

HP Performance Insight

For the HP-UX, Solaris, and Linux operating systems

Software Version: 5.41

Configuration Guide for Veritas Cluster

April 2011



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 Hewlett-Packard Development Company, L.P.

Trademark Notices

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 Scheemaeker, 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	VCS Implementation Prerequisites	11
	Hostname and IP Address	11
	Shared and Local Files	12
	VCS Agents.....	12
3	Configuring PI on VCS	14
	Task 1: Setting Failover Machines	15
	Task 2: Verifying System Resources and Machine Patch Levels.....	15
	Task 3: Synchronizing System Clocks	16
	Task 4: Setting the Logical Hostname.....	16
	Task 5: Installing a Database	18
	Sybase Database.....	18
	Oracle Database.....	19
	Task 6: Installing PI	20
	On Sybase Database.....	20
	On Oracle Database.....	20
	Task 7: Installing Report Packs	21
	Launch the Setup from the RNS CD.....	21
	Task 8: Editing trendtimer.sched	21
	Task 9: Completing the PI Configuration.....	22
	Task 10: Propagating Shared and Local Files	23
	Sybase Shared and Local Files	23
	Oracle Shared and Local Files.....	24
	PI Shared and Local Files	25
	PI and VCS Configuration Scripts.....	27
	The Configuration Script	27
	The Agent Script	28

4	Configuring PI to Function in a Cluster Setup.....	31
	Changing Hostname on a PI 5.41 Installation	32
	Changing Hostnames and IP Addresses in the Configuration Files and PI Tables....	33
	Installing a Patch in a VCS setup	38
5	Troubleshooting Veritas Cluster Scripts	39
	PI Resource Creation Scripts	39
	PI Agent Script.....	39
	PI Configuration Script.....	39
A	References.....	40
	Software Depots and Versions	40
	PI.....	40
	Sybase	40
	Oracle.....	40
	Report Packs.....	40
	Location for PI and Veritas Manuals	41
	Veritas.....	41
	PI.....	41

1 Introduction

This guide provides instructions on how to configure HP Performance Insight (PI) in a high availability environment using the Veritas Cluster Server (VCS) 5.0. 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 Veritas 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.

