# Oracle® VM Server for SPARC 3.5 Release Notes

**Part No: E80107**
July 2018

ORACLE

Oracle VM Server for SPARC 3.5 Release Notes

**Part No: E80107**



**Access to Oracle Support**

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info` or visit `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs` if you are hearing impaired.

**Référence: E80107**



**Accès aux services de support Oracle**

Les clients Oracle qui ont souscrit un contrat de support ont accès au support électronique via My Oracle Support. Pour plus d'informations, visitez le site `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info` ou le site `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs` si vous êtes malentendant.

# Contents

**Using This Documentation** ..... 7

**1 Oracle VM Server for SPARC 3.5 Release Notes** ..... 9
What's New in This Release ..... 10
Features That Depend on System Firmware, the Oracle Solaris OS, or Both ..... 12
Oracle VM Server for SPARC 3.5 System Requirements ..... 12
Deprecated and Removed Oracle VM Server for SPARC Features ..... 12
Known Issues ..... 13
Bugs Affecting the Oracle VM Server for SPARC Software ..... 13
Documentation Issues ..... 35
Resolved Issues ..... 37

# Using This Documentation

- **Overview** – Provides late-breaking information about the Oracle VM Server for SPARC 3.5 software, such as changes for this release and known bugs that affect the software.
- **Audience** – System administrators who manage virtualization on SPARC servers.
- **Required knowledge** – System administrators with working knowledge of UNIX systems and the Oracle Solaris operating system (Oracle Solaris OS).

## Product Documentation Library

Documentation and resources for this product and related products are available at http://www.oracle.com/technetwork/documentation/vm-sparc-194287.html.

## Feedback

Provide feedback about this documentation at http://www.oracle.com/goto/docfeedback.

♦ ♦ ♦ CHAPTER 1

# Oracle VM Server for SPARC 3.5 Release Notes

These release notes include information about issues that you might encounter if you are running the fully qualified versions of the Oracle Solaris OS, system firmware for a supported platform, and the Oracle VM Server for SPARC 3.5 software. If you are not running these qualified versions, you might encounter a larger set of issues.

---

**Note -** Ensure that you install and run the Oracle VM Server for SPARC 3.5 software with the fully qualified system firmware versions on the supported hardware platforms. All domains on the system must run the latest Support Repository Update (SRU) of an Oracle Solaris 11 OS. Guest domains can also run the latest patch for the Oracle Solaris 10 1/13 OS.

These release notes might include some known issues that exist with older versions of the software.

---

For information about the supported hardware and fully qualified software and system firmware, see Chapter 1, “System Requirements” in *Oracle VM Server for SPARC 3.5 Installation Guide*.

---

**Note -** Oracle VM Server for SPARC features are added and maintained on the supported hardware platforms listed in “Supported Platforms” in *Oracle VM Server for SPARC 3.5 Installation Guide*. However, new features will not be added and existing features will not be maintained on hardware platforms that have been removed from the list.

As a rule, new Oracle VM Server for SPARC features and functionality are made available for all price-listed, supported SPARC T-series servers, SPARC M-series servers, and SPARC S-series servers from Oracle and Fujitsu M12 servers and Fujitsu M10 servers at the time that the Oracle VM Server for SPARC software is released and not for SPARC-based servers that have already passed their last-order date.

---

**Note -** The features that are described in this book can be used with all of the supported system software and hardware platforms that are listed in *Oracle VM Server for SPARC 3.5 Installation Guide*. However, some features are only available on a subset of the supported system software and hardware platforms. For information about these exceptions, see "What's New in This Release" in *Oracle VM Server for SPARC 3.5 Release Notes* and What's New in Oracle VM Server for SPARC Software (`http://www.oracle.com/technetwork/server-storage/vm/documentation/sparc-whatsnew-330281.html`).

# What's New in This Release

For information about the features introduced in all versions of the Oracle VM Server for SPARC (Logical Domains) software, see What's New in Oracle VM Server for SPARC Software (`http://www.oracle.com/technetwork/server-storage/vm/documentation/sparc-whatsnew-330281.html`).

The major changes for the Oracle VM Server for SPARC 3.5 software are as follows. Note that each feature is available on all supported platforms unless otherwise stated. For the list of supported platforms, see "Supported Platforms" in *Oracle VM Server for SPARC 3.5 Installation Guide*.

- Improve Oracle VM Server for SPARC logging and add command logging. See "Logging Oracle VM Server for SPARC Events" in *Oracle VM Server for SPARC 3.5 Administration Guide*.
- Add `ldm start-domain -f` option to start an I/O domain if the root domains are unavailable. See "Starting a Domain" in *Oracle VM Server for SPARC 3.5 Administration Guide*.
- Specify an amount of time to wait for domains to stop after you issue an `ldm stop-domain` command. See "Stopping a Domain" in *Oracle VM Server for SPARC 3.5 Administration Guide*.
- Add features to the XML and XMPP interfaces to match the CLI. See *Oracle VM Server for SPARC 3.5 Developer's Guide*.
- Show blacklisted and evacuation-pending core and memory resources. See "Using FMA to Blacklist or Unconfigure Faulty Resources" in *Oracle VM Server for SPARC 3.5 Administration Guide*.
- Add ADI support across live migration. See "Migration Restrictions for Silicon Secured Memory Servers" in *Oracle VM Server for SPARC 3.5 Administration Guide*.
- Enable domain migration when SR-IOV or PCIe cards are present. See "Migrating a Domain That Has an SR-IOV Ethernet Virtual Function Assigned" in *Oracle VM Server for SPARC 3.5 Administration Guide*.

- Increase the success rate of domain migrations between machines whose memory layouts were previously incompatible. See "Migration Requirements for Memory" in *Oracle VM Server for SPARC 3.5 Administration Guide*.
- Ensure that migrated domains do not reappear after an SP reset. See "Saving Post-Migration SP Configurations Automatically" in *Oracle VM Server for SPARC 3.5 Administration Guide*.
- Support native CPU-family migration. See "Domain Migration Requirements for CPUs" in *Oracle VM Server for SPARC 3.5 Administration Guide*.
- Standardize the memory allocation size.
- Support the migration of named resources. See "Migrating a Domain That Uses Named Resources" in *Oracle VM Server for SPARC 3.5 Administration Guide*.
- Support SCSI devices that have no LUN0. See "Simulating a LUN0" in *Oracle VM Server for SPARC 3.5 Administration Guide*.
- Support the virtual switch relay. See "Using a Virtual Switch Relay" in *Oracle VM Server for SPARC 3.5 Administration Guide*.
- Better handle MAC address collisions. See "Detecting MAC Address Collisions" in *Oracle VM Server for SPARC 3.5 Administration Guide*.
- Support for updating alternate MAC addresses dynamically. See "Dynamically Updating Alternate MAC Addresses" in *Oracle VM Server for SPARC 3.5 Administration Guide*.
- Add support for up to 999 virtual networks. See "Virtual Network Device" in *Oracle VM Server for SPARC 3.5 Administration Guide*.
- Make MIB enhancements. See *Oracle VM Server for SPARC 3.5 Management Information Base User's Guide*.
- Make XMPP enhancements to the `ldm list-io` and `ldm list-netdev` commands. See the `ldm`(1M) man page.
- Update the OVM template utilities to provide automatic disk expansion. See "Oracle VM Server for SPARC Template Features" in *Oracle VM Server for SPARC 3.5 Developer's Guide*.
- Update the OVM template utilities to create an OpenStack image. See "Oracle VM Server for SPARC Template Features" in *Oracle VM Server for SPARC 3.5 Developer's Guide*.
- Add support for cross-CPU live migration on Fujitsu M12 servers. See *Fujitsu M12 Server Product Notes*.
- Add support for setting the maximum page size of a logical domain. See *Fujitsu M12 and Fujitsu M10/SPARC M10 System Operation and Administration Guide*.

  This feature is supported only on Fujitsu M12 servers and Fujitsu M10 servers.
- Bug fixes.

# Features That Depend on System Firmware, the Oracle Solaris OS, or Both

Some of the Oracle VM Server for SPARC 3.5 features are available only when the fully qualified system firmware and Oracle Solaris OS are installed. Note that the control domain must already be running the fully qualified Oracle Solaris OS.

