

# HP Matrix Operating Environment 7.2 Release Notes

## Abstract

This document describes the HP Matrix Operating Environment 7.2 release content and provides information about limitations, and major and minor issues with suggested actions.



© Copyright 2006–2013 Hewlett-Packard Development Company, L.P.

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.

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.

### Acknowledgments

HP-UX Release 10.20 and later and HP-UX Release 11.00 and later (in both 32 and 64-bit configurations) on all HP 9000 computers are Open Group UNIX 95 branded products.

UNIX is a registered trademark of The Open Group.

Java is a US trademark of Sun Microsystems, Inc.

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

Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates.

EMC, SYMMETRIX, and VMAX are either registered trademarks or trademarks of EMC Corporation in the United States and/or other countries.

### Warranty

HP will replace defective delivery media for a period of 90 days from the date of purchase. This warranty applies to all Insight Management products.

### Revision history

Document part number	Software version	Supported operating systems	Document edition	Publication date
T8671-91059	7.2.0	See the <a href="#">HP Insight Management 7.2 Support Matrix</a>	1	March 2013
T8671-91059	7.2.0	See the <a href="#">HP Insight Management 7.2 Support Matrix</a>	2	March 2013
T8671-91065	7.2.0	See the <a href="#">HP Insight Management 7.2 Support Matrix</a>	3	April 2013
T8671-91065a	7.2.0	See the <a href="#">HP Insight Management 7.2 Support Matrix</a>	4	October 2013

---

# Contents

1	Description.....	5
	Update recommendation.....	5
	Supersedes.....	5
	Operating systems and devices.....	5
	Enhancements.....	5
	Prerequisites.....	7
	Installation .....	7
2	Issues and suggested actions.....	15
	Matrix Operating Environment Online Help.....	15
	Matrix Operating Environment software.....	15
	Capacity Advisor.....	17
	Infrastructure orchestration.....	18
	Insight managed system setup wizard.....	27
	Instant Capacity Manager.....	27
	Logical server management.....	28
	Process Resource Manager .....	36
	Storage Provisioning Manager.....	37
	Managed node installation.....	38
	Matrix KVM Private Cloud.....	38
	HP Matrix Operating Environment infrastructure orchestration.....	44
3	Support and other resources.....	49
	Contacting HP.....	49
	Related information.....	50
4	Documentation feedback.....	53



---

# 1 Description

HP Matrix Operating Environment is advanced infrastructure management software that allows you to analyze and optimize both physical and virtual resources in the same way. Matrix OE gives you a powerful toolkit to accelerate complex technology projects and simplify daily operations. It enables you to plan capacity and power continuously, provision infrastructure consistently, and balance across physical and virtual resources. It also allows you to ensure cost-effective high availability for a larger set of applications.

**HP Matrix Operating Environment for ProLiant, HP Matrix Operating Environment for non-HP servers, and HP Matrix Operating Environment for HP-UX** are integrated by design, delivering capabilities greater than the sum of its parts. Matrix OE is based on proven HP technologies, such as HP Insight Control and HP Systems Insight Manager, the world's most popular platform management tool. Matrix OE takes advantage of HP Virtual Connect to bring the flexibility of virtualization to a physical environment.

- **Matrix Operating Environment for ProLiant and Matrix Operating Environment for non-HP servers** include capacity planning, advanced virtualization and logical server management, and Matrix OE infrastructure orchestration capabilities for automated design and provisioning.
- **Matrix Operating Environment for ProLiant** also includes Matrix OE recovery management for disaster recovery capabilities, Insight Control for advanced power management, performance management, and server deployment, and extends advanced virtualization management to HP server blades using HP Virtual Connect technology.
- **Matrix Operating Environment for HP-UX** includes capacity planning, advanced virtualization management, global workload management, and support for Integrity blades by HP Integrity Virtual Server Manager, HP-UX Workload Manager, and HP Insight Remote Support.

## Update recommendation

HP recommends that customers running 7.1 and earlier versions upgrade to version 7.2.

## Supersedes

The 7.2 version supersedes the 7.1 released versions.

## Operating systems and devices

Supported operating systems and devices are listed in the *HP Insight Management 7.2 Support Matrix*:

[Matrix OE Information Library](#)

## Enhancements

Matrix OE 7.2 software provides new capabilities and expanded platform support.

See the *HP Insight Management Support Matrix* at [Matrix OE Information Library](#) for more information.

## New features in Matrix Operating Environment

- VM hypervisor use of the Flat SAN technology to direct connect from the FlexFabric module to an HP 3PAR StoreServ Storage System support. The hypervisor use of Flat SAN is configured independent of Matrix. Matrix will not provision volumes for physical servers from the direct connect HP 3PAR StoreServ Storage System.
- Provisioning to an HP Matrix KVM Private Cloud

- Storage Provisioning Manager (SPM) Windows CMS only
  - HP Storage Provisioning Manager (SPM) Adapter for EMC Symmetrix VMAX
 

The SPM adapter for EMC Symmetrix VMAX storage supports storage operations for discovering arrays and reading pool and volume information, create/delete volumes, grow volume, and volume presentations on a given host using FC interface. This enables SPM to import EMC volumes into the storage catalog which can be used to fulfill storage requests for logical servers. For information about enabling and disabling the adapter, see the *HP Insight Management 7.2 Installation and Configuration Guide* at <http://www.hp.com/go/matrixoe/docs>.
  - Support for virtual domains
  - New SPM security groups (e.g., Matrix User Security Group)
- Data disk provisioning for HP Cloud Services servers
- User-defined properties in templates to facilitate external integration
- Organization lifecycle API and CLI extensions
- Bursting support for HP Matrix systems, as a source and a target
- Upgrade Matrix OE Operations Orchestration to version 9.06.0001 with Content Pack 9
- Extending bursting capabilities to partners
- Support for Insight Control server provisioning
- Enhancements to initial VM server import capabilities
- Assign service to another user
- Built-in OpenStack connector for Service Providers
- Support for 3PAR StoreServ 7000
- Per user template control
- Chargeback archival mechanism
- mxreconfig for CMS
- Support for HyperV non-HA VM with non-clustered RDM disk
- Support for Microsoft Internet Explorer 10 in Compatibility View
- New CMS support
  - Windows Server 2012 Standard
  - Windows Server 2012 Data Center
  - HP-UX 11iv3 (11.31) 1303
  - Support for VMware ESXi 5.0 Update 2
- New database support
  - Microsoft SQL Server 2008 R2 Express SP2
  - Microsoft SQL Server 2008 R2 Standard SP2
  - Microsoft SQL Server 2008 R2 Enterprise SP2
  - Microsoft SQL Server 2012 Standard Initial release and SP1
  - Microsoft SQL Server 2012 Enterprise Initial release and SP1

- Microsoft SQL Server 2012 Business Initial release and SP1
- Microsoft SQL Server 2012 Web Initial release and SP1 (HP SIM only)
- HP Virtual Connect Enterprise Manager 7.2
  - Support for Virtual Connect FW 3.75
- Support for HP BladeSystem C7000 Platinum Enclosure
- New managed system support
  - RHEL 5.9
  - KVM on RHEL 5.9
  - Microsoft Windows Server 2012 with Hyper-V
  - Microsoft Hyper-V Server 2012

## Prerequisites

- Hardware, firmware, and software requirements for this release are listed in the *HP Insight Management 7.2 Support Matrix*.
- Installation and upgrade prerequisites are documented in the *HP Insight Management 7.2 Preinstallation Worksheet*, *HP Insight Management 7.2 Installation and Configuration Guide*, and *HP Matrix Operating Environment Integrity CMS Installation and Configuration Guide*.

These documents are available from the [Matrix OE Information Library](#).

## Installation

Installation instructions are documented in the *HP Insight Management 7.2 Installation and Configuration Guide* and *HP Matrix Operating Environment Integrity CMS Installation and Configuration Guide*.

These documents are available from the [Matrix OE Information Library](#).

## Newer version available for download

### Matrix Operating Environment for HP-UX

The latest version of Matrix OE for Integrity is available from the following location:

<https://h20392.www2.hp.com/portal/swdepot/displayProductInfo.do?productNumber=vse>

### Matrix Operating Environment for ProLiant and non-HP servers

The latest version of Matrix OE for ProLiant and non-HP servers is available from the following location:

<https://h20392.www2.hp.com/portal/swdepot/displayProductInfo.do?productNumber=HPID>

## Accessing online help before licenses are installed

You can access the Matrix OE online help before you have installed licenses on the CMS by performing the following steps.

1. Select **Help**→**For HP Systems Insight Manager** from the top menu.
2. In your browser address bar (if enabled), change the URL from the first line to the second line:

`https://{CMS-FQDN}:50000/mxhelp/mxportal/en/help-welcome.html`

`https://{CMS-FQDN}:50000/mxhelp/vsemgmt/en/licenses.html`

where {CMS-FQDN} is the fully qualified domain name of the CMS.

3. Press Enter to display the Matrix OE license management topic. From there, you can select any of the help topics that apply to Matrix OE software. For help specific to an individual component of Matrix OE, select the **Related information** topic and choose one of the links in the section **Help for related products**.

## Back up your existing CMS configuration before upgrading the software

HP recommends making a backup of your pre-7.2 environment before upgrading the software. The upgrade process is not reversible. If, after you begin, you decide not to complete the upgrade to Matrix OE 7.2, a backup copy of your pre-7.2 environment is required to restore your CMS to the pre-upgrade state.

- 
- ❗ **IMPORTANT:** After the upgrade process has successfully completed, you must also perform a backup of the upgraded 7.2 configuration to restore your environment to a known good state in the event of a catastrophic failure. This is necessary to ensure that managed resources, if they were changed by the 7.2 CMS, will continue to be manageable if the CMS ever needs to be restored from backup.
- 

See the *Backing up and restoring HP Insight Software Central Management Server (Windows)* white paper at [Matrix OE Information Library](#) for detailed backup and restore instructions.

## Back up user-modified properties files before upgrade

During an upgrade, Matrix OE properties files from an earlier release are replaced by new properties files for the 7.2 release. If you made any modifications or customizations to the properties files, your changes are lost.

### Suggested action

Before beginning an upgrade to Matrix OE 7.2, save a copy of user-modified properties files or note any changes made to the files so that you can insert those modifications into the new versions of the files.

Following is a partial list of properties files that are commonly modified. Verify in your environment if you have modified these files before beginning an upgrade, and if you have modified them, save the files to a location that is not under the `..\Program Files\HP\Virtual Server Environment` directory tree.

```
C:\Program Files\HP\Virtual Server Environment\conf\esa.properties
```

After the upgrade has completed, but before the services are restarted or the CMS is rebooted, edit the new properties files and add any custom entries made to the previous version of the file.

## Re-scaling stored CPU utilization data in Capacity Advisor

A new copy of the Performance Index Database, which holds server performance benchmarks, is installed whenever you upgrade Matrix OE.

The Performance Index Database is used in pre-scaling the CPU utilization data that is stored for capacity planning. After a Matrix OE upgrade, some of the data must be recomputed and re-scaled. This process can take a long time if there is a large history of utilization data.

In Matrix OE 7.2, the recomputing task is performed after the rest of the Matrix OE upgrade is complete and the CMS is back online. This change was made to reduce the downtime of the CMS. Although new Capacity Advisor data is not collected while the recomputing process runs, the rest of Matrix OE is available.

To recompute the utilization data, follow these steps:

1. Wait until the Matrix OE 7.1 upgrade is complete.

The time of day to perform the recomputing task is flexible. Pick a time when Capacity Advisor data collection can be suspended.

2. Log in to the CMS console.
3. Navigate to one of the following locations, depending on the operating system:
  - Windows CMS: C:\Program Files\HP\Virtual Server Environment\bin\capupgrade
  - HP-UX CMS: /opt/vse/bin/capupgrade
4. Recompute the data:
  - Windows CMS: Run capupgrade.bat.
  - HP-UX CMS: Issue the capupgrade command. No parameters are required.

The recomputing process:

- Stops Capacity Advisor data collection.
- Determines the data to be recomputed.
- Copies the new Performance Index Database.
- Recomputes the pre-scaled data.
- Restarts the Capacity Advisor data collection service.

Recomputing can take several hours to run, depending on the number of systems and the amount of collected data. The text output for capupgrade provides an estimate of the time.

The example below shows capupgrade being run on a CMS with 20 servers.

```
C:\Program Files\HP\Virtual Server Environment\bin>capupgrade
==> For more information, please refer to
"C:\Program Files\HP\Virtual Server Environment\logs\capupgrade.log"
The Performance Index upgrade can take up to 3 minutes.
During this time, the HP Capacity Advisor Data Service
will be offline and data collection will be stopped.
The service will restart automatically at the end of the
capupgrade process.
Do you want to continue? [y\n] y
Progress: 1/20
Progress: 2/20
Progress: 3/20
Progress: 4/20
Progress: 5/20
...
```

---

**NOTE:** The capacity planning scenario editor will display a reminder message until capupgrade is run.

---

## Support for a rolling upgrade of a CMS federation

To decrease federation downtime, Matrix OE infrastructure orchestration 7.1 includes support for a rolling upgrade of a federated CMS environment. A rolling upgrade temporarily allows a federation containing CMSs running different versions of Matrix OE during the upgrade process.

- ① **IMPORTANT:** The primary CMS must be upgraded first.
- 

During a rolling upgrade, a CMS running Matrix OE 7.1 supports basic operations against a secondary CMS running 7.1. Not all 7.2 features are expected to work against a 7.1 CMS.

The following operations are *not* supported on a primary CMS running Matrix OE 7.2 with secondary CMSs running 7.1:

- Manual operating system provisioning of VM guests using resources discovered on a secondary CMS
- Provisioning a physical server group with auto-provisioned SPM volumes on a secondary CMS
- Changing the memory size or number of CPUs of an existing virtual machine
- Some chargeback operations, including detection of changes made outside of Matrix OE (and directly in VMware, Hyper-V, OA, and so on)

## Hyper-V VM hosts are moved to the IO unassigned pool after upgrade from 6.3

After an upgrade from Matrix OE 6.3 to 7.2, all Hyper-V VM hosts are automatically moved to the infrastructure orchestration Unassigned server pool. This issue does not occur after an upgrade from Matrix OE 7.0 to 7.2.

Manually move the Hyper-V VM hosts back to their original server pool in order to provision virtual servers to Hyper-V VM hosts.

## Deactivated logical servers are powered off after upgrade from 6.0 and 6.1

A Matrix OE infrastructure orchestration service that is deactivated before an upgrade from Matrix OE 6.2 or higher to 7.1 will contain logical servers that remain deactivated after the upgrade. Their state is not changed.

However, service that is deactivated before an upgrade from Matrix OE 6.0 or 6.1 to 7.2 will contain logical servers that are powered off after the upgrade. This changes the state of the logical server from paused to powered off.

After the upgrade, you can power on and deactivate the logical server, if desired.

## Check logical server management upgrade log file

When Matrix OE logical server management is upgraded, the pre-upgrade task completes even if there are exceptions or missing objects.

### **Suggested action**

After the upgrade is complete, check for important upgrade messages at `../Program Files/HP/Virtual Server Environment/logs/vselog.lsm.log`, and make any changes as necessary.

## Changing the CMS password

Systems Insight Manager includes a CMS reconfigure tool that allows you to make common reconfiguration changes to your CMS, including realigning the Insight Management service account credentials with the CMS operating system password after the operating system password is changed.

In earlier releases, manual steps were required for the Insight Management components to recognize a CMS password change and properly operate, which are detailed in the *Changing the Insight Software 6.0, 6.1 or 6.2 CMS password* white paper. This white paper does not apply to Matrix OE 7.2.

See the *HP SIM 7.2 User Guide* at [www.hp.com/go/foundationmgmt/docs](http://www.hp.com/go/foundationmgmt/docs) for information about using the CMS reconfigure tool.

## Enabling template access restrictions

HP Matrix Operating Environment infrastructure orchestration supports template access restrictions on a per user basis. This feature is disabled by default. If you do not modify the properties file, all users will see all templates.

To enable template access restrictions for users, edit the `C:\Program Files\HP\Matrix infrastructure orchestration\conf\hpio.properties` file, and set `template.access.restricted=true`. After access restrictions are enabled, only users assigned to a template are able to see and use the template.

---

**NOTE:** Irrespective of this setting, all users in an organization have access to all templates assigned to the organization. Template access controls for individual users in an organization are not supported.

---

## Operations Orchestration Studio installation and upgrade issues

### Pointing a remote Operations Orchestration Studio installation to Operations Orchestration Central installed on a CMS

A “Failed Login” error message is displayed when you try to log in to Operations Orchestration Studio 9.x pointing to a remote repository (for example, Operations Orchestration Central 9.x on the CMS).

To resolve this issue, perform the following steps:

1. Make a backup of the `rc_keystore` file on the CMS.
2. Copy the `rc_keystore` file installed on the CMS (`C:\Program Files\HP\Operations Orchestration\Central\conf\rc_keystore`) to the path on the remote Operations Orchestration Studio system (`C:\Program Files\HP\Operations Orchestration\Studio\conf\rc_keystore`).
3. Restart Operations Orchestration Studio.

### Certificate key mismatch after Operations Orchestration Studio upgrade

After upgrading Operations Orchestration Studio from version 7.51 to 9.x, Operations Orchestration Central and Studio might have different certificate keys. If Operations Orchestration Studio cannot access the URL `https://localhost:16443`.

To resolve this issue, perform the following steps:

1. Make a backup of the `rc_keystore` file on the CMS.
2. Copy the `rc_keystore` file installed on the CMS (`C:\Program Files\HP\Operations Orchestration\Central\conf\rc_keystore`) to the same path on the Operations Orchestration Studio system (`C:\Program Files\HP\Operations Orchestration\Studio\conf\rc_keystore`).
3. Restart Operations Orchestration Studio.

### Cannot download 64-bit Operations Orchestration Studio from infrastructure orchestration Designer

An icon in the Workflow dialog box in infrastructure orchestration Designer allows you to download and install the 32-bit version of Operations Orchestration Studio. If you download and attempt to install the 32-bit version on a 64-bit CMS, the following error is displayed: “Studio version: 9.00 does not match Central version: 2.0. Setup will exit.”

If you have a 64-bit CMS, install the 64-bit version of Operations Orchestration Studio, which is available on the Insight Software DVD #2 in the `/matrixio` folder.

### Red Hat Enterprise Linux 5.5 cannot be installed on a VM guest on any ESX/ESXi host on server blades BL465c G7 and BL685c G7

Red Hat Enterprise Linux 5.5 cannot be installed in a VM guest on an ESXi 4.1 host on BL465c G7 and BL685c G7 server blades.

To resolve this issue, install RHEL 5.5 on the host system, install a RHEL 5.4 guest and then upgrade the guest to RHEL 5.5, ensuring that kernel-2.6.18-194.3.1.el5 or higher is used.

For more information, see Red Hat's Knowledgebase article at <https://access.redhat.com/kb/docs/DOC-38013>.

## Managing Microsoft Windows servers without OpenSSH installed

Even though the product documentation states OpenSSH is required, Matrix OE supports Microsoft Windows environments not running OpenSSH. However, for Microsoft Windows systems without OpenSSH, there are several limitations with agent installation and troubleshooting.

For details on using Matrix OE with Microsoft Windows systems not running OpenSSH, see the *HP Insight Dynamics 6.0 software without OpenSSH on HP ProLiant and Integrity Servers Running Microsoft Windows* white paper, available at the following website:

[Matrix OE Information Library](#)

## Patch required to provision Hyper-V linked clone VM guests using HP Server Automation

If you are using HP Server Automation (SA) 9.0 or 9.10 on Red Hat Enterprise Linux 6 or SUSE Linux Enterprise Server 10 platforms and you want to provision Hyper-V linked clone VM guests, install the following version-specific core patches and hotfixes for each SA version.

For SA 9.0, download and install:

- SA 9.05 core patch from <http://support.openview.hp.com/selfsolve/patches>
- "Linked Clone errors for RHEL6 and Above" hot fix from [https://quixy.dev.hp.com/hotfix/index.php?HOTFIX\\_PROJECT=server\\_automation](https://quixy.dev.hp.com/hotfix/index.php?HOTFIX_PROJECT=server_automation)

For SA 9.1, download and install:

- SA 9.11 core patch from <http://support.openview.hp.com/selfsolve/patches>
- "Linked Clone errors for RHEL6 and Above" hot fix from [https://quixy.dev.hp.com/hotfix/index.php?HOTFIX\\_PROJECT=server\\_automation](https://quixy.dev.hp.com/hotfix/index.php?HOTFIX_PROJECT=server_automation)

## Patch required for proper management of Microsoft Windows Server 2008 with Hyper-V systems

A patch from Microsoft is required for managing Microsoft Windows Server 2008 with Hyper-V systems. Patches are installed on the Hyper-V hosts, and after you do so, you must install or update Hyper-V Integration Services from the updated host to each VM guest. Do this as a precaution because not all patches affect Integration Services on the guests.

Microsoft Windows Server 2008 SP2 with Hyper-V requires you to install the following patch that fixes the hang in host servers with many simultaneous virtual machine operations. The patch is located at <http://support.microsoft.com/kb/980081>.

See the *Managing Microsoft Windows Server 2008 Hyper-V with HP Insight Software* technology brief for more information about other patches that might be required for later editions of Microsoft Windows Server 2008 (R2):

[Matrix OE Information Library](#)

## Patch required for HP Integrity VM version 6.1 to deploy IO services containing shared guests

Patch PHSS\_43030 is required to deploy Matrix infrastructure orchestration services containing HP Integrity VM version 6.1 shared guests. Without this patch, Matrix infrastructure orchestration virtual provisioning fails with an insufficient memory error.

Download from <http://www.hp.com/go/hpsc>. Log on with an HP Passport. From the Download options list, select **Patch management** and enter the patch number in the search box. Click **Search**.

## Install Utilization Provider version A.01.08.07.00 for HP-UX 1203 Integrity VM host

Utilization Provider version A.01.08.07.00 is required for systems that have HP-UX vPars and Integrity VM 6.1. Download this version from the [HP Software Depot](#). Without this version, Utilization Provider displays excessive high system memory for HP-UX 1203 Integrity VM hosts even though there are no VMs hosted.

## VM kernel port name used for management traffic must be “Management Network”

In VMware ESXi, the default name of the Virtual Machine (VM) kernel port used for management traffic is “Management Network”. Do not modify this name in the vCenter client user interface. If this name is modified, the following will result.

- When trying to activate a logical server, the ESXi hosts appear as Rejected Target hosts with an error that the hosts are not part of vCenter, or that virtual machine management is unable to get vCenter information for the VM host.
- The association between the host and vCenter does not appear in Systems Insight Manager.

### Suggested action

Modify the **Network Label** of the VMKernel port used for management traffic to “Management Network” by performing the following steps.

1. Log into the vCenter using the vCenter client.
2. Select the **Hosts and Clusters** view from the menu **View**→**Inventory**.
3. Select the ESXi server from the tree displayed in the left pane.
4. Select the **Configuration** tab from the right pane.
5. From the **Hardware** section, select **Networking**.
6. Click **Properties** for the Virtual Switch where the VMKernel port used for management traffic is located. If there are multiple VMKernel ports located on multiple virtual switches and configured for management traffic, select the switch that has the VMKernel port that has the IP that is registered with HP Insight Control virtual machine management, or that will be used for registration with virtual machine management.
7. In the pop-up that displays the virtual switch properties, select the service console port and click the **Edit** button.
8. Change the Network Label to **Management Network** and click **OK**. Close the virtual switch properties pop-up window.

It can take a maximum of five minutes for the cache to be refreshed and the API to report the ESXi host.

## Managing a Microsoft Hyper-V (or any MSCS) cluster when an alias is assigned by DHCP

In Windows Server 2008 Failover Clusters, if the cluster alias is assigned using DHCP, the cluster alias PTR (reverse lookup) record in DNS will be incorrect, and will cause the cluster to be incorrectly discovered in Systems Insight Manager.

To correct this problem, delete the cluster objects from Systems Insight Manager, take the steps in the Microsoft Knowledge Base article at <http://support.microsoft.com/kb/969049> to modify the cluster DNS behavior, and then rediscover the cluster in Systems Insight Manager.

## User experience and application performance

Matrix OE is based on web client server technologies. Your perception of performance can be affected by many external sources, including the central management server’s health, network congestion and latency, and client side equipment and browser.

Every effort to balance performance and functionality has been made to the default settings to create the best user experience under multiple browsers with different server environments. Changing

these default settings can result in adverse application responsiveness, which is more noticeable in some browsers; for example, there are known memory issues in Internet Explorer 7 and 8 that can be exposed in large scale environments. No matter which browser you choose to use, HP recommends using the latest release and updates to improve your user experience and protect your environment from the latest security threats.

## Additional software dependencies

The following are additional dependencies related to Matrix OE for Integrity software.

- Instant Capacity Manager requires iCAP version 8.02 or higher installed on the HP-UX managed systems.
- Matrix OE for Integrity supports having multiple database instances hosted by Microsoft SQL Server. Refer to the Microsoft SQL Server documentation for instructions on how to set up a system to host multiple SQL Server database instances.

---

## 2 Issues and suggested actions

The issues and limitations in this release are listed in this chapter, organized by Matrix Operating Environment feature or functional area. If a feature or functional area has no limitations or issues, that category will not appear. The following categories are used:

Limitations	Limitations of the implemented feature or functional area
Major issues	Issues that may significantly affect functionality and usability
Minor issues	Issues that may be noticeable but do not have a significant impact on functionality or usability

## Matrix Operating Environment Online Help

### Limitations

When a search in the online help results in no hits, the system returns a truncated form of search terms that have many characters

For example, when there are no hits for *Orchestration*, your search returns no results for `orchestr`. This is expected behavior.

#### **Suggested action**

For more information, see the Known Issues section in the [HP Systems Insight Manager Release Notes](#).

## Matrix Operating Environment software

### Limitations

#### Supporting HP IO Accelerators

HP IO Accelerators are supported in the Matrix Operating Environment with limitations. The hypervisor version and Virtual Connect modules need to support the particular HP IO Accelerator. Matrix OE will not provision the hypervisor host to use the HP IO Accelerators (the installation and configuration of the hypervisor using these accelerators would be independent of Matrix OE). Matrix OE can be used to provision VMs to the hypervisor host. For information on HP IO Accelerators, see [http://h18004.www1.hp.com/products/storageworks/io\\_accelerator/index.html](http://h18004.www1.hp.com/products/storageworks/io_accelerator/index.html). Two whitepapers are available online that describe the setup and scenarios for using HP IO Accelerators. You can locate the *Insight Management and HP IO Accelerators (ioTurbine version 2.0.3.1) – Supported Environments and Scenarios White Paper* and the *Insight Management and HP IO Accelerators (ioTurbine version 2.1) – Supported Environments and Scenarios White Paper* in the Insight Management Information Library at <http://www.hp.com/go/insightmanagement/docs>.

#### Requirements for CMS identification changes

If any of the following key attributes of the CMS operating system need to be modified after Insight Management has been installed, use the CMS Reconfigure tool.

The key attributes are:

- hostname
- IP address associated with the hostname
- full name or domain (including the fully qualified domain name)
- DNS entries corresponding to the hostname and domain

If any of these attributes are changed after Insight Management has been installed and the CMS Reconfigure tool was not used, many features of the solution will fail to function properly. For more information on how to use CMS Reconfigure tool, see the “CMS Reconfigure Tool” section in the *HP Systems Insight Manager User Guide* available from the [Systems Insight Manager Information Library](#).

Consider using the Data Migration Tool to migrate existing data onto a newly installed CMS that will use the new hostname or IP address. For more information on how to perform this migration, see the *Data migration of an existing Microsoft Windows CMS to a new Insight Management system* white paper available from the [Matrix OE Information Library](#).

## Configuring Matrix OE with a remote Microsoft SQL Server database

Because Windows Single Sign On is used by Matrix OE to connect to the SQL Server database, the credentials associated with the user account executing a command or service are used to connect to the remote database. Therefore, user accounts used to execute Matrix OE software that accesses the database must satisfy the following conditions:

- The user credentials specified for the CMS access and database access during the installation should be the same, and be trusted domain credentials available to both the CMS and database systems.
- The user credentials specified during the installation for the database access should have access permissions to the database.

If you run a CLI command as a user that is not a domain user or does not have database access permission on the remote database, and the CLI command requires database access, the command will fail.

## Operations Orchestration intended usage in Matrix OE

A limited version of Operations Orchestration is included with Matrix OE. This version of Operations Orchestration is licensed for limited uses.

### Functionality

- Allows Matrix OE infrastructure orchestration to trigger Operations Orchestration workflows for pre- and post- provisioning tasks. Examples include Create Service/Add Server, Standby Server/Service, and Resume Server/Service.

Workflows in the limited version of Operations Orchestration included with Matrix OE may only be triggered from the Matrix OE infrastructure orchestration interface in relation to pre- and post- provisioning tasks. An upgrade to the full version of Operations Orchestration allows the launch of workflows from outside of the Matrix OE interface.

- The limited version of Operations Orchestration allows Matrix OE to be integrated with other management software. However, this requires additional services from HP consulting.
- The limited version of Operations Orchestration included within Matrix OE is licensed for one author. Upgrades to the full version of Operations Orchestration allow for multiple authors.
- The version of HP Operations Orchestration Central that is included in Matrix OE may only be used for flow troubleshooting. An upgrade to the full version of Operations Orchestration allows the use of Operations Orchestration Central to launch and schedule the launch of Operations Orchestration workflows from outside of Matrix OE/Insight Control.

### Content

The content supplied with the limited version is a subset of the content that is available in the full version of Operations Orchestration and is related to Infrastructure-as-a-Service workflows. An upgrade to the full version of Operations Orchestration provides additional content.

## Configuring Internet Explorer 10 Browser Mode

If you are using the runtime Internet Explorer Compatibility View, you might experience issues with Matrix OE components. Use the following instructions to set your runtime version.

1. Open Internet Explorer.
2. Press **F12**.
3. Click **Browser Mode** in the menu bar.
4. From the menu, click **Internet Explorer 8** or **Internet Explorer 9**.

## Capacity Advisor

### Limitations

#### Capacity Advisor does not work with VSP 6.2 hosts

Do not use the Scenario Editor (under the **Optimize**→**Capacity Advisor** menu) in Matrix OE 7.2 to model the behavior of VSP 6.2 hosts. There are issues in collecting data and in modeling the VSP 6.2 host and its VMs and vPars in the Scenario Editor.

With VSP 6.2, there can be a mix of VMs and vPars running on one host at the same time. In this case CapAd will model the vPars as VMs. This is correct for IO and memory, but there is a danger that spare CPU capacity in a vPar will be interpreted as being available to the VMs. CapAd may show there is enough capacity in the VSP host, but it could be that the VMs are not able to get all the CPU resources they need if the vPars are tying up more cores than they need.

Because of this modeling issue, CapAd 7.2 does not claim to support using VSP 6.2 vPars in the Scenario Editor. CapAd will collect utilization data from VSP VMs and vPars. You can view this data in the Profile Viewer. The Capacity Analysis dashboard will still work with VSP VMs and vPars.

#### Workaround

It is possible to work around this modeling issue. A VSP 6.2 host with VMs and vPars can be included in the Scenario Editor. You can also create a what-if VSP host by creating an Integrity VM host. What-if VMs and vPars can be created as VMs. Once you have built up your desired configuration, you need to check it for bottlenecks in two steps. First, check that none of the VMs, vPars, or the host are exceeding their utilization limits in any way. CapAd's Scenario Editor would show red utilization bars if you had an issue here. Then as a second step, possibly in a copy of this scenario, delete all the vPars and reduce the number of cores on the host by the total number of cores allocated to the vPars. In this new scenario, you again check that the VSP host does not have a CPU bottleneck.

---

**NOTE:** Do not use the **What-if Action**→**Automated Load Balance of VM Hosts...** or **What-if Action**→**Change Server to be VMs**→**Automated System Consolidation to VMs...** actions with VSP 6.2 hosts. There is a risk of over committing the resources of the VSP host.

---

Because CapAd does not collect data from VSP hosts in hybrid mode, VSP 6.2 hosts will not have any data collected for them. This will prevent CapAd from making an automatic estimate of the IO capacity of the host. You will have to manually enter your own estimate of the network and disk bandwidth capacity for some VSP 6.2 hosts that you bring into a scenario.

### Minor issues

#### Missing raw data for a time period

Capacity Analysis sometimes encounters missing data when it tries to compute statistics for a day, week, or month. When no data is collected for a day, the average CPU utilization for that day is stored as "unknown" and is displayed as a dash. If all of the data has been collected for a day, the average can be computed and stored. However, if data has been collected for only part of a

day, then Capacity Analysis can report the average of the data that was collected, or report that the average is unknown. In this release of Capacity Analysis, the rules for displaying a dash or a value based on partial data are not well defined and can vary from metric to metric.

### **Suggested action**

If you have an analysis query that is sensitive to the amount of available data, you can edit the query to add a filter expression for the "CPU % valid" (the CPU utilization percentage). You might want to display the CPU % value in any query results to show whether the results are full or partial. Additionally, you can change the analysis interval to a time period where all of the data is more likely to be available.

## Infrastructure orchestration

Infrastructure orchestration runs on an x86/x64 Windows CMS.

### Limitations

#### Cross-technology logical servers are not supported by infrastructure orchestration

Cross-technology logical servers are not supported by Matrix OE infrastructure orchestration. Only logical servers created or imported in Matrix OE logical server management, using **Create→Logical Server** or **Tools→Logical Servers→Import**, can be moved from physical to virtual or virtual to physical. A logical server created in infrastructure orchestration (by deploying a service template) cannot be managed as a cross-technology logical server. For example, if a physical logical server is created in infrastructure orchestration, it cannot be moved to a virtual machine using logical server management. Similarly, if a virtual logical server is created in infrastructure orchestration, it cannot be moved to a server with Virtual Connect using logical server management.

For more information about cross-technology logical servers, see the *HP Matrix Operating Environment Recovery Management User Guide* at [Matrix OE Information Library](#).

#### Number of networks provisioned to a logical server

The infrastructure orchestration designer does not restrict the number of networks that can be connected to a logical server group. In practice, only four networks can be connected to virtual logical servers, and 128 networks can be connected to physical logical servers designed to run on HP c-Class server blades. These practical limits will increase over time as Virtual Connect and hypervisors are enhanced.

If these limits are exceeded in an infrastructure template, provisioning operations using that template will fail.

### **Suggested action**

Limit the number of network connections to a logical server group to four for virtual logical servers and 128 for physical logical servers.

#### Manually configure the SAN boot path for ProLiant DL/ML servers before provisioning

The job "Deploy ProLiant System Configuration (BL BFS)," contained in the Insight Control server deployment software job folder deployed using infrastructure orchestration, cannot be used to configure ProLiant DL/ML target servers. The job configures ProLiant/Integrity BL target servers. If you attempt to configure ProLiant DL/ML servers using this job, the request will fail.

### **Suggested action**

Manually configure the SAN boot path before you use infrastructure orchestration to provision a ProLiant DL/ML target server on SAN boot LUN, as follows.

1. Create a server deployment job folder that does not contain the "Deploy ProLiant System Configuration (BL BFS)" job.
2. Change the Boot Controller Order (from ROM-Based Setup Utility) for the ProLiant DL/ML server to select Fibre Channel port (HBA) as the first in the boot order.
3. Use the HBA vendor's utility to configure the HBA's boot path to boot from the desired SAN boot LUN.
4. Power off the target server after the Boot Controller Order and the HBA's boot path have been configured correctly.

## Required HBA firmware and OS drivers to access multi-initiator N-Port ID Virtualization (NPIV) data disks

The HP blade multi-initiator NPIV feature allows you to configure multiple WWNs for a single HBA connected with a Virtual Connect Fibre Channel module. Although Matrix OE automates the configuration, certain versions of HBA firmware and drivers are required to access multi-initiator NPIV data disks configured for any target BL server.

### Suggested action

The HBA firmware and drivers shown for the following operating systems are required to support multi-initiator NPIV data disks.

- QLogic HBA (BIOS version 2.07 or higher)  
RHEL 5.3 and higher  
SLES 11 and higher  
SLES 10 SP3 and higher  
Windows 2003/2008: QLogic driver 9.1.7.29 or higher
- Emulex HBA (BIOS version 3.11a3 or higher)  
RHEL 5.4 or higher  
SLES 10 SP3 or higher  
SLES 11 in-box driver not yet available. Requires out-of-box Emulex driver version: 8.2.8.32  
Windows 2003/2008: Emulex driver 5.00.17.06-2 or higher

## Multi-initiator NPIV support on Integrity blades

Multi-initiator NPIV features used by Matrix OE and Virtual Connect Enterprise Manager are supported on Integrity blades using the QLogic FC controller with the HP-UX 11i v3 Update 8 (March 2011) and later release. Multi-initiator NPIV features are supported on Integrity blades with Emulex FC controllers with HP-UX 11i v3 March 2013 update and later release.

### Suggested action

On Integrity blades with Emulex controllers and/or versions of HP-UX earlier than 11i v3 March 2013 update release, create a single storage pool entry that combines a boot volume and the data volumes intended to be allocated to a single physical server. You can do this in one of the following ways:

- Pre-provision any number of these composite storage pool entries to enable fully automated allocation of SAN volumes to physical servers.
- Pre-provision storage pool entries with only boot volumes. An automated provisioning request pauses after the operating system is deployed if the infrastructure service template specifies data disks. You can modify the storage pool entry allocated to the server to add data volumes, and then continue the request.

## Required software versions for FlexFabric Integrity physical provisioning support

Matrix OE infrastructure orchestration requires the following minimum software versions to support HP FlexFabric for Integrity physical provisioning:

- Ignite-UX C.7.12.475 or higher
- HP-UX 11i v3 Update 8 (March 2011) or higher

### Install Microsoft Sysprep tool for customization of Windows 2003 guests on Hyper-V after installation

The Microsoft System Preparation (Sysprep) tool, which is required for customization of Windows 2003 guests on Microsoft Hyper-V, is not retained on the CMS after upgrading from previous versions of Insight Software. Without these tools, customization of Windows 2003 guests on Hyper-V results in failures.

Installation of the Sysprep tool is not required for customization of Windows 2008 and higher VM guests.

#### **Suggested action**

To enable customization of Windows 2003 guests on Hyper-V, install the Microsoft Sysprep tool on the CMS in the folder `C:\Program Files\HP\Insight Control virtual machine management\Sysprep\2003` after the installation or upgrade of Insight Management 7.1.

### Cannot power on a virtual server during provisioning if VM hardware version is 4 and a template contains more than three NICs

A request to power on an ESX virtual server may fail if the ESX VM hardware version is set to 4 and the virtual machine contains more than three NICs. Matrix OE infrastructure orchestration provisions virtual machines using ESX hardware version 4 by default, which allows a maximum of five PCI Controller devices (two SCSI controllers and three NICs) attached to a virtual machine. ESX hardware version 7 allows four SCSI controllers and ten NICs.

#### **Suggested action**

Use an ESX native template with a hardware version of 7, or reduce the number of NICs in the infrastructure orchestration template to three or fewer.

### Physical disks with multiple partitions are unsupported unless all but the first partition are excluded

If your environment contains a physical disk with multiple partitions (volumes with the same disk identifier) and you want to provision virtual machines on that disk, exclude all partitions except the first partition on the disk. A virtual machine created on a physical disk with two or more partitions without this exclusion is unsupported, and the service could fail to deploy with the error "Import failed for *logical-server*. VM host datastore not found."

#### **Suggested action**

Exclude all partitions except the first partition on the disk in the "VOLUMES TO EXCLUDE" section of the `hpio.properties` file.

For example, if there are three volumes (C:, D:, and E:) on a physical disk, exclude the D: and E: volumes to provision a logical server on the virtual machine created on the C: volume of the physical disk.

### Cluster shared volume does not appear after Microsoft Hyper-V VM provisioning with high availability enabled

If you use infrastructure orchestration to create and deploy a virtual provisioning template on a Microsoft Hyper-V VM host with the High Availability (HA) option enabled, and then view the virtual machine in the Failover Cluster Manager on the Hyper-V host, the cluster shared volume (CSV) is not displayed in the provisioned VM information.

However, if you manually deploy the virtual machine using the Microsoft Hyper-V Manager on the Hyper-V VM host, the Failover Cluster Manager correctly displays the disk resource and the virtual machine resources created.

### Suggested action

No action is needed; the highly available virtual machine disk was created correctly on the CSV.

### Cluster shared volume is not correctly recognized unless all VM hosts in the Hyper-V cluster are registered in Insight Control virtual machine management

When only one VM host in a Hyper-V cluster is managed by Insight Control virtual machine management, the following areas are impacted:

- Matrix infrastructure orchestration — IO treats a cluster shared volume (CSV) as a cluster disk, and allocates only one VM to the cluster disk. When trying to deploy multiple VMs, the following error displays:

```
Provisioning request has paused. Modify storage volumes on VM Hosts to satisfy the logical disk requirements. An IO Administrator has been notified.
```

- Matrix OE logical server management — Only the CSV and cluster disk are available on the VM host if the registered VM host ID is the VM host owner. If it is not the owner, the VM host is placed on the Rejected List and an error indicates that the storage is not available on the VM host.

### Suggested action

Add all VM hosts that are in the Hyper-V cluster to Matrix infrastructure orchestration, registering them with Insight Control virtual machine management as VM hosts. The VM hosts can be left in the Unassigned pool, or they can be used as regular VM hosts within IO.

### Cannot provision a server group on Microsoft Hyper-V R1 with linked clone and high availability enabled

If you attempt to provision a template containing multiple High Availability (HA)-enabled linked clones on a Hyper-V R1 server, or you attempt to add a server to a server group with a single HA-enabled linked clone on a Hyper-V R1 cluster disk, the request pauses indefinitely for storage provisioning. A linked clone group must be provisioned on a single datastore, but only one HA virtual machine can be provisioned to a Hyper-V cluster disk.

### Suggested action 1

1. Cancel the Create or Add Server request.
2. Modify the template to remove the HA requirement, or  
Modify the template to deselect linked clones and add a cluster disk for each virtual machine.
3. Resubmit the Create request, or  
For the Add Server request, delete the existing service and recreate it with the updated template. Take care to migrate data from the existing service to the new service before deleting the existing service.

### Suggested action 2

Upgrade to Microsoft Hyper-V R2, and specify a Cluster Shared Volume (CSV) instead of a cluster disk in the template.

### Microsoft SCVMM can be registered on one CMS

As part of the registration of Microsoft System Center Virtual Machine Manager (SCVMM), the web service component is copied from the CMS to the SCVMM host, which enables communication between the SCVMM server and the CMS. As a result, SCVMM can be registered on only one CMS. This means that in a federated CMS configuration, each CMS requires a different SCVMM host to manage Hyper-V hosts through SCVMM.

## Browsing to user interface requires default Systems Insight Manager port

HP Systems Insight Manager is installed with a default port of 50000 in the JBoss container. Browsing to the infrastructure orchestration user interfaces, for example, Designer, depends on this port value.

### Suggested action

Do not modify the Systems Insight Manager default port of 50000 in the JBoss container.

## Operations Orchestration Central and Studio contents are not localized to Japanese

In new installations of Matrix OE 7.2, Operations Orchestration Central and Studio screens, menus, and buttons are mostly localized to Japanese, and the contents (such as workflows and system properties) are not localized to Japanese when OO is installed in a Japanese environment.

If you are upgrading to Matrix OE 7.2, Operations Orchestration Central and Studio display the localization behavior of the version that was originally installed. Therefore, if you are upgrading from Insight Dynamics 6.3 to Matrix OE 7.2, the Operations Orchestration Central and Studio screens are displayed in English.

## Customizable allocation process deprecated

The customizable allocation process functionality in Matrix Infrastructure Orchestration has been deprecated in the 7.2 release. The functionality will be removed in a future version.

## Minor issues

### Delete service request leaves "clean-me" logical servers when used with Integrity servers if disk is not erased before delete service request

After successfully provisioning a physical Integrity server blade (using Ignite-UX, or through a secondary CMS in a federated CMS environment), there are manual steps you must perform to allow the disk to be erased during de-provisioning. Erasing the disk is part of a successful de-provisioning (delete request) for an Integrity service.

If the disk is not erased, the delete service request may fail and "clean me" logical servers will remain. In that case, a manual cleanup procedure is required. (See the **Suggested action** for the cleanup procedure.)

Before you start a delete service request from infrastructure orchestration:

1. Make sure that the Integrity service was successfully provisioned (the Integrity service is running).
2. Discover the Integrity blade where the Integrity service is running using its IP address from Systems Insight Manager. In a federated CMS environment, discover the Integrity blade on the primary CMS using its IP address from Systems Insight Manager. The Integrity node's identification information should contain the UUID and UUID (logical). You can discover the blade in either of the following ways:

- a. Use the Systems Insight Manager menu selection **Options**→**Discovery...**, or
- b. Use the Systems Insight Manager command: `mxnode -a <ip-address> -w`  
where `<ip-address>` is the IP address that was assigned to the Integrity blade after provisioning.

Add the root login of the Integrity blade using the Systems Insight Manager menu selection **Options**→**Security**→**Credentials**→**System Credentials...**

3. Set up an SSH connection between the CMS and the Integrity blade by entering the command: `mxagentconfig -a -n <ip-address> -u root -p <root password>`
4. If the IP address was discovered on the primary CMS in a federated CMS environment (in step 2 above), manually delete the managed node.

Use the Systems Insight Manager command: `mxnode -r <ip-address>`

## Suggested action

If you did not perform the preceding steps before you started a delete request for the Integrity service, the automatic delete service will not succeed completely. In that case, you must clean up the logical server associated with the Integrity blade and manually erase the disk. The logical server should be in an operable state.

1. Power on the server blade. (The blade was powered off when delete service request completed or failed.)
2. Perform steps 2 and 3 in the preceding section, if you did not do so before starting the delete service request.
3. To erase the disk, enter the following command from the CMS where the discovery was executed (in step 2 in the preceding section):

```
mxexec -t "Ignite Erase Disks" -n <ip-address>
```

4. If the IP address was discovered on the primary CMS in a federated CMS environment (in step 2 in the preceding section), manually delete the Systems Insight Manager node.  
Use the Systems Insight Manager command: `mxnode -r <ip-address>`

## Service creation fails for physical requests when Windows Server OS is selected

Matrix OE infrastructure orchestration service creation fails when a physical template with Windows Server OS is selected. This failure occurs when infrastructure orchestration times out waiting for Insight Control server deployment to complete installation of Insight Management Agents on the target server.

### Suggested action

When a request fails, infrastructure orchestration automatically retries the request and a new server is provisioned with the requested service. The original server is moved to the maintenance pool and can be cleaned and moved back to a valid server pool. If the failure is observed frequently, the following steps can be followed to disable the installation of Insight Management WBEM providers. Insight Management SNMP agents continue to provide management capabilities when WBEM providers are not installed.

1. Edit the `install.cmd` file, found under the `<Insight Software installation>\rdp\deployment server\hpfeatures\hpagents-ws\components` folder.
2. Remove the following two lines:

```
call :installsc wbem%PROCESSOR_ARCHITECTURE%  
if %cpreresult% geq 1 set /a cpreresults=cpreresults+1
```

## "Any Of" and "All Of" requirements in SPM are not displayed in infrastructure orchestration designer

Using SPM, you can define complex requirements with nested "Any of" (or) and "All of" (and) constructs. These requirements are correctly applied when provisioning storage in SPM, but infrastructure orchestration designer does not display these requirements in detail. The IO template architect may be allowed to enter values when defining storage that will cause storage provisioning to fail. For example, a combination of the "Any Of" or "All Of" requirements allows only a specific RAID level or specific size range. If the architect picks a RAID level or range that is not allowed in the template, storage provisioning will fail.

### Suggested action

When defining SPM storage templates, use the Description field to summarize the requirements (for example, if the template supports only RAID 1 and RAID 5). That description is visible to the architect in infrastructure orchestration designer. When using a specific SPM storage template in IO, be aware of the storage template's capacity, RAID, and tag requirements (included in the Description field). Do not specify values that conflict with these requirements or subsequent storage provisioning will fail.

## Provisioning a service may fail when storage volumes in a logical server group contain different tags

In an IO template, different storage tags can be assigned to the individual volumes in a logical server group. When storage is auto-generated, each volume is provisioned based on a number of defined requirements, including these tags. Within a logical server group, if the private data disk(s) contain different tags than the boot disk, the service will fail to provision.

### Suggested action

Use the same storage tags for the boot disk and private data disks in a logical server group. If separate tags are needed, separate storage pool entries can be manually defined for the boot disk and private data disks with those separate tags.

## Storage auto provisioning does not select a preconfigured data disk if the storage pool entry does not contain a boot disk

During storage provisioning, if a storage pool entry contains a volume that matches the requirements of an infrastructure orchestration service, the matching volume is used for the service.

If the IO template contains a boot disk and a data disk, and the storage pool entry contains only a boot disk, the boot disk is used and the data disk is automatically created. However, if the storage pool entry contains only a data disk, the volume is not used. Both the boot and data disk will be created.

### Suggested action

Add a boot disk to the storage pool entry containing the desired data disk, or allow both disks to be automatically created.

## Provisioning errors on Hyper-V cluster after CSV disk failure

If a cluster shared volume (CSV) is in a failed state (for example, if LUNs are not available or in a read-only state), Insight Control virtual machine management cannot obtain any information about the CSV. Subsequently, operations on the Hyper-V cluster (including provisioning) will fail, or all VM hosts in the cluster will be excluded from operations.

### Suggested action

1. Remove the failed CSV from the Hyper-V cluster and correct the issue that caused the failure.
2. Bring the CSV back online.
3. Manually refresh server resources in Matrix OE visualization by selecting **Tools**→**Logical Servers**→**Refresh...**
4. Retry the operation.

## NIC is not configured properly on a new VM after a successful RHEL deployment

Networking may not work on a provisioned ESX Linux virtual machine. This is due to a known VMware issue.

### Suggested action

For information about how to reconfigure your network on the provisioned VM, see [http://kb.vmware.com/selfservice/microsites/search.do?language=en\\_US&cmd=displayKC&externalId=2002767](http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2002767).

## VM cluster provisioning requires restart of VMware Management Agents for preferred hostname to be registered in DDNS

The VM cluster Operations Orchestration workflows enable Dynamic DNS integration by setting a VMware property. For this property to be activated, and the DDNS registration performed, the management network on the ESX host must be manually restarted.

### Suggested action

After provisioning an ESXi VM cluster, if full registration in a DDNS environment is required, restart the Management Agents on each VM host in the VMware cluster just created using the server console. You can alternatively use the vCenter VM host context menu to request that the host be re-booted. The fully qualified name of the host will then be registered.

### Set retry parameters for BL620c G7 and BL680c G7 blades

Matrix OE infrastructure orchestration may time out while waiting for the server to be recognized due to long server post times, particularly when using the BL620c G7 or the BL680c G7.

### Suggested action

To avoid this issue on any server:

1. Edit `C:\Program Files\HP\Matrix infrastructure orchestration\conf\hpio.properties` and set the `retry.count.ares.pxe` parameter to a value greater than 30.
2. Restart the HP Matrix infrastructure orchestration service.

### Some templates are not imported correctly and do not appear in Organization tab

Occasionally, some 6.3 infrastructure orchestration templates are not imported correctly. These templates appear in the infrastructure orchestration console Templates tab, but do not appear in the Organization tab Service Provider resource box. As a result, the templates cannot be assigned to an organization.

### Suggested action

Open the template in infrastructure orchestration designer, and "Save As" the template. It will then appear in the Organization tab.

### After upgrade, ioexec commands require installation path

The infrastructure orchestration Windows environment default installation path was changed from `C:\Program Files\HP\Insight Orchestration` to `C:\Program Files\HP\Matrix infrastructure orchestration`. In a new installation of Matrix OE 7.2, the `bin` folder for IO is added to the DOS path, enabling you to run `ioexec` commands from a command prompt in any folder on the CMS. However, after an upgrade from Insight Dynamics 6.3 to Matrix OE 7.2, `ioexec` commands can be run only from the `C:\Program Files\HP\Matrix infrastructure orchestration\bin` folder.

### Suggested action 1

Manually enter `C:\Program Files\HP\Matrix infrastructure orchestration\bin` on the command line before entering any `ioexec` commands.

### Suggested action 2

Reboot the CMS to automatically update the IO `bin` folder to the DOS path.

### Firefox 17 requires new method to accept self-signed certificate for access to IO

The first time you run IO in a Firefox 17 browser, a message "This Connection is Untrusted" might be displayed. Firefox 17 might need a new certificate.

### Suggested action

Do the following to assign a new certificate:

1. On the Firefox "This Connection is Untrusted" screen, click to expand **Technical Details** or **I Understand the Risks**.
2. Click **Add exception** or **or you can add an exception**.

If there is no option to add an exception:

- a. Open a new tab in Firefox.
  - b. Type the address for the IO self-service portal,  
`https://[your-CMS-name-or-IP]:51443/hpio/portal` or  
`https://[your-CMS-name-or-IP]:51443/hpio/portal.jsp`.
  - c. Firefox will display the “This connection is Untrusted” screen with the option to add an exception.
3. Firefox will display the Add Security Exception screen.
  4. Fill in the **Location** box, if it is not filled in by default.  
Type `https://[your-CMS-name-or-IP]:50000/`
  5. Click the **Permanently store this exception** checkbox.
  6. Click **Get certificate**.
  7. Click **Confirm Security Exception**.
  8. Return to the SIM tab and refresh the screen, or open a new tab and go to the SIM URL and start the IO.

### Sample ESA server files are changed after upgrade

If you are using Operations Orchestration workflow servers, and you upgrade from a previous release of Matrix OE to Matrix OE 7.2, the sample HP Extensible Server Adapter server files in `<Matrix infrastructure orchestration install dir>\esa_extensions\server` have been changed to comment out part of the XML file. This avoids confusion when you enable the ESA password in `esa.properties` configuration file, but it does not change the contents of the ESA server files. If you have modified the ESA server files, the sample files are not changed.

#### Suggested action

Note the changes to the server files after an upgrade.

### Templates that connect a physical disk to an Hyper-V HA virtual server group cannot be provisioned

Designer allows a physical disk to be connected to a High Availability Hyper-V virtual server group and allows you to set the Raw Device Mapping Type for the physical disk. However, provisioning a template with this configuration returns an error.

#### Suggested action

In infrastructure orchestration designer, delete the physical storage component and replace it with a virtual storage component, then retry the create service operation.

### Deleting the chargeback archive can cause the HPMOEMetering database transaction log to grow very large

Using the chargeback CLI `chargeback archive databasedelete` command can cause a large number of registers to be added to the transaction log, especially if the date specified by the `-d` qualifier is far in the past (a year or more). If the log grows large enough, the available space on the hard disk on which the database is running may run out of free space.

#### Suggested action

Before executing `archive databasedelete`, consult with your database administrator. If you have already executed the command, ask the administrator about performing a shrink operation on the transaction log.

### Matrix infrastructure orchestration may time out while concurrently provisioning Integrity VMs

When infrastructure orchestration attempts to provision Integrity VMs, the request may take longer than 60 seconds causing a time out.

## Suggestion actions

The timeout has to be increased from the default of 60 seconds:

1. Configure the timeout by adding the following entry to the `hpvmm.conf` file in `\Insight Control virtual machine management\bin` `VirtQueryTimeout=` (time in ms) > i.e. `xx000`. The default value of 60000 so change to something higher.
2. Restart Insight Control virtual machine management by selecting HP Insight Control virtual machine management, right-click and then click **Restart**.

## Insight managed system setup wizard

### Limitations

#### Privilege elevation configuration that requires a password by the privilege elevation tool is not supported

Insight managed system setup wizard supports privilege elevation only on a CMS that has privilege elevation configured so that no password is required by the privilege elevation tool (such as `su` and `sudo`). If you attempt to configure features when privilege elevation is turned on and requires a password, you may see failures in the wizard's execution output.

#### Suggested action

If a feature fails to configure correctly because of privilege elevation password requirements, configure the feature outside of the managed system setup wizard.

To turn off the requirement for a password, use the Systems Insight Manager

**Options**→**Security**→**Privilege Elevation** menu, and uncheck the check box labeled "A password is required for this privilege elevation tool." Note that systems managed by the CMS must be reconfigured to reflect this configuration change.

#### HP Integrated Lights-Out 3 (iLO3) support

HP Integrated Lights-Out 3 (iLO3) is supported in the current release of Insight managed system setup wizard. However, iLO3 for ProLiant requires firmware version 1.05 or higher.

#### Suggested action

Update iLO3 ProLiant firmware to version 1.05 or higher.

#### HP Integrity rx2620 is not supported

In this release, using the Insight managed system setup wizard to configure an HP Integrity rx2620 server is not supported.

#### Suggested action

Manually configure HP Integrity rx2620 servers.

## Instant Capacity Manager

### Limitations

#### Commands executed under privilege elevation may cause Instant Capacity Manager to monitor but not control managed nodes

Some HP-UX security settings cause output to be generated every time a user is logged in. This can occur, for example when `sudo` is used in Systems Insight Manager privilege elevation mode to execute commands on HP-UX managed nodes, which can generate apparently extraneous output to `stderr`. This output may be interpreted by some Instant Capacity (iCAP) Manager functions executed under privilege elevation as error output from those commands.

When this apparent error output occurs, the affected iCAP Manager functions fail. In this scenario, iCAP Manager can be used only to monitor, and not control, managed nodes configured in this way.

---

**NOTE:** HP Integrity Superdome 2 systems obtain monitoring data by executing commands using `sudo`. Monitoring data may not be available in this scenario.

---

### Suggested action

Configure HP-UX managed nodes so that no extraneous output is generated when the `sudo` command is executed.

For more information, see the following [HP-UX documentation](#):

- *HP-UX 11i Security Containment Administrator's Guide*
- *HP-UX System Administrator's Guide: Configuration Management* (HP-UX 11i v3)
- *HP-UX Servers and Workstations – HP-UX Standard Mode Security Extensions* (HP-UX 11i v2)
- the manpage `man 4 security` (particularly the discussion of `DISPLAY_LAST_LOGIN`)

## Logical server management

Logical server management runs on an x86/x64 Windows CMS.

### Limitations

#### Restrictions on using cross-technology logical servers

Cross-technology logical servers can be moved between ProLiant server blades with Virtual Connect to ESX and back, between Integrity server blades with Virtual Connect and Integrity VM and back, and between servers with Virtual Connect with unlike configurations. For more information, see the *HP Matrix Operating Environment Logical Server Management User Guide* and the *HP Matrix Operating Environment Recovery Management User Guide* at the [Matrix OE Information Library](#).

- Valid targets for activating or moving a cross-technology logical server:
  - ESX VM hosts
  - HP Integrity VM hosts (racks and server blades)
  - ProLiant c-Class servers with Virtual Connect
  - HP Integrity servers with Virtual Connect
- Storage requirements:
  - SAN volume presented to an ESX VM host (RDM disk) that can be added to a Virtual Connect profile.
  - One datastore shared on a SAN volume to store the HP Insight Control virtual machine management configuration file and RDM mapping file `[.vmdk]` on the ESX VM host. An iSCSI volume can be used for the HP Insight Control virtual machine management configuration file, but not the RDM mapping file.

- Not supported: SPM catalog storage volumes.  
In this release, catalog storage pool entries are supported on servers with Virtual Connect logical servers only. Catalog storage pool entries cannot be used with cross-technology logical servers that may be activated on an ESX or Integrity VM host.
- Integrity VM storage requirements:
  - SAN-based logical server storage is supported, including storage pool entries. File based storage is not supported with Integrity VM cross-technology logical servers.

---

**NOTE:** SAN volumes used for Integrity VM fluid cross-technology logical servers should not be presented to the Integrity VM host.

---

- ESX network requirements:
  - Single or multiple networks (limited to 4 by ESX)  
Not supported: VLAN networks
- Integrity VM network requirements:
  - Not supported: VLAN networks and Auto-Port Aggregation Link Aggregation configurations

### Performing operations on a logical server running as a VM host

Before performing any operations on a physical server running as a VM host, make sure that all jobs related to controlling the VM guests running on the VM host logical server have been completed. (A warning message is displayed only when you attempt to deactivate or delete virtual machines registered with Insight Control virtual machine management.)

Refer to the following table for the appropriate steps to prevent VM guests and VM logical servers from becoming inoperable.

VM host operation you want to perform	Steps to perform on VM guests	Steps to perform for the VM host <i>after</i> performing the steps on the VM guests
Move VC profile to another blade	Power off or perform a graceful shutdown of all VM guests	<ol style="list-style-type: none"> <li>1. Unregister the VM host from Insight Control virtual machine management. From HP SIM, select <b>Configure</b>→<b>Virtual Machine</b>→<b>Unregister Virtual Machine Host...</b></li> <li>2. Move the VM host logical server. Select <b>Tools</b>→<b>Logical Servers</b>→<b>Move</b>.</li> <li>3. Register the moved VM host using Insight Control virtual machine management. From HP SIM, select <b>Configure</b>→<b>Virtual Machine</b>→<b>Register Virtual Machine Host...</b></li> </ol>
Shut down VM host	Power off or perform a graceful shutdown of all VM guests	Shut down the VM host logical server.
Power off the VC blade	Power off or perform a graceful shutdown of all VM guests	<ol style="list-style-type: none"> <li>1. Perform a graceful shutdown of the VM host.</li> <li>2. Power off the VC blade where the VM host logical server resides. Select <b>Tools</b>→<b>Logical Servers</b>→<b>Power Off...</b></li> </ol>

VM host operation you want to perform	Steps to perform on VM guests	Steps to perform for the VM host after performing the steps on the VM guests
Delete the VM host logical server	<ul style="list-style-type: none"> <li>If you no longer need the VM guests, delete them. Otherwise, move all VM guests to a suitable host.</li> <li>If you have VM guests that are not logical servers, then use the appropriate tools to relocate the guests.</li> <li>If you have any inactive VM logical servers, reactivate them and then move them to a suitable host.</li> </ul>	<ol style="list-style-type: none"> <li>Unregister the VM host from Insight Control virtual machine management. From HP SIM, select <b>Configure</b>→<b>Virtual Machine</b>→<b>Unregister Virtual Machine Host...</b></li> <li>Delete the VM host logical server. Select <b>Delete</b>→<b>Delete Logical Server...</b></li> </ol>
Deactivate or unassign the VC profile	Power off or perform a graceful shutdown of all VM guests	Deactivate the VM host logical server. Select <b>Tools</b> → <b>Logical Servers</b> → <b>Deactivate</b> .
Activate or reassign the VC profile to a different blade	Power off or perform a graceful shutdown of all VM guests	<ol style="list-style-type: none"> <li>Unregister the VM host from Insight Control virtual machine management. From HP SIM, select <b>Configure</b>→<b>Virtual Machine</b>→<b>Unregister Virtual Machine Host...</b></li> <li>Activate the VM host logical server. Select <b>Tools</b>→<b>Logical Servers</b>→<b>Activate</b>.</li> <li>Register the VM host using Insight Control virtual machine management. From HP SIM, select <b>Configure</b>→<b>Virtual Machine</b>→<b>Register Virtual Machine Host...</b></li> </ol>

### Moving a logical server from a Microsoft Hyper-V R2 host to a Hyper-V R1 host is not supported

A limitation from Microsoft prevents you from moving or reactivating a virtual machine logical server from a Microsoft Hyper-V R2 host to a Hyper-V R1 host. An error message is displayed if you attempt the operation.

You can move or reactivate a virtual machine logical server from a Hyper-V R1 host to a Hyper-V R2 host.

**NOTE:** The Hyper-V host version is shown in the Platform column in the Available Target Hosts table on the activate or move logical server target screen. The host is a Hyper-V R1 host if the host version begins with 6.0 (for example, 6.0.6001 or 6.0.6002 in Windows 2008 Server – Hyper-V SP1 or Windows 2008 Server – Hyper-V SP2). The host is a Hyper-V R2 host if the host version begins 6.1 (for example, 6.1.7077 in Windows Server 2008 R2).

#### Suggested action

Be certain that you want to move or reactivate the logical server from Microsoft Hyper-V R1 to R2 before you complete the operation. After the operation completes, you cannot move or reactivate the logical server back to Hyper-V R1.

### Multiple datastores are not supported

All storage volumes for a virtual machine logical server must be part of the same datastore. (To import a virtual machine, the virtual machine must be using a single datastore.) File volumes are created on the same disk as the virtual machine's configuration file.

### Virtual machine datastore size and configuration filename limits

When creating a virtual machine logical server (ESX, Hyper-V, or Integrity VM), the maximum datastore size that you can specify is 2048 TB (2097151 GB). Additionally, the maximum number of allowed characters for a virtual machine configuration filename is 255. If you specify a size larger than the maximum, a warning message is logged and the entry is ignored.

VM technologies have individual configuration limits that may limit the VM disk size.

### Logical server management supports virtual machines created only on the first partition of a physical disk

If a physical disk has multiple partitions (volumes with the same disk identifier), only the first volume is recognized and supported by logical server management. For example, if a disk has volumes C:, D:, and E:, you can create and import a virtual machine logical server created only on the C: volume.

### Importing a Microsoft Hyper-V virtual machine with configuration file in the default location is not supported

A Hyper-V virtual machine whose configuration file is in the default location on the hard disk in `\ProgramData\Microsoft\Windows\Hyper-V` is not supported by logical server management. You cannot import a virtual machine guest as a logical server if the virtual machine's configuration file is in this location. There is a limitation in Microsoft Hyper-V that allows Insight Control virtual machine management to re-register the virtual machine if the configuration file is created in this directory. To prevent issues with reactivation, logical server management checks the path and rejects any virtual machines with a configuration file in the default location.

### Hyper-V cluster shared volumes created with GPT disk partitioning system have limited support

Hyper-V VM hosts configured with cluster shared volume (CSV) storage with a GUID partition table (GPT) disk partitioning system are supported with the following conditions:

- All hosts (cluster nodes) must contain the same LUN identifiers for the LUN
- LUN identifiers must be unique in the host if the LUNs are presented from different storage arrays

### iSCSI disk requirements for use with Microsoft Hyper-V virtual machine logical servers

To enable selection of a specific iSCSI datastore as the storage definition during logical server creation or modification, iSCSI disks associated with Microsoft Hyper-V hosts must be a cluster disk or a shared cluster volume, online, and contain a volume created with a drive letter.

With this limitation, only MBR (master boot record) formatted iSCSI cluster disks are supported.

### Microsoft Hyper-V virtual machine logical server boot disk is IDE and data disks are SCSI

When you create a Hyper-V virtual machine logical server, the boot disk (where Windows or Linux is installed) is always created using an IDE controller. Data disks are created using a SCSI controller. The IDE controller is required to allow the installation of the operating system using Deploy Tools or using images without additional drivers. If a virtual machine is created outside of Matrix OE visualization and imported as a logical server, it is possible that the boot disk controller type is SCSI. In this case, you must install the SCSI controller driver to allow installation of the operating system.

If a virtual machine is provisioned by Matrix OE infrastructure orchestration (using either a SCSI or an IDE controller) and the logical server is imported, the virtual machine has an operating system installed so the SCSI driver is not needed.

### One Microsoft Hyper-V target displayed when activating a logical server created on a cluster disk volume

A limitation in Microsoft Hyper-V R1 prevents the hypervisor from providing logical server management with required information about a Hyper-V cluster. The result is that only one target is shown when you attempt to activate a virtual machine logical server configured to be created on a cluster disk volume. In Hyper-V R1, a logical server that is high availability (HA)-enabled can be configured only on a cluster disk volume.

## Suggested action

Upgrade to Hyper-V R2. In R2, Microsoft corrected the problem so that the activation of a virtual machine logical server configured to be created on a cluster disk or a cluster shared volume lists all Hyper-V hosts in the cluster on the target screen.

## High Availability Hyper-V virtual machine with RDM disk can be activated only on cluster disk owner

When an HA Hyper-V logical server with a SAN-based storage pool entry (RDM disk) is activated, only the RDM disk host owner displays as a target on which to activate the logical server. It is not possible to enable HA for a VM with an RDM disk if the Hyper-V host that is hosting the VM is not the owner of the RDM disk.

## Some Virtual Connect profile capabilities are not supported by logical server management

Importing a server with Virtual Connect that implements the following features is not supported by Matrix OE logical server management. If the Virtual Connect profile is using these features, the error "The Server Profile cannot be imported because it contains the following capabilities unsupported by Logical Servers:" is displayed.

EXTENDED_FC	Server profile has two or more FC Connections using the same connection bay (with the same connectionBay attribute value), which identifies a server profile created for an HP Integrity BL860c)
EXTENDED_FCOE	Server profile has more than one Fibre Channel over Ethernet (FCoE) connection associated with the same interconnect bay (IO bay)
ISCSI	Server profile has at least one iSCSI connection
NAG	Server profile has an associated Network Access Group

## LUNs are not detected after activating a logical server using logical server management on an HP Integrity Virtual Machine host

A logical server, when activated directly on an HP Integrity Virtual Machine host using LSM, fails to detect the LUNs assigned to it at the virtual machine EFI prompt. Due to this limitation, the logical server will not boot up the operating system automatically even if it is assigned with a LUN that is pre-provisioned with an operating system.

## Suggested action

At the Integrity Virtual Machine EFI prompt, execute the following commands to see LUN mappings. When these steps are complete, the user can select the required LUN from the displayed device list to boot up or to install the operating system.

1. Check configurable controllers using the `drvcfg` command.

```
Shell> drvcfg
Configurable Components
Drv[2C] Ctrl[29] Lang[eng]
```

2. Enumerate all SCSI and FC LUNs using `drvcfg -s <drvnum> <ctrlnum>`. (`drvnum` and `ctrlnum` are displayed as part of step 1.) Press Y for confirmation on enumerating both SCSI and FC LUNs.

```
Shell> drvcfg -s 2C 29
Set Configuration Options
Drv[2C] Ctrl[29] Lang[eng]
HP AVIO Stor Driver Configuration
=====
Warning: Enumerating all SCSI or FC LUNs increases initialization times.
```

```
Enumerate all SCSI LUNs (Y/N)? [current setting: Y]: Y
```

```
Enumerate all FC LUNs (Y/N)? [current setting: Y]: Y
Drv[2C] Ctrl[29] Lang[eng] - Options set. Action Required is None
```

3. Reconnect drivers, refresh mapped drives and display available boot options using commands `reconnect -r` and `map -r`.

```
Shell> reconnect -r
ReconnectController(0,0,0) : Status = Success
```

```
Shell> map -r
Device mapping table
blk0 : Acpi(PNP0A03,0)/Pci(1|0)/Fibre(WWN50001FE1501A1F39,Lun4068000000000000)
blk1 : Acpi(PNP0A03,0)/Pci(1|0)/Fibre(WWN50001FE1501A1F3D,Lun4068000000000000)
```

## Applications that depend on IO hardware paths or legacy storage DSFs are not supported with Integrity cross-technology logical servers

HP-UX applications that have a dependency on I/O hardware paths or Legacy Storage Device Special Files (DSFs) might fail if these applications are running on Integrity cross-technology logical servers. This occurs because I/O hardware paths and Legacy Storage DSFs may change during cross-technology move operations. After a cross-technology logical server move operation, Networking I/O Instance Numbers and Storage Agile DSFs will not cause application failure.

## Major issues

A new Integrity VM logical server may be assigned MAC addresses and/or server WWNs belonging to previously unmanaged Integrity VM logical servers, resulting in duplicate MAC addresses and/or server WWNs

If you unmanage an Integrity VM logical server, then create a new Integrity VM logical server, MAC addresses and/or WWNs from the unmanaged logical server may be assigned to the new logical server resulting in duplicates. If this new logical server is then activated on the same Integrity VM host as the host containing the previously unmanaged Integrity VM logical server, an activation error occurs due to a conflict with a duplicate MAC address and/or server WWN assigned to the unmanaged Integrity VM logical server. Additionally, the previously unmanaged Integrity VM logical server cannot be imported again until the newly created logical server that introduced the duplicate MAC addresses and/or server WWNs is corrected or deleted.

### Suggested actions

- To prevent duplicate MACs and/or WWNs:
  - If possible, use the Delete action rather than the Unmanage action when removing an Integrity VM logical server.
  - If possible, re-import any previously unmanaged Integrity VM logical servers.
  - If you must unmanage an Integrity VM logical server, you may reserve its WWNs and MAC addresses to prevent them from being reassigned by using this procedure:
    1. Determine the WWNs of the unmanaged Integrity VM logical server by logging onto the VM host and using the `hpvmstatus` command:

```
hpvmstatus -P Name_of_Unmanaged_VM | grep npiv
```

Sample output:

```
hba avio_stor 0 1 npiv /dev/fclp0-0x5001438002A30041,0x5001438002A30044
```

Make a note of the server WWNs, converting them to lsmutil format. In this example, `0x5001438002A30041` and `0x5001438002A30044` become `50:01:43:80:02:A3:00:41` and `50:01:43:80:02:A3:00:44`.
    2. If using Storage Pool Entries (SPEs), delete the SPE containing the WWNs of the unmanaged Integrity VM logical server.

3. Reserve the WWNs that belonged to the unmanaged Integrity VM logical server by using the `lsmutil -reserve -WWN WWN` command (where **WWN** is the worldwide name that you made note of in Step 1. Issue this command once for each of the WWNs associated with the unmanaged Integrity VM logical server.
4. Create a placeholder logical server to prevent the MAC addresses from being assigned to new logical servers:

1. Create a new logical server template without specifying any storage entries in the storage configuration.
2. In the network selection page, click **Add Network Entries**.

---

**NOTE:** LSM shows the MAC addresses that were part of the unmanaged logical server.

---

3. Click **Next** and then **Finish** to create the logical server template.

This Integrity VM logical server template cannot be activated because it lacks storage entries, but the new template prevents the MAC addresses of the unmanaged logical server from being used again.

You may now create and activate Integrity VM logical servers.

- To correct duplicate MAC and/or WWNs:
  1. Delete the newly created Integrity VM logical server (after backing up its data, if necessary).
  2. Perform one of the procedures described above under *To prevent duplicate MACs and/or WWNs*.

You may now create and activate Integrity VM logical servers.

## HP Logical Server Automation service may run out of memory on large scale systems

On large scale systems, with large numbers of managed nodes and/or large numbers (1000+) of logical servers on a 64-bit CMS, the HP Logical Server Automation service may run out of memory and display out of memory errors in the `/logs/hp_lsa_service.log` or `/logs/vselog.lsm.log` file. You may see the following errors:

```
java.lang.OutOfMemoryError: GC overhead limit exceeded
java.lang.OutOfMemoryError: Java heap space
java.lang.OutOfMemoryError: PermGen space
```

For more information, see “Configuring HP Insight Management for large environments” in the *HP Insight Management Installation and Configuration Guide* at [Matrix OE Information Library](#).

### Suggested action 1

To correct the GC overhead limit and/or Java heap space errors, modify the configuration file to increase memory allocation:

1. Modify the configuration file (shown in the default location) at `C:\Program Files\HP\Virtual Server Environment\conf\hp_lsa_service.conf`.
2. Increase the value of `wrapper.java.maxmemory=1024`. HP suggests a value of 2048.
3. Restart the HP Logical Server Automation service.

### Suggested action 2

To correct the PermGen space error, modify the configuration file to increase permanent generation size:

1. Modify the configuration file (shown in the default location) at C:\Program Files\HP\Virtual Server Environment\conf\hp\_lsa\_service.conf.
2. Add the following line. (Change additional.2 to additional.3 if a line with additional.2 already exists; change additional.2 to additional.4 if lines with additional.2 and additional.3 already exist, and so on.)  

```
wrapper.java.additional.2=-XX:PermSize=32m -XX:MaxPermSize=128m
```
3. Restart the HP Logical Server Automation service.

## Minor issues

### Performing operations on logical servers outside of Matrix OE visualization may cause incorrect information to be displayed

If you use tools other than Matrix OE visualization to manage logical servers, modifications to the underlying resources managed by logical server management are not immediately reflected in the logical server.

For example, if you modify the power state of a managed resource outside of logical server management, the power state of the logical server is displayed incorrectly until logical server management resets its data.

#### Suggested action

Wait several minutes for the logical server to reset its data, or use **Tools**→**Logical Servers**→**Refresh...** to manually refresh server resources.

### Activating an Integrity VM virtual machine logical server with three or more vCPUs may fail with entitlement error in VM host

If you attempt to activate an Integrity VM logical server with three or more vCPUs on an Integrity VM host with 1596 MHz processor speed or more, and the entitlement value used to create the virtual machine on the VM host falls below the minimum 5% value, the activation may fail with the error "Value x.x for entitlement\_min is below minimum, setting to 5.0." This error is displayed on the VM host.

#### Suggested action

On the Integrity VM host, set the Integrity VM configuration variable GUESTMHZMIN to the nearest integer value that is five percent of the host processor speed in cycles. The following commands can be used to perform the GUESTMHZMIN variable setting on the Integrity VM host:

```
# typeset -i i j k
# i=$(hpvmstatus -sM | cut -d: -f3)
# j=i/20
# k=j*20
# if ((i>k)) ; then j=j+1 ; fi
# ch_rc -a -p GUESTMHZMIN=$j /etc/rc.config.d/hpvmconf
```

### LUNs not detected after activating a logical server using logical server management on an HP Integrity Virtual Machine host

A logical server, when activated directly on an HP Integrity Virtual Machine host using LSM, fails to detect the LUNs assigned to it at the virtual machine EFI prompt. Due to this limitation, the logical server will not boot up the operating system automatically even if it is assigned with a LUN that is pre-provisioned with an operating system.

#### Suggested action

At the Integrity Virtual Machine EFI prompt, execute the following commands to see LUN mappings. When these steps are complete, the user can select the required LUN from the displayed device list to boot up or to install the operating system.

1. 1. Check configurable controllers using 'drvcfg' command.

```
Shell> drvcfg
Configurable Components
Drv[2C] Ctrl[29] Lang[eng]
```

2. Enumerate all SCSI and FC LUNs using `drvcfg -s <drvnum> <ctrlnum>`. `drvnum` and `ctrlnum` are displayed as part of step 1. Press Y for confirmation on enumerating both SCSI and FC LUNs.

```
Shell> drvcfg -s 2C 29
Set Configuration Options
Drv[2C] Ctrl[29] Lang[eng]
HP AVIO Stor Driver Configuration
=====
Warning: Enumerating all SCSI or FC LUNs increases initialization times.
Enumerate all SCSI LUNs (Y/N)? [current setting: Y]: Y
Enumerate all FC LUNs (Y/N)? [current setting: Y]: Y
Drv[2C] Ctrl[29] Lang[eng] - Options set. Action Required is None
```

3. Reconnect drivers, refresh mapped drives and display available boot options using commands "reconnect -r" and "map -r"

```
Shell> map -r
Device mapping table
blk0 : Acpi (PNP0A03,0)/Pci(1|0)/Fibre (WWN50001FE1501A1F39, Lun4068000000000000)
blk1 : Acpi (PNP0A03,0)/Pci(1|0)/Fibre (WWN50001FE1501A1F3D, Lun4068000000000000)
```

## Moving a managed virtual machine guest using Microsoft SCVMM may change BIOS serial number of logical server

On rare occasions, the BIOS serial number of a VM guest managed by logical server management is changed if a file copy move operation is performed on the VM guest using Microsoft System Center Virtual Machine Manager (SCVMM). This may cause operations on the virtual machine logical server to be unsuccessful, or the virtual machine logical server may become inoperable.

### Suggested action

Unmanage the logical server using the Matrix OE visualization **Delete**→**Unmanage Logical Server...** menu selection, then import the virtual machine as a logical server using **Tools**→**Logical Servers**→**Import...**

## Viewing or editing the logical server management properties file may cause failure of the HP Logical Server Automation service

Viewing or editing the `C:\Program Files\HP\Virtual Server Environment\conf\lsa\lsa.properties` file may write lock the file while the HP Logical Server Automation service is running. This may cause a serious failure of the service.

### Suggested action

Use a text editor that does not lock the file, such as Emacs, Notepad, or TextPad.

## Process Resource Manager

### Limitations

#### Creating and modifying configuration files with privilege elevation enabled is not supported

If privilege elevation is enabled in Systems Insight Manager, you can monitor Process Resource Manager (PRM), but you may experience problems performing PRM operations. Specifically, creating and modifying PRM configuration files may report success even though the operations hang.

### Suggested action

Do not change PRM configuration files using the PRM user interface when privilege elevation is enabled.

# Storage Provisioning Manager

Storage Provisioning Manager runs on an x86/x64 Windows CMS.

## Limitations

### RAID level support restrictions for EMC VMAX arrays

Currently, only RAID 1 and RAID 5 volumes are supported on the EMC Symmetrix VMAX array concrete pool. However, a thin volume is supported for any given RAID level, provided EMC storage admin creates corresponding storage pool setting instances. For more information, contact your EMC Storage Administrator.

### Required version of Brocade Network Advisor and HP B-Series Network Advisor

Brocade Network Advisor 11.1.3 or 11.1.4 or HP B-Series Network Advisor 11.1.3 or 11.1.14 must be installed in order to manage Brocade-based SAN environments.

HP B-Series Network Advisor is available from [http://h18006.www1.hp.com/products/storageworks/dc\\_fabricmgr/index.html](http://h18006.www1.hp.com/products/storageworks/dc_fabricmgr/index.html). See the *HP Storage Provisioning Manager (SPM) User Guide* at [www.hp.com/go/matrixoe/docs](http://www.hp.com/go/matrixoe/docs) for additional details.

### Firmware for 3PAR StoreServ

3PAR StoreServ must have 3PAR OS 3.1.1 MU1 or higher installed in order to work properly with SPM. SPM will not allow 3PAR arrays with earlier versions to be imported into the SPM catalog.

### 3PAR StoreServ Storage enforces a 64 initiator WWN limit per controller port on the array

This limitation can impact Matrix OE-created servers that use SPM to provision its storage. For solutions to limit this restriction, see the *HP 3PAR StoreServ Storage best practices guide* at <http://h20195.www2.hp.com/V2/GetDocument.aspx?docname=4AA4-4524ENW&cc=us&lc=en>.

### HP 3PAR StoreServe Storage Zero Detect

SPM does not recognize the Zero Detect property of volumes on 3PAR StoreServ Storage. SPM 2.2 (included with Matrix 7.2) cannot create volumes with the Zero Detect property. SPM can import and use pre-provisioned volumes with the Zero Detect property, but it does not recognize the property value. SPM will continue to use volumes it has created after an administrator applies HP 3PAR Zero Detect in the HP 3PAR Management Console, but does not recognize if a volume contains a Zero Detect property value.

### 3PAR StoreServ Storage Peer Motion has limited support in SPM

If Peer Motion is used on an HP 3PAR StoreServ Storage System that is imported into an SPM catalog, a number of steps must be performed in order to ensure that SPM and Matrix OE can properly use the storage affected by Peer Motion. For details on these steps, see the SPM 3PAR StoreServ Storage Peer Motion Technical White Paper at <http://www.hp.com/go/matrixoe/docs>.

### 3PAR StoreServ Storage Peer Persistence is not supported in SPM

SPM 2.2 (provided in Matrix OE 7.2) does not support peer persistence. SPM 2.2 cannot import volumes that are part of a remote copy group, and cannot import two volumes that share the same worldwide name. For this reason, SPM cannot be used to manage volumes in a peer persistence configuration as noted in 7.2 documentation.

# Managed node installation

## Limitations

### Multipath I/O and Emulex FC HBA or Emulex FlexFabric Converged Network Adapter is not supported on Windows Server 2008 and Windows Server 2008 R2

Using multipath I/O during OS installation is not supported with an Emulex FC HBA or an Emulex FlexFabric Converged Network Adapter on any Matrix-supported server on Windows Server 2008 and Windows Server 2008 R2.

The following suggested actions enable you to install the OS single-path and then modify the storage definitions to create multiple-path access to the boot volume. After a successful OS installation, these instructions will define an additional physical path to the boot volume by:

- Modifying the storage definition to add an additional server HBA definition (port)
- Modifying the boot volume definition to add a second physical 'path' to its definition.

#### **Suggested action 1**

If you are defining a storage pool entry in Matrix OE logical server management:

1. Define the storage pool entry, adding multiple ports, and marking all non-boot disks as redundant. Keep the boot volume as a single path.
2. Assign the storage pool entry to a logical server.
3. Activate the logical server.
4. Provision the OS, and install the multipath IO drivers as appropriate.
5. Deactivate the logical server.
6. Modify the storage pool entry to make the boot volume redundant.
7. Reactivate the logical server.

#### **Suggested action 2**

If you are creating a template using Matrix OE infrastructure orchestration designer select **Redundant SAN paths to disk** when configuring the boot volume.

Beginning in Matrix OE 7.2, infrastructure orchestration executes steps to work around this issue by ensuring that the OS is single-path while the OS is being provisioned, and then modifies the storage definitions to create multipath access to the boot volume when OS provisioning is complete.

# Matrix KVM Private Cloud

## Limitations

### Matrix KVM Private Cloud is available in US English only

The current version of Matrix KVM Private Cloud is available in US English only.

#### **Suggestion action**

Specify only US English characters in the Matrix KVM Admin Console text fields.

### Manually changing the time or time zone in the KVM Admin Console appliance requires restarting services on the compute nodes

#### **Suggested action**

Changing the time or time zone in the appliance make all activated hosts unresponsive to provisioning and cloud instance life cycle requests. When this occurs, all activated hosts show a red "critical" icon on the Hosts screen. If the time or time zone was changed on the Settings screen in the Matrix KVM Admin Console appliance after hosts have been activated, perform the following steps to synchronize the time on the hosts with the appliance time.

1. Verify that each host is configured to have the appliance as the NTP server.
  - The `/etc/ntp.conf` has the appliance configured as an NTP server:  

```
server isc.mgmt.local #ISC
```
  - The `/etc/hosts` file has a private LAN IP address set for the appliance hostname `isc.mgmt.local`.  
For example: `192.0.2.0 isc.mgmt.local`
2. On each host that was previously activated and is now showing critical status, enter the following commands on the command line:  

```
service ntpd stop  
ntpdate isc.mgmt.local  
service ntpd start
```

If possible, protect against this circumstance by first synchronizing the appliance clock with an NTP server. See the *HP Matrix KVM Admin Console online help* for instructions on synchronizing the appliance with an NTP server.

### Images added through browser upload are dependent on browser size limitations

Uploading an image through your browser on the Add Image screen is limited to the following:

- Using the Add button: One image that is less than 2 GB using Internet Explorer or 4 GB using Mozilla Firefox
- Using the Add+ button: Five images less than 2 GB total size using Internet Explorer or 4 GB total size using Mozilla Firefox

#### Suggested action

1. Add an image by specifying a URL from a file server that is accessible from the host management subnet.
2. Verify the image by checking its size on the Images screen. If the size is 0.0 GB, the image is invalid or is not accessible. In this case, perform steps 2 and 3 below.

If you uploaded an image or images that exceed the size limitation and the screen displays "Loading..." for several minutes:

1. Close the browser window to cancel the upload operation.
2. In a new browser window, go to the Images screen and find the new, partially uploaded image by name. The status of the image will be Active.
3. From the Images screen, click the row of the image to be deleted.
4. Select **Actions**→**Delete**.

### Technology preview: Customizing instances using the OpenStack metadata service

Matrix 7.2 KVM Private Cloud includes a technology preview of instance of customization using the OpenStack metadata service. The technology preview allows early access use in non-production configurations. Future releases will contain full support.

Customizing an instance includes adding a hostname, instance type, public IP address, SSH keys, and custom user data to an instance. User-written scripts that add this information to an instance are embedded in an operating system image before an instance is deployed. These scripts are not standard in any OS images.

With this technology preview, you can write customized scripts that call the OpenStack metadata service. This service is available to users of any OS image. For Ubuntu images, you can optionally use cloud-init to handle early initialization of a cloud instance. Cloud-init is a client of the OpenStack metadata service. A cloud-init RPM package is available for RedHat and Fedora from the FedoraProject.org EPEL repository at <http://fedoraproject.org/wiki/EPEL>.

For more information about the OpenStack metadata service, refer to <http://docs.openstack.org/essex/openstack-compute/admin/content/metadata-service.html> and <http://docs.openstack.org/trunk/openstack-compute/admin/content/user-data.html>. For more information about cloud-init, refer to <https://launchpad.net/cloud-init> and <https://help.ubuntu.com/community/CloudInit>.

## Do not use the 10.x.x.x IP address range for both the fixed and floating networks

IP addresses in the 10.x.x.x range can be specified for the floating networks (host management subnet and cloud connectivity subnet) or for the fixed network (private subnet for instances), but not both.

### Suggested action

- If you specify IP addresses in the 10.x.x.x range for the host management subnet and cloud connectivity subnet, specify an IP address in the 172.x.x.x or 192.x.x.x range for the private subnet for instances for each VM host (compute node).
- If you specified 10.x.x.x IP address ranges for all subnets within the appliance and you have not activated any VM hosts, edit cloud networking on the appliance (from the Settings screen, select Edit Cloud Networking) and change the settings as described above.
- If you specified 10.x.x.x IP address ranges for all subnets within the appliance and you have activated VM hosts, open an SSH connection to each activated VM host and do the following:
  1. Append the following line to the end of the `/etc/nova/nova.conf` file:  
`fixed_range = <CIDR for your cloud connectivity subnet`  
For example: `10.1.50.0/24`
  2. Restart the nova-network service by entering the following command on the command line:  
`service openstack-nova-network restart`

## Minor Issues

### After an appliance reboot, communication is not restored with compute nodes

When the Matrix KVM Private Cloud appliance comes online after it reboots, it cannot communicate with VM hosts (compute nodes), and the VM provisioning stops.

### Suggested action

After an appliance reboot, connect to each VM host using ssh and execute the following commands to restore communication with the host:

```
service openstack-nova-compute restart
service openstack-nova-network restart
```

### Internet Explorer 8 displays “Out of Memory” popup if the Matrix KVM Admin Console Activity page is left open for an extended period of time

If your browser is Internet Explorer 8, you may see an “Out of Memory” popup if the Matrix KVM Admin Console Activity page is kept open for an extended period of time. The amount of time depends on the appliance memory allocation and load on the server, but in general this error can occur after approximately six days with 25 notifications displayed, or after 30 hours with 125 notifications displayed. If this error occurs, the web browser may become sluggish or unresponsive to user input.

### Suggested actions

1. Reload the page by pressing **CTRL-F5**.
2. Restart the browser.
3. To prevent the problem from reoccurring, upgrade to a newer browser version.

## Deactivating multiple hosts simultaneously may not free all floating IP addresses

If you deactivate more than one KVM host at a time, the system may not release the floating IP address associated with the host to the floating IP address pool. Therefore, those IP addresses will not be available for reuse.

### Suggested action

1. During initial cloud networking configuration on the Settings screen, allocate a large floating IP address pool in the Cloud connectivity subnet to allow for possible IP address leakage. HP recommends that you specify at least a /21 subnet (255.255.248.0) to ensure that there are sufficient IP addresses to allocate to virtual machines.
2. On the Hosts screen, deactivate only one host at a time.

## Deactivating a host may result in excessively large log files

When deactivating a VM host, the nova-network or nova-compute service may write large amounts of data to the log files in the `/var/log/nova` directory. If this occurs, the log files might exhaust the disk volume space. If the log files are written to a physical volume other than the boot disk, the deactivation will proceed normally. If the log files are written to the boot disk, this could cause a system crash.

### Suggested action

1. Before activation, configure the VM host so that log files are written to a data disk.
2. After deactivation, check the size of the log files in the `/var/log/nova` directory and delete the files if they are consuming too much space on the disk.

## Instance state remains as “Building” when there are no more floating IP addresses available

When a VM instance is booted and there are no more floating IP addresses available to attach to a VM instance, the state of the instance indefinitely remains as “Building” without reporting an error.

### Suggested action

To delete the instance from the VM host, perform the following steps:

1. Log on to the VM host console that was selected to host the VM instance.
2. Take the host offline by entering the following command:  

```
service openstack-nova-compute stop
```
3. On the Matrix KVM Admin Console Hosts screen, wait for the status of the VM host to change to offline (red status).
4. On the Instances screen, select the instance in the “Building” state and delete it. A warning displays about deleting the instance in this situation. Proceed with the deletion.
5. On the VM host console, enter the following command to bring the host back online:  

```
service openstack-nova-computer start
```
6. Free IP addresses by deleting instances that are no longer needed.

---

**NOTE:** IP addresses cannot be added after cloud networking is configured without resetting the cloud environment.

---

## VM instances cannot be customized if the host management subnet is 10.x.x.x

If you configure the Host management subnet to be in the 10.x.x.x range, activating a VM host causes the OpenStack metadata service to stop working. As a consequence, VM instances cannot be customized, which means that the hostname cannot be changed and the SSH public key cannot be injected into the created VM instance.

### Suggested action

1. During initial networking configuration on the Settings screen, specify a Host management subnet that is not in the 10.x.x.x range. For example, specify 192.0.2.0.
2. If a VM host has been activated, open an SSH connection to the activated VM host and restart the metadata service by executing the following command:

```
service openstack-nova-metadata-api restart
```

### Deployed instance does not boot and the error No available host can provide the specified resources displays

The error `No available host can provide the specified resources` displays even if the Hosts screen appears to show sufficient available resources causing the virtual machine instance not to boot.

The Hosts screen displays the percentage of resources in use and the total amount of resources. However, the actual available resources of a host are calculated by subtracting allocated resources (the virtual machine instances already provisioned to a host) from the capacity of the host.

If one or more virtual machine instances in a host are not consuming all allocated resources or are powered down, the host appears to have a high percentage of free resources, but the available resources are in fact already allocated to the powered down instances.

### Suggested action

Note the number of hosted virtual machine instances when evaluating whether resources on a particular host are available. If sufficient resources are not available, perform one of the following tasks:

- Add compute resources to the host.
- Free space on the host by deleting existing instances that are no longer needed.

### External yum repository can cause host activation to fail

An external yum repository, used for RPM-based Linux updates, that is enabled on a VM host can cause the first attempt to activate the host to time out and display “<Host> activation failed. Error: Failed to configure user accounts on remote host.”

### Suggested action

Retry activating the host. After the first failed attempt, subsequent attempts to activate the host will proceed as expected. To avoid this circumstance, disable all external yum repositories by logging into the VM host to be activated and setting `enabled=0` for the file `/etc/yum.repos.d/rhel.repo`.

### Adding a Cloud User account does not require specifying a public key, but a public key is required for provisioning

If a Cloud User account that is used as the Matrix integration service account has no key or an invalid key associated with it, provisioning from Matrix OE will not succeed. The Add User dialog will create a Cloud User without requiring a key.

### Suggested action

Specify a valid public key for the Cloud User account when the account is added, or modify an existing account and specify a valid key. Because hidden characters introduced during a copy/paste operation can change the key, take care to prevent special characters from being injected into the key when pasting it into the public key field.

## Filtering the Instances screen for “Resize pending” requires manual cleanup

On the Matrix KVM Admin Console Instances screen, filtering for instances in the “Resize pending” state adds the characters `state:Resize%20pending` to the search field. When this occurs, subsequent filtering for instances in any state does not return results.

### Suggested action

Erase all characters in the search field, then enter `state:Resize pending` in the search field.

## Sorting does not work as expected

On the Instance Types screen, clicking the numeric column might result in an alphabetic order, not numeric ordering.

On the Users screen, you cannot sort users by role, but you can sort by name.

### Suggested actions

1. Sort instance types by clicking a different column head.
2. Sort users by role by entering the role in the Search field.

## Activity screen does not always report failures to create instances and instance types

If you use the Matrix KVM Admin Console to create an instance type and the operations does not succeed, or you create an instance and the operation does not succeed because the OS image is larger than the disk size in the selected instance type, the Activity screen does not always indicate the failure or the reason.

### Suggested action

If the instance type does not appear on the Instance Types screen as expected, create the instance type again and specify a different name.

If the instance does not appear on the Instances screen as expected:

1. Open an SSH session to the VM host selected to create the instance.
2. Open the `/var/log/nova/compute.log` file.
3. Search the log file and confirm that it contains the string “Image Too Large”.
4. Create the instance again, and specify an instance type with larger disk size.

## Incorrect message is displayed when filtering or searching returns no results

If you filter or search for resources on a Matrix KVM Admin Console screen, the message is a hyperlink to create resources or take corrective action. The link is unnecessary because there are no resources to display based on the filter options.

### Suggested action

Ignore the message. Clear the search text and search for a known string that produces results, for example “source”.

## Intermittent timeouts may occur during virtual machine deployment

Deploying a virtual machine to a KVM host using Matrix infrastructure orchestration may occasionally time out with the error `The request for HP Matrix KVM Private Cloud server <server-name> has failed. ... Read timed out connecting to POST.` When this occurs, all hosts on the Matrix KVM Admin Console Hosts screen appear with an error status.

### Suggested action

1. Retry the create service request in infrastructure orchestration.
2. If the retry does not succeed, open an SSH connection to the VM host (compute node) and run the following commands from the command line:

```
service openstack-nova-compute restart
```

```
service openstack-nova-network restart
```

3. Retry the create service request again.

## HP Matrix Operating Environment infrastructure orchestration

### Enable Matrix infrastructure orchestration provisioning support of Linux VMs

By default, Matrix infrastructure orchestration does not provision a Linux VM on a Microsoft Hyper-V hypervisor. Matrix infrastructure orchestration can be set up to support provisioning of a Linux VM on a Microsoft Windows 2008 or 2012 Hyper-V hypervisor, but it does not personalize the VM.

#### Suggested action

To deploy Linux VMs on a Microsoft Windows 2008 or 2012 Hyper-V hypervisor, perform the following steps:

1. For Matrix infrastructure orchestration to deploy a Linux VM on Hyper-V, change the `hpio.properties` file:  

```
skip.linux.on.hyperv.template.personalization = true
```
2. To personalize the VM (for example, changing the host name, configuring the NIC for DHCP or static IP addresses, configuring DNS, etc.), use an Operations Orchestration workflow or perform the task manually using the Linux configuration tools.

### Choose fabrics when auto-generating storage through Matrix infrastructure orchestration

When auto-generating storage from Matrix infrastructure orchestration, if there is more than a single pair of fabrics to choose from, Matrix infrastructure orchestration determines the pairs to use. To set the preferred list of fabrics to use during the auto-generation of storage, add the `SPM_AUTO_PREFERRED_FABRICS` property to the `lsa.properties` file.

#### Suggested action

Add the following property to the `lsa.properties` file:

```
SPM_AUTO_PREFERRED_FABRICS=FabricA,FabricB,FabricC:FabricD
```

Fabrics separated by ":" are redundant pairings.

If the storage is not redundant, Matrix tries to find storage using the fabrics in the list in the order that they appear. If the storage is redundant, then Matrix only tries the redundant fabric pairings in the order that they appear.

In the above example, for non-redundant storage, Matrix tries FabricA, FabricB, FabricC and FabricD in that order. For redundant storage, Matrix tries the FabricC & FabricD pairing.

---

**NOTE:** The fabrics listed in this property are used by all auto-generated storage from Matrix infrastructure orchestration on that CMS. It cannot be set up to only be used by some users or some organizations. If storage is not found after going through the fabrics defined in this list, then all known fabrics are tried.

---

## Minor issues

### Provisioning of HP Integrity VMs with NPIV data disk fails

HP Integrity VMs with NPIV disks can be created through Matrix infrastructure orchestration. But the disks must be added after service creation through an Add disk operation. Any attempt to add NPIV data disks before service creation will fail and display the error message on the IO Requests screen:

```
Workflow execution failed for operation type EditLS: could not delete:
```

```
[com.hp.vse.lsa.datamodelV72.LsaFcConnection#44]
```

### Suggested action

1. Create the service from the IO template.
2. Power off the virtual server where the NPIV data disk will be added by selecting **Power off servers** from the Actions or Server Actions list in IO.
3. Add the required NPIV data disk to the virtual server by selecting **Add disks** from the Actions or Server Actions list in IO.
4. Power on the virtual server by selecting **Power on servers** from the Actions or Server Actions list in IO.

## Deployment of HPUX 11iv3 1203 or HPUX 11iv3 1209 on any Integrity i2 or i4 server blade using Matrix infrastructure orchestration with Ignite-UX may fail

When deploying HP-UX 11iv3 1203 or HP-UX 11iv3 1209 on a BL860c i2, BL870c i2, BL860c i4, BL 870c i4, or BL890c i4 server blade using Matrix infrastructure orchestration with Ignite-UX, the Create Service Request may fail and the server will not be provisioned. The step in the request that deploys the HP-UX OS fails and the following error displays:

```
Major, "Failed while deploying operating system for <your service name>. Cause: Logical server job (ID = <your logical server ID>) completed with a failure status. Failure: Deployment server does not have a record of the target system(s).."
```

### Suggested action

1. Erase the disk used by the target server for deployment.
  - a. Make sure all hosts are in the same zone then perform a LUN clean-up: Log in to the storage management server (EVA, 3PAR StoreServ Storage, etc.)
  - b. Unpresent the host associated with the vdisks and delete the vdisks.
  - c. Create the new vdisks with the required size and RAID level and present the vdisks to the respective hosts.
2. Deactivate the logical server associated with the target server.
  - a. From the HP Matrix Operating Environment menu, click **Tools**→**HP Matrix OE Visualization**.
  - b. Select **Logical Server** from the Perspective drop-down menu.
  - c. Select the logical server and click **Tools**→**Logical Servers**→**Deactivate...**
  - d. Monitor the status of the task by selecting **Report**→**Logical Server Job Status**.
  - e. Wait for the task to complete successfully before continuing to the next step.
3. Delete the logical server. Select the appropriate deactivated logical server and click **Delete**→**Delete Logical Server...**
4. Move the target server from the IO maintenance pool to the desired server pool.
  - a. From the HP Matrix Operating Environment menu, click **Tools**→**Infrastructure Orchestration**.
  - b. Click the Servers tab. Make sure the servers are not being used by checking the "Usage" column.
  - c. Move the servers under the "Maintenance" pool to the desired server pool by using the "Modify Pool" option.
5. Recreate the service request.

## Create or delete service request for an Integrity server fails due to an EFI disk read error

A create or delete service request for an Integrity server with an FCoE device may fail due to an EFI disk read error. A message displays in the Matrix infrastructure orchestration Requests screen that the request for the logical server has failed.

To determine if this error is caused by an EFI disk read error, open a console window on the host and look for a message similar to the following:

```
Automatic boot entry execution will start in 1 second .
Booting VC Pri. Fibre(0x5001438004C8E7B8,0x1000000000000)
Boot Failed. VC Pri. Fibre(0x5001438004C8E7B8,0x1000000000000)
Booting HP-UX Primary Boot:
0/0/0/3/0/0/2.0x5001438004c8e7b9.0x4001000000000000
Error - efi_disk_read returned (Device Error)
```

### Suggested action

If the EFI disk read error displays in the console, clean the logical server and retry the service request by performing the following steps:

1. Erase the disk used by the target server for deployment.
  - a. Make sure all hosts are in the same zone then perform a LUN clean-up: Log in to the storage management server (EVA, 3PAR StoreServ Storage, etc.)
  - b. Unpresent the host associated with the vdisks and delete the vdisks.
  - c. Create the new vdisks with the required size and RAID level and present the vdisks to the respective hosts.
2. Deactivate the logical server associated with the target server.
  - a. From the HP Matrix Operating Environment menu, click **Tools**→**HP Matrix OE Visualization**.
  - b. Select **Logical Server** from the Perspective drop-down menu.
  - c. Select the logical server.
  - d. Click **Tools**→**Logical Servers**→**Deactivate...** and click the **Deactivate** button to continue deactivating the logical server.
  - e. Monitor the status of the task by selecting **Report**→**Logical Server Job Status**.
  - f. Wait for the task to complete successfully before continuing to the next step.
3. Delete the logical server.
  - a. Select the appropriate deactivated logical server.
  - b. Click **Delete**→**Delete Logical Server...** and click the **Delete** button to continue deleting the logical server.
4. Move the target server from the IO maintenance pool to the desired server pool.
  - a. From the HP Matrix Operating Environment menu, click **Tools**→**Infrastructure Orchestration**.
  - b. Click the Servers tab. Make sure the servers are not being used by checking the “Usage” column.
  - c. Move the servers under the “Maintenance” pool to the desired server pool by using the “Modify Pool” option.
5. Recreate the service request.

## Matrix infrastructure orchestration provisioning failure on Integrity blades BL870c i2, BL890c i2, BL870c i4, and BL890c i4

Matrix infrastructure orchestration provisioning fails on Integrity server blades BL870c i2, BL890c i2, BL870c i4, and BL890c i4. The `serverboot.physical.wait.seconds` in the `hpio.properties` file and the `INTEGRITY_POWERON_WAITTIME` in the `lsa.properties` file are not set to the right values according to the blade models. The wait time consists of waiting for the Integrity blade to boot completely before infrastructure orchestration or Matrix OE logical server management executes the power off. This setting is required to prevent the potential QLogic HBA card EFI/FW from being corrupted while the Virtual Connect profile is being written to or removed from the NVRAM.

---

**NOTE:** The `hpio.properties` resides under `IOinstalledDir/conf`.

The `lsa.properties` resides under `VSEinstalledDir/conf/lsa`.

---

### Suggested action

1. Measure the time required for the power cycle of the Integrity blade being provisioned.

**NOTE:** The time varies based on the blade model, its memory size, and the number of QLogic cards. You can measure the power cycle time from the Integrity blade remote console using Onboard Administrator (OA) or VCEM.

---

2. Set the `serverboot.physical.wait.seconds` and `INTEGRITY_POWERON_WAITTIME` according to the time measured in step 1.

Use the following parameters as a minimum.

- For BL870c i2 and i4 blades:
  - Set the `serverboot.physical.wait.seconds` to 540 seconds.
  - Set the `INTEGRITY_POWERON_WAITTIME` to 540000 milliseconds.
- For BL890c i2 and i4 blades
  - Set the `serverboot.physical.wait.seconds` to 1080 seconds.
  - Set the `INTEGRITY_POWERON_WAITTIME` to 1080000 milliseconds.



---

## 3 Support and other resources

### Contacting HP

#### Information to collect before contacting HP

Be sure to have the following information available before you contact HP:

- Software product name
- Hardware product model number
- Operating system type and version
- Applicable error message
- Third-party hardware or software
- Technical support registration number (if applicable)

#### How to contact HP

Use the following methods to contact HP technical support:

- See the Contact HP Worldwide website for contact options:  
<http://www.hp.com/go/assistance>
- Use the `Contact hp` link on the HP Support Center website:  
<http://www.hp.com/go/hpsc>
- In the United States, call 1-800-334-5144 to contact HP by telephone. This service is available 24 hours a day, 7 days a week. For continuous quality improvement, conversations might be recorded or monitored.

#### Registering for software technical support and update service

HP CloudSystem Matrix includes as standard, three or one year of 24 x 7 HP Software Technical Support and Update Service and 24 x 7 four hour response HP Hardware Support Service. This service provides access to HP technical resources for assistance in resolving software implementation or operations problems.

The service also provides access to software updates and reference manuals in electronic form as they are made available from HP. Customers who purchase an electronic license are eligible for electronic updates.

With this service, Insight Management software customers benefit from expedited problem resolution as well as proactive notification and delivery of software updates. For more information about this service, see the following website:

<http://www.hp.com/services/insight>

Registration for this service takes place following online redemption of the license certificate.

#### How to use your software technical support and update service

As HP releases updates to software, the latest versions of the software and documentation are made available to you. The Software Updates and Licensing portal gives you access to software, documentation and license updates for products on your HP software support agreement.

You can access this portal from the HP Support Center:

<http://www.hp.com/go/hpsc>

After creating your profile and linking your support agreements to your profile, see the Software Updates and Licensing portal at <http://www.hp.com/go/hpsoftwareupdatesupport> to obtain software, documentation, and license updates.

## Warranty information

HP will replace defective delivery media for a period of 90 days from the date of purchase. This warranty applies to all Insight Management software products.

## HP authorized resellers

For the name of the nearest HP authorized reseller, see the following sources:

- In the United States, see the HP U.S. service locator web site:  
[http://www.hp.com/service\\_locator](http://www.hp.com/service_locator)
- In other locations, see the Contact HP worldwide web site:  
<http://welcome.hp.com/country/us/en/wwcontact.html>

## Security bulletin and alert policy for non-HP owned software components

Open source software (such as OpenSSL) or third-party software (such as Java) are sometimes included in HP products. HP discloses that the non-HP owned software components listed in the Matrix Operating Environment end user license agreement (EULA) are included with Matrix OE. To view the EULA, use a text editor to open the `/opt/vse/src/README` file on an HP-UX CMS, or the `<installation-directory>\src\README` file on a Windows CMS. (The default installation directory on a Windows CMS is `C:\Program Files\HP\Virtual Server Environment`, but this directory can be changed during installation.)

HP addresses security bulletins for the software components listed in the EULA with the same level of support afforded HP products.

HP is committed to reducing security defects and helping you mitigate the risks associated with security defects when they do occur. HP has a well defined process when a security defect is found that culminates with the publication of a security bulletin. The security bulletin provides you with a high level description of the problem and explains how to mitigate the security defect.

## Subscribing to security bulletins

To receive security information (bulletins and alerts) from HP:

1. Open a browser to the HP home page:  
<http://www.hp.com>
2. Click the **Support & Drivers** tab.
3. Click **Sign up: driver, support, & security alerts**, which appears under **Additional Resources** in the right navigation pane.
4. Select **Business & IT Professionals** to open the Subscriber's Choice web page.
5. Do one of the following:
  - Sign in if you are a registered customer.
  - Enter your e-mail address to sign-up now. Then, select the box next to **Driver and Support alerts** and click **Continue**.

## Related information

The latest versions of manuals and white papers for Matrix Operating Environment and related products can be downloaded from the HP Web. Documents for Matrix OE software can be found at [Matrix OE Information Library](#).

For more information about Matrix OE and related products and solutions, visit the following HP websites:

- HP Insight Management: <http://www.hp.com/go/insightmanagement>
- HP Matrix Operating Environment: <http://www.hp.com/go/matrixoe>



---

## 4 Documentation feedback

HP is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback ([docsfeedback@hp.com](mailto:docsfeedback@hp.com)). Include the document title and part number, version number, or the URL when submitting your feedback.