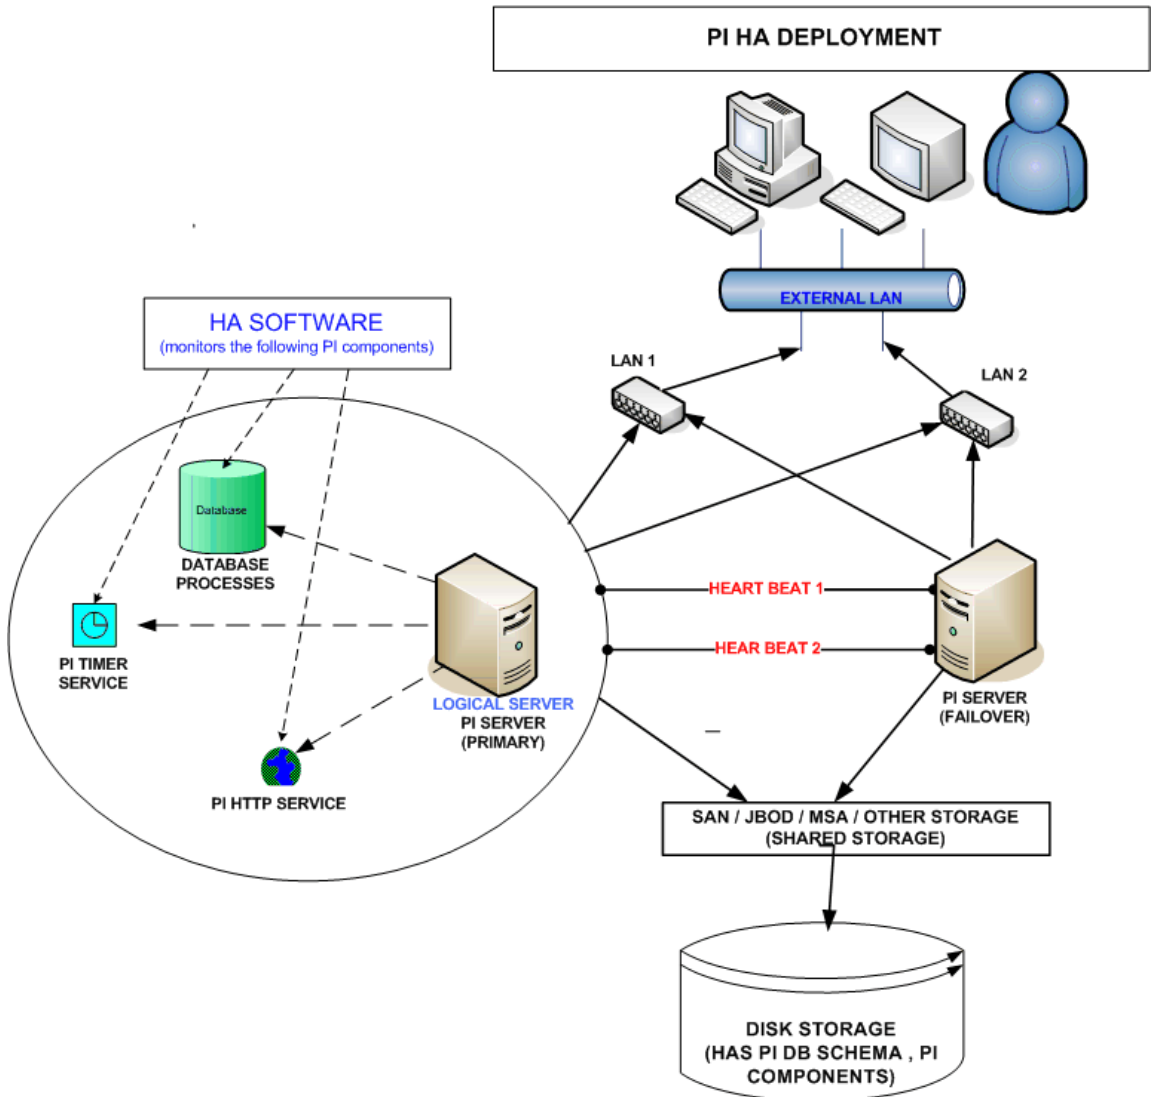
Configuration Considerations

This guide covers the Configuration of a two node cluster using a shared disk 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.

The following graphic provides an overview of the high availability setup:



2 VCS Implementation Prerequisites

The primary requirement for a VCS 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 (wherever applicable) 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 VCS:

- Assigning a logical hostname and associated IP address (wherever applicable)
- 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 (wherever applicable) 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 before installing the database or PI.

Shared and Local Files

In the VCS 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.

VCS Agents

VCS uses “agents” for monitoring the status and health of various resources, bringing them online, shutting them down in normal and emergency situations. VCS includes the agents required to handle common resources such as:

- Network Interface Cards – NIC Agent
- IP addresses – IP Agent
- Physical Disk drives – Disk Agent
- Logical Disk drives – NFS Agent, DiskGroup Agent and so on

Other resources also require similar type of monitoring and process management. For Sybase and Oracle databases, you can use the packages provided by Veritas. These packages include the agents as well.

Minimal-function agents for PI are available. These are: monitor, online, offline, and clean.

For other applications see the VCS Guides that provide the details for creating agents.

3 Configuring PI on VCS

The following variables/names are used in this chapter:

Name	Description
/NFS	The mount point for PI
/DB	The mount point for the database
failover1	The primary node of the failover cluster
logicalhost	The hostname assigned to the node that currently controls of the files and applications; also known as the “floater” or “virtual” hostname

Before configuring PI on the VCS cluster, ensure that all the nodes in the cluster can access the shared drives for both PI and the database (/NFS and /DB) through the Java UI Admin Console. However, only the system currently running the application must be able to read and write to the shared disks.

Configuring PI on VCS 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 the Logical Hostname*
- 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 VCS 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 of the machines in the cluster to a floating “virtual” IP address. This IP address shares the primary Network Interface Card (NIC) with the “real” IP address of each system.
- 4 Ensure that VCS is running. To verify, run the following command on both the nodes:

```
hastatus -sum
```

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:

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

See the xntpd (1M) manpage for details.

Task 4: Setting the Logical Hostname

Before installing the database and PI, set the logical hostname and IP address (wherever applicable) on the system that will function as the primary node.

If you face problems configuring the primary node to use the logical hostname and IP address, see *Installing PI into an Existing VCS Setup*.

Before setting the logical hostname:



Ensure that the IP assigned for the Veritas cluster is up and running.

