Oracle® VM Server for SPARC 3.3 Release Notes

Part No: E62362
October 2015

ORACLE

Pour plus d'informations sur l'engagement d'Oracle pour l'accessibilité à la documentation, visitez le site Web Oracle Accessibility Program, à l'adresse http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc.

Accès au support électronique

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

Preface 7

1 **Oracle VM Server for SPARC 3.3 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 3.3 System Requirements 12

Deprecated and Removed Oracle VM Server for SPARC Features 13

Known Issues 14

- Installation Issues 14
- General Issues 14
- Migration Issues 21
- SR-IOV Issues 28
- Bugs Affecting the Oracle VM Server for SPARC Software 31

Resolved Issues 60

Preface

- **Overview** – Includes information about this release of the Oracle VM Server for SPARC software, such as changes for this release, supported platforms, a matrix of required software and patches, and bugs affecting this software.
- **Audience** – System administrators who manage virtualization on SPARC servers
- **Required knowledge** – System administrators on these servers must have a 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 3.3 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 3.3 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 3.3 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 or the latest patch for the Oracle Solaris 10 1/13 OS.

For issues that you should address before you install the Oracle VM Server for SPARC 3.3 software, see "Installation Issues" on page 14.

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, "Oracle VM Server for SPARC 3.3 System Requirements," in *Oracle VM Server for SPARC 3.3 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 3.3 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 T-Series and M-Series servers from Oracle and Fujitsu M10 servers at the time that the Oracle VM Server for SPARC software is released and not for SPARC based systems 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 3.3 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" on page 10 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" on page 10
- "Features That Depend on System Firmware, the Oracle Solaris OS, Or Both" on page 12
- "Oracle VM Server for SPARC 3.3 System Requirements" on page 12
- "Deprecated and Removed Oracle VM Server for SPARC Features" on page 13
- "Known Issues" on page 14

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 3.3 software are as follows. Note that each feature is available on all supported platforms except UltraSPARC T2, UltraSPARC T2 Plus, and SPARC T3 unless otherwise stated. For the list of supported platforms, see "Supported Platforms" in *Oracle VM Server for SPARC 3.3 Installation Guide*.

- Support the virtual SCSI host bus adapter (vHBA) feature, which enables you to virtualize any type of SCSI device (such as disk, tape, CD, or DVD). The virtualized SCSI device is accessible from a guest domain.

  The vHBA feature is compatible with other I/O interfaces such as MPxIO multipathing, which enables a virtual LUN to have the same behavior as a physical LUN. vHBA also enables you to easily configure virtual SANs, which can contain an unbounded number of

SCSI devices. See "Introduction to Virtual SCSI Host Bus Adapters" in *Oracle VM Server for SPARC 3.3 Administration Guide* and the `ldm`(1M) man page.

This feature is supported on UltraSPARC T2, UltraSPARC T2 Plus, and SPARC T3 systems as well.

- Enable recovery mode by default. See "Recovering Domains After Detecting Faulty or Missing Resources" in *Oracle VM Server for SPARC 3.3 Administration Guide*.

  This feature is not supported on SPARC T4 systems.
- Add Oracle VM Templates commands to create, deploy, and configure Oracle VM templates for SPARC systems. See Chapter 18, "Using Oracle VM Server for SPARC Templates," in *Oracle VM Server for SPARC 3.3 Administration Guide*.

  This feature is supported on UltraSPARC T2, UltraSPARC T2 Plus, and SPARC T3 systems as well.
- Support whole-core dynamic resource management (DRM). See "Using Dynamic Resource Management" in *Oracle VM Server for SPARC 3.3 Administration Guide*.
- Add support for CPU socket constraints. See *Fujitsu M10/SPARC M10 Systems System Operation and Administration Guide*.

  This feature is supported only on Fujitsu M10 servers.
- Enhance the live migration feature to enable memory block splitting. Memory block splitting eliminates the requirement to have sufficient free contiguous memory regions on the target machine for each memory block that is assigned to the guest domain. The migration can split up memory blocks in the guest domain to fit into available free regions on the target machine.

