

# HP Matrix Operating Environment 7.4 Update 1 Release Notes

## Abstract

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

HP Part Number: 5900-4024a  
Published: May 2015  
Edition: 2



© Copyright 2006,2015 Hewlett-Packard Development Company, L.P.

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

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.

Links to third-party websites take you outside the HP website. HP has no control and is not responsible for the information outside HP.com.

### Acknowledgments

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

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

Microsoft® and Windows® are U.S. registered trademarks of Microsoft group of companies.

Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

Red Hat® is a registered trademark of Red Hat, Inc. in the United States and other countries.

VMware®, ESX®, and ESXi™ are the registered trademarks of VMware, Inc, in the United States and other jurisdictions.

### Warranty

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

### Revision history

Document part number	Publication date	Document edition	Edition Notes
5900-4024	April 2015	1	Initial publication
5900-4024a	May 2015	2	Updated Insight Control server provisioning support related information

---

# Contents

<b>1 Description</b> .....	<b>4</b>
Update recommendation.....	4
Supersedes.....	4
Operating systems and devices.....	4
Enhancements.....	4
Prerequisites.....	12
Installation .....	12
<b>2 Issues and suggested actions</b> .....	<b>21</b>
Matrix Operating Environment Online Help.....	21
Matrix Operating Environment software.....	21
Capacity Advisor.....	25
Matrix infrastructure orchestration.....	26
HP Matrix OE visualization.....	50
HP Matrix Operating Environment recovery management.....	60
HP Insight managed system setup wizard.....	60
HP Application Discovery.....	61
Instant Capacity Manager.....	62
Process Resource Manager .....	62
HP Storage Provisioning Manager.....	62
Managed node installation.....	69
HP Global Workload Manager (gWLM).....	70
HP Insight Control server provisioning.....	70
<b>3 Documentation notes</b> .....	<b>77</b>
Addendum.....	77
Corrections.....	77
<b>4 Support and other resources</b> .....	<b>79</b>
Contacting HP.....	79
Related information.....	81
<b>5 Documentation feedback</b> .....	<b>82</b>

---

# 1 Description

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

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

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

## Update recommendation

HP recommends that customers running 7.2.x, 7.3.x, and 7.4.x versions upgrade to version 7.4.1.

---

**NOTE:** If you are running on 7.2.x or 7.3.x versions, then you need to first upgrade to 7.4 version.

---

## Supersedes

The version supersedes the 7.2.x, 7.3.x, and 7.4 released versions.

## Operating systems and devices

Supported operating systems and devices are listed in the *HP Insight Management Support Matrix* located in the [HP Enterprise Information Library](#).

## Enhancements

Matrix OE 7.4 Update 1 software provides new capabilities and expanded platform support.

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

## New features in Matrix Operating Environment

- Matrix infrastructure orchestration
  - Change in the default value for `allow.skip.os.customization`

The default value of the `hpio.properties` `allow.skip.os.customization` property has changed to `true`, which enables display of the *Skip OS Customization* option on Matrix infrastructure orchestration designer dialog used to select the operating system. This option allows 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 OE, and becomes the responsibility of the architect to ensure that the VM is configured properly. This option to skip automated OS customization can be disabled by setting the property to `false` in the `hpio.properties` file.
  - Support for booting Gen8 (and later) servers using HP Intelligent Provisioning  
When using Insight Control server provisioning 7.4, Generation 8 (and later) servers are booted using HP Intelligent Provisioning, rather than PXE boot. The "PXE-less" boot offers shorter boot times and improved security. This assumes the default iLO configuration, which listens on port 443. If you are upgrading from a prior version of HP Server Automation, see the `readme` file in the `HP System Insight Manager/sa/bin` directory for the required procedures.

---

**NOTE:** Use of HP Intelligent Provisioning may not be supported by all server configurations or firmware versions. See the *HP Insight Management Support Matrix* for information about configurations that do not support its use.

To globally disable use of Intelligent Provisioning set the `enable.hp.intelligent.provisioning.boot` property to **False** in the `hpio.properties` file.

To disable its use on a per-server basis, perform the following procedure to delete all iLO users from the system:

  1. Navigate to the **Onboard Administrator**.
  2. Select the device bay.
  3. Select the **Web Administration** link.
  4. Select **Administration**→**User Administration**.
  5. Select **All local users**.
  6. Click **Delete**.
  7. Re-run the HP Systems Insight Management discovery task for the target system or enclosure.
  8. Disable any automatic discovery tasks associated with the system.

This disables the ability to logon directly to the iLO, but the Onboard Administrator can continue to perform operations and you can logon via the Onboard Administrator.

---
- Cloud provisioning support for HP CloudSystem 8.0, 8.1, and 8.1.2 Resource Pools  
Matrix infrastructure orchestration can provision and manage servers in CloudSystem 8 resource pools in the same manner that it does for other external cloud providers, such as HP Helion Public Cloud or Amazon EC2.

---

**NOTE:** Floating IP operations and the Power Off operation are not supported by Matrix OE, use CS8 interfaces to perform these operations.

---

- For improved security, support for legacy SSLv3 encryption is removed.
- Launch Matrix OE infrastructure orchestration from SIM without having to accept the self-signed certificate using a separate URL.

Starting the Matrix OE infrastructure orchestration console from within HP SIM no longer requires the workaround to first open the URL in a separate tab to accept the self-signed certificate. Matrix infrastructure orchestration now detects if the SSL certificate needed to access the console within SIM is accepted by the web browser. If not, it provides a button to load the certificate and then loads the console.

---

**NOTE:**

- The browser must be configured to enable pop-ups for the site.
  - Failure to accept the certificate results in the display of a blank page instead of the console.
- 

- Matrix OE visualization

- Support for FCoE Preferred port speed

Matrix OE and Virtual Connect 4.40 support the use of the HP Virtual Connect 16Gb 24-Port 16Gb Fibre Channel Module for c-Class BladeSystem. This module can support 16Gb FC HBAs when used in a Generation 3 c7000 chassis. The Matrix OE visualization **Add Storage Pool Entry** user interface provides a Custom (FCoE) speed choice, which can be used to specify the 16Gb speed.

- The Matrix OE visualization storage pool entry interface for SAN Catalog Storage Entry allows the selection of a specific SPM storage template from a prepopulated list. By default the Matrix Default Storage Template is used.

- Storage Provisioning Manager
  - Change in HP SPM (Storage Provisioning Manager) Volume Capacity requirement validation

As a thin provisioned volume is filled with data, the committed capacity can exceed the volume capacity by specific amounts as determined by the thin provisioning algorithms. SPM compared the user requested capacity with the committed capacity, and this resulted in the `Volume Capacity` requirement changing to `No Match` making the service *nonconformant*. To address this issue, the SPM logic in 7.4 Update 1 is changed to validate the requested capacity against the volume's capacity versus committed capacity to keep the SPM service *conformant* (even in situations where rapid growth of a thin provisioned volume has resulted in committed capacity exceeding the volume's capacity).

---

**NOTE:** The HP 3PAR StoreServ Storage System terminology for committed capacity is `Reserved User Size`, and the volume's capacity is `Virtual Size`. The thin provisioning algorithms can result in the `Reserved User Size` being slightly larger than the `Virtual Size`. For example, the SPM user may request a 2 GB volume and select a 5 GB volume. The user requested capacity is 2 GB, the volume's capacity is 5 GB, and the committed capacity of the thin provisioned volume will grow as data is written to the volume and may exceed 5 GB by a small amount.

---
- Matrix recovery management
  - Matrix recovery management messages that were logged in `mxdomainmgr.0.log` are now redirected to `logs\lsdt_web.log` in the Matrix recovery management installation directory.
  - Matrix recovery management now supports Microsoft Hyper-V logical switches.
  - Matrix recovery management now supports ESXi vNetwork Distributed Switch (dvSwitch, vDS).
  - Matrix recovery management now supports NT LAN Manager 2 (NTLMv2) Authentication.
  - Minor GUI enhancements:
    - Added log file guidance in error and warning messages.
    - Displays the site name instead of the default text `Local Site`.
- Enhanced Hyper-V Support
  - Support for Hyper-V Cluster RDM disks

The Matrix OE visualization interface for creating storage pool entries now provides an *Enable High Availability* check box for portability groups containing Hyper-V VM Hosts. This box must be checked when creating the storage to be used by Hyper-V HA VMs using clustered Disk-Pass Through disks. It should not be used for non-HA Hyper-V VMs. Matrix OE will ensure a logical server requiring Hyper-V HA will only use storage pool entries with

the *Enable High Availability* box checked (and that storage will not be used for non-HA VMs).

- HP SPM (Storage Provisioning Manager) can be used to provision raw LUNs for Hyper-V VMs (Disk Pass-Through configurations)
  - Storage pool entries must be manually created; Matrix infrastructure orchestration auto-generation of storage pool entries is not supported.
  - Masking and unmasking are not supported (the masking check box on the storage pool entry dialog is disabled).
  - When SPM is used to provision a hypervisor as a logical server or Matrix infrastructure orchestration service and SPM created the zoning, the zone name must be manually adjusted to remove SPM from the name; otherwise, VM host zoning will be impacted when VMs are deleted.
  - When a raw LUN is added to a logical server for an existing VM, by editing the storage pool entry to include additional volumes or adding a new storage pool entry, the volume will not be visible until the logical server is deactivated and then reactivated. This is not necessary when the raw LUN is added using Matrix infrastructure orchestration.
  - Growing a raw LUN provisioned through SPM is supported at the logical server level, but is not supported by the Matrix infrastructure orchestration **Edit Disk** operation.
  - SPM does not support 3PAR Flat SAN and therefore cannot be used to provision raw LUNs in this configuration.

- Administrators can now specify the LDAP distinguished name of the Organization Unit for Windows guests on both ESXi and Hyper-V VM hosts by specifying the following sysprep file entry: `MachineObjectOU = DNS-name, LDAP-path`.

---

**NOTE:** Microsoft currently does not support this setting with Windows 2008 guests. Previously, this was supported only for Hyper-V hosts.

---

- Administrators can now specify commands to be executed the next time a Microsoft Windows system is restarted by entering one or more entries of the following form in the sysprep file: `RunOnce=command`, where `command` is the command line to be executed.

- Enhanced Erase Disk functionality

When erasing disks during `service delete` using physical targets deployed via IC server provisioning 7.4.1 or SA 10.20 (Server Automation), the erase disk process now has enhanced functionality. The new functionality includes:

- All non-shared disks can now be erased using IC server provisioning 7.4.1 or SA 10.20.
- When deleting servers from a server group that still has servers in the group, shared disks will be removed automatically from the servers being deleted before the `erase disk` is performed, protecting from unintentional data loss.
- When deleting the final server of a server group (or deleting all servers via `delete-service`) shared disks will be erased.
- Reset of the target server's BIOS system configuration to factory default settings using the `rbsureset` utility.
- PXE is no longer required for Generation 8 (or later) servers, resulting in faster request completion, and full lifecycle PXE-less deployment. PXE is still used for

servers prior to Generation 8 servers (or when servers do not use HP Intelligent Provisioning).

- In environments, where an IC server provisioning or SA OS Build Plan named ProLiant HW-Erase Server does not exist (earlier versions of IC server provisioning/SA), the previous method of erasing disks is used (with related limitations about which disks are erased and the use of PXE).
- Support for HP Server Automation 10.20
- Storage support for:
  - 3PAR OS 3.2.1 MU2
  - HP P6000 Command View Software 10.3.6

---

**NOTE:** Matrix recovery management using CLX 5.01 does not support HP P6000 Command View Software version 10.2. CLX 5.01 is bundled with Matrix recovery management 7.4 Update 1, 7.4.0b, 7.3 Update 4, and 7.2 Update 4.

---

- HP B-Series SAN Network Advisor/Brocade Network Advisor (BNA) version 12.3.4.
- Hardware support for:
  - HP ProLiant ML110 Gen9
  - HP ProLiant XL170r Gen9
  - HP ProLiant XL190r Gen9
  - HP ProLiant WS460c Gen9
  - HP ProLiant XL250a Gen9
  - HP ProLiant DL120 Gen9
  - HP ProLiant DL80 Gen9
  - HP ProLiant DL60 Gen9
  - HP ProLiant ML150 Gen9
  - HP ProLiant ML10e v2
  - HP 16 Gb 24-port Fibre Channel Module

- Automated provisioning of ESXi clusters using Insight Control server provisioning or VMware Auto Deploy

The 7.3 Update 2 release introduced support for automated provisioning of ESXi clusters using Insight Control server provisioning in addition to VMware Auto Deploy. Insight Control server provisioning or Insight Control server deployment, in conjunction with the appropriate Operations Orchestration workflows can be used as an alternative to VMware Auto Deploy for automated deployment. See the *HP Matrix infrastructure orchestration: Automated ESXi Provisioning* white paper available at [HP Enterprise Information Library](#) for more information.

- Co-existence of Matrix OE with VMware vCenter Site Recovery Manager

The Matrix OE co-exists with VMware vCenter Site Recovery Manager (SRM). Testing has been done with ESXi 5.1 and 5.5 Update 1. See the “[Capacity Advisor](#)”

(page 25) section for a summary of related Capacity Advisor limitations. For more information on these procedures see the *Matrix Co-Existence with VMware vCenter Site Recovery Management (SRM)* white paper available at the [HP Enterprise Information Library](#).

- Co-existence of Matrix OE with Microsoft Hyper-V Replica  
The Matrix OE 7.4 Update 1 release provides co-existence with Microsoft Hyper-V Replica, a feature of Hyper-V that replicates all changes on a virtual machine to a counterpart virtual machine hosted by a different server. The following considerations apply:
  - HP SIM, Insight Control virtual machine management, and Matrix have various considerations for display of the Primary and Replica VMs and their associations to VM Hosts, logical servers, and Matrix infrastructure orchestration services.
  - Disk Pass-Through disks are not supported with Microsoft Hyper-V Replica.
  - Changes to the VM attributes, such as CPU or memory, are not applied to the replica VM.

See the *Matrix Operating Environment and Microsoft Hyper-V Replica* white paper at [HP Enterprise Information Library](#) for more details.

## New sections in the document

A new chapter **Documentation notes** has been added to summarize changes to Matrix Operating Environment documents that are not being revised in this release. This information will be included in future document revisions.

## Matrix OE support in VMware VXLAN environments

Matrix OE supports VMware VXLAN environments with specific prerequisites and considerations as outlined below. Correctly configured and used, Matrix OE can provision and import VMs in an environment using VMware VXLAN. Specifically, Matrix OE is able to provision infrastructure using virtual networks from vShield Manager, configured under dvSwitches. Customers can use IO to display networks from vCenter, configure those networks in IO, assign networks to organizations, and provision infrastructure using service templates.