- 1 Verify that both failover systems are on the same LAN with identical netmasks. For more information, type the following commands:

For Solaris or Linux

```
ifconfig -a
```


For HP-UX

lanscan

ifconfig <lan number>

- 2 Add logicalhost to DNS. To verify that logicalhost is recognized, type:

nslookup logicalhost

- 3 Disable the auto-startup under `/etc/rc*.d/` for those applications that you do NOT want to automatically start during restart. This will prevent any problems arising due to restarting with a different hostname.

To set the logical hostname, follow these steps:

On HP-UX systems (PARISC and Itanium)

- 1 At the command prompt, type **set_parms hostname**
- 2 Enter the logical hostname when prompted.
- 3 Restart the system.
- 4 Ping the system to verify if the logicalhost value is set correctly.
- 5 Type **hostname**. It should return logicalhost.

On Linux systems

- 1 Type **hostname** <logical hostname>
- 2 Ping the system to verify if the logicalhost value is set correctly.
- 3 Type **hostname**. It should return logicalhost.

On Solaris systems

- 1 Determine which system network files contain the physical node name.
Type:

```
cd /etc/
```

```
grep -il failover1 `find . -type f`
```

A list of files containing the machine name “failover1” appears.

Solaris 9.0

```
/etc/net/ticlts/hosts
```

```
/etc/net/ticots/hosts
```

```
/etc/net/ticotsord/hosts
```

```
/etc/nodename
```

```
/etc/hostname.*  
/etc/dumpadm.conf  
/etc/inet/ipnodes  
Solaris 10.0  
/etc/nodename  
/etc/hostname.*  
/etc/dumpadm.conf  
/etc/inet/ipnodes
```

- 2 Modify each file replacing “failover1” with “logicalhost”.



If “failover1” appears in the `/etc/hosts` file (sometimes linked symbolically to `./inet/hosts`) you need not edit it.

- 3 Verify the changes.
- 4 Restart the system
- 5 Ping the system to verify if the logicalhost value is set correctly.
- 6 Type `hostname`. It should return logicalhost.

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 VCS cluster.

Sybase Database

To install Sybase database, follow these steps:

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

```
mkdir -p /DB/sybase  
chmod 777 /DB/sybase
```
- 2 Verify the following:
 - a The primary node is set to return the “logicalhost” name. See the section [Setting the Logical Host Name](#).
 - b 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:
/DB/Sybase
- 6 Continue the installation per the instructions in the installation guide.
 - ▶ The name of the Sybase server (the DSQUERY value) defaults to <HOSTNAME_SYBASE> (the HOSTNAME being the logical hostname) during the installation.

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 each shared disk, create the Oracle target directory:

```
mkdir -p /DB/oracle
chmod 777 /DB/oracle
```
- 2 Verify the following:
 - a The primary node is set to return the “logicalhost” name. See the section [Setting the Logical Host Name](#).
 - b The shared disk is mounted.
- 3 Create a .profile file. Follow these steps,

```
umask 022

export ORACLE_SID=<oracle_sid>
export ORACLE_HOME=/DB/oracle
export ORACLE_BASE=/DB/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:/NFS/local/
bin
```

- 4 Follow the instructions provided in the *HP Performance Insight Installation and Upgrade Guide for Oracle – UNIX* to install Oracle 10.2.0.1 Enterprise Edition with Partitioning.
- 5 When prompted for the Installation Path, change the default value to `/DB/Oracle`. Continue the installation as per the Installation Guide.
- 6 Download and install the Oracle 10.2.0.4 patch. See the instructions in `README.html` available with the 10.2.0.4 patch.

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 VCS cluster

On Sybase Database

Ensure the following:

- The PI installation folder uses the shared drive. For example, `/NFS/OVPI`.
- The Sybase server name (DSQUERY) matches the one created during the Sybase installation.
- The Sybase host name uses the “logicalhost” name not the physical name of the machine.
- The Sybase directory is the location on the shared disk where you installed Sybase. For example, `/DB/Sybase`.

On Oracle Database

Ensure the following:

- The PI installation folder uses the shared drive. For example, `/NFS/OVPI`.
- The Oracle database name (SID) matches the name set during the Oracle database creation.