  This feature requires that you run the Oracle Solaris 11.3 OS on the guest domain being migrated.

  This feature is supported on UltraSPARC T2, UltraSPARC T2 Plus, and SPARC T3 systems as well.
- Enable autoreplacement of faulty CPUs with bound processes and faulty CPUs in single-core domains. See *Fujitsu M10/SPARC M10 Systems System Operation and Administration Guide*.

  This feature is supported only on Fujitsu M10 servers.
- Dynamically update the `net-dev` property value by using the `ldm set-vsw` command.
- Expand cross-CPU migration to include the SPARC T7 series servers and SPARC M7 series servers. See "Domain Migration Requirements for CPUs" in *Oracle VM Server for SPARC 3.3 Administration Guide* and "Cross-CPU Restrictions for Migration" in *Oracle VM Server for SPARC 3.3 Administration Guide*.
- Provide support for Ethernet over USB to enable the `ldmd` daemon to communicate with the SP and provide support for an `ldmd` connection check. See "How to Verify the ILOM Interconnect Configuration" in *Oracle VM Server for SPARC 3.3 Administration Guide*.

  This feature is supported only on SPARC T7 series servers and SPARC M7 series servers.

- Bug fixes.

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

Some of the Oracle VM Server for SPARC 3.3 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, "Oracle VM Server for SPARC 3.3 System Requirements," in *Oracle VM Server for SPARC 3.3 Installation Guide*. For Fujitsu M10 servers, see the latest *Fujitsu M10/SPARC M10 Systems Product Notes*.

For Oracle VM Server for SPARC 3.3, all new 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 3.3 features are not available:

- Virtual SCSI HBA, which requires the Oracle Solaris 11.3 OS
- Memory block splitting, which requires the Oracle Solaris 11.3 OS on the guest domain to be migrated

Oracle VM Server for SPARC 3.3 System Requirements

You can find information about the recommended and minimum software component versions to use with the Oracle VM Server for SPARC 3.3 release in Chapter 1, "Oracle VM Server for SPARC 3.3 System Requirements," in *Oracle VM Server for SPARC 3.3 Installation Guide*.

Deprecated and Removed Oracle VM Server for SPARC Features

The following Oracle VM Server for SPARC features have been removed from this release of the software:

- The Oracle VM Server for SPARC 3.3 produce no longer supports the running of the Oracle Solaris 10 OS in control domains. You can continue to run the Oracle Solaris 10 OS in guest domains, root domains, and I/O domains when using future releases. Oracle Solaris 10 patches will continue to be provided for Oracle VM Server for SPARC based on the documented support policy.
- Setting the `threading` property to manage single-thread workloads on SPARC T4 platforms has been deprecated previously. By default, domains are created for maximum throughput, and the Oracle Solaris OS automatically uses the critical thread API to optimize for single-thread workloads. This feature has been removed from the Oracle VM Server for SPARC 3.3 release.

  Before you install or upgrade to the Oracle VM Server for SPARC 3.3 software, ensure that all domains have the `threading` property set to `max-throughput`.
- Using the `-c` option of the `add-vcpu`, `set-vcpu`, and `rm-vcpu` subcommands to manage hard partitioning has been deprecated previously. Instead, use the `add-core`, `set-core`, or `rm-core` subcommand to assign whole cores. Also, use the `add-domain` or `set-domain` subcommand to specify the CPU cap (`max-cores`). This feature has been removed from the Oracle VM Server for SPARC 3.3 release.
- This Oracle VM Server for SPARC release no longer includes the `ldmconfig` utility as it runs only on Oracle Solaris 10 control domains. This utility will no longer be updated to address bug fix or enhancement requests.
- The network interface functionality of an Oracle VM Server for SPARC `vsw` driver was deprecated in Oracle Solaris 11.1. The Oracle VM Server for SPARC `vsw` driver continues to provide virtual network switching functionality for guest domains. See "Oracle Solaris 11 Networking Overview" in *Oracle VM Server for SPARC 3.3 Administration Guide*.
- The SR-IOV feature is no longer supported on a control domain that runs the Oracle Solaris 10 OS.