For information about the supported hardware, system firmware, and Oracle Solaris OS, see Chapter 1, "System Requirements" in *Oracle VM Server for SPARC 3.5 Installation Guide*. For Fujitsu M12 servers, see the latest *Fujitsu SPARC M12 Server Product Notes*. For Fujitsu M10 servers, see the latest *Fujitsu M10/SPARC M10 Server Product Notes*.

Generally, some Oracle VM Server for SPARC 3.5, features are available even if the system does not run the fully qualified system firmware.

If a guest domain, an I/O domain, or a root domain does not run the fully qualified Oracle Solaris OS version, the following Oracle VM Server for SPARC 3.5 features are not available:

- ADI migration
  - Requires at least Oracle Solaris 11.3 SRU 18
  - Requires any version of system firmware for at least the SPARC M7 series server or SPARC T7 series server
- SR-IOV Ethernet virtual function migration
  - Requires at least Oracle Solaris 11.3 SRU 9 on the domain to be migrated

# Oracle VM Server for SPARC 3.5 System Requirements

You can find information about the recommended and minimum software component versions to use with the Oracle VM Server for SPARC 3.5 release in Chapter 1, "System Requirements" in *Oracle VM Server for SPARC 3.5 Installation Guide*.

# Deprecated and Removed Oracle VM Server for SPARC Features

The following Oracle VM Server for SPARC feature is deprecated and will be removed from a future release of this software:

- Support for Logical Domains Manager based power management.

The following previously deprecated Oracle VM Server for SPARC feature will be removed from a future release of this software:

- Using the `ldm migrate-domain -p` *filename* command to initiate a non-interactive migration operation is deprecated. Instead, use SSL certificate-based authentication. See "Configuring SSL Certificates for Migration" in *Oracle VM Server for SPARC 3.5 Administration Guide*.

# Known Issues

This section contains general issues and specific bugs concerning the Oracle VM Server for SPARC 3.5 software.

## Bugs Affecting the Oracle VM Server for SPARC Software

This section summarizes the bugs that you might encounter when using this version of the software. The most recent bugs are described first. Workarounds and recovery procedures are specified, if available.

### Bugs Affecting the Oracle VM Server for SPARC 3.5 Software

#### SPARC M8 and SPARC T8 Series Servers: Oracle Solaris Might Experience a Kernel Panic or a Fatal Error When `cpu-arch=migration-class1` is Set on a Guest Domain

**Bug ID 27952673:** A SPARC M8 or SPARC T8 series server that runs the Oracle VM Server for SPARC 3.5, 3.5.0.1, or 3.5.0.2 software might experience a kernel panic or another fatal error condition if a guest domain has the `cpu-arch` property value set to `migration-class1`.

The panic might occur on any domain that has `cpu-arch=migration-class1` on the source system when the domain is migrated to a SPARC M8 or SPARC T8 series server.

In addition, a fatal error might occur on any domain that you create with `cpu-arch=migration-class1` on a SPARC M8 or SPARC T8 series server.

---

**Note -** No other SPARC servers are affected directly by this issue. However, a domain on another server that supports the `migration-class1` value of the `cpu-arch` property (at least the SPARC T4, SPARC M5, or SPARC S7 series server) becomes vulnerable if that domain is migrated to a SPARC M8 or a SPARC T8 series server.

To determine whether a guest domain is vulnerable to this issue, run the following command on the `primary` domain for each guest domain:

```
primary# ldm list -l domain-name | grep cpu-arch
```

A guest domain is vulnerable to this issue if the output shows `cpu-arch=migration-class1`.

---

**Cause:** The problem is triggered when the Oracle Solaris kernel or any user application attempts to reference a 2-Gbyte page on a SPARC M8 or SPARC T8 series server. The `cpu-arch=migration-class1` setting incorrectly permits 2-Gbyte page sizes on the underlying hardware platform, even though the SPARC M8 and SPARC T8 series servers do not support 2-Gbyte page sizes.

**Symptoms:** While the symptoms can vary, the Oracle Solaris kernel typically panics in the guest domain that has the `cpu-arch=migration-class1` setting. The associated panic string is `bad unexpected error from hypervisor call at TL 1`.

**Workaround:** Until an update is available, change the value of the `cpu-arch` property on all SPARC M8 and SPARC T8 series servers as follows:

- For a guest domain that you do not plan to live migrate or you plan to live migrate to another SPARC M8 or SPARC T8 series server, set `cpu-arch=native`.
- For a guest domain that you plan to live migrate to or from an older generation SPARC server, set `cpu-arch=generic`.

---

**Note -** Reboot the guest domain after you set the `cpu-arch` property value to make the change take effect.

---

### `ldmd` Crashes After a Failure to Remove Cores From a Domain

**Bug ID 26435797:** The `ldmd` daemon might dump core if a virtual CPU or CPU core removal operation fails. This failure might occur when all the CPUs in the target domain are bound or heavily loaded.

When this failure occurs, the `ldm remove-core` command might issue one of the following error messages:

```
Invalid response

Failed to receive version negotiation response from logical domain manager:
  Connection reset by peer
```

To perform the removal of virtual CPUs, you must unbind some of the CPUs from the target domain or lower the workload. Note that this problem does not affect virtual CPU removal operations on bound or unbound domains.

### `vsw-relay-mode` Behavior Reverts From `remote` to `local` When the Associated Service Domain Reboots

**Bug ID 26184111:** The `vsw-relay-mode` property is set on a virtual switch to enable reflective relay mode. This mode is not retained after a service domain reboot, so the state reverts to the default value of `local`.

**Workaround:** Run the following command on the virtual switch after the reboot that enables the `vsw-relay-mode` property:

```
primary# ldm set-vsw vsw-relay-mode=remote primary-vsw0
```

### Cross-CPU Migration Can Fail if Global Performance Counters are Enabled

**Bug ID 26047815:** In certain cross-CPU migration scenarios, a migration can fail with the following errors:

```
API group 0x20b v1.0 is not supported in the version of the firmware
 running on the target machine.
API group 0x214 v1.0 is not supported in the version of the firmware
 running on the target machine.
```

All of the following conditions must be present to encounter this problem:

- The domain has the `cpu-arch` property set to `generic` or `migration-class1`
- The domain has a `perf-counter` property setting that includes the `global` value
- The domain was booted on at least a SPARC M7 series server or a SPARC T7 series server
- The target machine is a platform released prior to the SPARC M7 series server or SPARC T7 series server

This problem occurs because a domain booted on at least a SPARC M7 series server or a SPARC T7 series server with a `perf-counter` property setting that includes the `global` value

will register platform-specific performance counter Hypervisor interfaces that do not exist on older platforms. As part of the migration, a check is performed to ensure that all the interfaces used by the domain are present on the target machine. When these SPARC M7 series server or SPARC T7 series server specific interfaces are detected, the migration is aborted.

**Workaround:** Do not set `perf-counter=global` if `cpu-arch` is note `native` and at least SPARC M7 series servers and SPARC T7 series servers are part of the migration pool.

## Virtual SCSI HBA Subsystem Does Not Support All SCSI Enclosure Services Devices

**Bug ID 25865708:** An SES device that is seen by the Oracle Solaris OS as a secondary function is an SES device type that cannot be supported by `vhba`. `vhba` can support an SES device whose device type has a value of `0xd` as specified in the `inq_dtype` field of the `INQUIRY` payload.

When the `vhba` binary in the guest domain attempts to initialize some SCSI enclosure services (SES) devices, `vhba` causes `scsi` to issue the following warning message:

```
... scsi: WARNING: scsi_enumeration_failed: vhba2 probe@w50080e51bfd32004,0,d
 enumeration failed during tran_tgt_init
```

The `,d` substring represents the `0xd` hexadecimal digit, which is the SCSI industry standard code for an SES device. The `,d` string indicates that this warning message is a result of an unsupported type of SES device.