### Prerequisites to configure the VMware VXLAN environment

1. Ensure that all of the ESX clusters which need to be part of the VXLAN are in the same VMware vSphere datacenter. Add all the clusters of the datacenter to VXLAN Networking.
2. All hosts in a cluster must be connected to a distributed switch to enable VXLAN networking (same VDS name).
3. Use vShield Manager to configure the VXLAN scope, networking, and port groups; create a single scope with all the clusters defined to a datacenter. Consult the VMware documentation for details.
4. When adding an ESX host to vCenter, HP recommends that you configure it using the FQDN and not the IP address.

### Considerations

- Matrix infrastructure orchestration ESXi cluster provisioning does not directly support postinstallation network configuration, like VXLAN. Any network configuration must

be accomplished either manually or by means of a custom kickstart file and enhanced server customization.

- Matrix infrastructure orchestration displays all the networks which are present in the vCenter datacenter (including VXLANs from different dvSwitches which are part of different VXLAN scopes or not part of the VXLAN network scope). Proper network selection is key, as noted in the following bullet.
- Ensure service templates are created using the correct networks to avoid connectivity issues. If a service is provisioned using a VXLAN, which is not part of the network scope, the VMs will have connectivity issues. Thus, the architect defining the service in Matrix infrastructure orchestration designer needs to choose networks carefully. If the administrator or user is customizing networks at the time the service is created, the network should be chosen carefully. It is possible to disable network customization if desired.
- ESX clusters/hosts must also be chosen carefully. Clusters which are not part of VXLAN scope can also see the port groups of the VXLAN from the vCenter. The VMs must be created using an ESX host which is part of VXLAN scope, otherwise the VMs created will have network connectivity issues.
- The VXLAN solution has only been tested for IPv4.

## Matrix OE support for Hyper-V NVGRE

HP Matrix Operating Environment (Matrix OE) 7.4 Update 1 supports Microsoft Hyper-V Network Virtualization using Generic Routing Encapsulation (NVGRE) environments with specific prerequisites and considerations. Matrix OE can provision virtual machines in an environment configured with Hyper-V NVGRE. Specifically, Matrix OE is able to provision Matrix infrastructure orchestration service templates using virtual networks which are created in Microsoft SCVMM and applied on Hyper-V VM Hosts. Virtual Subnets created in Microsoft SCVMM are listed in the Matrix infrastructure orchestration Networks tab. These networks can be configured in Matrix OE based on organization requirements and services can be provisioned using service templates.

### Prerequisites

- Matrix OE supports use of Hyper-V NVGRE networking (when used with SCVMM 2012 SP1 or later).
- Configure Microsoft SCVMM 2012 SP1 or SCVMM 2012 R2 with logical networks and logical switches.
- NVGRE functionality may not work on Microsoft Windows Server 2012 managed by SCVMM 2012 SP1. The administrator needs to install Microsoft update (KB2779768) on Windows Server 2012 when it is managed in SCVMM 2012 SP1.
- Hosts that need to be part of an NVGRE network need to be part of the same Hyper-V Cluster.
- VLAN tagging must be done at the switch level.

---

**NOTE:** Do not interlink VLAN.

---

### Considerations

- Matrix OE displays all the networks which are present in SCVMM (including Standard Switches and Logical Switches that are managed in SCVMM).
- Ensure service templates are created using the correct networks to avoid connectivity issues. If a service is provisioned using a VM subnet which is not part of the network

scope, the VMs will have connectivity issues. Thus, the architect defining the service template in Matrix infrastructure orchestration designer has to choose networks carefully. If the administrator or user is customizing networks at the time the service is created, the network must be chosen carefully. It is possible to disable network customization if desired.

- If used with SCVMM 2012, without the service pack, virtual switch creation is supported without the ability to create virtual networks.

See the *HP Matrix Operating Environment support of Microsoft Hyper-V NVGRE environments* white paper for more details. This document is available from the [HP Enterprise Information Library](#).

## Prerequisites

- Hardware, firmware, and software requirements for this release are listed in the *HP Insight Management Support Matrix*.
- Installation and upgrade prerequisites are documented in the *HP Insight Management Preinstallation Worksheet*, and the *HP Insight Management Installation and Configuration Guide*.
- HP Matrix Operating Environment support of VMware VXLAN environments is documented in the *HP Matrix Operating Environment support of VMware VXLAN environments* white paper.
- HP Matrix Operating Environment support of Microsoft Hyper-V NVGRE environments is documented in the *HP Matrix Operating Environment support of Microsoft Hyper-V NVGRE environments* white paper.

These documents are available from the [HP Enterprise Information Library](#).

## Installation

Installation instructions are documented in the *HP Insight Management Installation and Configuration Guide*.

This document is available from the [HP Enterprise Information Library](#).

## Newer version available for download

### Matrix Operating Environment for HP-UX

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

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

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

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

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

## Accessing online help before licenses are installed

To access the Matrix OE online help before installing licenses:

1. Select **Help**→**For HP Systems Insight Manager** from the top menu.
2. In your browser address bar (if enabled), change the URL from `https:// {CMS-FQDN} :50000/mxhelp/mxportal/en/help-welcome.html` to

<https://{CMS-FQDN}:50000/mxhelp/vsemgmt/en/licenses.html>  
where {CMS-FQDN} is the fully qualified domain name of the CMS.

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

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

---

- ❗ **IMPORTANT:** If you are upgrading from an earlier version of HP Insight Management and you have Insight Control performance management and HP Matrix infrastructure orchestration installed and configured on your highly available CMS with SQL Always On enabled, then you may have a failure in upgrade of Matrix infrastructure orchestration. To avoid this, perform the steps as listed under Suggested Action in section [“Failure during upgrade of Matrix infrastructure orchestration 7.2.x to 7.4”](#).
- 

HP recommends to back up the earlier version of the Matrix OE environment before upgrading the software. The upgrade process is not reversible. If, after you begin, you decide not to complete the upgrade to the current version of Matrix OE, a backup copy of your earlier version of the Matrix OE environment is required to restore your CMS to the preupgrade state.

---

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

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

---

## Back up user-modified properties files before upgrade

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

### Suggested action

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

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

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

Back up the SPM state folder before performing an upgrade: `C:\Program Files (x86)\HP\Storage Provisioning Manager\state\`

After the upgrade has completed, but before the services are restarted or the CMS is rebooted, merge changes contained in the saved configuration files and add any custom entries made to the previous version of the files.

The following are the `lsa.properties` properties for which the default values have changed. If you are merging changes to the saved version, you must update these properties with the listed values:

```
CACHE_EXPIRE_TIME_OUT=259200000
UNSUPPORTED_VCDG_CAPABILITIES=PARTIALLYSTACKEDDOMAIN

# ESX VERSION PROPERTIES
ESX_SUPPORTED_VERSIONS=4.0.*,4.1.*,5.0.*,5.1.*,5.5.*,6.0.*
ESXi_SUPPORTED_VERSIONS=4.0.*,4.1.*,5.0.*,5.1.*,5.5.*,6.0.*
ESXi_PATTERN=.*ESXi.*,.*Embedded.*, (?i).*VMware ESX( Server 3)?i.*
ESX_DEFAULT_MAX_HOST_VERSION=6.0
ESX_COMPATIBLE_VERSIONS=4.4.0:4.1:5.0:5.1:5.5:6.0,7:4.0:4.1:5.0:5.1:5.5:6.0,8:5.0:5.1:5.5:6.0,9:5.1:5.5:6.0,10:5.5:6.0,11:6.0
# Maximum vmVersion defined for ESX5.5.*
MAX_DEFAULT_ESX_VM_VERSION=9
ESX_VM_VERSION=4,7,8,9,10,11
```

The following properties are added to the `lsa.properties` file. If you are merging the changes with the saved version, add the following lines to the end of the file:

```
VMHOST_RESCAN_WAIT_INTERVAL=15000
# Supported SSL protocols
ENABLED_SSL_PROTOCOLS=TLSv1,TLSv1.1,TLSv1.2
# Supported SSL cipher suites
ENABLED_SSL_CIPHER_SUITES=TLS_RSA_WITH_AES_128_CBC_SHA,TLS_DHE_RSA_WITH_AES_128_CBC_SHA,TLS_RSA_WITH_AES_128_CBC_SHA256,
TLS_DHE_RSA_WITH_AES_128_CBC_SHA256
```

## Allocated storage may not be visible if SPP is not installed

To ensure that allocated storage is visible when deploying an OS, ensure the SPP is installed. For more information, see the SPP home page at the <http://www.hp.com/go/spp>. For more documentation about the product, see the SPP documentation page at [HP Enterprise Information Library](#).

## Support for a rolling upgrade of a CMS federation

To decrease federation downtime, Matrix OE infrastructure orchestration 7.4 Update 1 includes support for a rolling upgrade of a federated CMS environment. A rolling upgrade temporarily allows a federation containing CMSs running different versions of Matrix OE during the upgrade process.

---

❗ **IMPORTANT:** The primary CMS must be upgraded first.

---

During a rolling upgrade, a CMS running Matrix OE 7.4 Update 1 supports basic operations against a secondary CMS running 7.4. Not all 7.4 Update 1 features are expected to work against a 7.4 CMS.

The following operations are not supported on a primary CMS running Matrix OE 7.4 Update 1 with secondary CMSs running 7.4:

- Manual operating system provisioning of VM guests using resources discovered on a secondary CMS.
- Provisioning a physical server group with auto-provisioned SPM volumes on a secondary CMS.
- Some chargeback operations, including detection of changes made outside of Matrix OE (and directly in VMware, Hyper-V, OA, and so on).
- Booting Generation 8 (and later) servers using HP Intelligent Provisioning.
- Matrix infrastructure orchestration
  - Increased disk size of a virtual or physical disk
  - Detection of disk growth outside of Matrix
  - Improved storage tag granularity

- Supports Hyper-V Generation 2 VMs
- Supports Hyper-V logical switches
- Matrix OE visualization
  - Supports FCoE Preferred port speed.
  - Enhanced logical server storage resource relationships.
  - LSMUTIL enhancements.

## Check logical server management upgrade log file

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

### Suggested action

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

## Changing the CMS password

Use the `mxcpassword` command to change the CMS password. For more information on using the `mxcpassword` command see the *HP Systems Insight Manager User Guide* at <http://www.hp.com/go/insightmanagement>.

## HP Insight Control server provisioning

---

**NOTE:** Information related to HP Insight Control server provisioning 7.4.1 is available only in the English version of the *HP Matrix Operating Environment 7.4 Update 1 Release Notes* and not in other languages.

---

### Installation considerations

Each version of Matrix OE supports a specific release of the SPP. When using a version of IC server provisioning whose version number does not match the Matrix OE version, you will need to modify your Build Plans to use driver packages and SPP media that matches the version that Matrix OE supports.

The [Table 1 \(page 15\)](#) lists the SPP versions that go with each Matrix OE version, and the IC server provisioning versions that are supported with these.

**Table 1 Matrix OE and IC server provisioning supported versions**

Matrix OE version	Supported IC server provisioning version	Notes
7.2	7.2	Must use SPP version 2013.02.0
7.2.1	7.2.1	Must use SPP version 2013.02.0
7.2.2	7.2.2	Must use SPP version 2013.09B.0 or 2013.09C.0
7.3	7.2.2	Must use SPP version 2013.09B.0 or 2013.09C.0
7.3.1	7.3.1, 7.4.1	Must use SPP version 2014.02.0 or 2014.02b.0
7.3.2	7.3.1, 7.3.2, 7.4, 7.4.1	Must use SPP version 2014.06.0

**Table 1 Matrix OE and IC server provisioning supported versions** (continued)

Matrix OE version	Supported IC server provisioning version	Notes
7.4	7.4, 7.4.1	Must use SPP version 2014.09.0
7.4.1	7.4, 7.4.1	Must use SPP version 2015.03.0

**Table 2 Supported update paths for Insight Control server provisioning**

	To 7.2.1	To 7.2.2	To 7.3.1	To 7.3.1a*	To 7.3.2	To 7.3.2a*	To 7.4	To 7.4.0a*	To 7.4.1
<b>From 7.2</b>	Yes	Yes	No	No	No	No	No	No	No
<b>From 7.2.1*</b>	NA	Yes	No	No	No	No	No	No	No
<b>From 7.2.2</b>	NA	NA	Yes	No	No	No	No	No	No
<b>From 7.3.1</b>	NA	NA	NA	Yes	No	No	Yes	No	No
<b>From 7.3.1a*</b>	NA	NA	NA	NA	No	No	Yes	No	No
<b>From 7.3.2</b>	NA	NA	NA	NA	NA	Yes	Yes	No	No
<b>From 7.3.2a*</b>	NA	NA	NA	NA	NA	NA	Yes	No	No
<b>From 7.4</b>	NA	NA	NA	NA	NA	NA	NA	Yes	Yes
<b>From 7.4.0a*</b>	NA	NA	NA	NA	NA	NA	NA	NA	Yes

\* These versions of IC server provisioning are updates only and cannot be used to install as a new appliance. You must first install the correct older version and then apply the update.

ⓘ **IMPORTANT:** For Insight Control server provisioning issues that could impact Matrix Operating Environment, see the *HP Insight Control Release Notes*.

### Deployment server configuration does not support domain names

When installing Matrix infrastructure orchestration, the deployment server configuration screen does not have an explicit *Domain* field and does not support domain specification within the user name (For example, does not support the `domain\user` syntax). If Insight Control server deployment was installed with a domain user, complete the Matrix installation with the simple user name, and then use the `mxnodesecurity` command to add the domain to the user.

1. Remove the existing entry: `mxnodesecurity -r -p DSC_RDP -n <IP_ADDRESS>`.
2. Create a new entry for the domain user: `mxnodesecurity -a -p DSC_RDP -c <DOMAIN>\<USERNAME>:<PASSWORD>`.

## Operations Orchestration Studio installation and upgrade issues

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

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

To resolve this issue:

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

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

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

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

### Upgrading Operation Orchestration 10.01 content to 10.10

Operation orchestration projects created by OO Studio 10.01 are converted to OO 10.10 format by OO Studio 10.10, but the converted package of IO ESA flows supplied by Matrix infrastructure orchestration is not.

If any of the IO or ESA (Extensible Server Adapter) flows were edited under OO 10.01 Studio, follow these instructions to convert them to OO 10.10 format:

---

**NOTE:** These steps are needed only if IO or ESA flows were directly changed under OO Studio 10.01. If IO or ESA flows were used (dragged and dropped) by your flows, but not changed, replace the IO ESA flows with the new version provided by the HP Matrix infrastructure orchestration 7.4 installation.

---

1. Prepare `oo10-io_esa` project:
  - a. Create a backup of `c:\Users\<username>\.oo` folder. This folder stores the OO 10 Studio projects.
  - b. In the `oo10-io_esa` project folder, right click on OO Studio and select Show in Explorer.
  - c. Navigate to `Content/Configuration/System Accounts` folder.
  - d. Open `HpioSmtxCredentials.xml` and `HpioCmsCredentials.xml` files using a text editor.
  - e. Replace the token string of `HpioSmtxCredentials.xml` with the token found in the `HpioCmsCredentials.xml` file.
2. Upgrade Operations Orchestration Studio to version 10.10.
3. Open the studio, open all projects, and remove the old out-of-the-box content pack from the Dependencies pane and add a new one.
4. Fix any incompatibilities identified by diagnostic messages:

- a. Most problems are duplicated configuration items found on configuration tree of oo10-io\_esa project. To fix the issues, inspect all items in the `Configuration tree` folder. Click **Delete** and **Confirm**. OO Studio deletes the duplicated internal reference and the item will not appear as an error.
  - b. Run the OO Studio verification tool under **Tools/Validate Flows and Operations** menu to check.
5. Package the converted content pack.
  6. Deploy the converted content pack.

---

**NOTE:** The following white papers provide detailed information and considerations when upgrading the Operations Orchestration included with Matrix OE:

- *Upgrading HP Operations Orchestration included with Matrix OE to HP Operations Orchestration 9.00 (Full Version)*
  - *Upgrading HP Operations Orchestration included with Matrix OE to HP Operations Orchestration 10.x*
- 

## Managing Microsoft Windows servers without OpenSSH installed

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

For details on using Matrix OE with Microsoft Windows systems not running OpenSSH, see the *HP Matrix Operating Environment 7.0 software without OpenSSH on HP ProLiant and Integrity Servers Running Microsoft Windows* white paper, available at the following website:

[HP Enterprise Information Library](#)

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

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

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

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

[HP Enterprise Information Library](#)

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

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

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

### Suggested action

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

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

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

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

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

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

## User experience and application performance

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

Every effort to balance performance and functionality is made to the default settings to create the best user experience under multiple browsers with different server environments. Changing these default settings can result in adverse application responsiveness, which is more noticeable in some browsers; for example, there are

known memory issues in Internet Explorer 8 that can be exposed in large scale environments. No matter which browser you choose to use, HP recommends using the latest release and updates to improve your user experience and protect your environment from the latest security threats.

## Additional software dependencies

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

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

---

## 2 Issues and suggested actions

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

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

### Matrix Operating Environment Online Help

#### Limitations

When a search results in no hits, the system returns the stemmer rather than the search term in the A search of Performance returns a result of Perform panel

For example, the term *Orchestration* stems to *orchestr* and the results show as Your search returned no results for *orchestr*. This is expected behavior.

### Matrix Operating Environment software

#### Limitations

##### HP-UX support of IPv6

HP-UX and cloud servers do not support IPv6. See the *IPv6 Support in Insight Management* white paper at [HP Enterprise Information Library](#) for more information.

##### Supporting HP IO Accelerators

HP IO Accelerators are supported in the Matrix Operating Environment with limitations. The hypervisor version and Virtual Connect modules need to support the particular HP IO Accelerator. Matrix OE will not provision the hypervisor host to use the HP IO Accelerators (the installation and configuration of the hypervisor using these accelerators would be independent of Matrix OE). Matrix OE can be used to provision VMs to the hypervisor host. For information on HP IO Accelerators, see the [http://h18004.www1.hp.com/products/storageworks/io\\_accelerator/index.html](http://h18004.www1.hp.com/products/storageworks/io_accelerator/index.html). Three 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, the *Insight Management and HP IO Accelerators (ioTurbine version 2.1) — Supported Environments and Scenarios* white paper and the *HP Insight Management and IO Accelerators (ioTurbine version 2.2.0 or later) — Supported Environments and Scenarios* white paper at the [HP Enterprise Information Library](#).

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

As Windows Single Sign On is used by Matrix OE to connect to the SQL Server database, the credentials associated with the user account executing a command or service are used to connect to the remote database. Therefore, the user accounts that

are used to execute Matrix OE software that accesses the database must satisfy the following conditions:

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

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

## Operations Orchestration intended usage in Matrix OE

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

### Functionality

The limited version of Operations Orchestration in Matrix OE:

- Allows Matrix OE infrastructure orchestration to trigger Operations Orchestration workflows for pre and post provisioning tasks. Examples include Create Service/Add Server, Standby Server/Service, and Resume Server/Service. Workflows in the limited version of Operations Orchestration included with Matrix OE may only be triggered from the Matrix OE infrastructure orchestration interface in relation to pre- and post- provisioning tasks. An upgrade to the full version of Operations Orchestration allows the launch of workflows from outside of the Matrix OE interface.
- Allows Matrix OE to be integrated with other management software. However, this requires additional services from HP consulting.
- Is licensed for one author in OO Studio. Upgrades to the full version of Operations Orchestration allow for multiple authors.
- Is included in Matrix OE may only be used for flow troubleshooting. An upgrade to the full version of Operations Orchestration allows the use of Operations Orchestration Central to launch and schedule the launch of Operations Orchestration workflows from outside of Matrix OE/Insight Control.

### Content

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

## Provisioning of Hyper-V Generation 2 VMs is not supported by Insight Control server provisioning 7.4 or HP Server Automation

Generation 2 Hyper-V VMs must either be provisioned manually by using a Hyper-V template or with Insight Control server provisioning 7.4.1.

---

**NOTE:** The default Matrix OE visualization configuration, used by Matrix OE infrastructure orchestration, specifies the use of Generation 2 VMs when supporting hosts are present. This results in unexpected deployment failures. The configuration can be changed to always use Generation 1 VMs by editing the `Virtual Server Environment/conf/lisa.properties` file and changing the value of the `MAX_DEFAULT_HYPERV_VM_VERSION` property from 2 to 1. This changes only the default values; logical servers can be manually created with an explicit selection of Generation 2.

---

## Using SPM to provision RDM disks can cause VM Hosts to lose their zoning

When Matrix infrastructure orchestration or Matrix OE visualization with logical server management is used to provision a VM Host with disks provisioned using SPM (including automated zoning), use of SPM to provision RDM disks for a VM can result in the disks for the VM Host becoming unzoned and therefore is inaccessible to the VM Host. This potentially impacts all hosted VMs and the VM Host itself. This problem can occur with one CMS managing both VM Hosts and VM guests, or with one CMS managing the VM Host(s) and a second CMS managing the VM guests (but not with one storage volume managed by more than one CMS).

---

**NOTE:** This problem can occur only when the BNA managed switches are used and SPM is configured for automated zoning.

### Cause

SPM relies on the zone name to determine whether it created the zone and thus is responsible for cleaning it up. When a VM Host is initially created using SPM provisioned disks, SPM will create zones for the VM Host initiators. Later, when using SPM to provision RDM storage for VMs, these same zones are used. In the above scenario, when the VM RDM disks are removed, SPM may not realize that the zone is still being used by the VM Host itself and remove the zone when removing the storage.

### Suggested action

To prevent this from occurring, manually change the zone name for the VM Host and ensure that it does not include `SPM_` in the name. Do this prior to using SPM to provision RDM storage for VMs.

---

## VCM server profile assignment may fail for Gen9 servers

Matrix OE infrastructure orchestration service provisioning fails when it attempts to assign a VCEM profile to a Gen9 server. The VCEM log indicates the following error: Error during Assign server profile <profile> in domain <domain>. Details: VCEM received an unexpected error from the VC Manager.

### Suggested action

If this problem occurs, physically reset the server or use the Onboard Administrator to reset the server and then reapply the profile. If the problem persists, clear the server BIOS settings. To avoid this problem, change the value of the `GRACEFUL_POWEROFF` property from **False** to **True** in the `Virtual Server Environment/conf/lisa.properties` file.

## Null Pointer Exception Reported when OO Workflow Sends Email

A Null pointer exception is reported when an Operations Orchestration workflow attempts to send an email. This is because, the *HpioSmtxCredentials* System Account user name field has an entry, but the password field does not. This may be the case when the SMTP server does not require authentication.

### Suggested action

Ensure that both fields have values. Specify dummy data if necessary. These fields are not used when the SMTP server does not require authentication.

## Service creation fails during SA operating system deployment

Operating system deployment by HP SA (Server Automation) fails with the following diagnostics:

Failed while deploying operating system for service <server name>.

**Cause:** Logical server job (ID = <server ID>) completed with a failure status.

**Failure:** The requested operation is not allowed in the current life cycle state of the server according to the deployment server. One possible reason is that a node with the same MAC address or UUID exists on the deployment server.

### Suggested action

Upgrade to a version of HP Server Automation that supports HP Intelligent Provisioning or disable the use of HP Intelligent Provisioning by setting the *hp.intelligent.provisioning.boot* property to **False** in the *hpio.properties* file.

## 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 postinstallation 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:

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.

# Capacity Advisor

## Limitations

### Capacity Advisor provides support for the HP ProLiant WS460c workstations

Capacity Advisor is able to collect historical data from HP ProLiant WS460c after following the steps described in the *How to license HP ProLiant WS460c for Capacity Advisor* white paper located at the [HP Enterprise Information Library](#).

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

Capacity Advisor supports Linux servers, but does not support collecting data from KVM VMs that are 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 no drop-down menu to see 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.

### Capacity Advisor limitations with VSP

In the Scenario Editor, the only VSP systems available are VM Hosts that are only hosting VMs (See [Table 3](#)). In the Capacity Analysis dashboard the only VSP systems available are VM guest and vPars.

**Table 3 VSP systems availability for the Capacity Advisor Scenario Editor**

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 limitations with VMware SRM

With configurations that contain two CMSs, manual steps must be taken to capture the data collected while running on the secondary site and then to merge it with the data on the primary site. There are documented procedures available to minimize and potentially eliminate this from happening. For more information on these procedures see the *Matrix Co-Existence with VMware vCenter Site Recovery Management (SRM)* white paper available at the [HP Enterprise Information Library](#).

### Capacity Advisor limitations with Microsoft Hyper-V Replica environments

In environments using Matrix co-existing with Microsoft Hyper-V Replica, when DHCP is enabled, the Replica VM in Replica site gets its IP address from DHCP which is different from its IP address in the Primary site. In such a situation, the Replica VM must be manually discovered with the new IP address (DHCP assigned) in HP SIM. After this step, the Replica VM works fine with Capacity Advisor.

When DHCP is disabled and a static IP address is assigned to the VM in Primary site, the Replica VM gets the same IP address and is discovered automatically without any issue in HP SIM and works fine with Capacity Advisor.

## Minor issues

### Missing raw data for a time period

Capacity Analysis sometimes encounters missing data when it tries to compute statistics for a day, week, or month. When no data is collected for a day, the average CPU utilization for that day is stored as “unknown” and is displayed as a dash. If all of the data has been collected for a day, the average can be computed and stored. However, if data has been collected for only part of a day, then Capacity Analysis can report the average of the data that was collected, or report that the average is unknown. In this release of Capacity Analysis, the rules for displaying a dash or a value based on partial data are not well defined and can vary from metric to metric.

#### **Suggested action**

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

## Matrix infrastructure orchestration

Matrix infrastructure orchestration runs on an x64 Windows CMS.

### OS Customization may fail when a Matrix service template based on ESXi Linux VM template is customized to use a Hyper-V Windows VM template

HP SIM can be used to create a Hyper-V VM template from a powered off Hyper-V hosted Windows VM. If a Matrix infrastructure orchestration template, that is based on an ESXi Linux VM template, is later customized to use the Hyper-V VM template, service provisioning fails at the OS customization step. The software inventory listing for the Hyper-V VM template shows empty OS Name and OS Type fields.

#### **Suggested action**

Prior to using the SIM **Deploy**→**Virtual Machine**→**Templates**→**Create Virtual Machine Template...** command to create the Hyper-V VM template:

- Ensure that the Hyper-V hosted Windows VM is powered on.
- Ensure that Microsoft Hyper-V Integration Services is installed and running.

The resulting Hyper-V VM template will have a non-empty OS Name and OS Type listed in the software inventory and can be successfully provisioned.

### ESXi Linux VM customization fails when two network adapters have different DNS suffixes

If you create an ESXi Linux VM with two network adapters and those adapters have different DNS suffixes, the server will always join the AD Domain using the suffix of the primary network adapter.

#### **Suggested action**

Ensure the primary network adapter has appropriate DNS suffix (or have both network adapters use the same DNS suffix).

### VM network state is disconnected after ESXi Virtual provisioning

If the ESXi VM used to create a VM template has a network adapter with no connectivity, any VM created from that template will not have network connectivity. This is a VMware related issue.

### Suggested action

Ensure the network adapters for the VM are connected before creating the VM template.

## Matrix API and CLI add network interface operation do not support customization options

In version 8 of the Matrix API the `addNetworkInterfaceToLogicalServerGroup` operation defines an optional `customizeNetwork` parameter. This capability currently is not supported. The parameter should be omitted or a value of `false` specified. Specifying a value of `true` causes the request to fail.

### Suggested action

The CLI `addnic logicalservergroup` operation defines corresponding `-customize-operating-system` and `-O` options. These options must not be used as both causes the request to fail.

## Floating IP Address operations are not supported from Matrix OE with CloudSystem 8.0, 8.1, and 8.1.2 hosted servers

The Floating IP Address operations are not supported from Matrix with CS 8.0, 8.1, and 8.1.2 hosted servers.

### Suggested action

Such operations should be performed through the CloudSystem console.

## 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” \(page 27\)](#) for information on allowing Linux VMs on Hyper-V.

---

## Skipping automated OS customization

By default, after Matrix 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.

To skip automated OS customization:

### 1. 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.

---

**NOTE:** The per-server-group **Skip OS Customization** check box is enabled when `allow.skip.os.customization = true` in `hpio.properties` (this is the default value).

If you do not see the **Skip OS Customization** check box, edit the `hpio.properties` file to set `allow.skip.os.customization = true`.

---

2. Ensure appropriate OS customization is done (perhaps via an Operations Orchestration workflow attached to the service template or selecting an OS Build Plan modified to use enhanced server customization).

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

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

### Suggested action

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

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

Fabrics separated by ":" are redundant pairings.

For information on defining prefixes, see ["Define prefix used for host names on the storage array when provisioning storage"](#) (page 28).

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

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

---

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

---

## Define prefix used for host names on the storage array when provisioning storage

When provisioning storage through Matrix IO or Matrix OE, using SPM, a prefix can be defined for the host names that are created on the storage array. This can be defined by adding or modifying the property `SPM_HOSTNAME_PREFIX` within the `lsa.properties` file. This prefix can be five characters or less due to host name limitations on the storage arrays. If not selected, the default prefix used is "Matrix OE visualization". The rest of the host name will be the initiator WWN defined within the host definition to guarantee name uniqueness.

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

When deploying HP-UX 11iv3 on a BL860c i2, BL870c i2, BL860c i4, BL 870c i4, or BL890c i4 server blade using Matrix infrastructure orchestration with Ignite-UX, the

Create Service Request may fail and the server will not be provisioned. The step in the request that deploys the HP-UX OS fails and the following error displays:

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

### Suggested action

1. Erase the disk used by the target server for deployment.
  - a. Make sure all hosts are in the same zone then perform an LUN clean-up: Log in to the storage management server (HP P6000 Command View Software or HP 3PAR StoreServ Management Console, and so on).
  - b. Unpresent the host associated with the vdisks and delete the vdisks.
  - c. Create the new vdisks with the required size and RAID level and present the vdisks to the respective hosts.
2. Deactivate the logical server associated with the target server.
  - a. From the HP Matrix Operating Environment menu, click **Tools**→**HP Matrix OE Visualization**.
  - b. Select **Logical Server** from the **Perspective** drop-down menu.
  - c. Select the logical server and click **Tools**→**Logical Servers**→**Deactivate**.
  - d. Monitor the status of the task by selecting **Report**→**Logical Server Job Status**.
  - e. Wait for the task to complete successfully before continuing to the next step.
3. Delete the logical server. Select the appropriate deactivated logical server and click **Delete**→**Delete Logical Server**.
4. Move the target server from the IO maintenance pool to the desired server pool.
  - a. From the HP Matrix Operating Environment menu, click **Tools**→**Infrastructure Orchestration**.
  - b. Click the Servers tab. Make sure the servers are not being used by checking the "Usage" column.
  - c. Move the servers under the "Maintenance" pool to the desired server pool by using the "Modify Pool" option.
5. Recreate the service request.

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

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

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

```
Automatic boot entry execution will start in 1 second.  
Booting VC Pri. Fibre(0x5001438004C8E7B8,0x10000000000000)  
Boot Failed. VC Pri. Fibre(0x5001438004C8E7B8,0x100000000000000)  
Booting HP-UX Primary Boot:  
0/0/0/3/0/0/2.0x5001438004c8e7b9.0x4001000000000000  
Error - efi_disk_read returned (Device Error)
```

### Suggested action

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

1. Erase the disk used by the target server for deployment.
  - a. Make sure all hosts are in the same zone then perform an LUN clean-up: Log in to the storage management server (HP P6000 Command View Software or HP 3PAR StoreServ Management Console, and so on).
  - b. Unpresent the host associated with the vdisks and delete the vdisks.
  - c. Create the new vdisks with the required size and RAID level and present the vdisks to the respective hosts.
2. Deactivate the logical server associated with the target server.
  - a. From the HP Matrix Operating Environment menu, click **Tools**→**HP Matrix OE Visualization**.
  - b. Select **Logical Server** from the **Perspective** drop-down menu.
  - c. Select the logical server.
  - d. Click **Tools**→**Logical Servers**→**Deactivate**, and then click the **Deactivate** button to continue deactivating the logical server.
  - e. Monitor the status of the task by selecting **Report**→**Logical Server Job Status**.
  - f. Wait for the task to complete successfully before continuing to the next step.
3. Delete the logical server.
  - a. Select the appropriate deactivated logical server.
  - b. Click **Delete**→**Delete Logical Server**, and then click the **Delete** button to continue deleting the logical server.
4. Move the target server from the IO maintenance pool to the desired server pool.
  - a. From the HP Matrix Operating Environment menu, click **Tools**→**Infrastructure Orchestration**.
  - b. Click the Servers tab. Make sure the servers are not being used by checking the "Usage" column.
  - c. Move the servers under the "Maintenance" pool to the desired server pool by using the "Modify Pool" option.
5. Recreate the service request.

## Provisioning failure on Integrity blades BL870c i2, BL890c i2, BL870c i4, and BL890c i4

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

---

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

---

### Suggested action

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

---

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

---

2. Set the `serverboot.physical.wait.seconds` and `INTEGRITY_POWERON_WAITTIME` according to the time measured in step 1. Use the following parameters as a minimum.
  - For BL870c i2 and i4 blades:
    - Set the `serverboot.physical.wait.seconds` to 540 seconds.
    - Set the `INTEGRITY_POWERON_WAITTIME` to 540000 milliseconds.
  - For BL890c i2 and i4 blades:
    - Set the `serverboot.physical.wait.seconds` to 1080 seconds.
    - Set the `INTEGRITY_POWERON_WAITTIME` to 1080000 milliseconds.

## Some Virtual Connect profiles and domain group capabilities are not supported by specific VC blades or by the Matrix Operating Environment

Some Virtual Connect profiles and Domain Group capabilities are not supported. If you use Matrix infrastructure orchestration attempts to provision a service on a blade with an unsupported capability, the creation of the profile fails.

### Suggested Action

When adding servers to server pools, ensure to add only blades with VC capabilities that are supported by Matrix OE. Also, ensure to group VC blades with similar capabilities and connectivity within the same server pools. Different LOM and mezzanine cards could cause different VC blades to have different connectivity.

The following VC Domain Group capability is not supported by Matrix OE:

- Partially Stacked Domains.

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

- 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.

## Server Automation 9.x support

Matrix infrastructure orchestration 7.4 update 1 can be used with 9.x versions of Server Automation but requires additional files and configuration. For more information, see the `readme.txt` file in the `bin\SA` sub-directory of the SIM installation directory.

## Limitations

### Trunk network considerations

You cannot make a new network connection, if the selected network is already connected to a logical server group or a trunk attached to a logical server group. This limitation may not be enforced by all versions of Matrix infrastructure orchestration.

### Domain name discrepancies result in Matrix infrastructure orchestration VM hosts listed in the Unassigned server pool and failure to provision services to the host

If the Active Directory Domain name and DNS name of a VM Host (ESXi, Hyper-V or Integrity) do not match, VMM (Insight Control virtual machine management) will show the hostname (i.e. FQDN) as the DNS Name with the domain name appended. Matrix infrastructure orchestration will list the VM Host in the Unassigned server pool and will fail to provision services to the VM Host.

#### **Suggested action**

To resolve this issue, change the name in AD or DNS name, so that both names are the same. The VM Host will then appear in the appropriate server pool. An administrator can also manually move the VM Host from the Unassigned pool to the appropriate pool. The service provisioning request can then be repeated and can use the VM Host as appropriate.

### SCVMM shows duplicate VM entries in the SCVMM console

While associating a VM with a logical switch on SCVMM 2012 SP1 and SCVMM 2012 R2, the SCVMM console creates two VMs with same name and pointing to the same VM files.

This is an SCVMM console issue.

#### **Suggested action**

Ignore such duplicate instances created in SCVMM consoles. The Matrix Operating Environment shows only one correct VM instance.

### A SCVMM server cannot be managed by multiple CMSs

Attempts to register a SCVMM server with multiple CMSs may fail. Additionally, API calls may fail with unexpected errors with such configurations.

This is a VMM-related issue.

#### **Suggested action**

Do not register a SCVMM server with multiple CMSs.

### VMM upgrade from one version to another, makes the SCVMM VME status as "Disconnected" in VMM

If the user upgrades VMM from one version to another, VMM sets the VME status as "Disconnected".

#### **Suggested action**

Unregister and reregister the VME.

### SCVMM console will show refresh VM task in the logs folder

The SCVMM console shows some number of refresh VM tasks in the logs folder. Sometimes in Insight Control virtual machine management (VMM) you will notice differences in returned instances from WMI and PowerShell. Due to such discrepancies, VMM calls the refresh VM task explicitly in both places to maintain consistency.

### Suggested action

Ignore the refresh VM tasks in the SCVMM console.

## Network performance data is not collected in some Hyper-V 2012 R2 NIC teaming configurations

Hyper-V 2012 R2 NIC teaming enabled VMs that have not specified “**Allow management operating system to share this network adapter**” as part of Virtual Switch configuration fail to collect performance data for the network.

### Suggested action

This is a Microsoft issue. A WMI instance is not created when the **Allow management operating system to share this network adapter** option is not selected. You must select this option to collect performance data.

## Domain name not set when VMware Linux guest has a DNS suffix

If you specify a DNS suffix for a Linux VM residing on a VMware host, it fails to set the domain name. It sets the DNS suffix as the domain name.

This is a VMware related issue.

### Suggested action

Do not specify the DNS suffix while defining the Matrix infrastructure orchestration network. Manually set the DNS suffix for the Linux VM residing on the VMware host.

## 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 Matrix 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).

## 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 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 7.4 System* white paper available at the [HP Enterprise Information Library](#).

### 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 7.4 System* white paper available at the [HP Enterprise Information Library](#).

## 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:

---

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

---

- Integrity VM disks must 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. They cannot be spread across different data stores.
  - 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` properly from true to false. Setting this property to false also means that the users can specify storage volume names which are not part of the service template.

---

## Shared disks and some private data disks not automatically erased

By default, when an Insight Control server provisioning or Server Automation OS Build Plan named ProLiant HW - Erase Server exists it is used to erase disks according to the described behavior in the "Enhancements" (page 4) section under *Enhanced Erase Disk functionality*. When the preceding criteria is not met the suggested action remains valid.

### **Suggested action**

The administrator is responsible for ensuring the data security, with careful consideration of security implications. Erasing shared disks and private data disks defined in different storage pool entries can be done when necessary (after ensuring that all servers using those disks no longer require the information).

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

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

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

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

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

#### **Suggested action**

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

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

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

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

#### **Suggested action**

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

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

## Multi-initiator NPIV support on Integrity blades

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

---

**NOTE:** NPIV support has been added for Emulex CNA cards in HP-UX vPars and Integrity VM version 6.3. The host must have WEB1310 release (B.11.31.1310) or later versions of the Emulex FCoE driver installed. For details see, the *HP-UX vPars and Integrity VM V6.3 Release Notes* at <http://www.hp.com/go/hpux-hpvm-docs>.

---

### Suggested action

On Integrity blades with Emulex controllers and/or versions of HP-UX earlier than 11i v3 March 2013 update release or HP-UX vPars and Integrity VM prior to version 6.3, 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:

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

## Required software versions for FlexFabric Integrity physical provisioning support

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

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

## Microsoft Windows 2003 guest fails to join the domain

This is a Microsoft related issue, where the Microsoft Windows 2003 guest fails to join the domain at the first time but joins the domain after the VM is reset.

### Suggested action

To solve this issue, you must reset the VM on first failure.

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

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

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

### Suggested action

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

of Insight Management 7.2 Update 2, 7.2 Update 4, 7.3, 7.3 Update 1, 7.3 Update 2, 7.3, Update 4, 7.4, 7.4.0b, or 7.4 Update 1.

## Sysprep SkipRearm Consideration

If you anticipate running **sysprep** multiple times on a single system, you must follow the Microsoft recommendation of using the **SkipRearm** setting of the **Microsoft-Windows-Security-Licensing-SLC** component to postpone resetting of the activation clock. Failure to use this setting will result in the image becoming inoperable after the activation clock has been reset three times. This does not apply to activated KMS (Key Managed Service) clients, which do not limit the number of times the activation clock can be reset. For more information, see [http://technet.microsoft.com/en-us/library/cc766514\(v=ws.10\).aspx](http://technet.microsoft.com/en-us/library/cc766514(v=ws.10).aspx).

**sysprep** may be invoked as part of the Matrix infrastructure orchestration post-install OS customization process (used when automatic software deployment is used) and after Matrix infrastructure orchestration is used to add or remove a network interface.

## Microsoft Windows 2012 and Microsoft Windows 2012 R2 guest NIC customization fails for static IPv4 on Microsoft Windows 2012 R2 host

If you try to sysprep a Windows Server 2012 or Windows Server 2012 R2 running as a guest on a Windows 2012 R2 Hyper-V host, the provided `sysprep.xml` file from VMM does not set the IPv4 address properly. But, it sets the IPv6 address, and IPv4 DNS and gateway addresses.

### Suggested action

If the VM has a static IP set before sysprep is run, this issue does not occur. Hence, just setting a generic fake static IPv4 address prior to sysprep prevents this issue. If not, for Windows 8 Client and Windows Server 2012, **General Availability Update Rollup** needs to be installed in the VM before it is syspreped to prevent this issue. For more information on this see, the KB 2756872 (<http://support.microsoft.com/kb/2756872>) and link to download the updates (<http://www.microsoft.com/en-us/download/details.aspx?id=34943>). As part of customization, if you create a guest of Microsoft Windows 2012 on Hyper-V 2012 R2 Host and if you try to sysprep the guest using the command `C:\Windows\System32\Sysprep\Sysprep.exe` on the guest VM, it fails on reboot with the error as `vmbus.sys` couldn't be loaded. To resolve this issue, the above Windows 8 Client and Windows Server 2012 General Availability Update Rollup needs to be installed in the VM before it is syspreped.

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

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

### Suggested action

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

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

## High Availability Hyper-V VMs with pass-through RDM cluster disk considerations

High Availability Hyper-V VMs with nonclustered pass-through/RDM disks and non-High Availability Hyper-V VMs with clustered pass-through/RDM disks are not supported in this release and prior releases. Matrix OE supports only High Availability Hyper-V VMs with clustered pass-through/RDM disks and non-High Availability Hyper-V VMs with non-clustered pass-through/RDM disks.

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

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

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

### **Suggested action**

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

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

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

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

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

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

### **Suggested action**

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

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

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

### Suggested action 1

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

### Suggested action 2

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

## Microsoft SCVMM can be registered on only one CMS

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

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

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

### Suggested action

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

## Server power off operation is not supported for HP Cloud or next-generation HP CloudSystem resource pool hosted servers

In Matrix infrastructure orchestration, the **power off servers in group** operation is not supported for servers that are provisioned to HP CloudSystem 8. The request fails with an error indicating that the operation is not available in the cloud capacity pool.

If a server is powered off through the cloud provider's user interface, Matrix infrastructure orchestration does not recognize the resulting state. The state of such servers is listed as `Unknown`, which prevents Matrix infrastructure orchestration from performing any operation on these servers.

### Suggested action

Power the servers on or off from the cloud provider's user interface.

## Matrix OE does not support configurations other than Full Stacking

Virtual Connect 4.30 and later support Partially Stacked mode (Horizontal Stacking and Primary Slice Stacking), but Matrix OE does not support such configurations. If you attempt to use this unsupported configuration, the following symptoms are observed:

- Inoperable logical server appears in Matrix OE after changing the VC Stacking Mode from **Full Stacking** to **Horizontal Stacking** or **Primary Slice Stacking**.
- The portability group is not visible in Matrix OE after the creation or modification of the VC Domain Group using Partially Stacked support.

To change the stacking mode, perform the following in Virtual Connect Manager:

1. Turn on maintenance mode for the Virtual Connect Domain Group.
2. Change the VC Stacking Mode to **Full Stacking**.
3. Turn off maintenance mode for the Virtual Connect Domain Group.

**Suggested action**

Check if the VC Domain Group is configured with **Full Stacking**. Change the VC Stacking Mode to **Full Stacking**, if it is configured as **Horizontal Stacking** or **Primary Slice Stacking**.

## RHEL 7.0 Customization not supported

RHEL 7 can only be customized with the enhanced server customization feature when using Insight Control server provisioning.

## Growth of virtual or physical disks from Matrix is not supported for Integrity blades and HP-UX vPars and Integrity VMs

The ability to grow the size of a virtual or physical disk from Matrix (using Matrix infrastructure orchestration or Matrix OE visualization with logical server management) is not supported for Integrity blades and HP-UX vPars and Integrity VMs.

## Ability to add or remove networks from Integrity blades and HP-UX vPars and Integrity VMs is not supported

Addition or deletion of networks to or from server groups of already provisioned services using GUI, CLI or API is not supported for Integrity blades or HP-UX vPars and Integrity VMs.

## Major Issues

### After upgrade from 7.4 to 7.4.1 the Lifecycle operations in the DL service fails

The `esa.oo.admin.password` value is not correctly migrated during CMS upgrade from 7.4 to 7.4.1. Thus, after the CMS upgrade from 7.4 to 7.4.1, the Lifecycle operations in the DL service fails.

**Suggested action**

During an upgrade, Matrix Operating Environment properties files from an earlier release are replaced by new properties files for the 7.4 Update 1 release. Change the

following entries to identify different behavior for upgrade to full release and intermediate patches:

- For upgrade to minor version:

```
C:\Program Files\HP\Matrix infrastructure
Orchestration\conf\blade_models.properties

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

C:\Program Files\HP\Virtual Server
Environment\conf\lsa\hp_lsa_service.conf

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

C:\Program Files\HP\Virtual Server
Environment\conf\lsa\lsaclient.properties

C:\Program Files\HP\Virtual Server
Environment\conf\lsa\vc\vcController.properties

C:\Program Files\HP\Virtual Server
Environment\conf\lsa\vm\vmController.properties

C:\Program Files\HP\Virtual Server
Environment\conf\lsa\storage\storageController.properties
```

### [Linked clone provisioning of a VM template containing shared disks does not succeed](#)

During VM 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 Matrix 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 Matrix infrastructure orchestration designer and uncheck **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. This will cause the Matrix OE postinstall customization step to fail with an error stating *Error customizing Logical Server Logical server name realized by target UUID*.

#### **Suggested action**

Matrix OE environments can have the newly provisioned server join the domain by using enhanced server customization. For additional information on this, see the *HP Matrix Operating Environment Infrastructure Orchestration User Guide*.

---

**NOTE:** HP does not recommend using an OO workflow to join a newly provisioned server to the domain as it requires adding all computers as trusted hosts.

---

### Ignite fails to provision when SAN is configured with multiple paths

HP-UX deployment with Ignite may fail if multiple SAN paths are configured in active-passive mode. The installation of HP-UX operating system with physical servers might fail with the following error message:

```
The path to the root device
(0/0/0/5/0/0/0.0x207000c0ffda4ee1.0x4070000000000000) does not
support boot. You must choose a bootable path to the root disk.
Virtualized FibreChannel paths are not bootable.
```

This issue might be seen on environments, which have multiple SAN paths configured to the physical blade, and the paths are configured in active-passive mode.

#### **Suggested action**

Switch the multipath configuration to active-active mode. This sets all the SAN paths in active mode.

### Default gateway is not set on the provisioned VM for secondary NIC in dual NIC scenario

After successfully provisioning the VM, the default gateway is not set for the second NIC.

#### **Suggested action**

To access the IP addresses of secondary NIC, add the default gateway manually for the secondary NIC.

### Add disk operation with Hyper-V VM fails

If a Hyper-V VM has a DVD drive on an SCSI\IDE controller and there is no image associated with the drive, the HP Matrix infrastructure orchestration **Add disk** operation will fail with the following message: Cannot attach storage media to the controller because the specified location is in use.

#### **Suggested action**

Either attach the image path DVD to an IDE\SCSI controller or manually remove the DVD.

### Failure during upgrade of Matrix infrastructure orchestration 7.2.x to 7.4

If you are upgrading from HP Insight Management 7.2.x or later to HP Insight Management 7.4, and you have Insight Control performance management and HP Matrix infrastructure orchestration installed and configured on your highly available CMS with SQL Always On enabled, then this issue may be applicable to your environment. You may skip this section if you do not have HP Insight Control performance management and HP Matrix infrastructure orchestration installed and configured in your environment.

While upgrading in the high availability CMS environment with Always On enabled, the installer may result into following error or warning for HP Insight Control performance management and HP Matrix infrastructure orchestration:

- For HP Insight Control performance management  
SEVERE: The operation cannot be performed on database "PMP\_V5\_0" because it is involved in a database mirroring session or an availability group. Some operations are not

allowed on a database that is participating in a database mirroring session or in an availability group.

- For HP Matrix infrastructure orchestration  
Failed to set recovery mode to simple or auto close to false in HP Operations Orchestration database.

### **Suggested action**

Perform the following steps to avoid this error while upgrading the supported Insight Management stack in a high availability environment with Always On database.

1. Remove the HP Insight Control performance management (**PMP\_V5\_0**) and HP infrastructure orchestration (hpio, oo, and HPMOEMetering) databases, as applicable, from Always On Availability Group prior to the execution of upgrade from an earlier version of Insight Management to the latest version. If the removed databases exist in any of the replica servers, disconnect them from there to avoid any pending existing connections left for the removed databases in the Availability Group. Skipping this step will cause a failure when an attempt to add the removed databases back to the Availability database is made after the upgrade. For more information, see *Remove a Primary Database from an Availability Group (SQL Server)* at Microsoft MSDN Library.
2. Perform the upgrade on both nodes. Wait until the upgrade finishes successfully and then add the HP Insight Control performance management (**PMP\_V5\_0**) and HP infrastructure orchestration (hpio, oo, and HPMOEMetering) databases, as applicable, back to the Always On Availability Group using the *Add database wizard* in the Availability Group. For more information, see *Add a Database to an Availability Group (SQL Server)* at Microsoft MSDN Library.

### **Matrix infrastructure service requests fail during VM customization due to insufficient disk space**

Matrix service requests (such as Create Service or Add Server) fail during VM customization due to insufficient disk space. The request details show the attempt to customize the VM and then displays an error of powering on the logical server. It also displays the note to set the path to a disk with more storage space or to delete unnecessary files from the disk and then to try again.

### **Suggested action**

Free up space on the disk containing the VM template and retry the Matrix infrastructure orchestration service request.

## **Minor issues**

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

Provisioning of Integrity VMs with NPIV 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.

### Create Service request with Insight Control server provisioning fails as installation hangs at 60% when deploying SLES11 SP3 OS

The installation of SLES11 SP3 hangs on the target server with the following messages:  
In Matrix OE, the following error appears on the Matrix 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.
```

In the Insight Control server provisioning appliance console, the following error appears:  
Failed to integrate HP SA Agent at step 19.

#### Suggested Action

1. Perform manual cleanup, see the Manual clean-up process (physical) in the *HP Matrix Operating Environment Infrastructure Orchestration User Guide* at the [HP Enterprise Information Library](#). Refer to Manual cleanup for Insight Control server provisioning.
2. Recreate the service request in Matrix infrastructure orchestration.

### Create Service request with Insight Control server provisioning fails when deploying SLES11 SP3 OS with exit code 6 error on appliance console

Deployment of SLES11 SP3 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  
(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:  
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 the [HP Enterprise Information Library](#). Refer to Manual cleanup for 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 the [HP Enterprise Information Library](#). Refer to Manual cleanup for Insight Control server provisioning.
2. Recreate the service request in Matrix infrastructure orchestration.

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

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

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

Before you start a **Delete service request** from infrastructure orchestration:

1. Make sure that the Integrity service was successfully provisioned (the Integrity service is running).
2. Discover the Integrity blade where the Integrity service is running using its IP address from Systems Insight Manager. In a federated CMS environment, discover the Integrity blade on the primary CMS using its IP address from Systems Insight Manager. The Integrity node's identification information should contain the UUID and UUID (logical). You can discover the blade in either of the following ways:
  - a. Use the Systems Insight Manager menu selection **Options**→**Discovery...**, or
  - b. Use the Systems Insight Manager command: `mxnode -a <ip-address> -w`  
where *<ip-address>* is the IP address that was assigned to the Integrity blade after provisioning.  
Add the root login of the Integrity blade using the Systems Insight Manager menu selection **Options**→**Security**→**Credentials**→**System Credentials...**
3. Set up an SSH connection between the CMS and the Integrity blade by entering the command: `mxagentconfig -a -n <ip-address> -u root -p <root password>`
4. If the IP address was discovered on the primary CMS in a federated CMS environment (in step 2 above), manually delete the managed node.  
Use the Systems Insight Manager command: `mxnode -r <ip-address>`

### Suggested action

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

1. Power on the server blade. (The blade was powered off when delete service request completed or failed.)
2. Perform steps 2 and 3 in the preceding section, if you did not do so before starting the delete service request.

3. To erase the disk, enter the following command from the CMS where the discovery was executed (in step 2 in the preceding section):

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

4. If the IP address was discovered on the primary CMS in a federated CMS environment (in step 2 in the preceding section), manually delete the Systems Insight Manager node.

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

### Service creation fails for physical requests when Windows Server OS is selected when infrastructure orchestration times out waiting for Insight Control server deployment to complete installation of Insight Management Agents on the target server

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

#### Suggested action

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

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

```
call :installsc wbem%PROCESSOR_ARCHITECTURE%
if %cprresult% geq 1 set /a cprresults=cprresults+1
```

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

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

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

#### Suggested action

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

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

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

### Suggested action

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

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

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

### Suggested action

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

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

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

### Suggested action

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

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

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

### Suggested action

To avoid this issue on any server:

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

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

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

### **Suggested action**

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

### **Registration of SCVMM 2012 SP1 and SCVMM 2012 R2 server fails on a CMS**

Registration of SCVMM 2012 SP1 and SCVMM 2012 R2 server fails sometimes. This is an intermittent issue and observed several times on a CMS configured with Microsoft Windows Server 2012 or Microsoft Windows Server 2012 R2 operating system.

#### **Suggested action 1**

To solve this issue, retry registration of the SCVMM server.

#### **Suggested action 2**

Add a parameter `SCVMM_Retries` in `hpvmm.conf` file (located in `<IMinstallation>\Insight Control virtual machine management\bin`) and set the value to more than 5. Restart virtual machine management service and try registration again.

#### **Suggested action 3**

If it still fails to register SCVMM, then perform the following steps:

1. Check if the HP Matrix SCVMM Web Service is installed on the SCVMM server.
2. If the HP Matrix SCVMM Web Service is not installed, copy the `ScvmmWS2012.msi` file from `<IM installation>\Insight Control virtual machine management\bin` to the target SCVMM server.
3. Manually install the HP Matrix SCVMM Web Service.
4. Run the HP Matrix SCVMM Web Service on the SCVMM server and reregister the SCVMM server.

### **Guest having differential disk fails**

A VM guest with differential disk fails during the transport operation, customization, and control operation. The APIs that fetch information about disks display incorrect details because the **HP Insight Control virtual machine management** does not support differential disks.

#### **Suggested action**

To solve this issue, avoid using VMs which have differential disk attached.

### **SCVMM deployment fails for SCVMM template formed out of Generation 2 virtual machine**

When you create a template from a Generation 2 virtual machine, by generalizing the boot disk of VM boot disk in the process, and also copying the first boot device of the VM to the template verbatim, the template as its first boot device as an entry does not refer to any actual boot device for that template. This might not be a problem because when a VM is booted, it will go through the entries in the boot order one-by-one until the VM is booted. However, when you try to create a new VM out of this template, Insight Control virtual machine management will try to set the first boot device of the VM to the file entry as set in the template and will fail with error text: `VMM cannot find the device. or This device is not valid for a boot device.`

#### **Suggested action 1**

After creating a template from a VM, modify the first boot device of the template to a valid device.

## Suggested action 2

After the creation of VM and before taking the template, change the boot order and move up the virtual hard disk to first place in the boot order. Then take a template from SCVMM. This will resolve the SCVMM deployment problem for Generation 2 VMs.

### Host IP address field is populated with hostname instead of IP address

Logical Switch Detail is not available for host if Insight Control virtual machine management and SCVMM manage the host with different FQDNs. If VMM manages a host with one FQDN and in SCVMM you have added the same host with a different FQDN due to different DNS, the Logical Switch Detail is not populated.

#### Suggested action

To solve this issue, ensure that VMM and SCVMM register the host with the same FQDN.

As part of Disaster Recovery, VMware is returning the same UUID to VMM for both the live VM and the shadowed VM that is no longer accessible on protected site. (Disaster recovery could be triggered by using VMware vCenter SRM ).

This is a VMware issue.

#### Suggested action

If one or more hosts are disconnected or offline, do not bring them back online until ready to complete the SRM Recovery process. If one or more hosts cannot access their storage, these hosts must be shut down until the problem is resolved and the SRM Recovery process is ready to be initiated.

### Hyper-V provisioning fails on Microsoft Windows Server 2012 with Hyper-V and Microsoft Windows Server 2012 R2 with Hyper-V with timed out errors.

Hyper-V provisioning fails on Microsoft Windows Server 2012 with Hyper-V and Microsoft Windows Server 2012 R2 with Hyper-V with timed out errors. It sometimes, fails to register the deployed VM. The slow performance of the Hyper-V host delays the retrieval of the required data through WMI query.

#### Suggested action

To solve this issue perform the following steps:

1. Increase the CIM Client timed out period to **900000** using parameter `HyperVQueryTimeout` in the `<vmm-install>\bin\hpvmm.conf` and restart the VMM service.
2. Set the value of `serverTimeOut` to **900000** in `cimserver_planned.conf` file of `<MAPPER_INSTALLED_DIR>`, and then restart the Mapper.

### Error Adding CPU to ESXi VM without OS installed (CPU Hot Plug)

Matrix infrastructure orchestration does not allow you to add a CPU on a VM without an OS installed.

#### Suggested action

To update the CPU on a running VM, you should have an operating system that supports processor hot plug. A VM without an operating system cannot have the CPU allocation updated. Virtual Center does not allow CPU hot plug in a VM without an OS.

### Error Adding CPU to an ESXi VM with RedHat 6.4 installed (CPU Hot Plug)

An attempt to use Matrix infrastructure orchestration to increase the number of CPUs for an ESXi VM with RedHat 6.4 fails with the error: CPU hot plug is not supported for this virtual machine.

VMM/VMware returns CPU hot plug enable, but Enable CPU hot add only for this virtual machine is grayed out for the VM in vSphere. VMware and vSphere are not in sync with the flag. This issue occurs on the ESXi VM of ESXi 5.1 host and is not observed on ESXi 5.5 host.

#### **Suggested action**

Increase the number of CPUs for ESXi 5.1 VMs with RedHat 6.4 using Insight Control virtual machine management or vSphere directly. Matrix infrastructure orchestration can be used for other OSs within the ESXi 5.1 VM or Red HAT 6.4 in an ESXi 5.5 VM.

#### **RHEL 7.0 via SA 10.20 fails on Gen8 ProLiant with AMD-Vi (IOMMU) enabled**

When using SA 10.20 to provision RHEL 7.0 on a ProLiant BL465c Gen8 blade (AMD 6234 processor only), Operating system provisioning will fail when IOMMU-Vi is enabled.

#### **Suggested action**

The following are the suggested workarounds:

- If available, use a blade without the AMD 6234 (12 core) processor when provisioning a RHEL 7.0 system.
- If possible, disable the IOMMU-Vi feature within the blade before provisioning a RHEL 7.0 system on that blade.

## HP Matrix OE visualization

Matrix OE visualization runs on an x64 Windows CMS.

### Limitations

#### Raw LUNs added to VMs outside of Matrix not visible to Matrix

New raw LUNs added to a VM outside of Matrix (that is, through the hypervisor) are not discovered during Matrix refresh. The VM will have access to the raw LUN, but Matrix will not display that volume (and Matrix infrastructure orchestration will not charge for that volume).

#### **Suggested action**

Use Matrix to add the raw LUN to the VM via a storage pool entry (optionally using HP Storage Provisioning Manager).

### Restrictions on using cross-technology logical servers

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

The various restrictions on using cross-technology logical servers and their requirements are:

- Valid targets for activating or moving a cross-technology logical server:
  - ESX VM Hosts
  - HP Integrity VM Hosts (racks and server blades)

- ProLiant c-Class servers with Virtual Connect
- HP Integrity servers with Virtual Connect
- Storage requirements:
  - SAN volume presented to an ESX VM Host (RDM disk) that can be added to a Virtual Connect profile.
  - One or more shared datastores on a SAN Array to store the HP Insight Control virtual machine management configuration file and RDM mapping files [.vmdk] on the ESX VM Host. RDM mapping files can be on different SAN based shared datastores. An iSCSI volume can be used for the HP Insight Control virtual machine management configuration file, but not the RDM mapping file.
  - SPM catalog storage volumes:
 

With the 7.4 release, SPM catalog storage volume is supported and can be used with cross-technology logical servers that may be activated on ESXi or Integrity VM Hosts.
  - Integrity VM storage requirements:
    - SAN-based logical server storage is supported, including storage pool entries. File-based storage is not supported with Integrity VM cross-technology logical servers.

---

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

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

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

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

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

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

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

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

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

---

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

---

### Suggested action

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

### Multiple storage volume limitations

The 7.3 release added 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.
- A **Copy Move** operation for a VM using multiple data stores will result in all files on the chosen data store.

### Virtual machine datastore size and configuration filename limits

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

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

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

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

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

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

## Limited support for Hyper-V CSVs created

Hyper-V VM Hosts configured with CSV (cluster shared volume) storage with any type of partition are supported with the following conditions:

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

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

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

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

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

If a Generation 1 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. This is not necessary for Generation 2 VMs which include the SCSI drivers by default.

If a Hyper-V VM logical server is created with Matrix OE visualization, the boot disk is created with the default disk type for the appropriate VM Version (IDE for Generation 1 or SCSI for Generation 2), allowing the operating system installation using supported automation deployment solutions or using images without additional drivers.

If a Hyper-V VM is provisioned by Matrix infrastructure orchestration, IO may use either an SCSI or an IDE controller boot disk. If the Hyper-V VM is created using a VM Template, the controller will be the same as that in the VM Template. If the VM was created with manual OS deployment or using automated deployment solutions, the behavior is as described above for Matrix OE visualization.

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

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

### **Suggested action**

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

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

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

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

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

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

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

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

### Suggested action

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

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

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

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

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

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

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

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

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

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

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

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

## Inoperable logical server appears after changing the VC Stacking Mode from Full Stacking to Horizontal Stacking

Inoperable logical server appears after changing the VC Stacking Mode from **Full Stacking** to **Horizontal Stacking** or **Primary Slice Stacking**. Matrix OE does not support configurations other than Full Stacking.

### Suggested action

Do not use configurations other than **Full Stacking** (Virtual Connect 4.30 and later support Partially Stacked mode like **Horizontal Stacking** and **Primary Slice Stacking**, but Matrix OE does not support such configurations). To change the stacking mode perform the following in Virtual Connect Manager:

1. Turn on maintenance mode for the Virtual Connect Domain Group.
2. Change the VC Stacking Mode to **Full Stacking**.
3. Turn off maintenance mode for the Virtual Connect Domain Group.

## Unavailability of default portability group for VC Domain Group

The portability group is not visible in Matrix OE after the creation or modification of the VC Domain Group using Partially Stacked support (added in VC 4.30). Matrix OE supports only Full Stacking and does not support partially stacked domains.

### Suggested action

Check if the VC Domain Group is configured with **Full Stacking**. Change the VC Stacking Mode to **Full Stacking** if it is configured as **Horizontal Stacking** or **Primary Slice Stacking**. To change the stacking mode perform the following in Virtual Connect Manager:

1. Turn on maintenance mode for the Virtual Connect Domain Group.
2. Change the VC Stacking Mode to **Full Stacking**.
3. Turn off maintenance mode for the Virtual Connect Domain Group.

The portability group should be listed after the execution of this procedure.

## "Copy Logical Server (include all storage)" is not supported if using storage pool entries created by HP Storage Provisioning Manager

When using logical server copy, the option to include all storage is not supported if you use storage pool entries created by HP SPM (Storage Provisioning Manager). You can use *Copy Logical Storage (no storage)* and associate the resulting logical server with SPM-based storage pool entries you have created (for example, SAN Catalog Storage Entry). This ensures separate SPM storage services and storage volumes for the logical server created via the copy operation.

## Major issues

### Provisioning storage using SPM and automated zoning

Provisioning storage using SPM and automated zoning can fail because the principal switch in the network used for this storage is offline.

## Suggested action

Resynchronize the network in the SPM GUI, and then deactivate and reactivate the service involved.

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 are 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, reimport 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, delete the storage pool entry containing the WWNs of the unmanaged Integrity VM logical server.
    3. Reserve the WWNs that belonged to the unmanaged Integrity VM logical server by using the `lsmutil -reserve -WWN WWN` command (where **WWN** is the worldwide name that you made note of in Step 1). Issue this command once for each of the WWNs associated with the unmanaged Integrity VM logical server.

4. Create a placeholder logical server to prevent the MAC addresses from being assigned to new logical servers:
  1. Create a new logical server template without specifying any storage entries in the storage configuration.
  2. In the network selection page, click **Add Network Entries**.

---

**NOTE:** Matrix OE visualization shows the MAC addresses that were part of the unmanaged logical server.

---

3. Click **Next**, and then **Finish** to create the logical server template. This Integrity VM logical server template cannot be activated because it lacks storage entries, but the new template prevents the MAC addresses of the unmanaged logical server from being used again.

You may now create and activate Integrity VM logical servers.

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

You may now create and activate Integrity VM logical servers.

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

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

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

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

#### Suggested action 1

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

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

#### Suggested action 2

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

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

```
wrapper.java.additional.2=-XX:PermSize=32m -XX:MaxPermSize=128m
```

3. Restart the HP Logical Server Automation service.

## Minor issues

### VM logical server may not immediately activate when created using the logical server copy feature

When using Logical Server Copy (**Tools**→**Logical Servers**→**Copy**) to make a copy of a virtual logical server with storage, the newly created logical server may not immediately activate. It may fail activation with an error similar to the following:

```
The save Configuration operation on virtual machine configuration  
<vm path> failed: Failed to create VM! Invalid datastore path '[null]'.
```

#### Suggested action

Edit the newly created logical server and modify the datastore information in the logical server. Save the logical server and it will now activate successfully.

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

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

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

#### Suggested action

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

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

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

#### Suggested action

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

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

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

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

operations on the virtual machine logical server to be unsuccessful, or the virtual machine logical server may become inoperable.

#### **Suggested action**

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

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

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

#### **Suggested action**

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

#### Activation of an Integrity vPar based logical server may fail if base and floating memory configuration guidelines are not followed

The Integrity vPars base and float memory has to be specified as per specific guidelines. If you attempt to activate an Integrity vPar without required minimum amount of base or float memory, the LS activation fails and displays the following message: Unable to define guest. Unable to set memory. Please ensure sufficient base memory has been configured to allow the vPar to boot. There is insufficient free pool memory for this vPar or VM.

#### **Suggested action**

Follow the base and float guidelines specified in the *HP-UX vPars and Integrity VM Administrator Guide* at <http://www.hp.com/go/hpux-hpvm-docs>.

## HP Matrix Operating Environment recovery management

### Limitations

#### Federated Environment configuration requirements with Matrix recovery management

In a Federated Environment, Matrix recovery management is only supported on a primary CMS at both the primary and DR sites.

In other words, Matrix recovery management is not supported on a secondary CMS and all DR protected IO services and logical servers must be on the primary CMS. This can be controlled in Matrix infrastructure orchestration by ensuring server pools on the primary CMS are used when provisioning IO service templates that are recoverable. When manually defining logical servers, ensure they reside on the primary CMS if they are to be later DR protected with Matrix recovery management. For additional information, see the *HP Matrix Operating Environment Infrastructure Orchestration User Guide*.

## HP Insight managed system setup wizard

### Issues and suggested actions

[Select management features help page does not open when you click ? on Management Feature Details against HP Virtual Connect Enterprise Manager](#)

On the **Management Feature Details** page against the **HP Virtual Connect Enterprise Manager**, when you click the ? icon, the help page does not open.

### Suggested action

Perform the following steps:

1. Click the ? icon in the **Managed System Setup Wizard** window.
2. On the *HP Insight managed system setup wizard Online help* window, click **Contents**.
3. Select **Select management features** page to view the details.

## Limitations

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

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

### Suggested action

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

To turn off the requirement for a password, use the Systems Insight Manager **Options**→**Security**→**Privilege Elevation** menu, and uncheck the check box labeled "A password is required for this privilege elevation tool." Systems managed by the CMS must be reconfigured to reflect this configuration change.

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

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

### Suggested action

Update iLO3 ProLiant firmware to version 1.05 or higher.

## HP Application Discovery

### Issues and suggested actions

#### [Create a new application template help page does not open when you click What is this?](#)

On the **Application Discovery** window the **Create new AD application template** help page does not open when you click **What is this?** on **Processes not matched by this template (0)** label.

### Suggested action

Perform the following steps:

1. Click the ? icon on the **Application Discovery** window.
2. On the *HP Application Discovery Online help* window, click **Contents**.
3. Select **Managing application templates**→**Create a new application template** to create a new application template.

# Instant Capacity Manager

## Limitations

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

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

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

---

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

---

### Suggested action

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

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

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

# Process Resource Manager

## Limitations

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

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

### Suggested action

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

# HP Storage Provisioning Manager

Storage Provisioning Manager runs on an x64 Windows CMS.

## Limitations

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

Brocade Network Advisor or HP B-Series SAN Network Advisor versions 11.1.3, 11.1.4, 11.1.5, 12.1.4, 12.3.1, or 12.3.4, must be installed to manage Brocade-based SAN environments. No other versions of BNA are supported.

---

**NOTE:** It is recommended to use Brocade Network Advisor and HP B-Series SAN Network Advisor of version 12.1.4, 12.3.1 or 12.3.4. While using the SPM automated zoning capability with HP B-series/Brocade Network Advisor version 12.1.4, 12.3.1 or 12.3.4, you may disable Historical Data Collection to ensure BNA provides adequate performance for SPM. To disable Historical Data Collection from Brocade Network Advisor GUI, select *Monitor • Performance • Historical Data Collection • Disable All*.

---

SPM performance may be negatively impacted during refresh of resources when using BNA 11.x versions. For more information on how to improve refresh performance (for resources and services, across various versions), see the *HP Storage Provisioning Manager (SPM) User Guide*.

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

### SPM limitations on EMC Symmetrix VMAX arrays

For more operational details and information about configuring the SPM adapter for EMC Symmetrix VMAX, see the technical white paper on SPM adapter for EMC Symmetrix VMAX at the [HP Enterprise Information Library](#).

The SPM limitations on EMC Symmetrix VMAX arrays are:

- Creating a new volume with a given RAID requires preconfiguration of a storage pool setting instance by the EMC storage administrator who can create a volume of that RAID through EMC native tools (SMC or SYMCLI).
- Creating a thin volume for a given RAID requires creating 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. Thus, 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.
- 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 has not been tested with the OpenVMS host mode.
- For all supported host modes, the maximum number of LUNs that can be presented on a host initiator per EMC port is limited to 2000.

### Firmware for HP 3PAR StoreServ

HP 3PAR StoreServ must have a supported version of 3PAR OS installed to work properly with SPM. SPM will not allow 3PAR arrays with earlier versions to be imported into the SPM catalog. See the *HP Insight Management Support Matrix* at [HP Enterprise Information Library](#) for currently supported versions.

### HP 3PAR StoreServ Storage imposes an initiator WWN limit per controller port

SPM supports the feature 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, it prevents 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 HP 3PAR StoreServ Storage.

---

These limitations can impact servers provisioned by Matrix, using SPM for storage. For more information, see the *HP Matrix Operating Environment and HP Storage Provisioning Manager support for port groups: Storage controller port load balancing using SPM* white paper at the [HP Enterprise Information Library](#).

### HP 3PAR StoreServe 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 preprovisioned volumes with the Zero Detect property, but it does not recognize the property value. SPM will continue to use volumes it has created after an administrator applies HP 3PAR Zero Detect in the HP 3PAR Management Console, but does not recognize if a volume has the Zero Detect property value.

### HP 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 to ensure that SPM and Matrix OE can properly use the storage affected by Peer Motion. For details on these steps, see the *SPM 3PAR StoreServ Storage Peer Motion Technical White Paper* at the [HP Enterprise Information Library](#).

---

**NOTE:** The online migration from EVA to 3PAR is currently not supported.

---

### HP 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 worldwide name. For this reason, SPM cannot be used to manage volumes in a peer persistence configuration.

## Minor issues

### 3PAR OS 3.1.3 and later versions include new host personae that are supported by SPM

3PAR has added two new host personae with HP 3PAR OS 3.1.3 (also present in later versions, like 3.2.1). SPM has to support these two new personae and do the internal mapping of host mode strings accordingly. SPM has defined appropriate mappings based on the version of the 3PAR OS. 3PAR is also recommending a change in the personae used for Linux (for example, RHEL and SUSE must use value 2 Generic-ALUA versus value 1 Generic-1). Customers will need to manually add the new mappings by editing the SPM configuration file. After editing the file and restarting SPM, customers may find their services go nonconformant (for example, if the service was created using a previous mapping string and SPM now uses a different string based on the 3PAR OS version).

Changing the host mode of an existing service requires the SPM service to be reactivated (to pick up the new personae value) and in some cases the host OS may need to be rebooted to use the new personae for I/O (changing the Windows and HP-UX persona does not require a reboot, but the Linux persona changes do require reboot). For details on HP 3PAR StoreServ storage support, see the <http://www.hp.com/storage/spock> website. An HP Passport is required and the SAN Compatibility section allows you to view by your specific model, 3PAR OS, and server OS (including Host Mode information on the recommended personae).

For this reason, manual editing of the file is required to use the new values (rather than using them by default and having existing SPM services go nonconformant). The SPM configuration file can be manually edited to change the mapping strings (For example, to ensure the use of the new HP-UX 11iv3 or Windows personae for Matrix when working with 3PAR OS 3.1.3 or later). The SPM service must be restarted after the file edits. Such changes will make existing SPM services using the previous values become nonconformant until they are reactivated. The following steps will reactivate an SPM service:

1. Select the **Storage Services** view in the SPM GUI. The list of services will have some which are *nonconformant* (those with new personae values from the file versus from original creation).
2. Right click on each nonconformant SPM service and select *Configure Requirements*.

---

**NOTE:** There is no need to modify the requirements (SPM will have the new value from the configuration file edits). The Operating System value in each Host requirement must be the new value and will display as *Match (with Automation)*.

---

3. Click **Next**, then click **Finish**. SPM will reactivate the service and it should become *conformant*.

#### Suggested action

Consider the SPM behavior based on 3PAR OS version and determine if manual edits are appropriate for your environment. Typically, it is recommended to use the newer personae with later versions of 3PAR OS and the system OS (as noted in the 3PAR documentation referenced above). In the case of a fresh install of 7.4 (and upgrade to 7.4 Update 1), the previous personas used by the configuration file can (and must) be manually edited to take advantage of the new personae. In the case of an upgrade from a release prior to 7.4, no edits are required unless you wish to use the new personae (and are prepared to handle the nonconformant services). The steps above can be used to reactivate nonconformant SPM services after the file is edited. 3PAR

has introduced two new personae in HP 3PAR OS 3.1.3 and later releases as mentioned in the [Table 4 \(page 66\)](#).

**Table 4 New host personae supported by SPM**

Persona	ClientTypes	OtherClientTypeDescriptions
13	HPUX (11iV3)	HP-UX (11i v3)
15	Microsoft Windows	Windows 2003 should use value 1
	Microsoft Windows Server 2008	Generic Microsoft Windows
	Windows 2008/2008R2/2012	

For supporting the two new personae, SPM has added two new host mode strings, which are **HP-UX-11iV3** and **Windows**. In the 3PAR configuration file located under SPM state folder (such as `C:\Program Files\HP\Storage Provisioning Manager\state\3par_storage\Configuration.xml`), a new **OSPersonaMap** is added for HP 3PAR OS 3.1.3 and later releases. For backward compatibility (and to avoid existing SPM services going nonconformant), the new strings are not used unless manually edited into the file.

**NOTE:** An SPM restart is required after editing the configuration file.

Matrix continues to call SPM using the HP-UX personae, so customers wanting to use the new value for Matrix storage with 3PAR OS 3.1.3 and later will need to edit the HP-UX entry to be HP-UX-11iV3: `<OSPersona OS="HPUX" Persona="HP-UX-11iV3" />`

**NOTE:** This will cause existing services with the HP-UX personae to go nonconformant and require reactivation in SPM (using the steps noted above). This personae change is recommended for customers using HP-UX 11v3 only (it is not recommended for use with HP-UX 11iV2).

The mapping values for Windows, Windows 2008, and Windows 2012 remain unchanged for 3PAR OS 3.1.3 and later unless they are manually edited in the file. As noted, it is recommended to make these edits for Windows 2008 and Windows 2012. If the customers are still using Windows 2003, it is recommended to continue to use the old Generic-1 host mode (for more details on this, see the 3PAR documentation).