The following Oracle VM Server for SPARC features are deprecated in this release of the software and might be removed from the product in a future release:

- Using the Hybrid I/O feature is deprecated in favor of using the single-root I/O virtualization (SR-IOV) feature.
- The Netra Data Plane software suite, which includes the `vdpcc` and `vdpcs` virtual devices, is no longer supported by Oracle VM Server for SPARC.

- The `ldmp2v` utility will no longer be updated to address bug fix or enhancement requests. This utility will no longer be supported, but will continue to be included and documented as part of the Oracle VM Server for SPARC 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 3.3 Administration Guide*.
- The Logical Domains Manager auditing implementation is deprecated pending replacement and is turned off by default.

Known Issues

This section contains general issues and specific bugs concerning the Oracle VM Server for SPARC 3.3 software.

Installation Issues

threading Property Removal Requires That You Remove the threading=max-ipc Setting

The `threading` property has been removed from the Oracle VM Server for SPARC 3.3 release. Therefore, before you update your system to use the Oracle VM Server for SPARC 3.3 software, remove the `threading=max-ipc` property setting from all domains:

```
primary# ldm set-dom threading=max-throughput domain-name
```

If you do not remove this property setting, only one virtual CPU per core remains after the `ldmd` daemon is restarted and the whole-core constraint is lost.

Use the `ldm set-core` command to restore the whole-core constraint.

```
primary# ldm set-core n domain-name
```

General Issues

This section describes general known issues about this release of the Oracle VM Server for SPARC software that are broader than a specific bug number. Workarounds are provided where available.

After Canceling a Migration, `ldm` Commands That Are Run on the Target System Are Temporarily Unresponsive

If you cancel a live migration, the memory contents of the domain instance that is created on the target machine must be "scrubbed" by the hypervisor. This scrubbing process is performed for security reasons and must be complete before the memory can be returned to the pool of free memory. While this scrubbing is in progress, `ldm` commands become unresponsive. As a result, the Logical Domains Manager appears to be hung.

Recovery: You must wait for this scrubbing request to finish before you attempt to run other `ldm` commands. This process might take a long time. For example, a guest domain that has 500 Gbytes of memory might complete this process in up to 7 minutes on a SPARC T4 server or up to 25 minutes on a SPARC T3 server.

SPARC M5-32 and SPARC M6-32: Issue With Disks That Are Accessible Through Multiple Direct I/O Paths

When using the `ldm add-vcpu` command to assign CPUs to a domain, the Oracle Solaris OS might panic with the following message:

```
panic[cpu16]/thread=c4012102c860: mpo_cpu_add: Cannot read MD
```

This panic occurs if the following conditions exist:

- Additional DCUs have been assigned to a host
- The host is started by using a previously saved SP configuration that does not contain all the hardware that is assigned to the host

The target domain of the `ldm add-vcpu` operation is the domain that panics. The domain recovers with the additional CPUs when it reboots.

Workaround: Do not use configurations that are generated with fewer hardware resources than are assigned to the host.

To avoid the problem, do not add CPUs as described in the problem description. Or, perform the following steps:

1. Generate a new SP configuration after the DCUs have been added.

   For example, the following command creates a configuration called `new-config-more-dcus`:

   ```
   primary# ldm add-config new-config-more-dcus
   ```
2. Shutdown the domain.
3. Stop the host.

   ```
   -> stop /HOST
   ```
4. Start the host.

```
-> start /HOST
```

Destroying All Virtual Functions and Returning the Slots to the Root Domain Does Not Restore the Root Complex Resources

The resources on the root complex are not restored after you destroy all the virtual functions and return the slots to the root domain.

Recovery: Return all the virtual I/O resources that are associated with the root complex to its root domain.