- The Oracle Home directory is the location on the shared disk where you installed Oracle. For example, `/DB/Oracle`
- The location of the Oracle datafiles (in case you selected the option to allow the PI installation to create the necessary tablespaces) is on the shared disk. For example, `/DB/Oracle/dbs`
- The location for the collection cache is on the shared disk. For example, `/NFS/OVPI/collect`
- The location for the PI log files is on the shared disk. For example, `/NFS/OVPI/log`

Task 7: Installing Report Packs

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`. The report packs you choose will depend on your needs

See the Report pack documentation for details about extracting and installing report packs from the August 2009 Release Report Pack CD.

Launch the Setup from the RNS CD

- 1 Install the report packs only from the August 2009 Release Report Pack CD. To launch the RNS CD, as a root user, type:


```
./setup
```
- 2 In Package Manager, set the following values:
 - Installation Folder = `/NFS/OVPI/packages` (default)
 - Deploy Reports = **[checked]** (default)
 - Application Server Name: `logicalhost.<FQDN>`

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 the collectors (for example, `mw_collect`, `pa_collect`, and `ee_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.

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
```

Task 9: Completing the PI Configuration

To complete the basic PI configuration, follow these steps:

- 1 Add nodes (either through a node import via `node_manager` or with the SNMP discovery).
- 2 Complete type discovery.
- 3 Verify that collections have occurred successfully.



After completing the PI Configuration, change the hostname of the primary node back to the physical hostname. **This step is mandatory.**

See [Setting the Logical Hostname](#) for instructions on changing the hostname.

Task 10: 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 to Perform on Failover Nodes	Notes
System Files		
<code>/etc/group</code>	“groupadd sybase” to secondary nodes	Replicates “sybase” group to failover boxes
<code>/etc/passwd</code>	On secondary nodes: useradd sybase	Replicates “sybase” user to failover boxes
<code>/etc/shadow</code>	None if “useradd” command is used.	“useradd” will add a new user to both the passwd and shadow files. You need not make a manual entry to <code>/etc/shadow</code> .
<code>/etc/services</code>	On secondary nodes: Copy line containing the Sybase port. The service name changes per installation. Typically the port/protocol is 2052/tcp.	To add Sybase service name to <code>/etc/services</code> file.
Static File		
<code>/etc/init.d/Sybase</code>	Modify file name to prevent Sybase from restarting at boot time. For example, Sybase_DONTRUN Copy to secondary nodes	VCS will handle process startup. File is copied to failover machines in case a manual startup is necessary.

Oracle Shared and Local Files

File	Task to Perform on Failover Nodes	Notes
System Files		
/etc/group	“groupadd dba” to secondary nodes	Add group “dba” if that is the group created for the Oracle installation
/etc/passwd	“useradd oracle” to secondary nodes	Add user “oracle” in group “dba” with same home directory and shell
/etc/shadow	None if “useradd” command is used.	“useradd” will add a new user to both the passwd and shadow files. You need not make a manual entry to /etc/shadow.
/etc/services	Copy line containing the listener port to secondary nodes. The service name is usually “listener” with a port/protocol of 1521/tcp.	To add Oracle service name to /etc/services file.
Static Files		
~oracle/.profile	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.

PI Shared and Local Files

File	Task to Perform on Failover Nodes	Notes
System Files		
/etc/group	groupadd trendadm to each secondary node	Add group “trendadm”
/etc/passwd	useradd trendadm to each secondary node	Add user “trendadm”
/etc/shadow	None if “useradd” command is used.	“useradd” adds a new user to both the passwd and shadow files. You need not make a manual entry to /etc/shadow.
Static Files		
/etc/trend.conf	Copy file to secondary nodes.	Contains PI home directory
<p>Solaris and Linux</p> <ul style="list-style-type: none"> • /etc/init.d/ovpi_timer • /etc/init.d/ovpi_httpd <p>HP-UX</p> <ul style="list-style-type: none"> • /sbin/init.d/ovpi_timer • /sbin/init.d/ovpi_httpd 	<ol style="list-style-type: none"> 1 Modify file name to prevent PI from restarting at boot time For example, ovpi_<process>_DONTRUN 2 Copy to secondary nodes 	<p>VCS 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">• /etc/opt/OV/share/conf/snmpmib• /etc/opt/OV/share/conf/snmpmib.bin	<ul style="list-style-type: none">• Copy to shared disk• Create symlink on each failover node for /etc/opt/OV/share/conf 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)

/opt/OV/	Recursive copy to shared disk	Create symbolic links on failover nodes.
/var/opt/OV/	Recursive copy to shared disk	Create symbolic links on failover nodes.
/opt/perf (only on Solaris)	Recursive copy to shared disk	Create symbolic links on failover nodes.
/opt/dcelocal (only on Solaris)	Recursive copy to shared disk	Create symbolic links on failover nodes.

PI and VCS Configuration Scripts

To implement the failover of PI on VCS, you must create VCS “Service Group” and agent scripts to provide monitoring, startup, shutdown, and cleanup activities.

The following scripts are bundled with PI to simplify the configuration of PI failover on VCS:

- The configuration script: `PI_unix_vcsconfigure.ksh`
- The agent script: `Script.OVPI.VCS.StartStopMonitorClean.ksh`

The Configuration Script

The `PI_unix_vcsconfigure.ksh` script creates and configures Resource Types, Service Groups, Resources and Resource Dependencies, and configuring VCS to call the `Monitor_OVPI.ksh`, `Run_OVPI.ksh`, `Halt_OVPI.ksh`, and `Clean_OVPI.ksh` scripts.

Follow these steps:

- 1 Save the `PI_unix_vcsconfigure.ksh` script in the `/etc/VRTSvcs/conf/config/` directory only.
- 2 Edit the section `USER CUSTOMIZABLE PARAMETERS` (in the beginning of the script) to suit your environment. The instructions are provided within the script as comments.
- 3 Run the `PI_unix_vcsconfigure.ksh` script. Type:
`./PI_unix_vcsconfigure.ksh`

When you run the `PI_unix_vcsconfigure.ksh` script, it performs the following functions:

- Creates a service group, `PI_Resource_Group`. In VCS 5.0, you can access this group through the Symantec Veritas Cluster Manager (Java Console). You can use this console to perform administrative functions.
- Creates additional VCS resources for the NIC and IP associated with the service group.
- Adds the resources to the service group such that the four resources are linked together in a dependency tree as follows:

```

// resource dependency tree
//
//   group PI_Resource_Group
//   {
//   Application PI-application
//   {
//     Mount ora-mount-nfs
//     {
//       LVMVolumeGroup pi-volumegroup
//       {
//         IP pi-ip
//         {
//           NIC pi-nic
//         }
//       }
//     }
//   }
//   Mount pi-mount-pi
//   {
//     LVMVolumeGroup pi-volumegroup
//     {
//       IP pi-ip
//       {
//         NIC pi-nic
//       }
//     }
//   }
// }

```


- Enables and starts the resources on the primary node.
- Captures log messages in the following files:
 - /var/VRTSvcs/log/engine_A.log
 - /var/VRTSvcs/log/Application_*
- Informs the cluster about the four agent scripts.

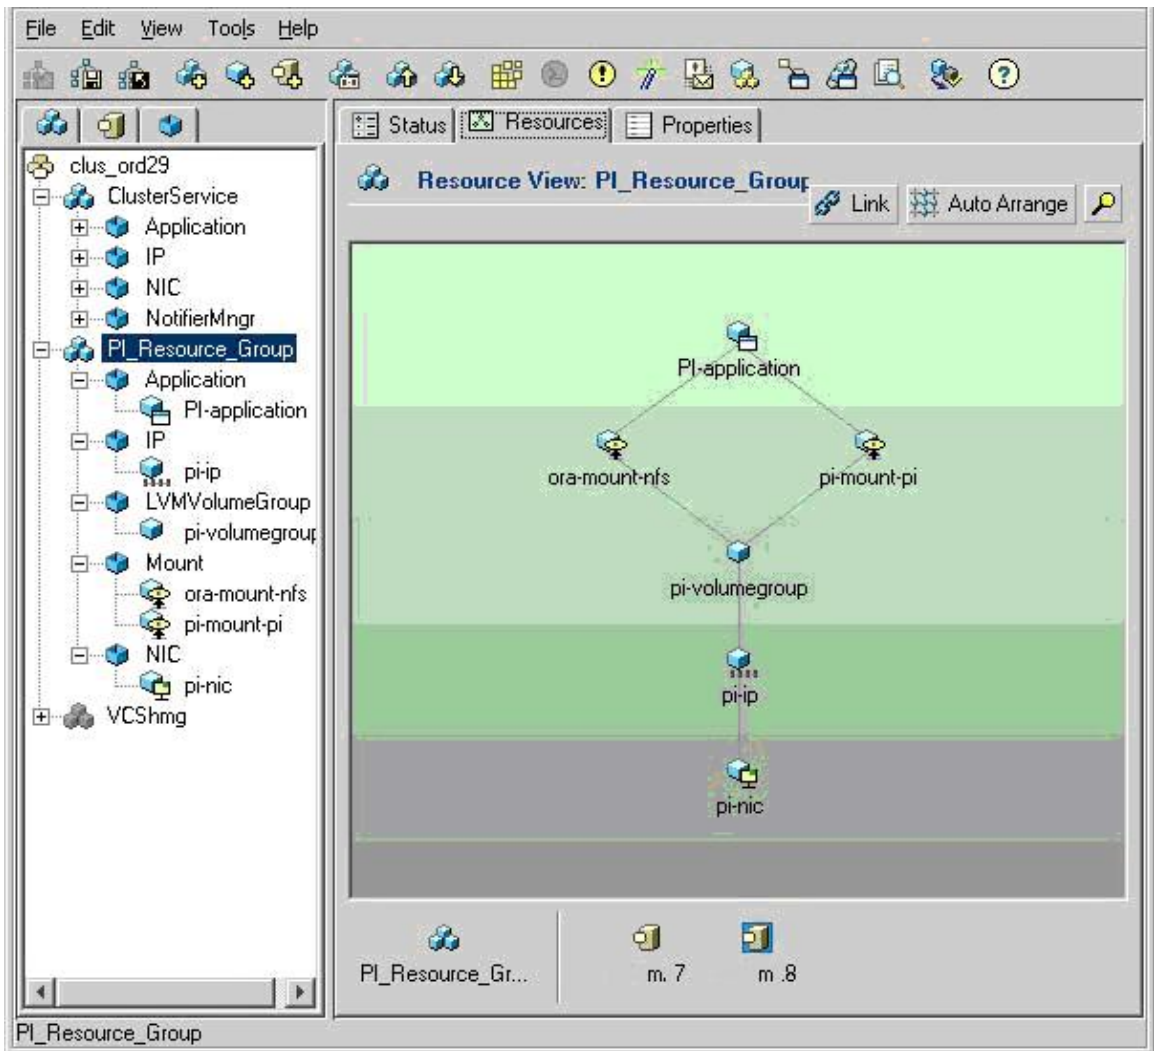
The Agent Script

Based on your environment, you must edit the agent script.

To configure the agent script, follow these steps:

- 1 Open the agent script `Script.OVPI.VCS.StartStopMonitorClean.ksh`.

- 2 Edit the section `USER CUSTOMIZABLE PARAMETERS` (in the beginning of the script) to suit your environment. The instructions are provided within the script as comments.
- 3 Scroll through the script to the `OVPI_Clean` procedure and include any other binary that you might want to remove.
- 4 Replicate the agent script into the following four scripts:
 - `Monitor_OVPI.ksh`
 - `Run_OVPI.ksh`
 - `Halt_OVPI.ksh`
 - `Clean_OVPI.ksh`
- 5 Create a folder on the primary node and place the four scripts in it
- 6 Create a folder with exactly the same name and directory structure on the secondary node and copy the four scripts in it.
 -  If you make a change in any of the four scripts, you must propagate the change to the other three scripts in that folder and then copy all the four scripts to the failover node.
- 7 Launch the Symantec Veritas Cluster Manager (Java Console) on the primary node.
- 8 Right-click **PI_Resource_Group** (in the left pane) and select **Online** → `<node_name>`. The Resource View appears as in the following figure:



The agent scripts log messages in the `VRTS_OVPI.log` file. This file is present in the `/tmp` directory.