```
<OSPersona OS="Windows" Persona="Generic-1" />
<OSPersona OS="Windows2008" Persona="Windows" />
<OSPersona OS="Windows2012" Persona="Windows" />
```

**NOTE:** This will cause existing services with the Generic-ALUA personae to go nonconformant and require reactivation in SPM using the steps outlined above, but typically does not require a Windows OS reboot to use the new Windows personae rather than Generic-ALUA (for more details on this, see the 3PAR documentation).

Similarly, the mapping values for Linux remain unchanged for 3PAR OS 3.1.3 and later unless they are manually edited in the file. 3PAR recommends use of Generic-ALUA for Linux operating systems (such as RHEL and SUSE) rather than the previous use of Generic-1 and an OS reboot is required to use the Generic-ALUA persona for Linux.

```
<OSPersona OS="Linux" Persona="Generic-ALUA" />
```

When SPM is upgraded from a version prior to 7.4, it retains the existing configuration file which did not have the new personae. In such situations, the configuration file must

be edited to add the version info and new personae as shown in the bold typeface below:

```
<OSPersonaMaps>
  <OSPersonaMap VersionMinimum="" VersionMaximum="3.1.2.219">
    <OSPersona OS="Windows" Persona="Generic-1" />
    <OSPersona OS="Windows2008" Persona="Generic-ALUA" />
    <OSPersona OS="Windows2012" Persona="Generic-legacy" />
    <OSPersona OS="Linux" Persona="Generic-1" />
    <OSPersona OS="HPUX" Persona="HP-UX" />
    <OSPersona OS="VMware" Persona="Generic-legacy" />
    <OSPersona OS="OpenVMS" Persona="Generic-1" />
  </OSPersonaMap>
  <OSPersonaMap VersionMinimum="3.1.2.220" VersionMaximum="3.1.2.999">
    <OSPersona OS="Windows" Persona="Generic-1" />
    <OSPersona OS="Windows2008" Persona="Generic-ALUA" />
    <OSPersona OS="Windows2012" Persona="Generic-ALUA" />
    <OSPersona OS="Linux" Persona="Generic-1" />
    <OSPersona OS="HPUX" Persona="HP-UX" />
    <OSPersona OS="VMware" Persona="VMware" />
    <OSPersona OS="OpenVMS" Persona="Generic-1" />
  </OSPersonaMap>
  <OSPersonaMap VersionMinimum="3.1.3" VersionMaximum="">
    <OSPersona OS="Windows" Persona="Generic-1" />
    <OSPersona OS="Windows2008" Persona="Windows" />
    <OSPersona OS="Windows2012" Persona="Windows" />
    <OSPersona OS="Linux" Persona="Generic-ALUA" />
    <OSPersona OS="HPUX" Persona="HP-UX" />
    <OSPersona OS="VMware" Persona="VMware" />
    <OSPersona OS="OpenVMS" Persona="Generic-1" />
    <OSPersona OS="HPUX(11iV3)" Persona="HP-UX-11iV3" />
  </OSPersonaMap>
```

---

**NOTE:**

- The above configuration uses the new Windows host mode for Windows 2008 and Windows 2012 and the Generic-ALUA mode for Linux. The HP-UX host mode is not changed. Existing SPM services using the older host modes will go nonconformant and it is necessary to reactivate those services using the SPM steps outlined above.
- When SPM is installed with 7.4 (perhaps for later upgrade to 7.4 Update 1), the configuration file retains the host mode string mapping of Windows as that of 3PAR OS 3.1.2 in the configuration file for 3PAR OS 3.1.3 and later (as shown above) for backward compatibility. As mentioned above, edits are required for Matrix OE to use the new HP-UX-11iV3 personae, or for the new Windows or Linux personae to be used by Matrix OE.

---

The results of those configuration file edits are shown in the bold typeface below:

```

<OSPersonaMaps>
  <OSPersonaMap VersionMinimum="" VersionMaximum="3.1.2.219">
    <OSPersona OS="Windows" Persona="Generic-1" />
    <OSPersona OS="Windows2008" Persona="Generic-ALUA" />
    <OSPersona OS="Windows2012" Persona="Generic-legacy" />
    <OSPersona OS="Linux" Persona="Generic-1" />
    <OSPersona OS="HPUX" Persona="HP-UX" />
    <OSPersona OS="VMware" Persona="Generic-legacy" />
    <OSPersona OS="OpenVMS" Persona="Generic-1" />
  </OSPersonaMap>
  <OSPersonaMap VersionMinimum="3.1.2.220" VersionMaximum="3.1.2.999">
    <OSPersona OS="Windows" Persona="Generic-1" />
    <OSPersona OS="Windows2008" Persona="Generic-ALUA" />
    <OSPersona OS="Windows2012" Persona="Generic-ALUA" />
    <OSPersona OS="Linux" Persona="Generic-1" />
    <OSPersona OS="HPUX" Persona="HP-UX" />
    <OSPersona OS="VMware" Persona="VMware" />
    <OSPersona OS="OpenVMS" Persona="Generic-1" />
  </OSPersonaMap>
  <OSPersonaMap VersionMinimum="3.1.3" VersionMaximum="">
    <OSPersona OS="Windows" Persona="Generic-1" />
    <OSPersona OS="Windows2008" Persona="Windows" />
    <OSPersona OS="Windows2012" Persona="Windows" />
    <OSPersona OS="Linux" Persona="Generic-ALUA" />
    <OSPersona OS="HPUX" Persona="HP-UX-11iV3" />
    <OSPersona OS="VMware" Persona="VMware" />
    <OSPersona OS="OpenVMS" Persona="Generic-1" />
    <OSPersona OS="HPUX(11iV3)" Persona="HP-UX-11iV3" />
  </OSPersonaMap>

```

---

**NOTE:**

- The above configuration uses the new Windows host mode for Windows 2008 and Windows 2012, for Linux, and the new HP-UX-11iV3 for all HP-UX. Existing SPM services using the older host modes will go nonconformant and it is necessary to reactivate those services using the SPM steps outlined above.
  - The new HP-UX-11iV3 and Windows personae do not apply to 3PAR OS versions prior to 3.1.3 and editing those values for earlier versions will lead to errors. Along with SPM service reactivation through Configure Requirements, changing the Linux persona may also require an OS reboot (for more details see, 3PAR documentation). For details on HP 3PAR StoreServ storage support, see the <http://www.hp.com/storage/spock> website. An HP Passport is required and the SAN Compatibility section allows you to view by your specific model, 3PAR OS, and server OS (including Host Mode information on the recommended personae).
-

After the 3PAR OS is upgraded to 3.2.1 MU1, if an array is imported into SPM, import fails

When an array is imported into SPM after HP 3PAR StoreServ Storage System is upgraded to 3PAR OS 3.2.1 MU1, the array results in an import failure and the server returns the following error:Unknown WBEM Client Error: Unexpected XML declaration. The XML declaration must be the first node in the document, and no white space characters are allowed to appear before it. Line 15 position 3. (See log for details.)

### Suggested action

Log in to 3PAR in CLI mode and execute the following commands:

1. `stopcim -f -x`
2. `startcim`
3. `showcim` [wait for "state" to become "active"]

After the server status changes to **active**, SPM can successfully import the array.

## Managed node installation

### Limitations

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

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

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

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

### Suggested action 1

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

Matrix 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 multiple-path access to the boot volume when OS provisioning is complete.

### Suggested action 2

If you are defining a storage pool entry in Matrix OE visualization:

1. Define the storage pool entry, adding multiple ports, and marking all nonboot 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.

# HP Global Workload Manager (gWLM)

## Known Issues

### gwlagent process crashes on undeployment of SRD with VSP workloads

gwlagent daemon crashes during undeployment of an SRD that has a VSP host and vPar guests (only vPars or combination of vPar and HPVM guests) as workloads of the SRD.

#### **Suggested action:**

Execute the following on the CMS for undeployment of the SRD: `gwlm undeploy --srd=<srd name> --force` Start the gwlagent on the VSP host.

### gWLM displays nPar compartment type during SRD creation for non iCAP complex

If a complex is non iCAP and part of a GiCAP group and if the user chooses to create SRD with any partition of the non-iCAP complex, then nPar will be displayed as a compartment type in the list of compartments during discovery.

#### **Suggested action**

The sub OS compartment types (pset/fss) can be chosen from the drop-down list of compartment types.

### SRDs defined when system was non-iCAP are unable to use TiCAP later

When an SRD (Shared Resource Domain) is defined on a non-iCAP system, and the system is later modified to be an iCAP system and TiCAP which can be used manually, the SRD does not make use of TiCAP even if the policy used has it enabled.

#### **Suggested action**

Perform the following steps:

1. If the SRD is deployed when the change from non-iCAP to iCAP occurred, click **Modify SRD**. You can see **TiCAP** option in the **Manage systems and workloads** tab. Click **OK**.
2. Redeploy the SRD.

## HP Insight Control server provisioning

### Known Issues

#### Network personalization of deployment NIC is limited to IPv4 and during OS installation

Since ESXi does not have an agent, you can personalize the deployment NIC only during an OS installation, and only with IPv4. Although the user interface allows you to enter data for other NICs, only the deployment NIC is configured. ESXi does not support post install network personalization.

#### RHEL 7.X VMs created to install via SA need to have at least 1.5 GB RAM

RHEL 7.x cannot be installed on VMs created using less than 1.5 GB of RAM. Trying to install RHEL 7.x via SA in VMs created with less than 1.5 GB will cause an error when SA OSBP tries to extract the SA agent in the target.

#### Configure Network issues with RHEL 7

When using the **Configure Network** feature on a RHEL 7 system, there are some restrictions like:

- You cannot set the domain name.
- If using KVM, the bridge interface BRO does not get configured.

### Suggested action

Set the domain name of the server and IP of the deployment NIC during OS installation time or manually.

## Linux installation may time out

If you are installing a lot of RPMs as part of your Linux installation, the Build Plan may time out on the `Monitor Installation` step because nothing is written to the installation log for an extended period of time, while the RPMs are installed.

### Suggested action

Increase the time, the `Monitor Installation` step waits for log activity as follows:

1. Make a copy of the Build Plan.
2. Edit your copy of the Build Plan.
3. Modify the parameters of the `Monitor Installation` step by adding a number after the log file name, which indicates the number of polls the step takes before timing out. The default is `200`, so increasing to `400` or more is recommended. Here is an example of the parameters in a Red Hat installation: `tmp/anaconda.log 400`.
4. Retry the installation with the new Build Plan.

## The Create Stub Partition step may fail for a Linux or ESXi installation when using 3PAR storage

The 3PAR 3.1.3 OS has a new ESX device type, which may result in a `Create Stub Partition` fail with exit code 1 error. Due to this error, an extra disk device is created, which does not represent an actual disk.

### Suggested action

Set the device type to RHEL (5.x, 6.x) on the 3PAR storage.

## Limitations

### Error while installing RHEL 6.6 on a target server booting from 3PAR storage

When performing an RHEL 6.6 boot from SAN OS installation on a 3PAR storage volume, the default `autopart` option which automatically creates the LVM partitions fails. This issue exists with LVM partition only.

### Suggested action

To complete the installation, use custom partitioning options when using LVM and create partition table sizes manually instead of using `autopart`. In the following example of a kickstart file, the `autopart` option is replaced with custom partitioning commands.

```
zerombr yes
clearpart --all --initlabel
part /boot --fstype=ext3 --size=512
## Uncomment below line if target is in UEFI mode ##
#part /boot/efi --fstype=efi --size=512
part / --fstype=ext3 --size=4096 --grow --maxsize=8192
part /var --fstype=ext3 --size=4096 --grow --maxsize=8192
part swap -recommended
```

---

**NOTE:** While using 3PAR storage, minimum boot partition size is 256 MB.

---

## Error while installing SLES 12 on a target server booting from 3PAR storage

When performing an SLES 12 boot from SAN OS installation on a 3PAR storage volume using default partitioning, the installation fails since it automatically creates an MBR partition of less than 256MB.

### Suggested action

Modify the custom partition settings in SLES12 autoyast configuration file and set the MBR partition size to at least 256MB as in the following example.

```
<partitioning config:type="list">
  <drive>
    <device>/dev/sddevice>/dev/sda>
    <partitions config:type="list">
      <partition>
        <mount>/boot</mount>
        <size>512mb</size>
      </partition>
      <partition>
        <mount>/</mount>
        <size>auto</size>
      </partition>
      <partition>
        <mount>swap</mount>
        <size>2gb</size>
      </partition>
    </partitions>
  </drive>
</partitioning>
```

## Installing SLES 11.3 to any Gen9 servers requires a Bootable Driver Kit (kISO)

To install SLES 11.3 on any ProLiant Gen9 servers, a special Bootable Driver Kit (previously called kISO) must be added to the SLES distribution on the media server.

### Suggested action

Perform the following steps to add Bootable Driver Kit to the SLES distribution on the media server:

1. Download the kit located at <http://drivers.suse.com/hp/HP-ProLiant-Gen9>.
2. Open the ISO.
3. On your media server, make a copy of the SLES 11.3 distribution and place it in a properly named folder. For example, `/Media/SLES-11SP3-kISO-x86_64`.

---

**NOTE:** You may want to keep the original copy as this kit is only used for Generation 9 servers.

---

4. In the new folder that you created in “step 3” (page 72), create an `addon` directory under new repository directory. For example, `Media/SLES-11SP3-kISO-x86_64/addon`.
5. Copy the kit contents to the `addon` directory.
6. Replace the original installation kernel and `initrd` files with the ones from the kit `addon` directory as follows:
  - Copy `addon/boot/x86_64/loader/initrd` to `boot/x86_64/loader/initrd`.
  - Copy `addon/boot/x86_64/loader/linux` to `boot/x86_64/loader/linux`.
7. In the appliance UI, find the `SLES11 SP3 Build Plan` and select **Save As** to create an OS Build Plan from the existing `SLES11 SP3 Build Plan` that can be edited.

8. Edit the `Set Media Source` step in the new Build Plan and change the parameter to point to the new directory.
9. Delete the `ProLiant Drivers for SLES 11 SP3 x64` step as the kit has all the necessary drivers.

### RHEL 7.0 may fail on servers with B140i Smart Array Controller

When installing RHEL 7.0 to servers with a B140i Smart Array Controller, an HP proprietary driver is required. However, there is a conflict between HP proprietary driver and the driver provided by the RHEL 7.0 distribution, which may cause issues during installation.

#### Suggested action

Create a new custom attribute, associated with either the server or the Build Plan named `kernel_arguments` and assign the value `blacklist=ahci`. This does not allow the driver within the RHEL 7.0 distribution to be used.

### HP Insight Control server provisioning cannot install on Generation 2 VMs in Hyper-V 2012 R2

Insight Control server provisioning cannot deploy to Generation 2 VMs hosted on Hyper-V 2012 R2 hosts.

---

**NOTE:** This limitation is applicable only to IC server provisioning 7.4. HP recommends use of IC server provisioning 7.4.1 which can deploy to Generation 2 VMs hosted on Hyper-V 2012 R2 hosts.

---

#### Suggested action

Use only Generation 1 VM guests. This can be done in Matrix by modifying the following property, `MAX_DEFAULT_HYPERV_VM_VERSION=1` within the directory `C:\Program Files\HP\Virtual Server Environment\conf\lsa`.

### HP Insight Control server provisioning does not currently support SLES 12 provisioning via FCoE

The Matrix Operating Environment uses Virtual Connect modules which use FCoE. FCoE is not supported by Insight Control server provisioning for SLES 12 deployment.

#### Suggested action

SLES 12 can be provisioned to a physical server using manual deployment.

### Windows Build Plan may fail on UEFI enabled servers running in Legacy Mode

A Windows scripted installation Build Plan may fail at the Install and boot into local WinPE step on any Gen9 or DL580Gen8 server configured in Legacy boot mode. The Build Plan fails with the error message: `Failed To Partition Disk for Windows: Cannot find on disk 0 a partition with label 'System'`.

#### Suggested action

To fix this, perform the following steps:

1. Make a copy of the Build Plan.
2. Edit your copy of the Build Plan.
3. Before the Install and boot into local WinPE step, add a `Prepare Disks on HP ProLiant Gen8` step.
4. Save the Build Plan and use that copy.

## Multiple limitations when installing RHEL 7 using FCoE

### Issue 1 — Disks may not appear to the OS Installer

Due to a timing issue, the FCoE disks may not appear to the RHEL 7.0 installer. When this happens, the Build Plan fails, and on the target servers console, you see that the installer has switched to interactive mode and asks for storage information.

#### Suggested action

Retry the Build Plan. The issue is intermittent and usually works for the second time.

### Issue 2 — Installation hangs or panics at final reboot

A RHEL 7.0 scripted installation Build Plan may fail at the final Wait for HP SA Agent step when using FCoE with certain HP FlexFabric/StoreFabric Adapters. The server may hang during the final reboot or may kernel panic when trying to boot the RHEL 7.0 production OS. This is due to a race condition when rebooting immediately after the installation is completed. The following adapters are affected when used in FCoE mode:

- HP FlexFabric 10Gb 2-port 536FLB Adapter
- HP FlexFabric 20Gb 2-port 630FLB Adapter
- HP FlexFabric 20Gb 2-port 630M Adapter
- HP FlexFabric 10Gb 2-port 534FLB Adapter
- HP FlexFabric 10Gb 2-port 534FLR Adapter
- HP FlexFabric 20Gb 2-port 630FLB Adapter
- HP StoreFabric CN1100R Dual Port Converged Network Adapter

#### Suggested action

Make a copy of the Build Plan and insert a `Sleep` script step with the parameter `--minutes=3` before the final `Reboot` step.

## Intermittent boot failures when installing SLES 11 SP3 with FCoE

Intermittent failures are seen when installing SLES 11 SP3 to FCoE devices. The RPM installation succeeds, but OS fails during the first reboot and the Build Plan fails with the error: `Step 22 of 22: Run OGFS Script 'Wait for HP SA Agent'`  
`Waiting for registration of agent running on production OS`  
`Failed To Wait for HP SA Agent: Timeout while waiting for agent`  
`Wait for HP SA Agent failed with exit code 120.`

#### Suggested action

Modify the disk partition section of the AutoYaST file as follows:

1. Use **Save as** to make copies of the `SLES 11 SP3` Build Plan and configuration file.
2. Edit the configuration file and uncomment the XML section that refers to the partitioning of the disk.
3. Edit the new Build Plan to use the newly modified configuration file.

## RHEL driver name change may cause problems when using different driver packages

In the `2014.09.0` RHEL driver packages delivered with the appliance, the name of the `Mellanox` driver has changed from that of previous driver releases. The HP provided RHEL kickstart files are updated to match the new driver name, however, if you choose to use an older kickstart file with the new driver packages, or if you choose to use a newer kickstart file with older driver packages, your Build Plans fail.

### Suggested action

Make sure the driver called out in the kickstart file matches the version of the driver package you are using.

- RHEL driver packages version 2014.09.0 — Driver name is `kmod-mlnx-en`
- RHEL driver packages prior to 2014.09.0 — Driver name is `kmod-mellanox-mlnx-en`

### Restrictions when using the HP FlexFabric 20Gb 2-port 630FLB, 650FLB, 650M or 630M Adapters

The following are some of the restrictions:

- Windows and Linux deployment is not supported with Intelligent Provisioning on Gen8 (or greater) servers. Windows and Linux deployment via PXE is supported.
- Deployment of SLES11 SP3 is not supported in iSCSI or FCoE configuration with these adapters.
- Deployment of Red Hat 6.x and Red Hat 7.x with FCoE configuration requires you to add the FCoE directive to the Kickstart file. The FCoE directive in the Kickstart file requires that the NIC attached to the SAN be specified.
  - `--nic=` (mandatory) — The name of the device to be activated.
  - `--dcb=` — Establish DCB (data center bridging) settings.
  - `--autovlan` — Discover VLANs automatically.

---

**NOTE:** At a minimum, you must specify the `--nic=` flag. For example, `fcoe --nic=eth0`. With Red Hat 7.x, the Ethernet device names have changed. For more information, see *Red Hat Enterprise Linux 7* documentation located at [https://access.redhat.com/documentation/en-US/Red\\_Hat\\_Enterprise\\_Linux/7/html/Networking\\_Guide/ch-Introduction\\_to\\_RHEL\\_Networking.html](https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Networking_Guide/ch-Introduction_to_RHEL_Networking.html).

---

### Broadcom CNA (converged network adapter) option cards may require additional FCoE initialization if the drive does not present using the “fcoe” directive in the Kickstart file

There is a timing issue that can sometimes prevent the LUN from being visible to the Broadcom CNA option card.

#### Suggested action

Successful deployment requires additional FCoE initialization in the kickstart file's `%pre` section. To do the initialization, add the following code to the `kickstart` file:

```
%pre
modprobe fcoe
fipvlan -c -s -a
sleep 50
fipvlan -c -s -a
%end
```

### BL465c Gen8 hangs during OS deployment if using Broadcom CNA mezzanine configured with FCoE configuration

OS deployment of Windows and Linux is not supported on the BL465c Gen8 if using Broadcom CNA mezzanine configured with FCoE configuration. The server hangs due to a platform issue and the OS deployment fails.

**Suggested action**

No work around exists.

## 3 Documentation notes

This chapter provides information about any one or all of the following:

- **Addendum**  
This section provides information on updates to Matrix Operating Environment documents that are not being revised in this release. This information will be included in future document revisions.
- **Corrections**  
This section provides the information meant to clarify certain inaccuracies that are required to be updated in Matrix Operating Environment documents that are not being revised in this release. This information will be included in future document revisions.

### Addendum

The following restriction as listed in the *Matrix Operating Environment 7.4.1 Infrastructure Orchestration Console and Designer Online Help Systems* and the *Matrix Operating Environment 7.4 Infrastructure Orchestration User Guide* is no longer applicable:

**Hyper-V VMs using logical switches cannot be DR protected using HP Matrix Operating Environment recovery management.**

Matrix recovery management now supports Microsoft Hyper-V logical switches.

### Corrections

The following are the corrections identified in this release:

#### Matrix Operating Environment 7.4 Infrastructure Orchestration User Guide

On page 43 of the *Matrix Operating Environment 7.4 Infrastructure Orchestration User Guide*, step 3 of the procedure for Joining a Windows Active Directory Domain included incorrect information. The corrected information appears below.

Define the following custom attributes for the OS Build Plan on the Insight Control server provisioning appliance. (See the Insight Control server provisioning documentation for more detailed information on using this step.)

The custom attributes can be defined for the OS Build Plan or as user-defined properties for the appropriate server group within the template in Matrix OE infrastructure orchestration designer.

Category	Example
DomainFQDN	mydomain.net # FQDN of your domain
DomainName	# NETBIOS name of your domain
DomainUser	Aduser # An AD user with permissions to join workstations to the domain
DomainPassword	supersecret #Password for DomainUser
EncryptedDomainPassword	76492d1116743f042341 #Encrypted password
Key	supersecretpasswordworks #Key used to generate above password

---

**NOTE:** Either *DomainPassword* or a combination of *EncryptedDomainPassword* with *Key* custom attributes need to be provided.

- *DomainPassword* is stored as plain text. You can delete it on completion of the Build Plan.
- *EncryptedDomainPassword* stores the encrypted password generated using the *Key* specified in the *Key Custom* attribute.

---

### Procedure 1 To generate an encrypted password

Use the following steps on a separate Windows system to generate an encrypted password using Windows PowerShell:

1. `$key="supersecretpasswordworks"`
2. `$password=read-host -assecurestring`
3. `<Type Password>`
4. `$encrypt=convertfrom-securestring $password -key $key.ToCharArray()`
5. `$encrypt | set-content c:\encrypted.txt #save password to file`

---

**NOTE:** The valid key can be any of 128/192/256 bits that is a byte array of 16/24/32 digits. In the preceding example, byte array to 24 digits is used. You must specify some key value in the *Key* custom attribute field to generate an encrypted password.

---

---

## 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:

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

#### How to contact HP

Use the following methods to contact HP technical support:

- See the Contact HP Worldwide website for contact options:  
<http://www.hp.com/go/assistance>
- Use the `Contact hp` link on the HP Support Center website:  
<http://www.hp.com/go/hpsc>
- In the United States, call 1-800-334-5144 to contact HP by telephone. This service is available 24 hours a day and 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, 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 website:  
[http://www.hp.com/service\\_locator](http://www.hp.com/service_locator)
- In other locations, see the Contact HP worldwide website:  
<http://welcome.hp.com/country/us/en/wwcontact.html>

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

Open source software (such as OpenSSL) or third-party softwares (such as Java) are sometimes included in HP products. HP discloses that the non-HP owned software components listed in the Matrix Operating Environment EULA (end user license agreement) are included with Matrix OE.

To view the EULA, use a text editor to open the `/opt/vse/src/README` file on an HP-UX CMS, or the `<installation-directory>\src\README` file on a Windows CMS. (The default installation directory on a Windows CMS is `C:\Program Files\HP\Virtual Server Environment`, but this directory can be changed during installation.)

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

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

## Subscribing to security bulletins

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

1. Open a browser to the HP home page:  
<http://www.hp.com>
2. Click the **Download Drivers** under Customer Support.
3. Click **Sign up: driver, support, & security alerts**, which appears under **Other support resources** in the center of the content pane.
4. Do one of the following:
  - Under **Subscribe today** enter your Email address, First and Last name, product information, and click **Subscribe**.
  - Under **Already subscribed?** sign in to customize your preferences and update your profile.

## Related information

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

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

- HP Insight Management documentation:
  - <http://www.hp.com/go/insightmanagement/docs>
- HP Insight Control documentation:
  - <http://www.hp.com/go/insightcontrol/docs>

---

## 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)). Include the document title and part number, version number, or the URL when submitting your feedback.