First, put the control domain in delayed reconfiguration..

```
primary# ldm start-reconf primary
```

Return all child PCIe slots to the root domain that owns the `pci_0` bus. Then, remove all of the child virtual functions on the `pci_0` bus and destroy them.

Finally, set `iov=off` for the `pci_0` bus and reboot the root domain.

```
primary# ldm set-io iov=off pci_0
primary# shutdown -y -g 10
```

Workaround: Set the `iov` option to `off` for the specific PCIe bus.

```
primary# ldm start-reconf primary
primary# ldm set-io iov=off pci_0
```

`init-system` Does Not Restore Named Core Constraints for Guest Domains From Saved XML Files

The `ldm init-system` command fails to restore the named CPU core constraints for guest domains from a saved XML file.

Workaround: Perform the following steps:

1. Create an XML file for the primary domain.

   ```
   # ldm ls-constraints -x primary > primary.xml
   ```
2. Create an XML file for the guest domain or domains.

   ```
   # ldm ls-constraints -x domain-name[,domain-name][,...] > guest.xml
   ```
3. Power cycle the system and boot a factory default configuration.
4. Apply the XML configuration to the `primary` domain.

   ```
   # ldm init-system -r -i primary.xml
   ```
5. Apply the XML configuration to the guest domain or domains.

   ```
   # ldm init-system -f -i guest.xml
   ```

Removing a Large Number of CPUs From a Domain Might Fail

You might see the following error message when you attempt to remove a large number of CPUs from a guest domain:

```
Request to remove cpu(s) sent, but no valid response received
VCPU(s) will remain allocated to the domain, but might
not be available to the guest OS
Resource modification failed
```

Workaround: Stop the guest domain before you remove more than 100 CPUs from the domain.

Newly Added NIU/XAUI Adapters Are Not Visible to the Host OS If Logical Domains Is Configured

When Logical Domains is configured on a system and you add another XAUI network card, the card is not visible after the machine has undergone a power cycle.

Recovery: To make the newly added XAUI visible in the control domain, perform the following steps:

1. Set and clear a dummy variable in the control domain.

   The following commands use a dummy variable called `fix-xaui`:

   ```
   # ldm set-var fix-xaui=yes primary
   # ldm rm-var fix-xaui primary
   ```

2. Save the modified configuration to the SP, replacing the current configuration.

   The following commands use a configuration name of `config1`:

   ```
   # ldm rm-spconfig config1
   # ldm add-spconfig config1
   ```

3. Perform a reconfiguration reboot of the control domain.

   ```
   # reboot -- -r
   ```

   At this time, you can configure the newly available network or networks for use by Logical Domains.

LSI SAS 2008 Cannot Be Added by Dynamic Bus or PCI-Box Hotplug Operations

If you attempt to remove a PCIe bus that hosts LSI SAS HBA devices, you cannot later add the devices by using dynamic bus or PCI-box hotplug operations.

In Certain Conditions, a Guest Domain's Solaris Volume Manager Configuration or Metadevices Can Be Lost

If a service domain is running a version of Oracle Solaris 10 OS prior to Oracle Solaris 10 1/13 OS and is exporting a physical disk slice as a virtual disk to a guest domain, then this virtual disk will appear in the guest domain with an inappropriate device ID. If that service domain is then upgraded to Oracle Solaris 10 1/13 OS, the physical disk slice exported as a virtual disk will appear in the guest domain with no device ID.

This removal of the device ID of the virtual disk can cause problems to applications attempting to reference the device ID of virtual disks. In particular, Solaris Volume Manager might be unable to find its configuration or to access its metadevices.

Workaround: After upgrading a service domain to Oracle Solaris 10 1/13 OS, if a guest domain is unable to find its Solaris Volume Manager configuration or its metadevices, perform the following procedure.

▼ How to Find a Guest Domain's Solaris Volume Manager Configuration or Metadevices

1. **Boot the guest domain.**

2. **Disable the `devid` feature of Solaris Volume Manager by adding the following lines to the `/kernel/drv/md.conf` file:**

   ```
   md_devid_destroy=1;
   md_keep_repl_state=1;
   ```

3. **Reboot the guest domain.**

   After the domain has booted, the Solaris Volume Manager configuration and metadevices should be available.

4. **Check the Solaris Volume Manager configuration and ensure that it is correct.**

5. **Re-enable the Solaris Volume Manager `devid` feature by removing from the `/kernel/drv/md.conf` file the two lines that you added in Step 2.**

6. **Reboot the guest domain.**

   During the reboot, you will see messages similar to this:

   ```
   NOTICE: mddb: unable to get devid for 'vdc', 0x10
   ```

   These messages are normal and do not report any problems.

Oracle Solaris Boot Disk Compatibility

Historically, the Oracle Solaris OS has been installed on a boot disk configured with an SMI VTOC disk label. Starting with the Oracle Solaris 11.1 OS, the OS is installed on a boot disk that is configured with an extensible firmware interface (EFI) GUID partition table (GPT) disk label by default. If the firmware does not support EFI, the disk is configured with an SMI VTOC disk label instead. This situation applies only to SPARC T4 servers that run at least system firmware version 8.4.0, to SPARC T5, SPARC M5, SPARC M6 servers that run at least system firmware version 9.1.0, and to Fujitsu M10 servers that run at least XCP version 2230.

The following servers cannot boot from a disk that has an EFI GPT disk label:

- UltraSPARC T2, UltraSPARC T2 Plus, and SPARC T3 servers no matter which system firmware version is used
- SPARC T4 servers that run system firmware versions prior to 8.4.0
- SPARC T5, SPARC M5, and SPARC M6 servers that run system firmware versions prior to 9.1.0
- Fujitsu M10 servers that run XCP versions prior to 2230

So, an Oracle Solaris 11.1 boot disk that is created on an up-to-date SPARC T4, SPARC T5, SPARC M5, SPARC M6, Fujitsu M10 server cannot be used on older servers or on servers that run older firmware.

This limitation restrains the ability to use either cold or live migration to move a domain from a recent server to an older server. This limitation also prevents you from using an EFI GPT boot disk image on an older server.

To determine whether an Oracle Solaris 11.1 boot disk is compatible with your server and its firmware, ensure that the Oracle Solaris 11.1 OS is installed on a disk that is configured with an SMI VTOC disk label.

To maintain backward compatibility with systems that run older firmware, use one of the following procedures. Otherwise, the boot disk uses the EFI GPT disk label by default. These procedures show how to ensure that the Oracle Solaris 11.1 OS is installed on a boot disk with an SMI VTOC disk label on a SPARC T4 server with at least system firmware version 8.4.0, on a SPARC T5, SPARC M5, or SPARC M6 server with at least system firmware version 9.1.0, and on a Fujitsu M10 server with at least XCP version 2230.

- **Solution 1:** Remove the `gpt` property so that the firmware does not report that it supports EFI.

  1. From the OpenBoot PROM prompt, disable automatic booting and reset the system to be installed.

     ```
     ok setenv auto-boot? false
     ok reset-all
     ```

     After the system resets, it returns to the `ok` prompt.

2. Change to the `/packages/disk-label` directory and remove the `gpt` property.

```
ok cd /packages/disk-label
ok " gpt" delete-property
```

3. Begin the Oracle Solaris 11.1 OS installation.

   For example, perform a network installation:

```
ok boot net - install
```

- **Solution 2:** Use the `format -e` command to write an SMI VTOC label on the disk to be installed with the Oracle Solaris 11.1 OS.

