewed via the ISF Adaptive Cluster link in the PMC. The actual values represent the average utilization of all the physical hosts in the cluster. The preconfigured utilization thresholds are provided on the same page as the actual utilization values so that the load on the cluster can easily be observed.

View the CPU and memory utilization of the cluster

1. Log on to the Platform Management Console as cluster administrator (Admin).
2. Click on the ISF Adaptive Cluster quick link on the dashboard.
3. Click on the VM Clusters tab.

A list of VM clusters displays.

Observe that the VM Cluster load increases and may increase above the configured low threshold as Symphony workload runs on VMs in each cluster. You can also try to run non-Symphony applications and observe the VM cluster load increasing above the configured high threshold. When CPU or memory utilization is above the configured high threshold, click on the VM cluster name. Observe that VMs are being closed.

Uninstall ISF Adaptive Cluster on Windows

On Windows, there is a separate msi entry point for "ISF Adaptive Cluster for Symphony" in Add/Remove programs. You can click Remove or double click the original msi package to uninstall the product.

The uninstaller will remove installed binaries and local conf files. The shared location is not affected. Also, if you already ran the post-install wizard, the new EGO IS service will not be removed. You must do a manual cleanup.

Manually remove these files under shared conf:

share_top\is

- eservice\esc\conf\services\IS.xml
- gui\conf\Catalina\localhost\isgui.xml
- gui\conf\pmcconf\pmc_conf_is.xml
- gui\conf\useracl\users_is.acl
- gui\conf\userconf\userconf.is.*.xml

Uninstall ISF Adaptive Cluster on Linux

To remove installed binaries and local files, run the following command:

```
rpm -e ISFAdaptiveClusterForSymphony-linux-1.1.0_XXXXX.rpm
```

The shared location is not affected. Also, if you already ran the post-install wizard, the new EGO IS service will not be removed. You must do a manual cleanup.

Manually remove these files under shared conf:

share_top/is

- eservice/esc/conf/services/IS.xml
- gui/conf/Catalina/localhost/isgui.xml
- gui/conf/pmconf/pmc_conf_is.xml
- gui/conf/useracl/users_is.acl
- gui/conf/userconf/userconf.is.*.xml

Troubleshooting

Log file location

Log files are an essential troubleshooting tool for production and test environments. By default, the IS Orchestrator log is located in IS_HOME/logs.

Component	Log file name	What it contains
IS Orchestrator	orch.log.hostname	Logs IS Orchestrator configuration. Logs retrieved results for vCenter and Symphony. Logs evaluation results for each monitoring cycle.
EGO IS service	run.log	Logs service startup messages.
PTM (PTM is the built-in framework of IS)	ptm.log	Logs PTM ERROR messages.

Log entry format

The standard format for log file entries is: "date [process] log_level class - description" where the date is expressed in YYYY-MM-DD hh:mm:ss.sss. For example:

```
2009-11-19 13:47:32,779 INFO PTM: Timers. 1
com.platfom.is.orchestrator.Adapters.PTMAadapter - Evaluation result: open
host num=0, close host num=0
```

Known issues

Workload is not dispatched to VMs

The Symphony workload VM is opened by IS Orchestrator, but Symphony application workload is not distributed to this VM by the Symphony scheduler. During the next scheduling cycle, the IS Orchestrator will close this VM because it is not allocated. Then the Symphony workload VM is opened and closed alternately, but workload is never dispatched to any VM by the Symphony scheduler. This problem is probably caused by the opened VM, which cannot satisfy the allocation request, for example:

- The quota of a consumer is reached. Try "egosh rg -l" to check the quota of the consumer.
- The client that submits the request is disconnected. Try "egosh client view" to check the client status.

FAQ

Why are no VMs open even though there is demand and the cluster utilization is below the thresholds?

It is possible that the resource groups being monitored contain overlapping resources. Check the log file `orch.log.<hostname>` to see if VMs are in their expected resource groups; if not, a misconfiguration probably caused resources to overlap between resource groups.

Why can't I add a new vCenter server to the cluster?

Adding a vCenter server in the PMC can fail when the DNS is unable to resolve the IP address or hostname of the vCenter server. To solve this problem, add the IP address and hostname for the vCenter server to the system `hosts` file in the `/etc` directory (Unix) and `C:\WINDOWS\system32\drivers\etc` (Windows).

What should I do if the IS service is still running on the original master host after failover?

This IS service has no impact other than it continues to write to the license audit log at each scheduling interval. It is suggested to terminate the service manually.

Why isn't the PMC displaying VMs?

Check that the VMware Tools package has been installed on the VM hosts. VMs will not be displayed if this package is not installed.