`vhba` can support an SES device that has a device type of `0xd` that is specified in the `inq_dtype` field of the `INQUIRY` payload:

```
# mdb -k
> ::vsan
vsan_t( 6400126e08c0 ) cfg-hdl(0) iport-path(/pci@300/pci@1/pci@0/pci@4/SUNW,emlxs@0,11/
fp@0,0)
    vsan_iport_t( 6400125b8710 )
        vsan_tport_t( 64001bf89718 ) tport_phys(w216000c0ff8089d5)
        vsan_lun_t( 640011aa65d0 ) lun(0) vlun-id(1127b) []

> 640011aa65d0::print vsan_lun_t vl_sd |::print struct scsi_device sd_inq |::print
 struct scsi_inquiry inq_dtype
inq_dtype = d
```

## `ldomHbaTable` Is Empty

**Bug ID 24393532:** The fix for bug ID 23591953 disabled both Oracle VM Server for SPARC Oracle VM Server for SPARC MIB monitoring, such as listing the Oracle VM

Server for SPARC MIB objects by using the `snmpwalk` command, and trap generation for the `ldomHbaTable` table. As a result, the Oracle VM Server for SPARC MIB `ldomHbaTable` table does not show contents.

```
primary# snmpwalk -v1 -c public localhost SUN-LDOM-MIB::ldomHbaTable
primary#
```

**Workaround:** Use the `ldm list-hba` command to view the HBA information.

### Inaccurate `Unable to Send Suspend Request` Error Reported During a Successful Domain Migration

**Bug ID 23206413:** In rare circumstances, a successful domain migration reports the following error:

```
Unable to send suspend request to domain domain-name
```

This issue occurs when the Logical Domains Manager detects an error while suspending the domain, and the Logical Domains Manager is able to recover and completes the migration. The exit status of the command is 0, reflecting the successful migration.

**Workaround:** Because the migration completes successfully, you can ignore the error message.

### Cold Migrating a Bound Domain With Many Virtual Devices Might Fail and Leave Two Bound Copies of the Domain

**Bug ID 23180427:** When cold migrating a bound domain that has a large number of virtual devices, the operation might fail with the following message in the SMF log:

```
warning: Timer expired: Failed to read feasibility response type (9) from target LDoms
 Manager
```

This failure indicates that the Logical Domains Manager running on the source machine timed out while waiting for the domain to be bound on the target machine. The chances of encountering this problem increases as the number of virtual devices in the migrating domain increases.

The timing of this failure results in a bound copy of the domain on both the source machine and the target machine. Do not start both copies of this domain. This action can cause data corruption because both domains reference the same virtual disk backends.

**Recovery:** After verifying that the copy of the migrated domain is correct on the target machine, manually unbind the copy of the domain on the source machine and destroy it.

### Migration Fails When the Target Machine Has Insufficient Free LDCs

**Bug ID 23031413:** When the target machine's control domain runs out of LDCs during a domain migration, the migration fails with no explanation and the following message is written to the SMF log:

```
warning: Failed to read feasibility response type (5) from target LDoms Manager
```

This error is issued when the domain being migrated fails to bind on the target machine. Note that the bind operation might fail for other reasons on the target machine, as well.

**Workaround:** For the migration to succeed, the number of LDCs must be reduced either in the domain being migrated or in the control domain of the target machine. You can reduce the number of LDCs by reducing the number of virtual devices being used by or being serviced by a domain. For more information about managing LDCs, see "Using Logical Domain Channels" in *Oracle VM Server for SPARC 3.5 Administration Guide*.

### `ovmtlibrary` Limits Disk Image File Name to 50 Characters

**Bug ID 23024583:** The `ovmtlibrary` command limits the disk image file name to 50 characters. The `ovmtlibrary` checks the `.ovf` file and compares the information in the `<ovf:References>` section with the actual file names of the decompressed disks.

An error is issued if the files are different or if the disk image file name is longer than 50 characters. For example:

```
# ovmtlibrary -c store -d "example" -q -o file:/template.ova -l /export/user1/
ovmtlibrary_example
event id is 3
ERROR: The actual disk image file name(s) or the actual number of disk
image(s) is different from OVF file: template.ovf
exit code: 1
```

The following example XML shows a disk image file name that is greater than 50 characters:

```
<ovf:References>
<ovf:File ovf:compression="gzip"
ovf:href="disk_image.ldoms3.5_build_s11_u3_sru15_01_kz_42G.img.gz"
```

```
ovf:id="ldoms3" ovf:size="6687633773"/>
</ovf:References>
```

**Workaround:** Limit the length of disk image file names to fewer than 50 characters.

### Virtual Network Devices Added to an Inactive Guest Domain Never Gets the Default `linkprop` Value

**Bug ID 22842188:** For `linkprop=phys-state` to be supported on a virtual network device, the Logical Domains Manager must be able to validate that the virtual switch to which the virtual network device is attached has a physical NIC that backs the virtual switch.

The Oracle VM Server for SPARC `netsvc` agent must be running on the guest domain so that the virtual switch can be queried.

If the guest domain is not active and cannot communicate with the agent in the domain that has the virtual network device's virtual switch, the virtual network device does not have `linkprop=phys-state` set.

**Workaround:** Only set `linkprop=phys-state` when the domain is active.

### `ldm set-vsw net-dev=` Fails When `linkprop=phys-state`

**Bug ID 22828100:** If a virtual switch has attached virtual network devices that have `linkprop=phys-state`, the virtual switch to which they are attached must have a valid backing NIC device specified by the `net-dev` property. The `net-dev` property value must be the name of a valid network device.

If this action is performed using `net-dev=`, the virtual switch still shows `linkprop=phys-state` even though the `net-dev` property value is not a valid NIC device.

**Workaround:** First, remove all the virtual network devices that are attached to the virtual switch, and then remove the virtual switch. Then, re-create the virtual switch with a valid `net-dev` backing device, and then re-create all the virtual network devices.

### A Domain That Has Socket Constraints Cannot Be Re-Created From an XML File

**Bug ID 21616429:** The Oracle VM Server for SPARC 3.3 software introduced socket support for Fujitsu M12 servers and Fujitsu M10 servers only.

Software running on Oracle SPARC servers and Oracle VM Server for SPARC versions older than 3.3 cannot re-create a domain with socket constraints from an XML file.

Attempting to re-create a domain with socket constraints from an XML file with an older version of the Oracle VM Server for SPARC software or on an Oracle SPARC server fails with the following message:

```
primary# ldm add-domain -i ovm3.3_socket_ovm11.xml
socket not a known resource
```

If Oracle VM Server for SPARC 3.2 is running on a Fujitsu M12 server or Fujitsu M10 server and you attempt to re-create a domain with socket constraints from an XML file, the command fails with various error messages, such as the following:

```
primary# ldm add-domain -i ovm3.3_socket_ovm11.xml
Unknown property: vcpus

primary# ldm add-domain -i ovm3.3_socket_ovm11.xml
perf-counters property not supported, platform does not have
performance register access capability, ignoring constraint setting.
```

**Workaround:** Edit the XML file to remove any sections that reference the `socket` resource type.

## Kernel Zones Block Live Migration of Guest Domains

**Bug ID 21289174:** On a SPARC server, a running kernel zone within an Oracle VM Server for SPARC domain will block live migration of the guest domain. The following error message is shown:

```
Guest suspension failed because Kernel Zones are active.
Stop Kernel Zones and retry.
```

**Workaround:** Choose one of the following workarounds:

- Stop running the kernel zone.

  **# zoneadm -z** *zonename* **shutdown**

- Suspend the kernel zone.

```
# zoneadm -z zonename suspend
```

- Perform a live migration of the kernel zone to another system before migrating the guest domain.

  See Chapter 3, “Migrating an Oracle Solaris Kernel Zone” in *Creating and Using Oracle Solaris Kernel Zones*.

### After Dropping Into `factory-default`, Recovery Mode Fails if the System Boots From a Different Device Than the One Booted in the Previously Active Configuration

**Bug ID 20425271:** While triggering a recovery after dropping into `factory-default`, recovery mode fails if the system boots from a different device than the one booted in the previously active configuration. This failure might occur if the active configuration uses a boot device other than the `factory-default` boot device.

---

**Note -** This problem applies to UltraSPARC T2, UltraSPARC T2 Plus, SPARC T3, and SPARC T4 series servers. This problem also applies to SPARC T5, SPARC M5, and SPARC M6 series servers that run a system firmware version prior to 9.5.3.

---

**Workaround:** Perform the following steps any time you want to save a new configuration to the SP:

1. Determine the full PCI path to the boot device for the `primary` domain.

   Use this path for the `ldm set-var` command in Step 4.

2. Remove any currently set `boot-device` property from the `primary` domain.

   Performing this step is necessary only if the `boot-device` property has a value set. If the property does not have a value set, an attempt to remove the `boot-device` property results in the `boot-device not found` message.

   ```
   primary# ldm rm-var boot-device primary
   ```

3. Save the current configuration to the SP.

   ```
   primary# ldm add-spconfig config-name
   ```

4. Explicitly set the `boot-device` property for the `primary` domain.

   ```
   primary# ldm set-var boot-device=value primary
   ```

   If you set the `boot-device` property after saving the configuration to the SP as described, the specified boot device is booted when recovery mode is triggered.

**Recovery:** If recovery mode has already failed as described, perform the following steps:

1. Explicitly set the boot device to the one used in the last running configuration.

   ```
   primary# ldm set-var boot-device=value primary
   ```

2. Reboot the `primary` domain.

   ```
   primary# reboot
   ```

   The reboot enables the recovery to proceed.

## Guest Domain `eeprom` Updates Are Lost if an `ldm add-spconfig` Operation Is Not Complete

**Bug ID 19932842:** An attempt to set an OBP variable from a guest domain might fail if you use the `eeprom` or the OBP command before one of the following commands is completed:

- `ldm add-spconfig`
- `ldm remove-spconfig`
- `ldm set-spconfig`
- `ldm bind`

This problem might occur when these commands take more than 15 seconds to complete.

```
# /usr/sbin/eeprom boot-file\=-k
promif_ldom_setprop: promif_ldom_setprop: ds response timeout
eeprom: OPROMSETOPT: Invalid argument
boot-file: invalid property
```

**Recovery:** Retry the `eeprom` or OBP command after the `ldm` operation has completed.

**Workaround:** Retry the `eeprom` or OBP command on the affected guest domain. You might be able to avoid the problem by using the `ldm set-var` command on the `primary` domain.

## Rebooting a Guest Domain With More Than 1000 Virtual Network Devices Results in a Panic

**Bug ID 19449221:** A domain can have no more than 999 virtual network devices (`vnets`).

**Workaround:** Limit the number of `vnets` on a domain to 999.

### Incorrect Device Path for Fibre Channel Virtual Functions in a Root Domain

**Bug ID 18001028:** In the root domain, the Oracle Solaris device path for a Fibre Channel virtual function is incorrect.

For example, the incorrect path name is `pci@380/pci@1/pci@0/pci@6/fibre-channel@0,2` while it should be `pci@380/pci@1/pci@0/pci@6/SUNW,emlxs@0,2`.

The `ldm list-io -l` output shows the correct device path for the Fibre Channel virtual functions.

**Workaround:** None.

### Oracle Solaris 11.3 SRU 12: `ssd` and `sd` Driver Functionality Is Merged for Fibre Channel Devices on SPARC Platforms

**Bug ID 17036795:** The Oracle Solaris 11.3 SRU 12 OS has merged the `ssd` and `sd` driver functionality for Fibre Channel devices on SPARC platforms.

This change affects device node names on the physical device path. The device node names change from `ssd@` to `disk@`. This change also affects device driver bindings from `ssd` to `sd`.

---

**Note -** Ensure that any application or client in the Oracle Solaris OS system that depends on these device node names or device driver bindings is adjusted.

---

This change is not enabled by default for Oracle Solaris 11.3 systems.

You must enable this change to perform live migrations of domains that use virtual HBA and Fibre Channel devices.

Before you enable this change, ensure that MPxIO is already enabled by running the `stmsboot -D fp -e` command.

Run the `format` command to determine whether MPxIO is enabled. When enabled, you should see `vhci` in device names. Alternatively, if the `mpathadm -list lu` output is empty, no MPxIO devices are enumerated.

Use the `beadm` command to create a new boot environment (BE). By using BEs, you can roll back easily to a previous boot environment if you experience unexpected problems.

Mount the BE and replace the `/etc/devices/inception_points` file with the `/etc/devices/inception_points.vhba` file. The `.vhba` file includes some feature flags to enable this change.

Finally, reboot after you activate the new BE.

```
# beadm create BE-name
# beadm mount BE-name /mnt
# cp /mnt/etc/devices/inception_points.vhba /mnt/etc/devices/inception_points
# beadm umount BE-name
# beadm activate BE-name
# reboot
```

After rebooting, use the `prtconf -D | grep driver | grep sd` command to verify the change.

If any disks use the `ssd` driver, there is a problem with the configuration.

You can also use the `mpathadm list lu` command to show multiple paths to the same disks if virtual HBA and the FibreChannel virtual function are both configured to see the same LUNs.

### Misleading Messages Shown for InfiniBand SR-IOV Remove Operations

**Bug ID 16979993:** An attempt to use a dynamic SR-IOV remove operation on an InfiniBand device results in confusing and inappropriate error messages.

Dynamic SR-IOV remove operations are not supported for InfiniBand devices.

**Workaround:** Remove InfiniBand virtual functions by performing one of the following procedures:

- "How to Remove an InfiniBand Virtual Function From an I/O Domain" in *Oracle VM Server for SPARC 3.5 Administration Guide*
- "How to Remove an InfiniBand Virtual Function From a Root Domain" in *Oracle VM Server for SPARC 3.5 Administration Guide*

### `ldm migrate -n` Should Fail When Performing a Cross-CPU Migration From SPARC T5, SPARC M5, or SPARC M6 Server to UltraSPARC T2 or SPARC T3 Server

**Bug ID 16864417:** The `ldm migrate -n` command does not report failure when attempting to migrate between a SPARC T5, SPARC M5, or SPARC M6 server and an UltraSPARC T2 or SPARC T3 server.

**Workaround:** None.

### Resilient I/O Domain Should Support PCI Device Configuration Changes After the Root Domain Is Rebooted

**Bug ID 16691046:** If virtual functions are assigned from the root domain, an I/O domain might fail to provide resiliency in the following hotplug situations:

- You add a root complex (PCIe bus) dynamically to the root domain, and then you create the virtual functions and assign them to the I/O domain.
- You hot-add an SR-IOV card to the root domain that owns the root complex, and then you create the virtual functions and assign them to the I/O domain.
- You replace or add any PCIe card to an empty slot (either through hotplug or when the root domain is down) on the root complex that is owned by the root domain. This root domain provides virtual functions from the root complex to the I/O domain.

**Workaround:** Perform one of the following steps:

- If the root complex already provides virtual functions to the I/O domain and you add, remove, or replace any PCIe card on that root complex (through hotplug or when the root domain is down), you must reboot both the root domain and the I/O domain.
- If the root complex does not have virtual functions currently assigned to the I/O domain and you add an SR-IOV card or any other PCIe card to the root complex, you must stop the root domain to add the PCIe card. After the root domain reboots, you can assign virtual functions from that root complex to the I/O domain.
- If you want to add a new PCIe bus to the root domain and then create and assign virtual functions from that bus to the I/O domain, perform one of the following steps and then reboot the root domain:
  - Add the bus during a delayed reconfiguration
  - Add the bus dynamically

### Guest Domains in Transition State After Reboot of the `primary` Domain

**Bug ID 16659506:** A guest domain is in transition state (`t`) after a reboot of the `primary` domain. This problem arises when a large number of virtual functions are configured on the system.

**Workaround:** To avoid this problem, retry the OBP disk boot command several times to avoid a boot from the network.

Perform the following steps on each domain:

1. Access the console of the domain.

   ```
   primary# telnet localhost 5000
   ```

2. Set the `boot-device` property.

   ```
   ok> setenv boot-device disk disk disk disk disk disk disk disk disk disk net
   ```

   The number of `disk` entries that you specify as the value of the `boot-device` property depends on the number of virtual functions that are configured on the system. On smaller systems, you might be able to include fewer instances of `disk` in the property value.

3. Verify that the `boot-device` property is set correctly by using the `printenv`.

   ```
   ok> printenv
   ```

4. Return to the `primary` domain console.
5. Repeat Steps 1-4 for each domain on the system.
6. Reboot the `primary` domain.

   ```
   primary# shutdown -i6 -g0 -y
   ```

### `WARNING: ddi_intr_alloc: cannot fit into interrupt pool` Means That Interrupt Supply Is Exhausted While Attaching I/O Device Drivers

**Bug ID 16284767:** This warning on the Oracle Solaris console means the interrupt supply was exhausted while attaching I/O device drivers:

```
WARNING: ddi_intr_alloc: cannot fit into interrupt pool
```

This limitation applies only to the supported SPARC systems prior to the SPARC M7 series servers and SPARC T7 series servers.

The hardware provides a finite number of interrupts, so Oracle Solaris limits how many each device can use. A default limit is designed to match the needs of typical system configurations, however this limit may need adjustment for certain system configurations.