1. Write an SMI VTOC label on the disk.

   For example, select the `label` option and specify the SMI label:

```
# format -e c1d0
format> label
[0] SMI Label
[1] EFI Label
Specify Label type[1]: 0
```

2. Configure the disk with a slice 0 and slice 2 that cover the entire disk.

   The disk should have no other partitions. For example:

```
format> partition

partition> print
Current partition table (unnamed):
Total disk cylinders available: 14087 + 2 (reserved cylinders)

Part      Tag    Flag     Cylinders         Size            Blocks
  0       root    wm       0 - 14086      136.71GB    (14087/0/0) 286698624
  1 unassigned    wu       0                0         (0/0/0)             0
  2     backup    wu       0 - 14086      136.71GB    (14087/0/0) 286698624
  3 unassigned    wm       0                0         (0/0/0)             0
  4 unassigned    wm       0                0         (0/0/0)             0
  5 unassigned    wm       0                0         (0/0/0)             0
  6 unassigned    wm       0                0         (0/0/0)             0
  7 unassigned    wm       0                0         (0/0/0)             0
```

3. Re-write the SMI VTOC disk label.

```
partition> label
[0] SMI Label
[1] EFI Label
Specify Label type[0]: 0
Ready to label disk, continue? y
```

4. Configure your Oracle Solaris Automatic Installer (AI) to install the Oracle Solaris OS on slice 0 of the boot disk.

   Change the `<disk>` excerpt in the AI manifest as follows:

```
<target>
   <disk whole_disk="true">
        <disk_keyword key="boot_disk"/>
        <slice name="0" in_zpool="rpool"/>
   </disk>
```

```
[...]
</target>
```

5. Perform the installation of the Oracle Solaris 11.1 OS.

Sometimes a Block of Dynamically Added Memory Can Be Dynamically Removed Only as a Whole

Due to the way in which the Oracle Solaris OS handles the metadata for managing dynamically added memory, you might later be able to remove only the entire block of memory that was previously dynamically added rather than a proper subset of that memory.

This situation could occur if a domain with a small memory size is dynamically grown to a much larger size, as shown in the following example.

```
primary# ldm list ldom1
NAME  STATE FLAGS   CONS VCPU MEMORY UTIL UPTIME
ldom1 active -n--   5000 2    2G     0.4% 23h

primary# ldm add-mem 16G ldom1

primary# ldm rm-mem 8G ldom1
Memory removal failed because all of the memory is in use.

primary# ldm rm-mem 16G ldom1

primary# ldm list ldom1
NAME  STATE FLAGS   CONS VCPU MEMORY UTIL UPTIME
ldom1 active -n--   5000 2    2G     0.4% 23h
```

Workaround: Use the `ldm add-mem` command to sequentially add memory in smaller chunks rather than in chunks larger than you might want to remove in the future.

Recovery: Perform one of the following actions:

- Stop the domain, remove the memory, and then restart the domain.
- Reboot the domain, which causes the Oracle Solaris OS to reallocate its memory management metadata such that the previously added memory can now be removed dynamically in smaller chunks.

Migration Issues

Migration of a Domain Between SPARC T7 Series Servers With Fragmented Memory Might Cause `ldmd` to Crash

Bug ID 21554591: During a live migration, the `ldmd` service on the target machine might dump core then restart.

This problem might occur when the memory on the domain to be migrated is highly fragmented into multiple memory segments and the target machine's free memory layout is not compatible. The problem is more likely to occur if you use memory DR to remove memory from the domain prior to live migration.

The stack trace of the core dump is similar to the following:

```
restore_lgpg_mblk+0x398(17bbc88, 16c39c8, 80000000, 80000000, 0, 40000000)
rgrp_restore_lgpg+0x39c(0, 0, 1733948, 1711598, 0, 20000000)
mem_allocate_real+0x92c(0, 20000000, ffbff868, 13aec88, 80808080, 373cd8)
affinity_bind_resources+0x9f4(17bbc88, ffbff948, 13aec88, 3a10c000, 3a10c000, 1010101)
mem_bind_real+0x468(17bbc88, ffbff9d4, 13aec88, 3a10c000, 3a10c000, 1010101)
mem_bind_real_check+0xf4(17bbc88, 12ee338, 13aec88, 0, 376468, ff29fd80)
mig_tgt_bound_feasibility_check+0x168(164be08, ff000000, ff, 1, 0, 0)
i_tgt_do_feasibility_check+0x168(164be08, 0, 12390, 1, f960d244, ffffff)
sequence+0x4a4(0, ff000000, ff322a40, 1, f960d244, ffffff)
main+0xb54(5, ffbffc64, ffbffc7c, f960a900, 0, ff320200)
_start+0x108(0, 0, 0, 0, 0, 370b60)
```