4 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 (wherever applicable) 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 VCS 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 (wherever applicable) 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

The instructions provided in this section are for a Standalone PI server on which PI is installed. However, you can customize the steps for a distributed environment too.

You will need to change the hostname in the following two scenarios:

- You have to change the name and IP address (wherever applicable) of a non-clustered PI server.
- You have to install/upgrade or apply a patch to a High Availability PI server (in a cluster setup) running under VCS.

In this scenario, the clustered server will host (at least) two IP addresses and hostnames:

- The “physical” or “real” hostname with its associated IP address: You can get this hostname by typing `uname -n` or `hostname` at the command prompt.
- The “logical” or “virtual” or “floating” hostname with the associated IP address: This is the virtual machine name that can be shifted from one clustered server to another thereby implementing High Availability (HA).

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.

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

Make the following replaceable values are used in the code given in the instructions. You must make the 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.

Solaris and Linux:

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

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

HP-UX:

```
/sbin/init.d/ovpi_httpd stop
```

```
/sbin/init.d/ovpi_timer stop
```

- 2 Start the database.

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:

Solaris and Linux – Type `/etc/init.d/Sybase start`

HP-UX – Type `/sbin/init.d/Sybase start`

- 3 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>"
```

- 4 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

```

- 5 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
```

- 6 Stop the database.