Specifically, the limit may need adjustment if the system is partitioned into multiple logical domains and if too many I/O devices are assigned to any guest domain. Oracle VM Server for SPARC divides the total interrupts into smaller sets given to guest domains. If too many I/O devices are assigned to a guest domain, its supply might be too small to give each device the default limit of interrupts. Thus, it exhausts its supply before it completely attaches all the drivers.

Some drivers provide an optional callback routine which allows Oracle Solaris to automatically adjust their interrupts. The default limit does not apply to these drivers.

**Workaround:** Use the `::irmpools` and `::irmreqs` MDB macros to determine how interrupts are used. The `::irmpools` macro shows the overall supply of interrupts divided into pools. The `::irmreqs` macro shows which devices are mapped to each pool. For each device, `::irmreqs` shows whether the default limit is enforced by an optional callback routine, how many interrupts each driver requested, and how many interrupts the driver is given.

The macros do not show information about drivers that failed to attach. However, the information that is shown helps calculate the extent to which you can adjust the default limit. Any device that uses more than one interrupt without providing a callback routine can be forced to use fewer interrupts by adjusting the default limit. Reducing the default limit below the amount that is used by such a device results in freeing of interrupts for use by other devices.

To adjust the default limit, set the `ddi_msix_alloc_limit` property to a value from `1` to `8` in the `/etc/system` file. Then, reboot the system for the change to take effect.

To maximize performance, start by assigning larger values and decrease the values in small increments until the system boots successfully without any warnings. Use the `::irmpools` and `::irmreqs` macros to measure the adjustment's impact on all attached drivers.

For example, suppose the following warnings are issued while booting the Oracle Solaris OS in a guest domain:

```
WARNING: emlxs3: interrupt pool too full.
WARNING: ddi_intr_alloc: cannot fit into interrupt pool
```

The `::irmpools` and `::irmreqs` macros show the following information:

```
# echo "::irmpools" | mdb -k
ADDR             OWNER   TYPE   SIZE  REQUESTED  RESERVED
00000400016be970 px#0    MSI/X  36    36         36

# echo "00000400016be970::irmreqs" | mdb -k
ADDR             OWNER   TYPE   CALLBACK NINTRS NREQ NAVAIL
00001000143acaa8 emlxs#0 MSI-X  No       32     8    8
00001000170199f8 emlxs#1 MSI-X  No       32     8    8
```

```
000010001400ca28 emlxs#2 MSI-X  No       32     8    8
0000100016151328 igb#3   MSI-X  No       10     3    3
0000100019549d30 igb#2   MSI-X  No       10     3    3
0000040000e0f878 igb#1   MSI-X  No       10     3    3
000010001955a5c8 igb#0   MSI-X  No       10     3    3
```

The default limit in this example is eight interrupts per device, which is not enough interrupts to accommodate the attachment of the final `emlxs3` device to the system. Assuming that all `emlxs` instances behave in the same way, `emlxs3` probably requested 8 interrupts.

By subtracting the 12 interrupts used by all of the `igb` devices from the total pool size of 36 interrupts, 24 interrupts are available for the `emlxs` devices. Dividing the 24 interrupts by 4 suggests that 6 interrupts per device would enable all `emlxs` devices to attach with equal performance. So, the following adjustment is added to the `/etc/system` file:

```
set ddi_msix_alloc_limit = 6
```

When the system successfully boots without warnings, the `::irmpools` and `::irmreqs` macros show the following updated information:

```
# echo "::irmpools" | mdb -k
ADDR             OWNER   TYPE   SIZE  REQUESTED  RESERVED
00000400018ca868 px#0    MSI/X  36    36         36

# echo "00000400018ca868::irmreqs" | mdb -k
ADDR             OWNER   TYPE   CALLBACK NINTRS NREQ NAVAIL
0000100016143218 emlxs#0 MSI-X  No       32     8    6
0000100014269920 emlxs#1 MSI-X  No       32     8    6
000010001540be30 emlxs#2 MSI-X  No       32     8    6
00001000140cbe10 emlxs#3 MSI-X  No       32     8    6
00001000141210c0 igb#3   MSI-X  No       10     3    3
0000100017549d38 igb#2   MSI-X  No       10     3    3
0000040001ceac40 igb#1   MSI-X  No       10     3    3
000010001acc3480 igb#0   MSI-X  No       10     3    3
```

## SPARC T5-8 Server: Uptime Data Shows a Value of 0 for Some `ldm` List Commands

**Bug ID 16068376:** On a SPARC T5-8 server with approximately 128 domains, some `ldm` commands such as `ldm list` might show `0` seconds as the uptime for all domains.

**Workaround:** Log in to the domain and use the `uptime` command to determine the domain's uptime.

### `ldm list -o status` on Control Domain Reports Incorrect Migration Progress

**Bug ID 15819714:** In rare circumstances, the `ldm list -o status` command reports an incorrect completion percentage when used to observe the status of a migration on a control domain.

This problem has no impact on the domain that is being migrated or on the `ldmd` daemons on the source or target control domains.

**Workaround:** Run the `ldm list -o status` command on the other control domain that is involved in the migration to observe the progress.

### `ldm init-system` Command Might Not Correctly Restore a Domain Configuration on Which Physical I/O Changes Have Been Made

**Bug ID 15783031:** You might experience problems when you use the `ldm init-system` command to restore a domain configuration that has used direct I/O or SR-IOV operations.

A problem arises if one or more of the following operations have been performed on the configuration to be restored:

- A slot has been removed from a bus that is still owned by the `primary` domain.
- A virtual function has been created from a physical function that is owned by the `primary` domain.
- A virtual function has been assigned to the `primary` domain, to other guest domains, or to both.
- A root complex has been removed from the `primary` domain and assigned to a guest domain, and that root complex is used as the basis for further I/O virtualization operations.

  In other words, you created a non-`primary` root domain and performed any of the previous operations.

If you have performed any of the previous actions, perform the workaround shown in Oracle VM Server for SPARC PCIe Direct I/O and SR-IOV Features (Doc ID 1325454.1) (`https://support.oracle.com/epmos/faces/SearchDocDisplay?amp;_adf.ctrl-state=10c69raljg_77&_afrLoop=506200315473090`).

### Guest Domain Panics When Running the `cputrack` Command During a Migration to a SPARC T4 Server

**Bug ID 15776123:** If the `cputrack` command is run on a guest domain while that domain is migrated to a SPARC T4 server, the guest domain might panic on the target machine after it has been migrated.

**Workaround:** Do not run the `cputrack` command during the migration of a guest domain to a SPARC T4 server.

### Limit the Maximum Number of Virtual Functions That Can Be Assigned to a Domain

**Bug ID 15775637:** An I/O domain has a limit on the number of interrupt resources that are available per root complex.

On SPARC T3 and SPARC T4 servers, the limit is approximately 63 MSI/X vectors. Each `igb` virtual function uses three interrupts. The `ixgbe` virtual function uses two interrupts.

If you assign a large number of virtual functions to a domain, the domain runs out of system resources to support these devices. You might see messages similar to the following:

```
WARNING: ixgbevf32: interrupt pool too full.
WARNING: ddi_intr_alloc: cannot fit into interrupt pool
```

### Trying to Connect to Guest Domain Console While It Is Being Bound Might Cause Input to Be Blocked

**Bug ID 15771384:** A domain's guest console might freeze if repeated attempts are made to connect to the console before and during the time the console is bound. For example, this might occur if you use an automated script to grab the console as a domain is being migrated onto the machine.

**Workaround:** To unfreeze console, perform the following commands on the domain that hosts the domain's console concentrator (usually the control domain):

```
primary# svcadm disable vntsd
primary# svcadm enable vntsd
```

### `ldm remove-io` of PCIe Cards That Have PCIe-to-PCI Bridges Should Be Disallowed

**Bug ID 15761509:** Use only the PCIe cards that support the Direct I/O (DIO) feature, which are listed in this support document (`https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&doctype=REFERENCE&id=1325454.1`).

---

**Note -** The direct I/O feature is deprecated starting with the SPARC T7 series servers and the SPARC M7 series servers.

---

**Workaround:** Use the `ldm add-io` command to add the card to the `primary` domain again.

### Live Migration of a Domain That Depends on an Inactive Master Domain on the Target Machine Causes `ldmd` to Fault With a Segmentation Fault

**Bug ID 15701865:** If you attempt a live migration of a domain that depends on an inactive domain on the target machine, the `ldmd` daemon faults with a segmentation fault and crashes. The `ldmd` daemon is restarted automatically, but the migration is aborted.

**Workaround:** Perform one of the following actions before you attempt the live migration:

- Remove the guest dependency from the domain to be migrated.
- Start the master domain on the target machine.

### DRM and `ldm list` Output Shows a Different Number of Virtual CPUs Than Are Actually in the Guest Domain

**Bug ID 15701853:** A `No response` message might appear in the Oracle VM Server for SPARC log when a loaded domain's DRM policy expires after the CPU count has been substantially reduced. The `ldm list` output shows that more CPU resources are allocated to the domain than is shown in the `psrinfo` output.

**Workaround:** Use the `ldm set-vcpu` command to reset the number of CPUs on the domain to the value that is shown in the `psrinfo` output.

### Simultaneous Migration Operations in “Opposite Direction” Might Cause `ldm` to Hang

**Bug ID 15696986:** If two `ldm migrate` commands are issued between the same two systems simultaneously in the “opposite direction,” the two commands might hang and never complete. An opposite direction situation occurs when you simultaneously start a migration on machine A to machine B and a migration on machine B to machine A.

The hang occurs even if the migration processes are initiated as dry runs by using the `-n` option. When this problem occurs, all other `ldm` commands might hang.

**Recovery:** Restart the Logical Domains Manager on both the source machine and the target machine:

```
primary# svcadm restart ldmd
```

**Workaround:** None.

### SPARC T3-1 Server: Issue With Disks That Are Accessible Through Multiple Direct I/O Paths

**Bug ID 15668368:** A SPARC T3-1 server can be installed with dual-ported disks, which can be accessed by two different direct I/O devices. In this case, assigning these two direct I/O devices to different domains can cause the disks to be used by both domains and affect each other based on the actual usage of those disks.

**Workaround:** Do not assign direct I/O devices that have access to the same set of disks to different I/O domains. To determine whether you have dual-ported disks on a SPARC T3-1 server, run the following command on the SP:

```
-> show /SYS/SASBP
```

If the output includes the following `fru_description` value, the corresponding system has dual-ported disks:

```
fru_description = BD,SAS2,16DSK,LOUISE
```

If dual disks are found to be present in the system, ensure that both of the following direct I/O devices are always assigned to the same domain:

```
pci@400/pci@1/pci@0/pci@4  /SYS/MB/SASHBA0
pci@400/pci@2/pci@0/pci@4  /SYS/MB/SASHBA1
```

### Using the `ldm stop -a` Command on Domains in a Master-Slave Relationship Leaves the Slave With the `stopping` Flag Set

**Bug ID 15664666:** When a reset dependency is created, an `ldm stop -a` command might result in a domain with a reset dependency being restarted instead of only stopped.

**Workaround:** First, issue the `ldm stop` command to the master domain. Then, issue the `ldm stop` command to the slave domain. If the initial stop of the slave domain results in a failure, issue the `ldm stop -f` command to the slave domain.

### Dynamically Removing All the Cryptographic Units From a Domain Causes SSH to Terminate

**Bug ID 15600969:** If all the hardware cryptographic units are dynamically removed from a running domain, the cryptographic framework fails to seamlessly switch to the software cryptographic providers, and kills all the `ssh` connections.

This issue only applies to UltraSPARC T2, UltraSPARC T2 Plus and SPARC T3 servers.

**Recovery:** Re-establish the `ssh` connections after all the cryptographic units are removed from the domain.

**Workaround:** Set `UseOpenSSLEngine=no` in the `/etc/ssh/sshd_config` file on the server side, and run the `svcadm restart ssh` command.

All `ssh` connections will no longer use the hardware cryptographic units (and thus not benefit from the associated performance improvements), and `ssh` connections will not be disconnected when the cryptographic units are removed.

### The Logical Domains Manager Does Not Start if the Machine Is Not Networked and an NIS Client Is Running

**Bug ID 15518409:** If you do not have a network configured on your machine and have a Network Information Services (NIS) client running, the Logical Domains Manager will not start on your system.

**Workaround:** Disable the NIS client on your non-networked machine:

```
# svcadm disable nis/client
```

### Cannot Connect to Migrated Domain's Console Unless `vntsd` Is Restarted

**Bug ID 15513998:** Occasionally, after a domain has been migrated, it is not possible to connect to the console for that domain.

Note that this problem occurs when the migrated domain is running an OS version older than Oracle Solaris 11.3.

**Workaround:** Restart the `vntsd` SMF service to enable connections to the console:

```
# svcadm restart vntsd
```

**Note -** This command will disconnect all active console connections.

### Simultaneous Net Installation of Multiple Domains Fails When in a Common Console Group

**Bug ID 15453968:** Simultaneous net installation of multiple guest domains fails on systems that have a common console group.

**Workaround:** Only net-install on guest domains that each have their own console group. This failure is seen only on domains with a common console group shared among multiple net-installing domains.

### Behavior of the `ldm stop-domain` Command Can Be Confusing

**Bug ID 15368170:** In some cases, the behavior of the `ldm stop-domain` command is confusing.

```
# ldm stop-domain -f domain-name
```

If the domain is at the kernel module debugger, `kmdb(1)`, prompt, then the `ldm stop-domain` command fails with the following error message:

`LDom <domain-name> stop notification failed`

# Documentation Issues

This section contains documentation issues and errors that have been found too late to resolve for the Oracle VM Server for SPARC 3.5 release.

---

**Note -** The changes described in the following documentation errata have been made to the English version of *Oracle VM Server for SPARC 3.5 Reference Manual* on OTN.

These changes are not reflected in the man pages delivered with the Oracle VM Server for SPARC 3.5 software product or in the Japanese version of *Oracle VM Server for SPARC 3.5 Reference Manual* on OTN.

---

## `ldmd`(1M): Missing Description of the `ldmd/migration_adi_legacy_compat` SMF Property

The `ldmd`(1M) man page is missing the following description of the `ldmd/migration_adi_legacy_compat` SMF property:

`ldmd/migration_adi_legacy_compat`

Specifies whether to permit a domain migration between servers that support Silicon Secured Memory (SSM) even if one of the machines does not have support for the migration of Application Data Integrity (ADI) version information that is introduced in Oracle VM Server for SPARC 3.5.

If both the source machine and the target machine are running the latest versions of the Oracle VM Server for SPARC software, you do not need to use this SMF property.

---

**Caution -** If you intend to perform a domain migration on your servers that support SSM, it is best that they run at least the Oracle VM Server for SPARC 3.5 software. If this is not possible, take extreme caution when using the `ldmd/migration_adi_legacy_compat` SMF property. Improper use of this property can result in undefined application behavior if ADI is in use in the domain being migrated.

---

By default, the property value is `false`, which prevents a domain migration unless both the source machine and the target machine support SSM and run the required version of the Oracle VM Server for SPARC software. This property has no effect on servers that do not support SSM.

When the value is `true`, the domain migration proceeds without support for the migration of ADI version information.

So, if either the source machine or target machine runs a version of the Oracle VM Server for SPARC software that is older than 3.5, which does not support the migration of ADI version information, the migration is permitted.

Only set the `ldmd/migration_adi_legacy_compat` SMF property value to `true` if both the following circumstances are true:

- You cannot upgrade both the source machine and target machine to a version of the Oracle VM Server for SPARC software that supports the migration of ADI version information
- You know for certain that ADI versioning is not in use within the domain to be migrated

Setting this property to `true` permits migrations where ADI version information is not transferred to the target machine. This situation can result in undefined application behavior if ADI is in use in the domain being migrated.

The `ldmd/migration_adi_legacy_compat` SMF property is not recognized by Oracle VM Server for SPARC versions older than 3.5. Use of this property is applicable only on a source machine or a target machine is running at least Oracle VM Server for SPARC 3.5.

## `ldm`(1M): Updated Description of the `set-domain` Subcommand and the `-i` Option

The `ldm`(1M) man page includes the following updates:

- The first paragraph now reads as follows:

  > The `set-domain` subcommand enables you to modify properties such as `boot-policy`, `mac-addr`, `hostid`, `failure-policy`, `extended-mapin-space`, `master`, and `max-cores` for a domain. You *cannot* use this command to update resources.

- The description of the `-i` now reads as follows:

  > `-i` *file* specifies the XML configuration file to use in setting the properties of the logical domain.
  >
  > *Only* the `ldom_info` nodes specified in the XML file are parsed. Resource nodes, such as `vcpu`, `mau`, and `memory`, are ignored.
  >
  > If the `hostid` property in the XML file is already in use, the `ldm set-domain -i` command fails with the following error:

`Hostid` *host-ID* `is already in use`

Before you re-run the `ldm set-domain -i` command, remove the `hostid` entry from the XML file.

### `ldm`(1M) Incorrectly References the Command History Buffer

The `ldm`(1M) man page incorrectly refers to a command history buffer that you can view by using the `ldm list-history` command.

The first and second paragraphs of the Command History section have been updated with the following paragraphs:

> Use the `ldm list-history` command to view the Oracle VM Server for SPARC command history log. This log captures `ldm` commands and commands that are issued through the XMPP interface. By default, the number of commands shown by the `ldm list-history` command is ten.
>
> To change the number of commands output by the `ldm list-history` command, use the `ldm set-logctl` command to set the `history` property value. If you set `history=0`, the saving of command history is disabled. You can re-enable this feature by setting the `history` property to a non-zero value.

The description of the `history` property in the Control Logging Operations section has been updated as follows:

> `history=`*num* specifies the number of commands output by the `ldm list-history` command. Setting the value to 0 disables the saving of command history.

The description of the `-a` option in the View Logging Capabilities section has been updated as follows:

> `-a` shows the logging capability values for all logging types and the number of commands output by the `ldm list-history` command.

## Resolved Issues

The following enhancement requests and bugs have been fixed for the Oracle VM Server for SPARC 3.5 software release:

| | |
|---|---|
| 15527921 | Migrate explicit console group/port bindings |
| 15639066 | ldm stop ldg1 &; CTRL-C; ldm stop -f ldg1; leads to ldmd assert |
| 15697510 | Memory leak in create_basic_vsw_node() |
| 15719675 | RFE: Add an ldm history subcommand |
| 15787709 | Unify static (boot-time) and dynamic memory allocation alignment |
| 15789213 | Remove support for very old HV revisions from Logical Domains Manager |
| 15794303 | Migration of a domain with SR-IOV Ethernet virtual functions |
| 15797943 | Provide way to specify how long the 'ldm stop' command should wait |
| 15809245 | RFE: PCIe bus names need to use NAC names instead of the alias |
| 15811686 | Migration error message for an invalid password is inaccurate |
| 15814148 | XML list-bindings command should have an "extended" output option |
| 15816196 | Migrated guest domain reappears if followed by SP reset |
| 15821739 | Vestige of the XML v2 interface in mdstore.c |
| 16524266 | The list-io CLI command should be available over XMPP interface |
| 16922190 | Propagate ldc_mapin_base and ldc_mapin_size directly from factory-default HV MD |
| 17401528 | Logical Domains Manager XML interface does not report CPU usage for physical assignment |
| 17401564 | XML Interface does not report any error when trying to create a duplicate VDSDEV |
| 17812407 | Logical Domains Manager build with unnecessary rpaths |
| 18111544 | Two guest domains have same MAC and IP addresses, and are functioning |

| | |
|---|---|
| 18167985 | MIB vnetTable is missing some attributes such as link prop, pvlan, maxbw, etc |
| 18320689 | No argument check in code to send command args from ldm to ldmd |
| 18375880 | Affinity algorithm mis-allocates cpus in multiple node config |
| 18827424 | The list-netdev and list-netstat CLI commands should be available over XMPP interface |
| 19167766 | database_ldom_add_net_client() & other vnet funcs have too many parameters |
| 19917454 | Rethink ldmd's use of libssl |
| 19932503 | HV MBLOCK FOR PHYSICAL RC IS NOT CORRECTLY LINKED |
| 19934101 | ADI tags should be preserved across a live migration |
| 19944379 | Support mem-#ra-bits property for determining RA limits |
| 20085077 | ldm start-reconf/cancel-reconf loop causes ldmd abort and hv_mblock exhaustion |
| 20662802 | A common function to handle the enumerated string and boolean string |
| 20731165 | "No PRI node interleave value" warning on system with one latency group |
| 20766195 | Migration of a domain using named core resources |
| 20828870 | ldmd logging needs a lot of improvement |
| 21201719 | ldom-mgr: EOL use of alloca |
| 21274667 | Add -f option to ldm start to start I/O domain if root domain(s) are down |
| 21354984 | Migration fails with 'auxiliary connection failure' error |
| 21445956 | Delayed reconfig messages are not consistent between core and cpu operations |
| 21563292 | Migration should use effective LPS, not platform LPS |
| 21780022 | Live migration occasionally fails with Oracle VM Server for SPARC 3.3 |

| | |
|---|---|
| 21895478 | Oracle VM Server for SPARC MIB IOBusTable needs updates to include new/added information |
| 22175515 | ldm list-domain -e should be available over XMPP interface |
| 22197942 | Change in OBP variables not always reflected correctly by eeprom command |
| 22204673 | ls-bindings (vsw and network) could use prtvec utilities functions to simplify the implementation of managing temporary lists of ptrs |
| 22377138 | ldm ls-spconfig should take optional specific config arg |
| 22529020 | Fix/enable Autosave optimization from 18746688 |
| 22556221 | CPU DR error handling can assert in mmu_info_assign(): mmu == cmmu |
| 22577014 | CMI domains can't be recovered correctly in recovery mode |
| 22597094 | ldomVswInterVnetLink should differentiate between on/auto and off/auto |
| 22956865 | ldmd linked with libadimalloc dumps core with add-vcpu commands |
| 23025823 | Restoration of I/O devices fails on addboard with 2 or more RCs on root domain |
| 23026660 | Remove support for HVFRAG_TYPE_RKEY_TABLE |
| 23048855 | Check in no_mem_cmp() incorrectly calculates affinity |
| 23074031 | After upgrade from Oracle VM Server for SPARC 3.1 to 3.3 vdisk, vdsdev and autosave configs are lost |
| 23108061 | ldmd dumps core on bind after a bind failure |
| 23119242 | Guest only shows 999 vnets after adding 2000 vnets |
| 23144895 | MIB ldomSPConfigTable did not list all the sp-configs except factory-default |
| 23154856 | Logical Domains Manager crashes during recovery mode with ADI malloc enabled |

| | |
|---|---|
| 23254423 | Support native migration within the same CPU family |
| 23260980 | Logical Domains Manager ptrvec improvements |
| 23282766 | XML interface and MIB should show [degraded] tag for degraded config |
| 23284476 | Remove both instances of "SPARC-SN" from mig_hwcaps.c |
| 23292662 | Changes made to the STR_MATCH macro have broken PAPSAT's main menu |
| 23488114 | p2v: warning: name redefined by pragma redefine_extname declared static: link |
| 23571820 | Dynamic Reconfiguration (add and remove) of named core resources |
| 23572679 | Need a command or API to list blacklisted resources |
| 23587909 | Update ldm I/O CLI commands to support SRIOV Ethernet Live migration |
| 23587949 | Logical Domains Manager support for Reflective Relay |
| 23587961 | Logical Domains Manager support for dynamic alt-mac-address |
| 23698097 | ldm ls-constraints lose the configured mpgroup function for unbound domains |
| 23700874 | Failed set-vcpu attempt on primary removes the whole core constraint |
| 23725678 | Remove system domain support |
| 23732652 | physical-bindings=core not getting removed cleanly |
| 23751486 | ldmd dumps core at complete_pending_evac_memory subsequent to row fault |
| 23756620 | Live Migration support for Fujitsu SPARC M12 servers |
| 23756677 | Fujitsu SPARC M12 server support for PPAR DR and Large Pages |
| 23756776 | Fujitsu SPARC M12 server Group Demap Support |
| 23757096 | Failed to recover CMI domains with init-system |

