

HP Performance Insight

For the Solaris operating system

Software Version: 5.41

Configuration Guide for the Sun Cluster

July 2012



Legal Notices

Warranty

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notices

© Copyright 2009, 2011-2012 Hewlett-Packard Development Company, L.P.

Trademark Notices

Windows® and Windows Server™ 2003 are U.S. registered trademarks of Microsoft® Corp.

UNIX® is a registered trademark of The Open Group.

Oracle and Java are registered trademarks of Oracle Corporation and/or its affiliates.

Acknowledgements

This product includes Xerces XML Java Parser software, which is Copyright (c) 1999 The Apache Software Foundation. All rights reserved.

This product includes JDOM XML Java Parser software, which is Copyright (C) 2000-2003 Jason Hunter & Brett McLaughlin. All rights reserved.

This product includes JClass software, which is (c) Copyright 1997, KL GROUP INC. ALL RIGHTS RESERVED.

This product includes J2TablePrinter software, which is © Copyright 2001, Wildcrest Associates (<http://www.wildcrest.com>)

This product includes Xalan XSLT Processor software, which is Copyright (c) 1999 The Apache Software Foundation. All rights reserved.

This product includes EXPAT XML C Processor software, which is Copyright (c) 1998, 1999, 2000 Thai Open Source Software Center Ltd and Clark Cooper Copyright (c) 2001, 2002 Expat maintainers.

This product includes Apache SOAP software, which is Copyright (c) 1999 The Apache Software Foundation. All rights reserved.

This product includes O'Reilly Servlet Package software, which is Copyright (C) 2001-2002 by Jason Hunter, jhunter_AT_servlets.com. All rights reserved.

This product includes HTTPClient Package software, which is Copyright (C) 1991, 1999 Free Software Foundation, Inc. 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA.

This product includes Perl software, which is Copyright 1989-2002, Larry Wall. All rights reserved.

This product includes Skin Look And Feel software, which is Copyright (c) 2000-2002 L2FProd.com. All rights reserved.

This product includes nanoXML software, which is Copyright (C) 2000 Marc De Scheemaecker, All Rights Reserved.

This product includes Sixlegs PNG software, which is Copyright (C) 1998, 1999, 2001 Chris Nokleberg

This product includes cURL & libcURL software, which is Copyright (c) 1996 - 2006, Daniel Stenberg, <daniel@haxx.se>. All rights reserved.

This product includes Quartz - Enterprise Job Scheduler software, which is Copyright 2004-2005 OpenSymphony

This product includes Free DCE software, which is (c) Copyright 1994 OPEN SOFTWARE FOUNDATION, INC., (c) Copyright 1994 HEWLETT-PACKARD COMPANY, (c) Copyright 1994 DIGITAL EQUIPMENT CORPORATION, Copyright (C) 1989, 1991 Free Software Foundation, Inc. 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

This product includes DCE Threads software, which is Copyright (C) 1995, 1996 Michael T. Peterson

This product includes Jboss software, which is Copyright 2006 Red Hat, Inc. All rights reserved.

This product includes org.apache.commons software developed by the Apache Software Foundation (<http://www.apache.org/>).

Documentation Updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent updates or to verify that you are using the most recent edition of a document, go to:

<http://h20230.www2.hp.com/selfsolve/manuals>

This site requires that you register for an HP Passport and sign in. To register for an HP Passport ID, go to:

<http://h20229.www2.hp.com/passport-registration.html>

Or click the **New users - please register** link on the HP Passport login page.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HP sales representative for details.

Support

Visit the HP Software Support web site at:

www.hp.com/go/hpsoftwaresupport

This web site provides contact information and details about the products, services, and support that HP Software offers.

HP Software online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support web site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HP support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and sign in. Many also require an active support contract. To register for an HP Passport ID, go to:

<http://h20229.www2.hp.com/passport-registration.html>

To find more information about support access levels, go to:

http://h20230.www2.hp.com/new_access_levels.jsp

Contents

1	Introduction	9
2	Sun Cluster Implementation Pre-requisites.....	10
	Hostname and IP Address	10
	Shared and Local Files	11
	Sun Cluster Data Services.....	11
3	Configuring PI on Sun Cluster.....	13
	Task 1: Setting Failover Machines	13
	Task 2: Verifying System Resources and Machine Patch Levels.....	14
	Task 3: Synchronizing System Clocks	15
	Task 4: Setting Global Disk Partitions and Shared Disk for Quorum Device	15
	Task 5: Installing a Database	15
	Sybase Database.....	15
	Oracle Database.....	16
	Task 6: Installing PI	17
	On Sybase Database.....	17
	On Oracle Database.....	17
	Task 7: Installing Report Packs	18
	Install the Report Packs from the RNS CD	18
	Task 8: Editing trendtimer.sched	19
	Configuring PI to function in a Cluster Setup.....	21
	Changing Hostname on a PI 5.41 Installation	22
	Changing Hostnames and IP Addresses in the Configuration Files and PI Tables ...	22
	Task 9: Propagating Shared and Local Files	27
	Sybase Shared and Local Files	27
	Oracle Shared and Local Files	28
	PI Shared and Local Files	29

4	Configuring PI on Sun Cluster.....	31
5	Installing a Patch in a Sun Cluster Setup	34
6	Troubleshooting Sun Cluster	36
	Verifying the Cluster Configuration	36
	Viewing Sun Cluster Logs	36
	Viewing PI Resource Group Log Files	36
A	References.....	37
	Software Depots and Versions	37
	PI.....	37
	Sybase	37
	Oracle.....	37
	Report Packs.....	37
	Location for PI and Sun Cluster Manuals	38
	Sun Cluster.....	38
	PI.....	38

1 Introduction

This guide provides instructions on how to configure HP Performance Insight in a high availability environment using the Sun Cluster version 3.2. Configuring PI in high availability environment improves its availability when used as a mission critical application.

▶ This guide does not cover how to set up a Sun Cluster.

High availability (HA), as used in this guide, implies that single points of failure (SPOFs) are eliminated from the environment. Examples of SPOFs include System Processing Unit (SPU), disks and disk controllers, LAN interface cards and cables, and power connection.

These potential SPOFs are removed by clustering the SPUs, mirroring or using RAID technology, providing redundant LAN interface cards, and attaching UPSs to the system. Clustering also facilitates operating system and application upgrades. HA solutions, however, cannot protect against failures caused by defects in applications and OS panics.

Configuration Considerations

This guide covers the Configuration of a two node cluster using a shared disk connected to each node through SCSI cable (and not through NFS) with PI 5.41 installed in a standalone setup (All PI components installed on a single machine.)

▶ The scripts provided in this guide can also be used in PI setups where the Performance Manager and database server are installed on a system separate from the Web Access Server/Web Application Server.

The database can be either Sybase or Oracle.

2 Sun Cluster Implementation Prerequisites

The primary requirement for a Sun Cluster implementation is that in case of a failover, all the designated processes are initiated seamlessly on the secondary (failover) system and the LAN connection is moved.

To provide access to the application processes regardless of the physical system on which they run, you must assign a logical hostname and associated IP address to the one and only physical system currently running. Though both the primary and secondary systems access the same shared disks, it is never at the same time. Only the system currently running the application must read and write to the shared disks.

There are, therefore, two core considerations when implementing high availability with Sun Cluster:

- Assigning a logical hostname and associated IP address
- Appropriate handling of files both shared and local



OVPI supports active-passive cluster type. At any point of time only one node is active and the other is a standby. It is a failover cluster support.

Hostname and IP Address

For a simple system failover implementation there must be two or more servers that are each capable of hosting a unique “floating” hostname and associated IP address that are not associated with a physical system (that is a “logical” hostname and IP address).

Only one of the failover systems will host this name and IP address at any given time. During a system failover, the hostname and IP is “transferred” from the failed server to the failover system. The applications, however, will always access the same hostname and IP.

Both Sybase and Oracle create some files that store the information about the hostname and IP address. PI also stores the values of hostname and IP

during installation for use at run time. You must, therefore, set the logical name after installing the database or PI.

Shared and Local Files

In the Sun Cluster high availability setup, the PI installation files are placed under a specified directory on the shared disk. However, some of the files can not be placed on the shared disk and require special handling. These files are grouped under the following categories:

- Static files – These files are placed on the primary system during installation. These files do not change. Do *ONE* of the following:
 - Copy these files to each of the secondary systems.
 - Move these files to the shared disk and create symbolic links, which reference the file locations on the shared disk, on all failover systems.
- Dynamic files – These files are modified by the application during normal use. You can move these files to the shared disk and create symbolic links, which reference the file locations on the shared disk, on all failover systems.
- System files – These files are shared with other HP Software applications. These files are, for example, `/etc/passwd`, `/etc/group`.
You must perform similar actions on these files (for example, add group, add user) on each of the failover systems.
- Files shared with other HP Software applications – Example for such a file is `/opt/OV`. You must handle these files like static files.

Sun Cluster Data Services

Sun Cluster uses Data Services for managing resources like:

- Network Interface Cards (NIC)
- IP addresses
- Physical disk drives
- Logical disk drives
- Databases such as Oracle or Sybase
- Applications such as PI

For each Resource Type known to Sun Cluster, there must be a Data Service agent for monitoring the status and health of the resource, for bringing the resource online, and for shutting down the resource. The data service agents for NICs, IP addresses, physical disk drives, and logical disk drives are bundled with the Sun Cluster. For Oracle and Sybase databases, Sun Cluster provides separate packages that include the Data Service agents.

PI comes bundled with its own set of Data Service agent.

You can create your own data service agents in the following manner:

- Use the Resource Management API (RMAPI) and the Data Service Development Library (DSDL) libraries provided by Sun Microsystems. See the accompanying documentation for step by step instructions on building data Services on top of these libraries using ksh or C language.
- Use the Generic Data Service (GDS) package - SUNW.gds provided by Sun Microsystems. This package performs the task of integrating with the Sun Cluster Resource Manager. You only have to provide the Start, Stop, and Probe scripts along with the configuration parameters.

The instructions provided in this guide are based on the GDS solution and will help you manage PI and the underlying database as separate resources with the former depending on the latter.

You can, however, manage the Oracle or Sybase database using any other Data Service as well. To do this, you will have to alter the scripts and configuration details provided later in this guide.

3 Configuring PI on Sun Cluster

The following variables/names are used in this chapter:

Name	Description
/NFS	File system on the shared disk to be shared by the two-node cluster
failover1	The primary node of the failover cluster
failover2	The secondary node of the failover cluster to which control transfers in a failover event
logicalhost	The hostname assigned to the node that currently controls of the files and applications; also known as the “floater” or “virtual” hostname

Configuring PI on Sun Cluster involves the following tasks:

- Task 1: *Setting Failover Machines*
- Task 2: *Verifying System Resources and Machine Patch Levels*
- Task 3: *Synchronizing System Clocks*
- Task 4: *Setting Global Disk Partitions and Shared Disk for Quorum Device*
- Task 5: *Installing a Database*
- Task 6: *Installing PI*
- Task 7: *Installing Report Packs*
- Task 8: *Editing trendtimer.sched*
- Task 9: *Completing the PI Configuration*
- Task 10: *Propagating Shared and Local Files*

Task 1: Setting Failover Machines

To set failover machines, follow these steps:

- 1 In a two node cluster, set one system to function as the primary node and one system to function as secondary node and install the Sun Cluster software on both.

Ensure that both the primary and secondary nodes have similar system resources.

This is because when you install the database and PI; the database tuning parameters are set based on the system resources available on the machine on which it is installed. In case there is a failover and the secondary system has significantly less resources than the primary system, the parameters used to initialize and tune the database might not only be sub-optimal but could also prevent the database from being started.



If your hardware availability is limited, you might want to install the database and PI on the less powerful machine to avoid incompatible settings in case of a failover. However, it is best practice to use the more powerful machine as the primary system.

- 2 Set up a shared disk.
- 3 Set one floating “virtual” IP address.
This IP address shares the primary Network Interface Card (NIC) with the “real” IP address of each system.
- 4 Obtain a NIC and IP address to mount the PI resource group.

Task 2: Verifying System Resources and Machine Patch Levels

Follow these steps:

- 1 Verify that each failover node has adequate resources (for example, RAM, swap, kernel settings) and patch levels for the operating system and Java before installing the database and PI.
- 2 See the pre-installation checklist provided in the *HP Performance Insight Installation Guide for Unix* before installing PI.
- 3 See the Oracle Installation Guide for your operating system before installing the Oracle database.

Task 3: Synchronizing System Clocks

PI has many time-dependent, time-critical processes. Therefore, it is important that each machine in the cluster uses the same source for time synchronization to keep data collection, aggregation, reporting, and logging correct and consistent.

You can use the Network Time Protocol Daemon (xntpd) to keep the system's time-of-day in agreement with Internet standard time servers.

For example, to set up the xntpd for Solaris, follow these steps:

- Copy `/etc/inet/ntp.server` to `/etc/inet/ntp.conf`
- Add a line to `ntp.conf` that identifies the local time server providing synchronization
- Launch the xntpd daemon. Type:
`/etc/init.d/xntpd start`.

See the xntpd (1M) manpage for details.

Task 4: Setting Global Disk Partitions and Shared Disk for Quorum Device

To complete the installation and configuration of a Sun Cluster installation, you must set up:

- Global Disk Partitions that reside either on the local hard disk of each failover node or on a common shared external drive.
- Quorum Device that is an external disk drive physically connected to all failover nodes. This cannot be an NFS mounted disk.

Task 5: Installing a Database

You can customize the database installation based on your environment.

You must install the database on the primary node in the Sun Cluster.

Sybase Database

To install Sybase database, follow these steps:

- 1 On the shared disk, create the Sybase target directory:

```
mkdir -p /NFS/Sybase
```

- 2 Verify the shared disk is mounted
 - 3 Mount the PI DVD as per the instructions in the *HP Performance Insight Installation and Upgrade Guide for Sybase – UNIX* and launch the setup.
 - 4 Select **Sybase 15.0.2**
 - 5 In the Sybase settings window change the default Installation Path to:
`/NFS/Sybase`
This path is symbolically linked to the external shared disk.
 - 6 Continue the installation per the instructions in the installation guide.
- For details on installing and configuring Sybase database for PI, see the *HP Performance Insight Installation and Upgrade Guide for Sybase – UNIX*.

Oracle Database

To install Oracle database, follow these steps:

- 1 On shared disk, create the Oracle target directory:

```
mkdir -p /NFS/oracle
```
- 2 Verify the shared disk is mounted
- 3 Add the appropriate environment variables for the oracle user.
- 4 Create a `.profile` file. Type these commands:

```
umask 022  
export ORACLE_SID=scora  
export ORACLE_HOME=/NFS/oracle  
export ORACLE_BASE=/NFS/oracle  
export ORACLE_OWNER=oracle  
export LD_LIBRARY_PATH=$ORACLE_HOME/lib:$LD_LIBRARY_PATH  
export CLASSPATH=$ORACLE_BASE/oraInventory:$ORACLE_HOME/jlib:$CLASSPATH  
export PATH=$PATH:$ORACLE_HOME/bin:$ORACLE_HOME:/mnt/HADisk/local/bin
```
- 5 Create necessary directories and set ownership. Type:

```
chown oracle:dba ~oracle/.profile /NFS/oracle
```



```
mkdir -p /NFS/local/bin; chmod 777 /NFS/local  
/NFS/local/bin
```

- 6 Follow the instructions provided in the *HP Performance Insight Installation and Upgrade Guide for Oracle – UNIX* to install Oracle 10.2.0.1/11.2.0.1 Enterprise Edition with Partitioning.
- 7 Download and install the Oracle 10.2.0.4 patch. See the instructions in `README.html` available with the 10.2.0.4 patch.

Oracle installation is done by the “oracle” OS user, therefore, the Oracle Universal Installer installs Oracle on the shared drive as specified the `.profile` file.

To install Oracle database and subsequent patches faster, do not create a database for PI during the installation of Oracle. Launch the `$ORACLE_HOME/bin/dbca` user interface after installing the 10.2.0.4 patch install to create a database for PI.

Task 6: Installing PI

The following are the prerequisites for installing PI in a cluster environment. You can customize the PI installation based on your environment. For detailed installation steps, see the *HP Performance Insight Installation and Upgrade Guide-UNIX* for the respective database.

You must install PI on the primary node in the Sun Cluster

On Sybase Database

Ensure the following:

- The PI installation folder must use the shared drive. For example, `/NFS/OVPI`.
- The Sybase server name (DSQUERY) must match the one created during the Sybase installation
- The Sybase directory must be the location on the shared disk into which Sybase was installed. For example, `/NFS/Sybase`.

On Oracle Database

Ensure the following:

- The PI destination folder must use the shared drive. For example, `/NFS/OVPI`.
- The Oracle database name (SID) must match the name set during the Oracle database creation
- The Oracle Home directory must be the location on the shared disk in which Oracle was installed. For example, `/NFS/oracle`.
- The location of the Oracle data files (in case you select the option to allow the PI installation to create the necessary tablespaces) must be on the shared disk. For example, `/NFS/oracle/dbs`.
- The location for the collection cache must be on the shared disk. For example, `/NFS/OVPI/collect`.
- The location for the PI log files must be on the shared disk. For example, `/NFS/OVPI/log`.

Task 7: Installing Report Packs

You can install the report packs based on your requirements. For PI 5.41 you must install the report packs only from the August 2009 Release Report Pack CD.

Installing report packs from the primary node, will place the report pack files on the shared disk at `DPIPE_HOME/packages`. `DPIPE_HOME` is the environment variable that identifies the installation directory of PI. In this guide `DPIPE_HOME` is `/NFS/OVPI`.

Install the Report Packs from the RNS CD

To install the report packs in the appropriate directories, you must install the report packs from the report pack CD only. Follow these steps:

- 1 Before starting the installation, verify that the destination folder for the packages references the PI location on the shared disk. For example, `/NFS/OVPI/packages`
- 2 Instantiate the report packs and datapipes on PI through the PI Package Manager. In Package Manager, set the following values:
 - Destination Folder = `/NFS/OVPI/packages` (default)
 - Deploy Reports = **[checked]** (default)

Task 8: Editing trendtimer.sched

To collect data using the Report Packs installed on a PI/HA node you must edit the configuration file: `$DPIPE_HOME/lib/trendtimer.sched`. This file contains specifications on how `trendtimer` invokes `mw_collect`.

Follow these steps:

- 1 Open the file `$DPIPE_HOME/lib/trendtimer.sched`.
- 2 Edit the default command-line arguments for all `mw_collect` entries to include the `-H <logical hostname>` option along with the `-n` option. The `-H` option along with the `-n` option gives the flexibility to specify the logical name in case of HA and ensures that only the specified nodes get polled.



Edit any other collector (for example, `pa_collect` and `ee_collect`) that you might be using in a similar manner.

For example,

OLD

```
5 - - {DPIPE_HOME}/bin/mw_collect -n -i 5 -K 1
. . . . .
24:00+1:00 - - {DPIPE_HOME}/bin/mw_collect -n -i 1440 -K 1
```

NEW

```
5 - - {DPIPE_HOME}/bin/mw_collect -n -H <logical hostname>
-i 5 -K 1
. . . . .
24:00+1:00 - - {DPIPE_HOME}/bin/mw_collect -n
-H <logical hostname> -i 1440 -K 1
```


Configuring PI to function in a Cluster Setup

In a scenario where both the PI and database installations have a physical hostname and IP address instead of a logical hostname and IP address, follow these steps to change to logical hostname and IP address (to make them a part of the Sun Cluster cluster):

- 1 Install the underlying database for PI. See, the section *Installing a Database*.
- 2 Install PI. See, the section *Installing PI*.

Now, the PI installation is usable with the physical name/IP on the clustered node on which it is installed. However, it is not usable with the logical name/IP.

- 3 Replace all occurrences of the physical name and IP address with the logical name and IP address. See the section *Changing Hostname on a PI 5.41 Installation*. Replace the “old hostname” and “new hostname,” with the “physical hostname” and “logical hostname,” respectively.

Similarly, replace “Old IP address” and “new IP address” with appropriate values for “physical IP address” and “logical IP address,” respectively.

Changing Hostname on a PI 5.41 Installation

When you install the complete PI (including report packs) on one of the clustered servers in an HA scenario, the configuration of both PI and the underlying database will be based on the “physical” hostname and the IP address of the clustered server on which PI is installed.

To reconfigure PI and the database to run on the logical/virtual/ floating hostname and IP address, follow these instructions:

Make sure the following replaceable values are used in the code given in the instructions. Make necessary substitutions as per your environment.

`<OLD_HOST>` – The “old” hostname. In the HA case, this will be the "physical" hostname obtained when you run `uname -n`.

`<OLD_IPADDR>` – The “old” IP address associated with `<OLD_HOST>`.

`<NEW_HOST>` – The “new” hostname. In the HA case, this will be the logical/virtual/floating hostname.

`<NEW_IPADDR>` – The “new” IP address associated with `<NEW_HOST>`.

Changing Hostnames and IP Addresses in the Configuration Files and PI Tables

Follow these steps:

- 1 Stop PI if it is running.

```
/etc/init.d/ovpi_httpd stop
```

```
/etc/init.d/ovpi_timer stop
```

- 2 Start the database. For Oracle, the Listener is not required to be running.

Oracle:

- a From the `<ORACLE_HOME>/bin` directory, type `sqlplus`

- b Type the username sys as `sysdba`

- c Type the appropriate password

- d Type `startup`

Sybase:

Type `/etc/init.d/Sybase start`

- 4 Connect to the database using sqlplus or isql as appropriate:

Oracle:

Type, `sqlplus dsi_dpipe/<Password>`

Sybase:

Type, `su - trendadm -c "isql -Udsi_dpipe -P <Password>"`

- 5 Update the DSI_SERVER table.

Oracle:

```
SELECT NAME, HOST_NAME, HOST_ADDRESS FROM DSI_SERVER;
```

```
UPDATE DSI_SERVER SET NAME = '<NEW_HOST>'
```

```
WHERE NAME = '<OLD_HOST>';
```

```
UPDATE DSI_SERVER SET HOST_NAME = '<NEW_HOST>'
```

```
WHERE HOST_NAME = '<OLD_HOST>';
```

```
UPDATE DSI_SERVER SET HOST_ADDRESS = '<NEW_IPADDR>'
```

```
WHERE HOST_ADDRESS = '<OLD_IPADDR>';
```

```
SELECT NAME, HOST_NAME, HOST_ADDRESS FROM DSI_SERVER;
```

Sybase:

```
SELECT name, host_name, host_address FROM dsi_server
```

```
go
```

```
UPDATE dsi_server SET name = '<NEW_HOST>''
```

```
WHERE name = '<OLD_HOST>'
```

```
go
```

```
UPDATE dsi_server SET host_name = '<NEW_HOST>''
```

```
WHERE host_name = '<OLD_HOST>'
```

```
go
```

```
UPDATE dsi_server SET host_address = '<NEW_IPADDR>'
```

```
WHERE host_address = '<OLD_IPADDR>'
```

```
go
```

```
SELECT name, host_name, host_address FROM dsi_server
```

```
Go
```

- 6 Update the DSI_INSTALLED_DATAPIPE table. The NAME and HOST_NAME fields appear both with and without the fully qualified domain name appended. If you are using NAME and HOSTNAME field without the domain name appended (and vice-versa) then use the same form when updating these values.



If more than one row is selected by any of the following sql “updates,” modify the update statement with a more-specific constraint clause to avoid uniqueness constraint violations.

Oracle:

```
SELECT * FROM DSI_INSTALLED_DATAPIPE;
UPDATE DSI_INSTALLED_DATAPIPE SET NAME = '<NEW_HOST>'
WHERE NAME = '<OLD_HOST>';
UPDATE DSI_INSTALLED_DATAPIPE
SET HOST_NAME = '<NEW_HOST>.<FULL_DOMAIN>'
WHERE HOST_NAME = '<OLD_HOST>.<FULL_DOMAIN>';
UPDATE DSI_INSTALLED_DATAPIPE
SET HOST_ADDRESS = '<NEW_IPADDR>'
WHERE HOST_ADDRESS = '<OLD_IPADDR>';
SELECT * FROM DSI_INSTALLED_DATAPIPE;
```

Sybase:

```
SELECT * FROM dsi_installed_datapipe
go
UPDATE dsi_installed_datapipe SET name = '<NEW_HOST>'
WHERE name = '<OLD_HOST>'
go
UPDATE dsi_installed_datapipe
SET host_name = '<NEW_HOST>.<FULL_DOMAIN>'
WHERE host_name = '<OLD_HOST>.<FULL_DOMAIN>'
go
UPDATE dsi_installed_datapipe
SET host_address = '<NEW_IPADDR>}'
```



```

WHERE host_address = '<OLD_IPADDR>'
go
SELECT * FROM dsi_installed_datapipe
Go

```

7 Stop the database.

Oracle:

- b From the <ORACLE_HOME>/bin directory, type **sqlplus**
- e Type the username sys as **sysdba**
- f Type the appropriate password
- g Type **shutdown**

Sybase:

Type **/etc/init.d/Sybase stop**

8 Make the following changes in the database configuration file:

Oracle:

- a Type, **cd \$ORACLE_HOME/network/admin/**
- b Type, **vi listener.ora tnsnames.ora**
- c Replace the old system name with the new system name in files similar to:


```

ADDRESS = (PROTOCOL = TCP)(HOST = .....

```

or

```

snmp.longname.listener = listener_.....

```
- d Save the file. If intelligent agent is installed edit the **snmp_ro.ora** file too.

Sybase:

- a Type **cd ~sybase**
- b Backup the old file and type:


```

vi interfaces

```
- c In the **master** and **query** lines replace the physical hostname with the logical hostname.

9 Edit the PI configuration files, type:

cd \$DPIPE_HOME/data

- a Open the `systems.xml` file. Type,

vi systems.xml



Backup the `systems.xml` file before making any changes to it. If this file is corrupted PI will not work.

- b Replace all occurrences of the old hostname and IP address with the new hostname and IP address:

Example: `<Name><OLD_HOST </Name>`

`<HostName><OLD_HOST>.<FULL_DOMAIN></HostName>`

`<IPAddress><OLD_IPADDR></IPAddress>`

`<Host><OLD_HOST>.<FULL_DOMAIN></Host>`

Oracle:

`<JdbcString>jdbc:oracle:thin:@`

`<OLD_HOST>:1521:mcsgora</JdbcString>`

`<OdbcString>DSN=PI_ORACLE;SID=mcsgora;PORTNUMBER=1521
;HOSTNAME=<OLD_HOST></OdbcString>`

Sybase:

`<JdbcString>jdbc:sybase:Tds:`

`<OLD_HOST>.{FULL_DOMAIN}:2052</JdbcString>`

`<OdbcString>DSN=PI_SYBASE;DB=dpipe_db;
NA=<OLD_HOST>.{FULL_DOMAIN},2052</OdbcString>`

- c Open the `systems.xml` file. Type,

vi config.prp



Backup the `config.prp` file before making any changes to it. If this file is corrupted PI will not work.

- d Replace all occurrences of the old hostname and IP address with the new hostname and IP address:

Example: `database.host=<OLD_HOST>`

`appserver.host=<OLD_HOST>`

`localhost=<OLD_HOST>`

`database_schema.host=<OLD_HOST>`

`server.host=<OLD_HOST>`

Task 9: Propagating Shared and Local Files

Complete the following tasks to ensure that “local” files are available in case of a failover. Both PI and the database (Oracle or Sybase) are installed on the shared disk. Some of these changes must be replicated to the other failover machines (secondary nodes) or disabled on the local system as follows:

Sybase Shared and Local Files

File	Task	Notes
System Files		
/etc/group	On secondary nodes, type: groupadd sybase	Replicates “sybase” group to failover boxes
/etc/passwd	On secondary nodes, type: useradd sybase	Replicates “sybase” user to failover boxes
Static File		
/etc/init.d/Sybase	Copy to the secondary nodes in a directory other than /etc/init.d	Sun cluster will handle the Sybase process startup. File is copied to failover machines in case a manual startup is necessary.

Oracle Shared and Local Files

File	Task	Notes
System Files		
<code>/etc/group</code>	On secondary nodes, type: groupadd dba	Add group “dba” if that is the group created for the Oracle installation
<code>/etc/passwd</code>	On secondary nodes, type: useradd oracle	Add user “oracle” in group “dba” with same home directory and shell
Static Files		
<code>~oracle/.profile</code>	Copy this file to the ORACLE_HOME folder on the shared drive.	Verify that the oracle user’s .profile has the appropriate Oracle environment variables before copying. At a minimum, it should include entries for ORACLE_SID, ORACLE_BASE and ORACLE_HOME.
<code>/etc/oratab</code>	Copy this to all failover machines.	
<code>/var/opt/oracle/oraInst.loc</code>	Copy file to secondary nodes.	
<code>/opt/ORCLfmap</code>	Recursively copy this directory to secondary nodes.	There are 3 subdirectories with a total of 6 files to be copied.

PI Shared and Local Files

File	Task	Notes
System Files		
/etc/group	On secondary nodes, type: groupadd trendadm	Add group “trendadm”
/etc/passwd	On secondary nodes, type: useradd trendadm	Add user “trendadm”
Static Files		
/etc/trend.conf	Copy file to secondary nodes.	Contains PI home directory
<ul style="list-style-type: none"> • /etc/init.d/OVPI_timer • /etc/init.d/OVPI_httpd <p>Move these file from /etc/init.d to another location to prevent PI processes from starting automatically on restart.</p>	Copy to the secondary node in a directory other than /etc/init.d	<p>Sun Cluster handles process startup.</p> <p>File is copied to failover machines if a manual startup is necessary.</p>

Dynamic Files

<ul style="list-style-type: none">• <code>/etc/opt/OV/share/conf/snmpmib</code>• <code>/etc/opt/OV/share/conf/snmpmib.bin</code>	<ul style="list-style-type: none">• Copy to shared disk• Create symbolic link on each failover node for <code>/etc/opt/OV/share/conf</code> to new location on shared disk	Stores MIBs loaded via the PI MIB Browser. Must be shared so that updates are available to any failover machine
---	---	---

Shared Files (Static)

<code>/opt/OV/</code>	Recursive copy to shared disk	Create symbolic links on failover nodes.
<code>/var/opt/OV/</code>	Recursive copy to shared disk	Create symbolic links on failover nodes.
<code>/opt/dcelocal</code>	Recursive copy to shared disk	Create symbolic links on failover nodes.

4 Configuring PI on Sun Cluster

The Data Service agent script `Script.OVPI.SUN.StartStopMonitor.ksh` performs the monitoring, startup, and shutdown activities for PI and the database.

To configure PI failover, follow these steps:

- 1 Create the following subdirectories and place the individual Data Service agent scripts in them:

- `$DS_AGENT_PATHPREFIX/ovpiServerDS/Run_OVPI.ksh`
- `$DS_AGENT_PATHPREFIX/ovpiServerDS/Monitor_OVPI.ksh`
- `$DS_AGENT_PATHPREFIX/ovpiServerDS/Halt_OVPI.ksh`
- `$DS_AGENT_PATHPREFIX/ovpi<dbtype>DS/Run_OVPI.ksh`
- `$DS_AGENT_PATHPREFIX/ovpi<dbtype>DS/Monitor_OVPI.ksh`
- `$DS_AGENT_PATHPREFIX/ovpi<dbtype>DS/Halt_OVPI.ksh`

In this instance, `<dbtype>` is either Oracle or Sybase.

- `$DS_AGENT_PATHPREFIX/Script.OVPI.SUN.StartStopMonitor.ksh`

The `Script.OVPI.SUN.StartStopMonitor.ksh` script is symbolically linked to the other agent scripts.

Manually copy `Script.OVPI.SUN.StartStopMonitor.ksh` to other failover nodes maintaining the same directory structure and symbolic links.

- 2 Register Sun Cluster's Generic Data Service Resource, `SUNW.gds`. Type, **# clresourcetype register SUNW.gds**
- 3 Create a new Sun Cluster Resource Group named **ovpiRG<ID>** where `<ID>` is a unique user defined ID number.
- 4 Create two new Sun Cluster Resources, one each for the PI server and database:
 - **ovpiSRVres<ID>**
In this instance, `<ID>` is a unique user defined ID number.
 - **ovpi<db>res<ID>**

In this instance, `<ID>` is a unique user defined ID number and `<db>` is either **ORA** or **SYB**.

- 5 Link the three resources (created in step 3 and step 4) together in a dependency tree as follows:

OVPI_SERVER_RESOURCE - depends on ---

DB_RESOURCE - which in turn depends on ---

LOGICAL_HOSTNAME_RESOURCE

- 6 Start the Resource Group on the primary node.
- 7 On each failover node, enable Sun Cluster Data Services to log messages to `/var/adm/messages`. Follow these steps :

- a Edit the file `/etc/syslog.conf` to add the following lines:

***.info /var/adm/messages**

***.debug /var/adm/messages**



Only use tabs – not blank spaces – in these lines. Adding these lines enables `/bin/logger` to put all messages of type “info” and “debug” into `/var/adm/messages`.

- b Run the command **kill -HUP `cat /var/run/syslog.pid`**

This enables the `/usr/sbin/syslogd` daemon to re-read its configuration file.

See the man pages for `syslogd(1m)` and `logger(1)` for more information.

- 8 Place the main Network Interface Card (NIC) on each failover node into its IP MultiPathing (IPMP) group.

If you already have the primary NIC in an IPMP Group, the groupname appears when you run **ifconfig -a**. For example:

```
bge0: flags= .....
```

```
inet ...
```

```
groupname ...
```

```
ether ....
```

If the groupname does not appear on running **ifconfig -a**, choose a IPMP Group name. For example, `ovpiIPMPG`.

- 9 Run the following (for **bge0**, substitute **hme0** or **eri0** or other interface name as required):


```
ifconfig bge0 group ovpiIPMPG
```

- 10 Run **ifconfig -a** again. The following line must appear:

```
groupname ovpiIPMPG
```

Make this change permanent so that it does not change on restart. To do this you can do either of the following:

- Create a 1-line script and save it in the `/etc/rc.*/` directory.
- Open the `/etc/init.d/network` file and append the following line:

```
/sbin/ifconfig bge0 group ovpiIPMPG
```

- 11 Sometimes error messages like the following may appear when you make changes in the Sun Cluster configuration

```
failover1 - Could not validate the settings in  
/etc/nsswitch.conf. It is recommended that the settings for  
host lookup consult `files` before a name server.
```

To stop these error messages from appearing, open the `/etc/nsswitch.conf` file and edit the following lines:

```
OLD: group: compat ...
```

```
hosts: cluster dns ...
```

```
NEW: group: files compat ...
```

```
hosts: files cluster dns ...
```

- 12 Restart the cluster for the configuration changes to come in effect.

5 Installing a Patch in a Sun Cluster Setup

To apply a patch in a Sun Cluster setup, follow these steps:

- 1 Shut down the PI resource and the DB resource by stopping the PI package. Type: **clresourcegroup offline**
<PI_Resource_Group_Name>
Do *not* shutdown Sun Cluster.
- 2 Edit `Script.OVPI.SUN.StartStopMonitor.ksh` script to set `DEBUG_SERVICES=1` instead of the default 0 and restart the PI resource group. Type:
clresourcegroup online *<PI_Resource_Group_Name>*
Setting `DEBUG_SERVICES=1` will not start, stop, or monitor PI and DB instead it will start, stop, or monitor logical IP address and shared disk.
- 3 Start the database manually.
This is **mandatory** for any PI upgrade or patch installation.
- 4 Back up your system before installing the patch and ensure that you have proper administrative privileges.
- 5 Run the upgrade or patch setup.
- 6 Follow the instructions in the wizard for the upgrade or patch installation.
- 7 Check the log file for any upgrade or patch errors.
- 8 If the upgrade or patch needs a report pack upgrade, upgrade the report pack to the latest version.
- 9 Stop the PI package. Type: **clresourcegroup offline**
<PI_Resource_Group_Name>
- 10 Edit the `Script.OVPI.SUN.StartStopMonitor.ksh` script to reset `DEBUG_SERVICES` to the default value 0.
- 11 Restart the PI package. Type: **clresourcegroup online**
<PI_Resource_Group_Name>

6 Troubleshooting Sun Cluster

Verifying the Cluster Configuration

To verify, if the Sun Cluster is configured appropriately, type the following command at the command prompt.

```
# sccheck -v1 -h phys-schost-1,phys-schost-2 -o  
/var/cluster/sccheck/myReports
```

Viewing Sun Cluster Logs

To view the messages logged in the `commandlog` file for the current week (one screen at a time), follow these steps;

- 1 Type the following command.

```
# more /var/cluster/logs/commandlog
```
- 2 At a different shell prompt, run the following command:

```
tail -f /var/adm/messages
```

When the `commandlog` file fills up, it rolls over to `/var/adm/messages.0`, `messages.1`, and so on.

Viewing PI Resource Group Log Files

To view the PI Resource Group log files, type the following command;

```
/var/cluster/logs/DS/<PI_Resource_Group_Name>/
```

A References

Software Depots and Versions

PI

The latest version of PI is PI 5.41. You can upgrade to PI 5.41 or install PI 5.41 from the product DVD.

Sybase

Use the Sybase version bundled in the PI 5.41 product DVD. The version is 15.0.2.

To verify the version, log in to `isql` and run:

```
select @@version
```

Oracle

PI 5.41 currently supports only Oracle version 10.2.0.4/11.2.0.1. To verify the version of Oracle, connect to `sqlplus` and run:

```
select * from v$version
```

Report Packs

Use the August 2009 Release Report Pack CD with PI 5.41.

Location for PI and Sun Cluster Manuals

Sun Cluster

You can download the latest versions of all the manuals from the following location:

<http://docs.sun.com/app/docs/prod/sun.cluster32>

- Sun cluster system administration guide for Solaris Operating System

<http://dlc.sun.com/pdf/819-2971/819-2971.pdf>

- Sun Cluster Data Services Developer's Guide for Solaris Operating System

<http://dlc.sun.com/pdf/819-2972/819-2972.pdf>

PI

You can download PI related guides from the following location:

<http://h20230.www2.hp.com/selfsolve/manuals>

Log in to the site using your HP Passport id and select Performance Insight from the product list.