Oracle:

- a From the Oracle home/bin directory, type **sqlplus**

- b Type the username sys as **sysdba**
- c Type the appropriate password
- d Type **shutdown**

Sybase:

Solaris and Linux – Type **/etc/init.d/Sybase stop**

HP-UX – Type **/sbin/init.d/Sybase stop**

- 7 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.

- 8 Edit the PI configuration files. Follow these steps:

- a Type: **cd \$DPIPE_HOME/data**
- b 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.

- c 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:vcsora</JdbcString>

<OdbcString>DSN=PI_ORACLE;SID=vcsora;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>
```

- d Open the `config.prp` 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.

- e 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>
```

- 9 In a non-HA scenario, change the system name and restart the system. In an HA scenario, use the HA software to enable the Floating/Logical IP address if it is not already enabled.
- 10 Start the database (including the Oracle Listener if applicable). Verify that it is running successfully.
- 11 Start PI and verify that it is running successfully.

Installing a Patch in a VCS setup

Before installing any patch in a VCS environment, you must take care of the following:

- The machine on which you want to install the patch must be set to a logical hostname.
- If the patch changes the local files: Propagate these changes to all other failover systems or move these files to the shared disk and provide a symbolic link to all the failover systems.

To apply a patch, in a VCS setup, follow these steps:

- 1 Stop the PI resource group on all the systems using the Cluster Manager console.
- 2 Change the name of the PI daemons from `ovpi_<process>_DONTRUN` to `ovpi_timer` and `ovpi_httpd`.
- 3 Change the physical hostname and IP address (wherever applicable) of the primary node to the logical hostname and IP address (wherever applicable). See the section, *Setting the Logical Hostname* for instructions on changing hostname and IP address (wherever applicable).
- 4 Install the patch.
- 5 Change the logical hostname and IP address (wherever applicable) back to the physical hostname and IP address (wherever applicable).
- 6 Rename the PI daemons to `ovpi_<process>_DONTRUN`.
- 7 Restart the system.

5 Troubleshooting Veritas Cluster Scripts

PI Resource Creation Scripts

To troubleshoot the PI resource creation scripts, follow these steps:

- 1 Open the `/var/opt/OV/log/vcs_resource_configuration.log` file.
- 2 Check whether the following attributes are correct:
 - Mount points
 - Virtual host names
 - Directory names
 - File names
 - Agent script names
- 3 Check if the file permissions are appropriate.

PI Agent Script

To troubleshoot the `Script.OVPI.VCS.StartStopMonitorClean.ksh` script, follow these steps:

- 1 Check the `/var/VRTSvcs/log` folder. This folder contains the logs for all resources. For each resource a log file of the name `<ResourceName>_A.log` is created. In this instance, A stands for agent.
- 2 Open the respective log file to check for error messages.

PI Configuration Script

To troubleshoot the `PI_unix_vcsconfigure.ksh` script, run the following command to check the syntax of the configuration.

```
#hacf -verify /etc/VRTSvcs/conf/config
```

A References

Software Depots and Versions

PI

The latest version of PI is PI 5.41. You can 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. 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.

Location for PI and Veritas Manuals

Veritas

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

http://www.symantec.com/business/support/all_products.jsp

- VCS 5.0 Installation Guide:

ftp://ftp.support.veritas.com/pub/support/products/ClusterServer_UNIX/283868.pdf

- VCS 5.0 User's Guide:

ftp://ftp.support.veritas.com/pub/support/products/ClusterServer_UNIX/283869.pdf

- VCS 5.0 Agent Developer's Guide:

ftp://ftp.support.veritas.com/pub/support/products/ClusterServer_UNIX/283870.pdf

- VCS 5.0 Bundled Agents Guide:

ftp://ftp.support.veritas.com/pub/support/products/ClusterServer_UNIX/283871.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

We appreciate your feedback!

If an email client is configured on this system, by default an email window opens when you click on the bookmark “Comments”.

In case you do not have the email client configured, copy the information below to a web mail client, and send this email to **docfeedback@hp.com**

Product name:

Document title:

Version number:

Feedback:

