

# HP CloudSystem Matrix 7.3 Release Notes

## Abstract

This document describes the HP CloudSystem Matrix 7.3 release content and provides information about limitations and major and minor issues with suggested actions.

HP Part Number: 754815-001  
Published: December 2013  
Edition: 1



## Notices

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.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

## Acknowledgments

Adobe ® is a trademark of Adobe Systems Incorporated. Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

## Revision history

Published	Changes	Document version
December 2013	<b>HP CloudSystem Matrix 7.3 release</b> Adds the following: <ul style="list-style-type: none"><li>• Microsoft Windows Server 2012 R2 Standard</li><li>• Microsoft Windows Server 2012 R2 Data Center</li><li>• OpenVMS 8.4</li><li>• Microsoft Internet Explorer 11</li><li>• Mozilla Firefox 24 ESR</li><li>• Insight Control software 7.3 enhancements</li><li>• Matrix OE support of high availability (HA) Hyper-V VMs with RDM or pass-through disks (with constraints)</li><li>• Matrix OE API enhancements</li></ul>	754815-001

---

# Contents

1	Description.....	4
	Update recommendation.....	4
	Supersedes.....	4
	Operating systems and devices.....	4
	Enhancements.....	4
	Fixes.....	5
	Prerequisites.....	5
2	Installation and configuration.....	6
	.....	6
3	Issues and suggested actions.....	8
	HP Insight managed system setup wizard.....	8
	HP Insight Control software.....	8
	HP Matrix Operating Environment software.....	14
	Managed node installation.....	41
	Other issues and suggested actions.....	41
4	Support and other resources.....	43
	Contacting HP.....	43
	Related information.....	44
	Typographic conventions.....	44
5	Documentation feedback.....	46

---

# 1 Description

HP CloudSystem Matrix is designed to simplify the deployment of infrastructure, applications, and cloud services by delivering IT capacity through pools of readily deployed resources. The goal of the Matrix Operating Environment (Matrix OE) is to accelerate provisioning, optimize IT capacity across physical and virtual environments, and to ensure predictable delivery and service levels.

CloudSystem Matrix is built on the modular, standards-based HP BladeSystem and its proven innovations, including Virtual Connect (VC) and the Matrix OE. With its open, extensible approach, CloudSystem Matrix runs any application workload and provisions the entire infrastructure, including servers, storage, and networking, within minutes. CloudSystem Matrix is optimized for HP ProLiant and HP Integrity server blades, HP storage, and HP networking but also supports third-party x86 servers, networking, and storage. CloudSystem Matrix is deeply integrated with virtualization technologies such as VMware and Microsoft, through the integration of virtual machine (VM) template libraries.

## Matrix versioning

HP CloudSystem Matrix versions are indicated by major.minor.update.revision. For example, version 7.2.0.1 is Matrix 7.2 Revision 1. Revisions are only for firmware, driver, and/or agent changes to the CloudSystem Matrix recipe.

## Update recommendation

- Customers should update to CloudSystem Matrix 7.3 if they want the enhancements provided by 7.3.
- All upgrades to HP CloudSystem Matrix 7.3 or later from versions prior to CloudSystem Matrix 7.3 must be performed by HP Technology Services.

## Supersedes

The 7.3 version supersedes the 7.2, 7.2.1, or 7.2.2 released versions.

## Operating systems and devices

Supported operating systems and devices are listed in the *HP CloudSystem Matrix 7.3 Compatibility Chart v7.3.0.0* available at:

[www.hp.com/go/matrixcompatibility](http://www.hp.com/go/matrixcompatibility)

## Enhancements

The HP CloudSystem Matrix 7.3 solution contains a requirement for a specific set of firmware versions for most of the important hardware components in the solution. The *HP CloudSystem Matrix 7.3 Compatibility Chart v7.3.0.0* specifies the supported hardware components and defines the required firmware versions.

## Hardware and firmware enhancements

The HP CloudSystem Matrix 7.3 firmware set contains the following list of notable hardware and firmware additions:

- HP MSA 2040 iSCSI
- HP MSA 2040 SAS
- WS460c Gen8

## Software enhancements

Some of the software enhancements to Insight Management 7.3 include:

### New operating systems, databases, and hypervisors

Adds the following:

- Microsoft Windows Server 2012 R2 Standard
- Microsoft Windows Server 2012 R2 Data Center

---

**NOTE:** Insight Control server provisioning does not provide OS Build Plans for Windows Server 2012 R2. See the *Insight Control server provisioning How to Create an OS Build Plan for Installing Windows 2012 R2* white paper for more information.

---

- Microsoft Hyper-V Server 2012 R2
- OpenVMS 8.4

### Other software enhancements

Adds the following:

- System Center 2012 R2 Virtual Machine Manager (SCVMM)
- Microsoft Internet Explorer 11
- Mozilla Firefox 24 ESR
- Enhancements with HP Storage Provisioning Manager version 2.3
- Enhancements to Matrix Operating Environment (OE) provisioning capabilities
- Matrix OE support of high availability (HA) Hyper-V VMs with RDM or pass-through disks (for various constraints, see *HP Matrix Operating Environment 7.3 Release Notes*)
- Matrix OE API enhancements (see *HP Matrix Operating Environment 7.3 Release Notes*)

## Fixes

The objectives of this release are to support the latest HP hardware platforms, support the latest revisions of already supported operating systems and hypervisors, add new functionality, and fix software issues.

Software updates improve performance in large scale environments.

## Prerequisites

- Hardware, firmware, and software requirements for this release are listed in the *HP CloudSystem Matrix 7.3 Compatibility Chart v7.3.0.0*.
- Installation and upgrade prerequisites are documented here.

CloudSystem Matrix documentation is available at:

[www.hp.com/go/matrixcompatibility](http://www.hp.com/go/matrixcompatibility)

---

## 2 Installation and configuration

### Install Microsoft hotfix for Windows 2008 R2 or Windows 2008 R2 SP1

Microsoft KB2577795 is required to fix a socket leak which causes network communication issues. The hotfix does not come with Windows updates and must be installed manually on the CMS. Install the hotfix at <http://support.microsoft.com/default.aspx?scid=kb;EN-US;2577795>.

### Patches required for Microsoft Windows Server 2008 with Hyper-V systems

Install required patches on the Hyper-V hosts, and then install or update Hyper-V Integration Services from the updated host to each VM guest as a precaution. Not all patches affect Integration Services on the guests.

- For Microsoft Windows Server 2008 SP2 with Hyper-V, to fix a hang issue for host servers running several simultaneous VM operations, install the patch at <http://support.microsoft.com/kb/980081>.
- If a Windows 2008 or 2008 R2 Hyper-V server will be running any version of Windows 2012, install the patch located at <http://support.microsoft.com/kb/2744129>.

For more information about other patches that may be required for later editions of Microsoft Windows Server 2008 R2, see the *Managing Microsoft Windows Server 2008 Hyper-V with HP Insight Software* technology brief at:

<http://www.hp.com/go/matrixoe/docs>

### Install and configure HP 3PAR StoreServ Storage

HP 3PAR StoreServ Storage must be installed, configured, and fully operational before initiating an HP CloudSystem Matrix Starter Kit Implementation Service. Use the HP 3PAR StoreServ Storage Software Installation and Startup Service, included with the storage device, to complete the required HP 3PAR StoreServ Storage installation and configuration.

### Manually update Integrity I/O firmware if an update is needed before OS installation

The HP SUM Integrity CloudSystem Matrix bundle version 7.2.2.0 supports online updates of I/O firmware on Integrity server blades only. If an I/O firmware update is required before operating system installation, the I/O firmware must be manually updated using EFI packages.

### Installing the HP Insight Control server provisioning RDP Migration Utility

The RDP Migration Utility within the Insight Control server provisioning download package enables data and target server migration from Insight Control server deployment to Insight Control server provisioning. The RDP Migration Utility is run from the RDP Deployment Server.

For installation instructions, see the *Data Migration from Insight Control server deployment to Insight Control server provisioning* white paper at <http://www.hp.com/go/insightcontrol/docs>.

### Installing WAIK 3.1 for HP Insight Control server deployment

Insight Control server deployment requires WAIK 3.1. To install this version:

1. Uninstall WAIK 1.1.
2. Install WAIK 3.0 from the Microsoft download center at [www.microsoft.com](http://www.microsoft.com).
3. Install the WAIK supplement, from [www.microsoft.com](http://www.microsoft.com), to update WAIK to 3.1.

For more information, see the *Windows Automated Installation Kit for Windows 7 Readme* at [http://technet.microsoft.com/en-us/library/dd349350\(v=WS.10\).aspx#SP1](http://technet.microsoft.com/en-us/library/dd349350(v=WS.10).aspx#SP1).

## Updating the Insight Control server provisioning appliance

---

- △ **CAUTION:** Do not attempt to update your appliance without first reviewing and then carefully following the update instructions in the *HP Insight Control Server Provisioning 7.2 Update 2 Installation Guide*. It is critical to follow all of the instructions as they are described. Some portions of the update process require special procedures to ensure a successful update.
-

---

## 3 Issues and suggested actions

Issues and limitations of this release are listed below, organized by HP Insight Management product or functional area. The following categories are used:

Limitations	Limitations of the implemented functions and features of this release
Major issues	Issues that may significantly affect functionality and usability in this release
Minor issues	Issues that may be noticeable but do not have a significant impact on functionality or usability

### HP Insight managed system setup wizard

#### Limitations

Privilege elevation configurations that require a password by the privilege elevation tool are not supported

The Insight managed system setup wizard only supports privilege elevation on a CMS that has privilege elevation configured so that a password is not required by the tool (such as `su` and `sudo`). If you attempt to configure features while privilege elevation is turned on and a password is required, you may see failures in the wizard's execution output.

#### Suggested action

If a feature cannot be configured 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 **A password is required for this privilege elevation tool** checkbox. Systems managed by the CMS must be reconfigured to reflect this configuration change.

### HP Insight Control software

For complete release notes for HP Insight Control software, see the *HP Insight Control 7.3 Release Notes*. Note that Insight Control server provisioning 7.2 Update 2 is the supported version for the CloudSystem 7.3 release. Release notes for Insight Control server provisioning 7.2 Update 2 can be found in both the *HP Insight Control 7.3 Release Notes* and the *HP Insight Control 7.2 Update 2 Release Notes*.

### HP Insight Control power management

#### Limitations

[DCPC works only on the primary address of a system](#)

Data Center Power Control (DCPC) works only on the primary address of a system. It cannot detect a system with a non-primary address.

#### Suggested action

Provide the primary IP address of target node while creating a rule in DCPC.

[Power delivery devices should not utilize a resolvable hostname](#)

If a power delivery device is created with a name that is a resolvable hostname, HP Systems Insight Manager (HP SIM) discovery may change the type to match that of the hostname. When this occurs, power management analysis views that attempt to display the power delivery device may report the following error message:

```
Circuit {Name} does not exist.
```



### Suggested action

Correct the system type of the device to be a power delivery device:

1. In the Systems Insight Manager Search box, enter the device name.
2. On the System page for the device, click the device name link.
3. Click the **Tools & Links** tab.
4. Click **Edit System Properties**.
5. Set the **System type:** to **Power Delivery Device**.
6. Be sure that the **Prevent the Discovery, Identification and Data Collection processes from changing these system properties** option is selected.
7. Click **OK**.

### Major issues

[Scheduled tasks for HP Power Regulator or Power Cap are not listed in the “Scheduled tasks” list or are not editable](#)

If you scheduled a Power Regulator or Power Cap task using Power Management 6.0 or earlier and then upgraded your system to Power Management 7.2.x, the task may not be listed in the Scheduled tasks list or may result in failure when you try to edit the task.

#### Suggested action

Delete the existing Power Regulator or Power Cap task and schedule a new task.

### Minor issues

[Racks and data center do not have power data after discovery](#)

Unmanaged nodes (racks and data center) do not have power data after discovery.

#### Suggested action

For the unmanaged node to receive power data, do one of the following:

- Calculate power history by selecting **Reports**→**Power Management**→**Calculate Power History**.  
or
- Click **Refresh Data** on the Display Power/Thermal data page to refresh the data for the unmanaged node.

### Web browser crashes while creating or editing a data center or rack

The web browser crashes while creating or editing a data center or rack. The crash information dialog in Mozilla Firefox (about:crashes) may contain the signature:

Caused by a plugin: Shockwave Flash (npswf32.dll)

#### **Suggested action**

Download and install the supported version of Adobe Flash player from: <http://get.adobe.com/flashplayer>.

### Unable to shut down the hosts and guests in either a forced or graceful manner

VMware ESXi/vSphere versions do not support SSH capability.

#### **Suggested action**

Shut down the guests or hosts manually.

### Inconsistent system details may appear when multiple open browser instances share the same user session

If you have multiple browser windows open to view different systems and these browser instances share the same browser session, inconsistent system detail data may be displayed inside some of the browser instances.

#### **Suggested action**

Refresh the browser window that displays inconsistent system details. Start and sign in to HP SIM using a separate browser session:

- For Microsoft Internet Explorer 8, click **File->New Session** to start the new browser session.
- For Mozilla Firefox, start multiple browser sessions using separate profiles. For example, `firefox.exe -no-remote -P profile_name`.

### Graphs and reports do not display when the time difference between the CMS and OA is more than 24 hours

Insight Control power management cannot display power, utilization, and temperature graphs if the time difference between the CMS and managed node is more than 24 hours, even though data collection status reports are successful. Power data is stored based on the time set on the managed node.

#### **Suggested action**

If necessary, update the date and time on the CMS and managed nodes to be consistent with the respective time zones.

## HP Insight Control server migration

### Limitations

Agent launch from the CMS to a source server with ", :, or \ characters in the password is not supported

Windows source agent deployment does not support Administrator passwords that contain the ", :, or \ characters.

### Major issues

Operating system fails to boot on a destination server after migration when logical volumes are not erased on destination disks

After a completed Insight Control server migration, the OS may not boot on the destination server if the server is repurposed and the following criteria are met:

1. The primary storage controller is reconfigured from RAID mode to non-RAID mode without erasing the existing logical volumes.
2. The server is designated as a destination server by Insight Control server migration.

### **Suggested action**

Do one of the following actions:

- Configure the controller in RAID mode, erase all logical volumes, reconfigure the controller into non-RAID mode, and then run the migration again.
- Configure the server to use the RAID mode of the controller and then run the migration again.

## HP Insight Control server deployment

### Limitations

#### Deployment of SLES 11 SP2/KVM might fail if kISO is used

The deployment of SUSE Linux Enterprise Server (SLES) 11 SP2/KVM may fail if kISO is used along with one of the following option cards:

- HP QMH2672 16Gb FC HBA for BladeSystem c-Class
- HP Ethernet 10Gb 2-port 561FLR-T Adapter
- HP Ethernet 10Gb 2-port 561T Adapter
- SN1000Q 16Gb 1P HBA's (With Single & Double port)

### **Suggested action**

To resolve this issue:

1. Edit Run Script - Configure Boot Environment to pass the kISO parameter.
2. Remove # (comment character) in line `#export kiso=no` in the script.

### Major issue

#### During PXE regeneration, an error occurs when booting to WinPE targets

During PXE regeneration, the `Exiting Emulex PXE ROM` error occurs when booting to WinPE targets if Altiris PXE Manager Services are not re-started in the proper order.

### **Suggested action**

PXE services must be restarted if there is a change in drivers or policy settings that requires a PXE regeneration. To restart in the proper order:

1. Restart Altiris PXE Manager Service.
2. Regenerate PXE Images using PXE Configuration Utility.
3. Restart Altiris PXE Config helper service.

## HP Insight Control server provisioning

For additional issues and suggested actions, see the Troubleshooting index in the *Insight Control server provisioning Online Help*.

### Initial setup

#### Console browser may display an empty Servers screen after first time network setup

After entering network information on the console during the initial appliance setup, click **OK**. The Applying network settings screen displays initially, but then the Servers screen displays a message stating that servers cannot be retrieved. This behavior is expected and can be ignored. Do not attempt to perform any operations on the appliance. Wait one to two minutes for the appliance to reboot automatically.

## Limitations

### Server provisioning appliance restart does not log users out

When a password is compromised, some administrators might use an appliance restart in an attempt to log out all users. However, on restart of the IC server provisioning appliance via the Settings page, users are not logged out, despite a message indicating that all logged-in users will be logged out.

#### Suggested action

The workaround is to delete the user whose password is compromised. This will result in the user being logged out without having to restart the server provisioning appliance. An administrator can then re-add the user and assign a new password.

### On Hyper-V, server provisioning appliance cannot be configured with 32GB of memory or more

When setting VM appliance resources on a Hyper-V host, you cannot set the memory of the appliance to 32GB or higher, as this will cause a kernel panic.

#### Suggested action

Configure your appliance with less than 32GB of memory. The optimal configuration for the appliance at scale is 30GB memory and 8 CPUs. The VMware appliance does not have this restriction.

## Major issues

### Linux SPP fails after step 2 for SPP 2013.09b

When installing the Service Pack for ProLiant (SPP) version 2013.09b on Linux operating systems using the ProLiant SW - Install Linux SPP Build Plan, the plan fails at the `hpqlgc` QLogic component package with exit code 255 and the `hpsum_log.txt` displays the following error:

```
No install was run, due to failed dependencies
```

The failure occurs at the dependency check and therefore, no driver or online firmware updates will be attempted for any component in the SPP.

For Matrix OE customers, perform the following before running the ProLiant SW - Install Linux SPP Build Plan:

#### Suggested action

From the IC server provisioning media server:

1. Remove the rpms that are failing on the dependency from the 2013.09b SPP directory (`Media\spp\2013.09b\hp\swpackages`). The following table lists the rpm files to remove:

Linux OS version	The hpqlgc files to remove from the SPP 2013.09b directory
Red Hat Linux 5.9	<code>kmod-hpqlgc-nx_nic-4.0.590.6.1_3.0.13_0.27-1.rhel5u9.x86_64.rpm</code>
Red Hat Linux 6.3	<code>kmod-hpqlgc-nx_nic-4.0.590.6.1_3.0.13_0.27-1.rhel6u4.x86_64.rpm</code>
Red Hat Linux 6.4	<code>kmod-hpqlgc-nx_nic-4.0.590.6.1_3.0.13_0.27-1.rhel6u3.x86_64.rpm</code>
SUSE Linux 11 SP2	<code>hpqlgc-nx_nic-kmp-default-sles11sp2.x86_64.rpm</code> <code>hpqlgc-nx_nic-kmp-xen-sles11sp2.x86_64.rpm</code>
SUSE Linux 11 SP3	<code>hpqlgc-nx_nic-kmp-default-sles11sp3.x86_64.rpm</code> <code>hpqlgc-nx_nic-kmp-xen-sles11sp3.x86_64.rpm</code>

2. To install the correct version of these RPM files, copy the files from the 2013.02 SPP directory at `\Media\spp\2013.02\hp\swpackages` and paste them to the 2013.09 SPP directory at `\Media\spp\2013.09b\hp\swpackages`.

For non-Matrix OE customers, see the *HP Insight Control 7.2 Update 2 Release Notes* for other suggested actions.

### BL460c Gen8 server fails to boot from SAN

When running the ProLiant HW – Enable Boot from SAN OS Build Plan on a BL460c Gen8 server, the local array controller is not disabled, and the SAN controller is not set as the first boot controller within the BIOS. Despite this, the OS Build Plan completes successfully and no error is displayed.

#### **Suggested action**

You must manually disable the local Smart Array controller and set the SAN controller as the first boot controller within the BIOS.

## Minor issues

### After updating the Insight Control server provisioning appliance, the Settings page may display a DHCP error

If the Insight Control server provisioning appliance is not providing DHCP service at the start of an update (such as when using an external DHCP server), the following message appears on the Settings page after the update is complete:

The DHCP settings do not conform. The network specification is missing.

#### **Suggested action**

Navigate to the Edit DHCP page, and select **None** for the **Service provided by appliance** setting.

### After updating the Insight Control server provisioning appliance, WinPE may indicate that it is not installed

After updating the Insight Control server provisioning appliance from 7.2 or 7.2.1, WinPE may indicate `Not installed` after the update completes, even when WinPE was installed before the update.

#### **Suggested action**

Create a new WinPE PXE image using the IC server provisioning WinPE PXE Image Creation utility from the 7.2.2 release and upload it to the appliance.

### Add server action may produce two server entries if deploying to blade servers in an enclosure with Virtual Connect

If using Insight Control server provisioning to deploy to blade servers in an enclosure with Virtual Connect, two server entries may be produced when using the **Add server** function. One is the normal server entry, and one looks like `ILOHOST_X.X.X.X`. This occurs because a Virtual Connect profile is not fully applied to the server until it is booted at least once, and the **Add server** operation queries the iLO before that first boot happens.

#### **Suggested action**

To correct the problem when it happens, delete the server entry with the name `ILOHOST_X.X.X.X` once the second server entry appears.

To avoid the problem, choose one of the two following actions:

- Add the server to the appliance by PXE booting the server, rather than using the **Add server** action. Servers discovered via PXE booting have their iLO added automatically after the server boots, so the profile is applied before discovery. See the *HP Insight Control Server Provisioning 7.2.2 Administrator Guide* chapter on Adding servers for more information.
- After applying the server profile, power on the server, wait for it to complete its power on self test, and select **Add server**. This way, the profile is applied before the iLO is queried.

An error may occur during a Red Hat Enterprise Linux installation if installing RHEL 6.3 or 6.4 on a server with more than one disk

When you attempt to install Red Hat 6.3 or 6.4 on a server with more than one disk on which the Erase Server Build Plan was previously run, the following error may appear on the target server's console during the installation process:

```
This device may need to be reinitialized. REINITIALIZING WILL CAUSE ALL DATA TO BE LOST!
```

This error occurs because the Erase Server Build Plan erased all data from the disks and the Red Hat installer cannot determine the proper formatting.

### **Suggested action**

Add `zerombr` to your `kickstart` file. See the Red Hat documentation for information about `zerombr` before using it, so you can familiarize yourself with what it does.

The PXE boot menu at the target server console incorrectly states that the `Linux6-64-ogfs` option is based on RHEL 6.3 for an appliance that was upgraded to 7.2.2

The `Linux6-64-ogfs` option is actually based on RHEL 6.4 x64.

### **OS Build Plan may fail if the Media Server user password contains special characters**

If the password used in the Media Server Settings contains any of the following special characters: '@' (an at sign), '#' (a pound sign), '/' (a forward slash), ':' (a colon), '\\' (two backward slashes), or '?' (a question mark); then, the OS Build Plan fails with the following error:

```
Issues and suggested actions 15 Failed To Set Media Source: NET use Z: \\123@xxx.xxx.xx.xx\deployment2/USER:xyz /persistent:no ***** failed with exit code 2.
```

### **Suggested action**

Do not use these special characters in the media server user password.

## **HP Matrix Operating Environment software**

For complete release notes for the HP Matrix Operating Environment software, see the *HP Matrix Operating Environment 7.3 Release Notes*.

## **Limitations**

### **Matrix OE support for Generation-2 VMs on Hyper-V 2012 R2**

Matrix OE does not support importing Generation-2 VMs on Hyper-V 2012 R2. Attempts to import Generation-2 VMs result in the following error message:

```
Unable to import the specified system with UUID VM_UUID.VM_NAME.
```

For Matrix infrastructure orchestration, Generation-2 VMs will not be available in the list of VMs that can be imported.

These VMs will also not be supported for the Import VM functionality. If you try to import, the following message will be displayed:

```
Import failed for VM <Virtual Machine UUID that is available on HPSIM>. Check that the VM Host is configured appropriately with HP Insight Control virtual machine management on the CMS, then refresh logical server virtual machine resources.
```

---

**NOTE:** The Matrix Operating Environment will create Generation-1 Hyper-V VMs.

---

## HP IO Accelerators limited support

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 white papers 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>.

## Managing Microsoft Windows servers without OpenSSH installed

Even though the product documentation states that 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 Matrix Operating Environment 7.0 software without OpenSSH on HP ProLiant and Integrity Servers Running Microsoft Windows* white paper, available at:

<http://www.hp.com/go/insightmanagement/docs>

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

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

- The user credentials specified for CMS access and database access during installation must match, and be trusted domain credentials available to both the CMS and database systems.
- The user credentials specified during installation for database access must have “Create Database” access permissions. The user must have “write” access to the log files and “read/write” access after installation.

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

## Using Operations Orchestration in Matrix OE

A version of HP Operations Orchestration (OO) that is licensed for limited use is included with Matrix OE. For more information, see the *HP Matrix Operating Environment Infrastructure Orchestration User Guide* at: <http://www.hp.com/go/matrixoe/docs>

## Pointing a remote OO Studio installation to OO Central installed on the CMS

A “Failed Login” error message may appear when logging in to an instance of OO Studio 9.x that points to a remote repository, such as OO Central 9.x installed on the CMS.

### Suggested action

1. Back up the `rc_keystore` file on the CMS.
2. Copy the `rc_keystore` file installed on the CMS at `C:\Program Files\HP\Operations Orchestration\Central\conf\rc_keystore`.
3. Paste the `rc_keystore` file on the remote OO Studio system at `C:\Program Files\HP\Operations Orchestration\Studio\conf\rc_keystore`.
4. Restart OO Studio.

# HP Capacity Advisor

## Limitations

### Capacity Advisor provides partial support for the HP ProLiant WS460c Gen8 workstation

Capacity Advisor does not support the “Customize system from hardware catalog” feature for this workstation. The system template will not be available in the **Edit Scenario**→**Create System** screen. Planning with scenarios that contain a WS460c Gen8 is not supported as it can result in some inaccuracy on planning results. Capacity Advisor is able to collect historical data from ProLiant WS460c Gen8 after following the steps described in the *How to license HP ProLiant WS460c Gen8 Blade Workstations for Capacity Advisor* white paper located at <http://www.hp.com/go/matrixoe/docs>.

### Capacity Advisor does not collect CPU utilization data from VSP 6.1 or 6.2 hosts

When using the Profile Viewer to view a Virtualization Services Platform (VSP) host, the CPU data is not available. CPU data is only available for VSP VMs and vPars. In the Scenario Editor, the only VSP systems available are VM hosts that are only hosting VMs as shown in the table below. In the Capacity Analysis dashboard, the only VSP systems available are VM guest and vPars.

System	Available to Capacity Advisor Scenario Editor
VSP Hosts that only host VMs	Yes
VSP Hosts that host VPARs	No
VSP VM guests	No <sup>1</sup>
VSP vPar guests	No

<sup>1</sup> VM guests are not individually available to the Scenario Editor. Adding a VSP Host to a Scenario Editor will also add all VM guests that are hosted by the VSP Host.

### Capacity Advisor does not support collecting data from KVM VMs

Capacity Advisor supports Linux servers, but does not support collecting data from KVM VMs running on a Linux server. If a KVM host is added to HP SIM and licensed to Matrix Operating Environment, then Capacity Advisor collects data from the KVM host. In the Profile Viewer, the utilization data of a KVM host displays but there is not a drop-down menu to view its VMs. In the Scenario Editor, a KVM host is displayed as a hypervisor with no VMs so it appears as an idle host. A KVM host is not visible in the Capacity Analysis Dashboard.

## Minor issue

### Capacity Advisor does not recognize a new OS installed on a managed node

After installing a new OS on a node previously discovered and licensed for Matrix OE, Capacity Advisor does not update the current OS information, if it has already started collecting data. To resolve this issue:

1. Export the `profile.txt` and execute the following commands at the command prompt:  

```
cd C:\Program Files\HP\Virtual Server Environment\bin  
capprofile -x nameofSystem > nameofSystem.txt
```
2. Modify the file `nameofSystem.txt` by opening the saved file and editing the following fields:  
OS and OSName  
WINNT - Windows, Hyper-V  
LINUX - Linux  
HYPERVISOR - ESX



HPUX – HP-UX

OPENVMS – OpenVMS

3. Import `profile.txt` and execute the following commands at the command prompt:  

```
cd C:\Program Files\HP\Virtual Server Environment\bin
capprofile -i -o -S nameofSystem nameofSystem.txt
```
4. Delete the system in HP SIM by performing the following tasks:
  - a. In **System and Event Collections** in the HP SIM console, select **All System**.
  - b. Select the targeted system and click **Delete**.
5. Remove the profile by executing the following commands at the command prompt:  

```
cd C:\Program Files\HP\Virtual Server Environment\bin
capprofile -r nameofSystem
```
6. Rediscover the system in HP SIM.

## HP Matrix Operating Environment infrastructure orchestration

### Limitations

#### High Availability Hyper-V VMs with RDM/pass-through disks violating various constraints from Hyper-V results in error

Service provisioning or logical server VM creation or reactivation fails with a job message similar to the following:

```
The operation has failed. (The start function on virtual machine
vsels_name [VM info and XML path] failed. )
```

Matrix OE supports High Availability (HA) Hyper-V VMs with RDM/pass-through disks, with various constraints from Hyper-V. Violating these constraints can result in failures. Only HA VMs can have pass-through cluster disks attached. In addition, HA VMs can only have such disks attached. The cluster disks must be in maintenance mode prior to provisioning and activation.

#### Suggested action

If the error occurs when provisioning an HA VM using one or more pass-through cluster disks:

1. Using Microsoft Failover Cluster Manager, select **More Actions**→**Turn On Maintenance Mode** for the cluster disks.
2. Repeat the provisioning operation such as submitting a service create or creating a logical server again.
3. Using Microsoft Failover Cluster Manager select **More Actions**→**Turn Off Maintenance Mode** once the VM is provisioned and the OS is deployed.

If the error occurs when reactivating the HP VM using one or more pass-through cluster disks:

1. Using Microsoft Failover Cluster Manager, select **More Actions**→**Turn On Maintenance Mode**.
2. Repeat the provisioning operation again.
3. Using Microsoft Failover Cluster Manager select **More Actions**→**Turn Off Maintenance Mode** once the VM is reactivated.

#### Migrating from one CMS to another using Operations Orchestration

If you are using Matrix with Operations Orchestration (OO) workflows, you may experience workflow failures after migrating from one CMS to another CMS. These failures have two sets of possible causes and actions:

## 1. Duplication of values

The Run ID values from one CMS are not migrated via the DMT (Data Migration Tool) when moving from one CMS to another. The new CMS will start the values over rather than resuming from a previous value, resulting in duplicate values.

### Suggested action

Run the SQL script provided in the *Data Migration of an Existing Microsoft Windows CMS to a New Insight Management 7.3 System* white paper available at <http://www.hp.com/go/matrixoe/docs>.

## 2. Lack of necessary information

The logic of the OO workflow is dependent on remote information or some state information that is not available or appropriate on the new CMS. If you are using the Custom Hostnames feature, for example, and have modified the sample OO workflows to support your specific server naming strategy, the migration tool will not handle workflow-customized server-naming strategies. If your workflow stored state information on the CMS, it was probably not migrated to the new CMS.

### Suggested action

Modify the OO workflows on the new CMS as appropriate. You must ensure that the workflows and any external resources are manually updated to run properly on the new CMS (for example, to provide necessary connectivity information to a remote server with a database of custom hostnames, or to provide suitable state information based on previous work done on the prior CMS). For more information, see the *Data Migration of an Existing Microsoft Windows CMS to a New Insight Management 7.3 System* white paper available at <http://www.hp.com/go/matrixoe/docs>.

## IO/OO Load Balancer workflows not able to process large HAProxy configuration requests

The Create Logical Load Balancer Groups workflow targeting the Load Balancer HAProxy may fail with the following error message:

```
The execution of the workflow Create a Logical Load Balancer Groups has failed.
```

If this error message appears, the HAProxy configuration file is too big and the infrastructure orchestration/OO Load Balancer workflows are not able to process this request. To create many services that use this load balancer, HP recommends that you use the F5 BIG-IP Load Balancer. For more information, see the F5 - BIG-IP Local Traffic Manager website at [http://h17007.www1.hp.com/one/alliance/f5/big-ip.htm?jumpid=reg\\_r1002\\_usen\\_c001\\_title\\_r0001](http://h17007.www1.hp.com/one/alliance/f5/big-ip.htm?jumpid=reg_r1002_usen_c001_title_r0001).

## VM disks using different datastores

Designer allows boot and data disks connected to Virtual Server Groups to use different storage volume names. This feature allows VM boot and data disks to be allocated to different storage volumes with the following limitations:

- Integrity VM disks must all use the same storage volume names or be blank.
- Cloud server disks do not support storage volume names and must be blank.
- Hyper-V VM Linked Clone disks must have the same storage volume names or be blank. Hyper-V VM Linked Clone cannot be spread across different datastores.
- Microsoft does not support VM Linked Clone using multiple datastores.

---

**NOTE:** VM logical servers manually created from Matrix OE visualization can also specify VM disks on different data stores.

---

## Cross-technology logical servers are not supported by 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 VM using logical server management. Similarly, if a virtual logical server is created in infrastructure orchestration, it cannot be moved to a server with VC using logical server management.

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

## Limiting the 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 increase over time as VC and hypervisors are enhanced.

If the limits are exceeded in an infrastructure template, provisioning operations using that template 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.

## Multi-initiator NPIV support on Integrity server blades

Multi-initiator NPIV features used by Matrix OE and HP Virtual Connect Enterprise Manager (VCEM) are supported on Integrity server blades using:

- QLogic FC (Fibre Channel) controllers with the HP-UX 11i v3 Update 8 (March 2011) and later release.
- Emulex FC controllers with HP-UX 11i v3 March 2013 update and later release.

### Suggested action

On Integrity server 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.

## Installing the Microsoft Sysprep tool for post-installation customization of Windows 2003 guests on Hyper-V

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

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

### Suggested action

After installing or upgrading to Insight Management 7.2, 7.2 Update 1 or Update 2, install the Microsoft Sysprep tool on the CMS in the C:\Program Files\HP\Insight Control virtual machine management\Sysprep\2003 folder.

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

Matrix infrastructure orchestration provisions VMs 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 VM. 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 Matrix infrastructure orchestration template to three or fewer.

### CSV does not appear after Hyper-V VM provisioning with HA enabled

If you use Matrix infrastructure orchestration to create and deploy a virtual provisioning template on a Hyper-V VM host with the high availability (HA) option enabled, and then view the VM 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 VM using the Microsoft Hyper-V Manager on the Hyper-V VM host, the Failover Cluster Manager correctly displays the disk resource and the VM resources created.

#### **Suggested action**

No action is needed. The HA VM disk was created correctly by infrastructure orchestration on the CSV.

### CSV 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 CSV as a cluster disk, and allocates only one VM to the cluster disk. When trying to deploy multiple VMs, the following error appears:  
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 — The CSV and cluster disk are only available on the VM host if the registered VM host ID is the VM host owner. Otherwise, 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 in the Hyper-V cluster to Matrix infrastructure orchestration and register 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 Hyper-V R1 with linked clone and HA enabled

If you attempt to provision a template containing multiple 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 VM 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 VM.
3. Resubmit the Create request, or  
For the Add Server request, delete the existing service and recreate it with the updated template.  
Migrate data from the existing service to the new service before deleting the existing service.

### Suggested action 2

Upgrade to Hyper-V R2, and specify a CSV instead of a cluster disk in the template.

### Registering Microsoft SCVMM on a single 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, enabling communication between the two servers. As a result, SCVMM can be registered on only one CMS. In a federated CMS configuration, each CMS requires a separate SCVMM host to manage Hyper-V hosts through SCVMM. This requirement differs from VMware vCenter which can be registered to more than one CMS. For more information about registering SCVMM, see the *HP Matrix Operating Environment Infrastructure Orchestration User Guide* available at:

<http://www.hp.com/go/matrixoe/docs>.

### Browsing to the user interface requires the default HP SIM port

HP SIM is installed with a default port of 50000 in the JBoss container. Browsing to the Matrix infrastructure orchestration user interfaces, such as infrastructure orchestration designer, requires this port.

### Suggested action

Do not modify the HP SIM default port of 50000 in the JBoss container.

### Enabling Matrix infrastructure orchestration provisioning support of Linux VMs on Hyper-V

**NOTE:** The `hpio.properties` value `skip.linux.on.hyperv.template.personalization` allows Linux VMs on Hyper-V, but it is deprecated. See “[Skipping automated OS customization](#)” (page 21) for information on allowing Linux VMs on Hyper-V.

### Skipping automated OS customization

By default, after infrastructure orchestration deploys an OS to a server, it automatically customizes (personalizes) the OS by setting the host name, IP addresses, netmask, default gateway, and other attributes as appropriate. As of the 7.3 release, automatic OS customization can be skipped by performing the actions noted below. These actions allow the service template architect to design service templates that use OS images intended to be customized outside of Matrix OE, such as when deploying Linux VMs onto Hyper-V. When the skip customization feature is used, customization of the deployed VM is no longer performed by Matrix, and becomes the responsibility of the architect to ensure that the VM is configured properly.

#### 1. Enable the Skip OS Customization check box

Set `allow.skip.os.customization = true` in `hpio.properties`. This enables a per-server-group Skip OS Customization check box to appear in Designer.

#### 2. Select the Skip OS Customization check box

In Designer, navigate to **Configure Server Group**→**Software**→**1. Select Operating System**. Select the **Skip OS Customization** check box.

#### 3. Ensure appropriate OS customization is done (perhaps via an Operations Orchestration workflow attached to the service template).

## Selecting fabrics when auto-generating storage pool entries through Matrix infrastructure orchestration

When auto-generating storage pool entries from Matrix infrastructure orchestration, if more than a single pair of fabrics exist, Matrix infrastructure orchestration determines which pair to use.

### Suggested action

To set the preferred list of fabrics to use during the auto-generation of storage pool entries, edit the `lsa.properties` file and add the

`SPM_AUTO_PREFERRED_FABRICS=FabricA,FabricB,FabricC:FabricD` property where FabricA, FabricB, and so on are replaced with user-defined fabric names.

Fabrics separated by a colon (:) are redundant pairings.

If the storage is not redundant, Matrix infrastructure orchestration searches for storage pool entries using the specified fabrics in the order listed. If the storage is redundant, Matrix infrastructure orchestration only searches for storage pool entries using the specified redundant fabric pairings in the order listed.

In the above example:

- For non-redundant storage, IO searches FabricA, FabricB, FabricC and FabricD in order.
- For redundant storage, IO searches the FabricC/FabricD pairing.
- If no storage pool entry is found after searching the fabrics defined in this list, IO searches all known fabrics based on the pre-defined rules.

---

**NOTE:** The fabrics listed in this property are used by all auto-generation of storage pool entries from Matrix infrastructure orchestration on the CMS. This property cannot be configured at the user or organization level.

---

## Major issues

### Linked clone provisioning of a VM template containing shared disks does not succeed

During ESX server provisioning, errors similar to the following are displayed:

```
Error taking a snapshot on virtual Logical Server nnnnn. Virtual machine is configured to use a device that prevents the snapshot operation: Device is a SCSI controller engaged in bus-sharing.
```

The VM template created from an infrastructure orchestration template specifies linked clone servers and deploys software as an embedded shared disk. This combination is not supported.

#### Suggested action 1

Edit the template in Designer and clear **Deploy as Linked Servers**.

#### Suggested action 2

Select other software to deploy without an embedded shared disk, and then retry the Create Service operation.

### Provisioning of Matrix OE service fails in customization

Provisioning of Matrix OE services can fail in customization if Insight Control server provisioning or SA deployment is attempting to join the server to an Active Directory during the OS installation. Doing so will cause the Matrix OE post-install customization step to fail with the following error:

```
Error customizing Logical Server <Logical server name> realized by <target UUID>.
```

#### Suggested action

Create an OO workflow to join the newly provisioned server to the domain.

---

**NOTE:** When utilizing the PowerShell PSEXEC command, the OO RSJRAS service should be run as administrator.

---

## Minor issues

### Create Service request with Insight Control server provisioning fails when deploying SLES11 SP3 OS

The installation of SLES11 SP3 hangs at 60% or fails with an exit code 6 error on the appliance console.

- In Matrix OE, one of the following error appears on the infrastructure orchestration Requests screen and the failed server is moved to the maintenance pool:  

```
Failed while deploying operating system for service. Cause: Logical Server job has failed. Logical server job completed with a failure status. Failure: At least one OS or software deployment has failed. Check the Deployment Service to diagnose the details
```

or

```
Task for Logical Server has failed. Logical server job (logicalservername_FQDN) completed with a failure status. Failure: At least one OS or software deployment has failed
```
- In the Insight Control server provisioning appliance console, one of the following error appears:  

```
Failed to integrate HP SA Agent at Step 19
```

or

```
Wait for HPSA Agent failed with exit code 6 at step 21
```

#### Suggested action

1. Perform manual cleanup, see the Manual clean-up process (physical) in the *HP Matrix Operating Environment Infrastructure Orchestration User Guide* at <http://www.hp.com/go/matrixoe/docs> and perform the steps Manual cleanup for HP Insight Control server provisioning.
2. Recreate the service request in Matrix infrastructure orchestration.

### Create Service request with Insight Control server provisioning fails when deploying RHEL 5.9 OS

The installation of RHEL 5.9 OS fails with the following messages:

- In Matrix OE, the following error appears on the infrastructure orchestration Requests screen and the failed server is moved to the maintenance pool:  

```
Task for Logical Server has failed. Logical server job (ID = logicalservername_FQDN) completed with a failure status. Failure: At least one OS or software deployment has failed
```
- In the Insight Control server provisioning appliance console, the following error appears:  

```
Monitor Installation with exit code 1 at step 19
```

The target server remains in LinuxPE in maintenance server mode.

#### Suggested action

1. Perform manual cleanup, see the Manual clean-up process (physical) in the *HP Matrix Operating Environment Infrastructure Orchestration User Guide* at <http://www.hp.com/go/matrixoe/docs> and perform the steps Manual cleanup for HP Insight Control server provisioning.
2. Recreate the service request in Matrix infrastructure orchestration.

### Provisioning of Integrity VMs with NPV disks with storage pool entries containing both boot and data volumes may fail

Provisioning of Integrity VMs with NPV disks may fail if the storage pool entry selected has both boot and data entries. The following error appears:

```
Failed while deploying operating system for service <service-name>.
Cause: Logical server job (ID = <Logical-server-job-id>) completed with
a failure status. Failure: An unexpected error has occurred. Check
Alc1_0.0.log and HP SIM log files under HP SIM logs directory for
additional details
```

#### **Suggested action**

When creating a service from a template containing physical boot and data disks, create a separate storage pool entry for each boot disk and a separate entry for each data disk. Do not create a single entry for both boot and data disks.

### 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 disks contain different tags than the boot disk, the service fails to provision.

#### **Suggested action**

Use the same storage tags for the boot disk and private data disks in a logical server group. If different tags are needed, manually define separate storage pool entries for the boot disk and private data disks and include the required tags.

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

If a 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 information about the CSV. Subsequently, operations on the Hyper-V cluster (including provisioning) fail, or all VM hosts in the cluster are 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 failed operation.

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

Matrix infrastructure orchestration designer allows you to connect a physical disk to a Hyper-V HA virtual server group and set the Raw Device Mapping Type for the physical disk. However, provisioning a template with this configuration results in an error.

#### **Suggested action**

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

### NIC is not configured properly on a new VM after a successful Linux guest deployment in an ESX/ESXi/vSphere environment

Networking may not work on a provisioned Linux VM on an ESX/ESXi/vSphere host due to a known VMware issue that impacts all versions of Linux in an ESX/ESXi/vSphere environment.

#### **Suggested action**



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).

### Setting retry parameters for BL620c G7 and BL680c G7 server blades

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

#### Suggested action

Increase the retry parameter:

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

### Deployment of HP-UX 11i v3 1203 or HP-UX 11i v3 1209 on an Integrity i2 or i4 server blade using Matrix infrastructure orchestration with Ignite-UX may fail

When deploying HP-UX 11i v3 1203 or HP-UX 11i v3 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 is not provisioned. The request cannot deploy the HP-UX OS and the following error appears:

```
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. Be sure that all hosts are in the same zone, and then perform a LUN clean-up using the storage management server (such as EVA or 3PAR StoreServ Storage).
  - b. Unpresent the host associated with the vdisks and delete the vdisks.
  - c. Create the new vdisks with the required sizes and RAID levels 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. Select the 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 and verify that the Usage column indicates that the servers are not in use.
  - 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 Service or Delete Service request for an Integrity server fails due to an EFI disk read error

A Create Service 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 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:

1. Erase the disk used by the target server for deployment.
  - a. Be sure that all hosts are in the same zone, and then perform a LUN clean-up using the storage management server (such as HP P6000 Command View Software or 3PAR StoreServ Storage).
  - b. Unpresent the host associated with the vdisks and delete the vdisks.
  - c. Create the new vdisks with the required sizes and RAID levels 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. Select the 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 and verify that the Usage column indicates that the servers are not in use.
  - 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 fails on Integrity BL870c i2, BL890c i2, BL870c i4, and BL890c i4 server blades](#)

The `serverboot.physical.wait.seconds` parameter in the `IOinstalledDir/conf/hpio.properties` file and the `INTEGRITY_POWERON_WAITTIME` parameter in the `VSEinstalledDir/conf/lisa/lisa.properties` file are not set correctly according to the server blade models. The wait time specified in these parameters allows the Integrity server blade to boot completely before infrastructure orchestration or Matrix OE logical server management executes a power off command. Correct values are required to prevent a QLogic HBA card EFI from being corrupted while the VC profile is written to or removed from the NVRAM.

### Suggested action

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

---

**NOTE:** The time varies based on the server blade model, memory size, and number of QLogic cards. You can measure the power cycle time from the Integrity server blade remote console using 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:

- BL870c i2 and i4 server blades:
  - `serverboot.physical.wait.seconds` — 540 seconds
  - `INTEGRITY_POWERON_WAITTIME` — 540000 milliseconds
- BL890c i2 and i4 server blades:
  - `serverboot.physical.wait.seconds` — 1080 seconds.
  - `INTEGRITY_POWERON_WAITTIME` — 1080000 milliseconds.

### Firefox 17 requires a self-signed certificate for access to Matrix infrastructure orchestration

The first time you run Matrix infrastructure orchestration in a Firefox 17 browser, the “This Connection is Untrusted” message may appear, indicating that Firefox requires a new certificate.

#### Suggested action

To assign a new certificate:

1. On the Firefox “This Connection is Untrusted” screen, click **Technical Details** or **I Understand the Risks**.
2. Click **Add exception**.

If Add exception does not appear:

- a. Open a new tab in Firefox.
  - b. Enter 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. The “This Connection is Untrusted” screen appears and includes the Add an exception option.
3. The Add Security Exception screen appears. If the **Location** box is not populated, enter `https://[your-CMS-name-or-IP]:50000/`.
  4. Click the **Permanently store this exception** checkbox.
  5. Click **Get certificate**.
  6. Click **Confirm Security Exception**.
  7. Return to the SIM tab and refresh the screen, or open a new tab, navigate to HP SIM, and then restart Matrix infrastructure orchestration.

### OO Central and OO Studio contents are not localized to Japanese

In a Japanese environment, new installations of Matrix OE 7.2.2 include screens, menus, and buttons that are mostly localized to Japanese, but the contents of OO Central and OO Studio (such as workflows and system properties) are not localized to Japanese.

If you are upgrading to Matrix OE 7.2.2, OO Central and OO Studio display the localization behavior of the original version. Therefore, if you are upgrading from Insight Dynamics 6.3 to Matrix OE 7.2.2, the OO Central and OO Studio screens display in English.

### Deleting the chargeback archive can cause the 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 HPMOEMetering database transaction log, especially if the date specified by the `-d` qualifier is a year or more. If the log grows large enough, the hard disk on which the database is running may run out of free space.

## Suggested action

Before executing `chargeback 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.

## HP Matrix OE logical server management

### Limitations

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

Cross-technology logical servers can be moved between ProLiant server blades with VC to ESX/ESXi/vSphere and back, between Integrity server blades with VC to Integrity VM and back, and between servers with VC 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 [www.hp.com/go/matrixoe/docs](http://www.hp.com/go/matrixoe/docs).

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

Before performing any operations on a physical server running as a VM host, be sure that all jobs related to controlling the VM guests running on the VM host logical server are complete. Then, use the appropriate steps in the following table 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 the VC profile to another server 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 the VM host.	Power off or perform a graceful shutdown of all VM guests.	Shut down the VM host logical server.
Power off the VC server 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 server blade where the VM host logical server resides. Select <b>Tools</b>→<b>Logical Servers</b>→<b>Power Off...</b></li> </ol>
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>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. 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

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
		<b>Tools→Logical Servers→Deactivate.</b>
Activate or reassign the VC profile to a different server 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→Virtual Machine→Unregister Virtual Machine Host...</b></li> <li>2. Activate the VM host logical server. Select <b>Tools→Logical Servers→Activate.</b></li> <li>3. Register the VM host using Insight Control virtual machine management. From HP SIM, select <b>Configure→Virtual Machine→Register Virtual Machine Host...</b></li> </ol>

### Multiple storage volume limitations

The 7.3 release adds support for VMware and Hyper-V virtual machines using multiple data stores. The specification of storage volume names for VM disks is restricted in the following cases:

- Integrity VM disks must either specify the same storage volume name or not specify a name at all.
- Cloud server disks must not specify a storage volume name.
- Hyper-V VM linked clone disks must either specify the same storage volume name or not specify a name at all.
- Microsoft does not support VM Linked Clone using multiple data stores.

**NOTE:** By default, specifying storage volume names when adding a disk to a VM is restricted to administrators. To enable all users to specify storage volume names when adding a disk to a VM, change the value of the `hpio.properties user.disk.add.storage.volume.names.restricted` property from true to false.

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

A Hyper-V VM with a configuration file (.xml) in the default location on the hard disk (`\ProgramData\Microsoft Windows\Hyper-V`) is not supported by logical server management. You cannot import a VM guest as a logical server if the VM 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 VM if the configuration file is created in this directory. To prevent issues with reactivation, logical server management checks the path and rejects any VMs with a configuration file in the default location.

### Limited support for Hyper-V CSVs created with GPT disk partitioning system

Hyper-V VM hosts configured with 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 Hyper-V VM logical servers

To enable the selection of a specific iSCSI datastore as the storage definition during logical server creation or modification, iSCSI disks associated with Hyper-V hosts must:

- Be a cluster disk or a shared cluster volume.
- Be online.
- Contain a volume created with a drive letter.

With these limitations, only master boot record-formatted iSCSI cluster disks are supported.

### VM datastore size and configuration filename limits

When creating a VM logical server (ESX/ESXi/vSphere, 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 VM configuration filename is 255. If you specify a size larger than the maximum, a warning is logged and the entry is ignored.

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

### Hyper-V VM imported as a logical server may require a SCSI controller driver

If a Hyper-V VM is created outside of Matrix OE 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 Hyper-V VM logical server is created with Matrix OE visualization, the boot disk is created with an IDE controller. The IDE controller allows the installation of the operating system using Deploy Tools or using images without additional drivers.

If a Hyper-V VM is provisioned by Matrix infrastructure orchestration, IO may use either a SCSI or an IDE controller boot disk. If the VM was created using a SCSI boot disk, IO obtains the SCSI driver for the SCSI controller from the VM template.

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

There is a limitation in Hyper-V R1 that prevents the hypervisor from providing logical server management with required information about a Hyper-V cluster. This causes only one target to appear when you attempt to activate a VM logical server configured to be created on a cluster disk volume. In Hyper-V R1, a logical server that is HA-enabled can be configured on a cluster disk volume only.

#### **Suggested action**

Upgrade to Hyper-V R2. In Hyper-V R2, the activation of a VM logical server configured to be created on a cluster disk or a CSV lists all Hyper-V hosts in the cluster on the target screen.

### HA Hyper-V VM with an RDM disk can be activated on cluster disk owner only

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 VC profile capabilities are not supported by logical server management

Importing a server with a VC profile that uses unsupported features results in the following error:

The Server Profile cannot be imported because it contains the following capabilities unsupported by Logical Servers:

The following VC profile features are not supported by Matrix OE logical server management:

EXTENDEDFC	The server profile includes 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.
EXTENDEDFCOE	The server profile includes more than one FCoE connection associated with the same interconnect bay (IO bay).
ISCSI	The server profile includes at least one iSCSI connection.
NAG	The server profile is associated with a Network Access Group.

### Applications that depend on I/O 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) may fail on Integrity cross-technology logical servers. This occurs because I/O hardware paths and Legacy Storage DSFs can change during cross-technology Move operations. After a cross-technology logical server Move operation, Networking I/O Instance Numbers and Storage Agile DSFs do 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/fc1p0-0x5001438002A30041,0x5001438002A30044
```

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

---

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

---

- c. 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.

## Minor issues

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

If you attempt to activate an Integrity VM logical server with three or more vCPUs on an Integrity VM host with a 1596 Mhz or more processor speed, and the entitlement value used to create the VM on the VM host falls below the minimum 5% value, the activation may fail with the following error on the VM host:

```
"Value x.x for entitlement_min is below minimum, setting to 5.0."
```

#### Suggested action

To set the Integrity VM configuration variable `GUESTMHZMIN` on the Integrity VM host to the nearest integer value that is 5% of the host processor speed in cycles, use the following commands:

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

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

When activated directly on an Integrity VM host using Matrix OE logical server management, a logical server fails to detect its assigned LUNs at the VM EFI prompt. The logical server does not boot the operating system automatically, even if it is assigned with a LUN that is pre-provisioned with an operating system.

#### Suggested action



View the LUN mappings, and then select the LUN in the device list that is required to boot or install the operating system. To view the LUN mappings, perform the following steps from the Integrity VM EFI prompt:

1. To check the configurable controllers, run the `drvcfg` command.

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

2. To enumerate all SCSI and FC LUNs, run `drvcfg -s <drvnum> <ctrlnum>`. The `drvnum` and `ctrlnum` parameters display in the output from step 1. Press **Y** when prompted for confirmation.

```
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. To reconnect drivers, refresh mapped drives and display available boot options, run the `reconnect -r` and `map -r` commands.

```
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)
```

### Move or Activate operation may fail on a VC server blade imported into a P2V portability group

Before CloudSystem Matrix 7.1 Update 2, Matrix OE did not support a Move or Activate operation on a VC server blade that was imported into a portability group containing a mix of VC server blades and VM hosts. In 7.1 Update 2, this feature is supported. However, if the import occurred in a version before 7.1 Update 2 and a Move or Activate operation is attempted after upgrading to 7.1 Update 2, the operation fails with a message similar to the following:

```
Placement service failed to get targets because <starting VC blade>
cannot be cast to <possible VC blade or VM target>.
```

Matrix OE displays the same message even though there are no VMs in the Move or Activate operation.

The error does not occur:

- If the resources are created through Matrix OE within the portability group.
- If the resources are imported into a logical server portability group containing only VC server blades such as a VC-based logical server imported into a VC only portability group.
- If the VC server blade is imported into the portability group in the 7.1 Update 2 release.

### Suggested action

1. Unmanage the logical server, select **Delete**→**Unmanage Logical Server...**
2. Re-import the logical server into the desired portability group, which can be a portability group including both VC server blades and VM hosts.
3. Perform the Move or Activate operation again.

## HP Matrix Operating Environment recovery management

### Limitations

In a federated environment, Matrix recovery management is only supported on a primary-federated CMS at both the primary and DR sites. In other words, Matrix recovery management is not supported on a secondary-federated CMS.

## Storage Provisioning Manager

SPM runs on an x86/x64 Windows CMS. iSCSI is not supported by SPM.

### Limitations

#### Brocade Network Advisor software package 12.0.2

BNA software package 12.0.2 is not supported with SPM version 2.3.

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

Two features in SPM 2.3 can be used to limit the number of initiator WWNs that may be connected to a storage controller port. The white list capability for HP 3PAR StoreServ storage allows a storage administrator to control the controller ports used by SPM, thus preventing SPM from using controller ports that are reserved for other uses. The port group feature gives a storage administrator more detailed control over how SPM uses controller ports for all storage.

When adding ports to a storage system for which a white list is configured, if any new ports are to be used by SPM and Matrix OE, they must be added to the white list, and the storage system must be resynchronized in SPM.

---

**NOTE:** The white list capability can only be used for 3PAR.

---

These limitations can impact Matrix OE-created servers that use SPM to provision its storage.

#### 3PAR StoreServ Storage Zero Detect

SPM does not recognize the Zero Detect property of volumes on HP 3PAR StoreServ Storage. SPM cannot create volumes with the Zero Detect property. SPM can import and use pre-provisioned volumes with the Zero Detect property, but does not recognize the property value. SPM continues 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 more information, 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 does not support Peer Persistence. SPM cannot import volumes that are part of a remote copy group, and cannot import two volumes that share the same world wide name.

#### SPM limitations on EMC Symmetrix VMAX arrays

For more information about configuring the SPM adapter for EMC Symmetrix VMAX, see the *Insight Management 7.2 Update 1 Release Notes* at <http://www.hp.com/go/matrixoe/docs>. For more

operational details, see the technical white paper on SPM adapter for EMC Symmetrix VMAX at <http://www.hp.com/go/matrixoe/docs>.

- Creating a new volume with a given RAID requires pre-configuration of a storage pool setting instance by the EMC storage administrator. Generally, this can be done by creating a volume of that RAID through EMC native tools (SMC or SYMCLI).
- Creating a thin volume for a given RAID requires creation of a thin pool for that RAID, if it does not already exist, through EMC native tools (SMC or SYMCLI) by the EMC storage administrator.
- SPM defines thick and thin volumes based on the capacity requirements. If a new volume is to be created with committed capacity set to 100%, then a thick volume is expected to be created by the adapter. If the committed capacity is set to 0% (or to any value less than 100%), then a thin volume is expected to be created by the adapter. This behavior is expected irrespective of the storage pool being sent in the request. Hence the adapter implementation overrides the way EMC creates the thick and thin volumes:
  - When SPM requests a thick volume on a thin pool, the adapter will create a thin volume on the given thin pool.
  - When SPM requests a thin volume on a concrete pool, the adapter will create a thick volume on the concrete pool.
- Growing a thin volume is not supported in EMC Symmetrix VMAX.
- The SPM adapter for EMC Symmetrix VMAX supports the host modes Linux, Windows, Windows 2008, Windows 2012, VMware, and HP-UX. It is has not been tested with the OpenVMS host mode.
- For all supported host modes, the maximum number of LUNs that can be presented to on a host initiator per EMC port is limited to 2000.

## Minor issues

### Windows 2008 may fail on servers with Emulex LPe1205 HBA with multipath

When an attempt is made to deploy Windows 2008 on a server that contains Emulex LPe1205 HBA, the deployment may fail if more than one storage controller port is zoned for visibility. The server fails to boot from SAN and continuously reboots.

#### **Suggested action**

Since SPM automated zoning provides visibility to all controller ports on the given fabric, it is necessary to set SPM for manual zoning on the given fabric/network.

1. Create an XML file for the unmanaged network and set the state to Zoned.
2. From SPM, choose **Edit Network** and check **Allow External Processing**.
3. Contact the storage administrator to manually zone each initiator WWN to a single controller port. The controller port used in the storage pool entry for the volume also needs to be zoned.
4. Edit the SPM XML file representing the fabric/network to include the zone set(s).

After the next refresh or a resynchronization of the network, the Matrix workflow can be resumed. For more information, see the *HP Storage Provisioning Manager User Guide* and the *Faster storage provisioning in the HP Matrix Operating Environment: Use of the HP Storage Provisioning Manager storage catalog with controlled storage operations* available at <http://www.hp.com/go/matrixoe/docs>.

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

### 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—Adds one image that is less than 2 GB using Internet Explorer or 4 GB using Mozilla Firefox
- Using the **Add+** button—Adds up to 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 not accessible. In this case, perform steps 2 through 4 below.

If you uploaded an image or images that exceed the size limitation and the screen displays “Loading...” for several minutes, perform the following steps:

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 is Active.
3. From the Images screen, click the row of the image to be deleted.
4. Select **Actions**→**Delete**.

## Major issues

### RHEL 6.4 with KVM OS Build Plan fails with error

The default Insight Control server provisioning RHEL 6.4 KVM OS build plan configures the “br0” device, causing post-installation customization to fail with the following error:

```
Error customizing Logical Server <Logical server name> realized by  
<target UUID>.
```

### Suggested action

To prevent this, edit the kickstart file referenced in the OS Build Plan as follows. The “br0” device configuration is located in the section delineated between the comments “BEGIN KVM BRIDGE CONFIGURATION” and “END KVM BRIDGE CONFIGURATION.”

1. Remove the “br0” device configuration by removing the section delineated between the comments “BEGIN KVM BRIDGE CONFIGURATION” and “END KVM BRIDGE CONFIGURATION.”
2. In the %post section of the kickstart file, add the `chkconfig NetworkManager off` command as shown below:

```
%post --log /tmp/icsp-networkconfig.log  
chkconfig NetworkManager off
```

Deployment using the above methods will result in error messages appearing in the Insight Control server provisioning jobs. Specifically, you will see BRDC write personalization task output that contains entries similar to the following:

```
BRDC read personalization task for server <server number>
```

Invalid eth name vibr0

These messages are expected and can be ignored.

## Minor issues

### Clear browser cache if upgrade is unsuccessful

If an error occurs during the upgrade to Matrix 7.2 Update 1 KVM Private Cloud, you must clear your browser's cache to remove partially installed files before retrying the upgrade procedure.

#### Suggested action

1. Log out of the appliance and close the browser.
2. Clear your browser's cache:
  - In Internet Explorer, use **Tools**→**Delete Browsing History**.
  - In Firefox, use **Tools**→**Clear Recent History**.
3. Log back into the appliance and click **Install update** from the Settings->Update Appliance screen to install the patch kit. Reinstalling the pre-update patch kit is not necessary.

### Internet Explorer 8 displays "Out of Memory" if the Matrix KVM Admin Console Activity page is left open

Internet Explorer 8 may display an "Out of Memory" message 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 6 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 action

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.

### Intermittent timeouts may occur during VM deployment

Deploying a VM 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 display an error status.

#### Suggested action

1. Retry the create service request in Matrix infrastructure orchestration.
2. If the retry does not succeed, perform the following steps:
  - a. Restart the rabbitmq-server service on the management host by entering the following command from the command line:

```
service rabbitmq-server restart
```
  - b. 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.

### Cloud connectivity exclusion ranges cannot include spaces

After an upgrade to Matrix 7.2 Update 1 KVM Private Cloud or an appliance reboot, spaces inserted between IP addresses in the cloud connectivity exclusion range on the Edit Cloud Networking screen cause the portion of the range after the space to not appear. This does not affect IP address assignment; it is an error on the screen only. However, if you want to edit the range, you are not able to see the current range.

#### Suggested action

When specifying cloud connectivity exclusion ranges on the Edit Cloud Networking screen, do not insert a space between IP addresses. For example, specify the following, without a space after the comma:

```
192.0.2.1-127,192.0.2.135-254
```

If a space was inserted before a reboot or upgrade and the exclusion range is not correctly displayed, manually edit `/ci/data/isc-network-info` to remove any existing spaces.

### Error deploying a KVM private cloud may create an additional, incomplete VM instance

Deploying a KVM private cloud using Matrix OE may, in rare situations, return the following error in the Matrix infrastructure orchestration Requests Details screen:

```
Task for Logical Server <logical server name> has failed. Logical server job (ID = <server ID>) completed with a failure status. Failure: Could not create HP Matrix KVM Private Cloud server <logical server name> in capacity pool isc: 413 Request Entity Too Large The body of your request was too large for this server.
```

If this error appears, the service was created successfully. However, the Matrix KVM Admin Console Instances screen shows one VM in addition to the number of VMs requested in the service template. This instance has an error status.

#### Suggested action

In the Matrix KVM Admin Console, delete the extra instance that was created with an error status. If you subsequently attempt to delete the infrastructure service without first deleting this instance, the delete service request fails.

### OS type of virtual machine instances are “not set” after restart

After restarting the Matrix KVM Private Cloud appliance, the OS type of existing virtual machine instances is displayed as “not set” on the Instances screen.

#### Suggested action

Expand the details of the instance and note the image name. From the Images screen, search for the image name. The OS type of the image is shown, which is also the OS type of the instance.

### Deployed instance does not boot and an error appears

A VM instance did not boot, even though the Hosts screen appears to show sufficient available resources. The following error appears:

```
No available host can provide the specified resources
```

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 VM instances already provisioned to a host) from the capacity of the host.

If one or more VM 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 already allocated to the powered down instances.

#### Suggested action

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

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

### Deleting a large number of virtual machine instances from Matrix infrastructure orchestration occasionally times out

Attempting to continuously delete a large number (more than 25) of VM instances from Matrix infrastructure orchestration may not succeed. When this occurs, you may see an error in IO similar to the following:

```
ERROR: Reached TIMEOUT attempting to delete Nova instance with serverId:
xxxx-xxxx Failed to stop server. Instance in task_state deleting. Cannot
stop while the instance is in this state.
```

This error occurs when the VM host on which the instances were running times out and is disconnected from the appliance.

#### **Suggested action**

1. Restart services on the VM host (compute node) by opening an SSH connection to the host and entering the following command from the command line:  

```
service openstack-nova-compute restart
```
2. Retry the delete instance operation from Matrix infrastructure orchestration.

### External yum repository causes host activation to fail

An external yum repository (used for RPM-based Linux updates) that is enabled on a VM host may cause the first host activation 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 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 `/etc/yum.repos.d/rhel.repo` file.

### 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. Subsequent filtering for instances in any state does not return results.

#### **Suggested action**

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

### Sorting does not work as expected

On the Instance Types screen, clicking the numeric column may result in alphabetic ordering. On the Users screen, you cannot sort users by role, but you can sort by name.

#### **Suggested action**

- On the Instance Types screen, sort instance types by clicking a different column heading.
- On the Users screen, sort users by role by entering a role in the Search field.

### Incorrect hyperlink message displays when filtering or searching returns no results

When filtering or searching for resources on a Matrix KVM Admin Console screen and no results are found, an incorrect message displays with 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".

## Managed node installation

### Limitations

#### Multipath I/O and Emulex FC HBA or Emulex FlexFabric Converged Network Adapter is not supported on Windows Server 2008, Windows Server 2008 R2, and Windows 2012

Using multipath I/O during OS installation is not supported with an Emulex FC HBA or an Emulex FlexFabric Converged Network Adapter on any CloudSystem Matrix-supported server on Windows Server 2008, Windows Server 2008 R2, and Windows 2012.

The following suggested action enables 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 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 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.

#### Suggested action 2

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

1. Define the storage pool entry. Add multiple ports and mark 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.

## Other issues and suggested actions

### Minor issues

#### RHEL 5.9 OS deployment may fail on ProLiant servers configured with SAN multipath

For RHEL 5.9 deployments in a multi-path environment, refer to the deployment software documentation for how to enable multi-path.

---

**NOTE:** If Storage Provisioning Manager 2.2 is involved in the zoning of the LUN, SPM zones all connected controllers with the server WWN. If there are multiple target ports on the storage solution, then multipath must be configured for any RHEL 5.9 deployment with the LUN provisioned by SPM.

---

## HP SIM security certificate is not imported

A problem occurs when an Insight Remote Support Advanced A.05.80 installation is supporting an HP-UX Mission Critical contract and the HP SIM security certificate is 2048 bits. System Management Homepage (SMH) does not automatically import the HP SIM 2048-bit certificate causing communication with HP-UX Advanced Configuration Collections to fail. In HP SIM 7.2, the security certificate is 2048 bits by default, but this issue also occurs if a previous HP SIM 7.x setup changed the certificate to 2048 bits.

### Suggested action

Upgrade to the latest HP-UX SMH version (A.3.2.6/A.3.2.2/A.2.2.9.3.1) from <https://h20392.www2.hp.com/portal/swdepot/displayProductInfo.do?productNumber=SysMgmtWeb>.

## Managing an 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 is incorrect, causing the cluster to be discovered incorrectly in HP SIM.

To correct this problem, delete the cluster objects from HP SIM, follow 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 HP SIM.

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

In VMware ESXi/vSphere, the default name of the 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:

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

### Suggested action

Modify the **Network Label** of the VMKernel port used for management traffic to "Management Network":

1. Log into vCenter using the vCenter client.
2. Select the **Hosts and Clusters** view from the **View**→**Inventory** menu.
3. Select the **ESXi server** from the tree in the left pane.
4. Select the **Configuration** tab in the right pane.
5. In 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 with the IP that is registered with Insight Control virtual machine management, or that will be used for registration with Insight Control virtual machine management.
7. In the virtual switch properties pop-up window, select the service console port, and then click **Edit**.
8. Change the Network Label to "**Management Network**," and click **OK**. Close the virtual switch properties pop-up window.

It can take up to five minutes for the cache to refresh and the API to report the ESXi/vSphere host.

---

## 4 Support and other resources

### Contacting HP

#### Information to collect before contacting HP

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

- HP CloudSystem Matrix Starter Kit or Expansion Kit HP BladeSystem c7000 Platinum Enclosure serial number and/or SAID if applicable
- 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)

- 
- ❗ **IMPORTANT:** Be sure to mention that this is a HP CloudSystem Matrix configuration when you call for support. Each HP CloudSystem Matrix Starter Kit or Expansion Kit HP BladeSystem c7000 Platinum Enclosure serial number identifies it as a HP CloudSystem Matrix installation.
- 

#### 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>

## Related information

The latest versions of manuals and white papers for HP CloudSystem Matrix and related products can be downloaded from the web at [www.hp.com/go/assistance](http://www.hp.com/go/assistance).

HP CloudSystem Matrix documentation refers to Matrix Operating Environment documents, HP Server Automation (SA), and HP Cloud Service Automation (CSA).

For Matrix Operating Environment documents, see the Insight Management documentation: <http://www.hp.com/go/matrixoe/docs>.

For HP Server Automation (SA) and HP Cloud Service Automation (CSA), search the SSO portal to retrieve relevant documentation:

1. Navigate to the SSO portal (<http://support.openview.hp.com/selfsolve/manuals>).
2. Login to HP Passport, if necessary. If you do not already have an HP Passport account, you will need to create one.
3. In the Product menu, select **Server Automation or Cloud Service Automation**.
4. In the Product version menu, select the most current version.
5. In the Operating system menu, select the relevant operating system.
6. Click **Search**.

## Typographic conventions

This document uses the following typographical conventions:

<i>Book title</i>	The title of a book. On the web, this can be a hyperlink to the book itself.
Command	A command name or command phrase, for example <code>ls -a</code> .
Filename	The name of a file or the path to a file location.
Computer output	Information displayed by the computer.
<b>Ctrl-x</b>	A key sequence that indicates you must hold down the keyboard key labeled <b>Ctrl</b> while you press the letter <code>x</code> .

ENVIRONMENT VARIABLE	The name of an environment variable, for example, <code>PATH</code> .
<b>Key</b>	The name of a keyboard key. <b>Return</b> and <b>Enter</b> both refer to the same key.
<b>Term</b>	A term or phrase that is defined in the body text of the document, not in a glossary.
<b>User input</b>	Indicates commands and text that you type exactly as shown.
<i>&lt;Replaceable&gt;</i>	The name of a placeholder that you replace with an actual value.
[ ]	In command syntax statements, these characters enclose optional content.
{ }	In command syntax statements, these characters enclose required content.
	The character that separates items in a linear list of choices.
...	Indicates that the preceding element can be repeated one or more times.
WARNING	An alert that calls attention to important information that, if not understood or followed, results in personal injury.
CAUTION	An alert that calls attention to important information that, if not understood or followed, results in data loss, data corruption, or damage to hardware or software.
IMPORTANT	An alert that calls attention to essential information.
NOTE	An alert that contains additional or supplementary information.
TIP	An alert that provides helpful information.

---

## 5 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)).CMS