| | |
|---|---|
| 23761597 | ldmd crashes with --disable-migration-comp option |
| 23762236 | Blacklist retire operations should not update autosave |
| 23763000 | Assertion failure on CMI-enabled Fujitsu SPARC M12 server 2 socket system |
| 24310184 | ldmd assert hit in mem_unconfigure_blacklisted_range() when stop -f run |
| 24341689 | ldmd crashes with --migration-pagesize 8192 |
| 24343698 | Dynamic reconfiguration (add and remove) of named memory resources |
| 24354353 | "ldm rm-spconfig" cannot remove autosave for last created config |
| 24428051 | ldm ls-io -p -l listing error in data fields |
| 24788913 | developer/opensolaris/ldoms has dependency on gcc version about to be removed |
| 24469054 | No Fujitsu M10 PM request once memory util drops below self refresh power state |
| 24513690 | list-io non-parseable output has bus and domain in wrong column |
| 24517671 | ldmd should not use OpenSSL PKCS#11 engine by default on Oracle Solaris |
| 24579103 | ldm set-io usage msg for virtual function should show the user-assigned name option |
| 24604492 | ldm add/rm-io <NAC name of BUS> doesn't work |
| 24608425 | Effective Page Size miscalculation caused unexpected migration failure on Fujitsu M10 |
| 24615854 | ldm set-io to assign username to Fibre Channel virtual functions reports false error |
| 24468983 | deleteboard with unbind=shutdown causes inconsistent configuration |
| 24671677 | ldmd dumps core while trying to recover bootsets from the system controller |

| | |
|---|---|
| 24682710 | Cleanup the warning message for degraded config after blacklist retire operation |
| 24717093 | ldm ls-spconfig Failed After SP Failover |
| 24718083 | PM policy via snmp cannot be set back the budget once SNMP disconnected |
| 24810649 | Command logging needs to include XML initiated commands |
| 24819893 | socket command caused core dump with database saved from another configuration |
| 24819898 | ldmd dumped core on system set socket constraints by Board DR addboard |
| 24819902 | deleteboard with ratio mode may remove extra vcpus from domains |
| 24819912 | Named resource dynamic reconfig may dump core in delayed reconfig |
| 24826440 | ovmtdeploy is failing with error "Failed to expand device" |
| 24849679 | Migration of domains using ADI requires Logical Domains Manager support |
| 24919193 | Minor memory leak in libds_chan_create_rsp() |
| 24947310 | ldmd dumps core while live migrating |
| 24965576 | User Assigned Name for a virtual function missing in ldoms-constraints |
| 25034173 | Provide mechanism to disable Power Management via SMF in non-debug bits |
| 25037061 | ldmd core dump in umem_do_abort |
| 25070368 | Reflective relay does not go along with aggregation |
| 25084643 | Introduce a CLI/XML interface for migration of named resources |
| 25084669 | Migration of a domain using named memory resources |
| 25084688 | Logical Domains Manager crashes while migrating a bound guest |

| | |
|---|---|
| 25140841 | ldm commands hang after injecting a memory UE followed by ldm ls-devices -B |
| 25165458 | I/O domain virtual functions failed to resume after primary domain reboot, policy ignore |
| 25189906 | Update diskio to better support sparse files |
| 25240380 | Improve error reporting for ADI-related migration errors |
| 25240482 | Multiple password files can be passed to a 'migrate' command |
| 25290866 | ldm migrate should log better error message on network io timeout |
| 25316565 | Memory leaked from mig_src_warm_init() |
| 25340760 | Migration of a domain using named memory resources takes partial mappings |
| 25341049 | Oracle Solaris panic during VTS Component stress memory test and power test |
| 25347257 | Logical Domains Manager should support less than 1500 MTU in vnet/vsw |
| 25378217 | ldm set-vsw vid=xxx accepts negative values (but leaves VID unchanged) |
| 25381139 | The string "name" is used ambiguously in some SR-IOV commands |
| 25388410 | Function prototype dr_mem.h:dr_mem_qry_f does not match dr_mem.c:dr_mem_query |
| 25390105 | ldmd dumped core when ldm ls-constraint was executed if evacuated virtual function exists |
| 25394230 | Change some PPAR DR debug messages to info messages |
| 25394451 | MAC collision check for all MAC addresses when it should not |
| 25416007 | Logical Domains Manager Logging minor fixes |
| 25420434 | ldm init-system does not recreate vdsdevs that were created with the (-f) flag |

| | |
|---|---|
| 25422195 | ldm should check if mode=[sc] is consistent across LM |
| 25429988 | ldm should check vnet's vid != vsw's dvid |
| 25438641 | ldmd --logctl command line parsing issue |
| 25496410 | ls-constraints -x output is incorrect boot-policy n/a |
| 25498880 | Migrating a domain from 3.4 to 3.5 fails on ADI capable platforms |
| 25511830 | ldmd dumps core at domain reboot after PRI update caused by PPAR DR |
| 25511942 | recreate_named_cids not freeing memory |
| 25512076 | IOV pciedev code has memory free issues |
| 25535305 | Target check of memory congruence call missing a domain pointer |
| 25559758 | Logical Domains Manager to send blacklist replay on ILOM interconnect restore |
| 25560010 | dr_cpu_remap support needed for Fujitsu M10 servers |
| 25577939 | OpenBoot PROM variable updates should be blocked if migration is in progress |
| 25585009 | ldm set-vsw -q doesn't work properly - no msg but non-zero exit value |
| 25600483 | Memory leaked from xml_v3_named_map() in migration |
| 25600818 | Memory leaked from mig_tgt_warm_recv_state() in migration |
| 25606866 | Memory leak in 'ldm set-vsw vid=xxx' |
| 25631877 | Segmentation fault in iov_mig_tgt_add_vfs() during migration dry run |
| 25690054 | Domain MAC collision check is missing in migrate_domain_create() |
| 25690079 | Initialization of cli_conn_t is inconsistent, error prone, and risky |
| 25712656 | Logical Domains variable update should print an error if varconfig-update is unavailable |
| 25721740 | ldmd abort between DR CPU PRE and POST commands can cause core remap to fail |

| | |
|---|---|
| 25721810 | ldmd wastes huge amount of memory during deleteboard sequence |
| 25743217 | Need to support a missing 'hwcap-list' MD property |
| 25766479 | Oracle VM Server for SPARC 3.5_b22 CLI: "ldm add-vsw" not showing "Usage" |
| 25772308 | Error message when connection to ILOM is down hints at OS specific solution |
| 25774540 | Named Resource Migration allows invalid mblockmap |
| 25789968 | Calls to getopt_long() cannot rely on optind = 0 |
| 25790237 | Primary domain keys are not persisted correctly after a spconfig update |
| 25802105 | sparc64-cpu module for SPARC64_CLASS1 migration group |
| 25852557 | Named resource migration crash in i_tgt_chk_ra_pa_congruence() |
| 25861447 | testDynAltMacUpd failed in updating alt-mac-address |
| 25861735 | Minor input validation error in ldm stop -t sec |
| 25861951 | ldmd should update physical function info when root domain transitions to support Fibre Channel SR-IOV |
| 25871519 | Effective page size calculated twice when a domain is in delayed reconfig |
| 25891463 | The 'force' option for the migrate command is not propagated to target |
| 25926495 | Cleanup the degraded config messaging after blacklist evacuation |
| 25954547 | ldm init-system failed to restore vsw when backend is link aggregation |
| 25976734 | deleteboard may fail if the system has a number of I/O devices |
| 25976819 | CPU autoreplacement remaps faulted strand even when drd is not ready |
| 25976856 | PRE cancel may be issued inappropriately during crash recovery |
| 25987495 | Need informative message for native migration from 3.4 stack to 3.5 stack |

| | |
|---|---|
| 26001877 | iov_seq_work() leaks memory |
| 26004742 | ldmpower fails for users with LDoms Power Mgmt Observability RBAC profile |
| 26021889 | ovmtlibrary uses /usr/sfw/bin/wget which is being retired by PSARC/2017/103 |
| 26035583 | ldmd core with mem free issue when ds_netsvc calls send_data_request_to_agent() |
| 26052693 | OS panic happens after PPAR DR deleteboard because of incorrect mem-latency-grp |
| 26086941 | cpu_del() should handle CPU_DR_OP_FAILED and return false |
| 26138962 | Migration virtual function removal and dirty page tracking do not mix well |
| 26166680 | MALLOC_CHECK_=2 reports free of invalid ptr in xml_call_cli_cmd/ xml_v3_list_hba |
| 26169909 | deleteboard with unbind=shutdown may dump core if memory remap failed |
| 26176351 | CPU chip is not power managed if pm-rm registration is delayed |
| 26235395 | ldmd goes into maintenance: Failed to allocate a LDC channel while parsing MDs |
| 26315289 | ldmd goes into maintenance: Failed to allocate a LDC channel while parsing MDs |
| 26325102 | Oracle VM Server for SPARC core dumps due to illegal attempt to set no_reset_flag on primary |