When this problem occurs, the guest domain continues to run. If the `ldmd` services restarts successfully, no further recovery is needed.

If the `ldmd` service fails to restart and goes into maintenance mode due to Bug 21569507, you must perform a power cycle of the host or applicable physical domain before you can restart `ldmd`.

Workaround: Stop and unbind the guest domain and then perform a cold migration. Do not use memory DR to remove memory from the guest domains to be migrated.

Kernel Zones Block Live Migration of Guest Domains

Bug ID 21289174: On a SPARC system, 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*.

Cross-CPU Live Migrations Between SPARC T7 Series Servers and SPARC M7 Series Servers and Older Platforms Require at Least the Oracle VM Server for SPARC 3.2 Software on the Source Machine and the Target Machine

Bug ID 20606773: Cross-CPU live migrations between a SPARC T7 series server or a SPARC M7 series server and an older platform require that you run at least Oracle VM Server for SPARC 3.2 software on the source and target machines.

For example, live migration between a SPARC T5 system and a SPARC T7 series server requires that at least Oracle VM Server for SPARC 3.2 software is installed on the SPARC T5 system.

Domain Migration Might Fail Even Though Sufficient Memory in a Valid Layout Is Available on the Target System

Bug ID 20453206: A migration operation might fail even if sufficient memory in a valid layout is available on the target system. Memory DR operations might make it more difficult to migrate a guest domain.

Workaround: None.

Oracle Solaris 10 Guest Domains That Have Only One Virtual CPU Assigned Might Panic During a Live Migration

Bug ID 17285751: Migrating an Oracle Solaris 10 guest domain that has only one virtual CPU assigned to it might cause a panic on the guest domain in the function `pg_cmt_cpu_fini()`.

Workaround: Assign at least two virtual CPUs to the guest domain before you perform the live migration. For example, use the `ldm add-vcpu` *number-of-virtual-CPUs domain-name* command to increase the number of virtual CPUs assigned to the guest domain.

Domain Migrations From SPARC T4 Systems That Run System Firmware 8.3 to SPARC T5, SPARC M5, or SPARC M6 Systems Are Erroneously Permitted

Bug ID 17027275: Domain migrations from SPARC T4 systems that run system firmware 8.3 to SPARC T5, SPARC M5, or SPARC M6 systems are not permitted. Although the migration succeeds, a subsequent memory DR operation causes a panic.

Workaround: Update the system firmware on the SPARC T4 system to version 8.4. See the workaround for "Guest Domain Panics at `lgrp_lineage_add(mutex_enter: bad mutex, lp=10351178)`" on page 47.

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

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 machine and an UltraSPARC T2 or SPARC T3 machine.

Workaround: None.

`ldm list -o status` on Target Control Domain Reports Bogus Migration Progress

Bug ID 15819714: In rare circumstances, the `ldm list -o status` command reports a bogus 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.

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

Bug ID 15776123: If the `cputrack` command is run on a guest domain while that domain is migrated to a SPARC T4 system, 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 system.

Guest Domain That Uses Cross-CPU Migration Reports Random Uptimes After the Migration Completes

Bug ID 15775055: After a domain is migrated between two machines that have different CPU frequencies, the uptime reports by the `ldm list` command might be incorrect. These incorrect results occur because uptime is calculated relative to the `STICK` frequency of the machine on which the domain runs. If the `STICK` frequency differs between the source and target machines, the uptime appears to be scaled incorrectly.

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

The uptime reported and shown by the guest domain itself is correct. Also, any accounting that is performed by the Oracle Solaris OS in the guest domain is correct.

`nxge` Panics When Migrating a Guest Domain That Has Hybrid I/O and Virtual I/O Virtual Network Devices

Bug ID 15710957: When a heavily loaded guest domain has a hybrid I/O configuration and you attempt to migrate it, you might see an `nxge` panic.

Workaround: Add the following line to the `/etc/system` file on the `primary` domain and on any service domain that is part of the hybrid I/O configuration for the domain:

```
set vsw:vsw_hio_max_cleanup_retries = 0x200
```

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 the domain on the target machine restarts. Although you can still perform a migration, it will not be a live migration.

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 Fails to Restore the Default Number of Virtual CPUs for a Migrated Domain When the Policy Is Removed or Expired

Bug ID 15701853: After you perform a domain migration while a DRM policy is in effect, if the DRM policy expires or is removed from the migrated domain, DRM fails to restore the original number of virtual CPUs to the domain.

Workaround: If a domain is migrated while a DRM policy is active and the DRM policy is subsequently expired or removed, reset the number of virtual CPUs. Use the `ldm set-vcpu` command to set the number of virtual CPUs to the original value on the domain.

Migration Failure Reason Not Reported When the System MAC Address Clashes With Another MAC Address

Bug ID 15699763: A domain cannot be migrated if it contains a duplicate MAC address. Typically, when a migration fails for this reason, the failure message shows the duplicate MAC address. However in rare circumstances, this failure message might not report the duplicate MAC address.

```
# ldm migrate ldg2 system2
Target Password:
Domain Migration of LDom ldg2 failed
```

Workaround: Ensure that the MAC addresses on the target machine are unique.

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.

Workaround: None.

Migration of a Domain That Has an Enabled Default DRM Policy Results in a Target Domain Being Assigned All Available CPUs

Bug ID 15655513: Following the migration of an active domain, CPU utilization in the migrated domain can increase dramatically for a short period of time. If a dynamic resource management (DRM) policy is in effect for the domain at the time of the migration, the Logical Domains Manager might begin to add CPUs. In particular, if the `vcpu-max` and `attack` properties were not specified when the policy was added, the default value of `unlimited` causes all the unbound CPUs in the target machine to be added to the migrated domain.

Recovery: No recovery is necessary. After the CPU utilization drops below the upper limit that is specified by the DRM policy, the Logical Domains Manager automatically removes the CPUs.

Explicit Console Group and Port Bindings Are Not Migrated

Bug ID 15527921: During a migration, any explicitly assigned console group and port are ignored, and a console with default properties is created for the target domain. This console is created using the target domain name as the console group and using any available port on the first virtual console concentrator (`vcc`) device in the control domain. If there is a conflict with the default group name, the migration fails.

Recovery: To restore the explicit console properties following a migration, unbind the target domain and manually set the desired properties using the `ldm set-vcons` command.

Migration Can Fail to Bind Memory Even If the Target Has Enough Available

Bug ID 15523120: In certain situations, a migration fails and `ldmd` reports that it was not possible to bind the memory needed for the source domain. This situation can occur even if the total amount of available memory on the target machine is greater than the amount of memory being used by the source domain.

This failure occurs because migrating the specific memory ranges in use by the source domain requires that compatible memory ranges are available on the target as well. When no such compatible memory range is found for any memory range in the source, the migration cannot proceed. See "Migration Requirements for Memory" in *Oracle VM Server for SPARC 3.3 Administration Guide*.

Recovery: If this condition is encountered, you might be able to migrate the domain if you modify the memory usage on the target machine. To do this, unbind any bound or active logical domain on the target.

Use the `ldm list-devices -a mem` command to see what memory is available and how it is used. You might also need to reduce the amount of memory that is assigned to another domain.

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.

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

```
# svcadm restart vntsd
```

Note – This command will disconnect all active console connections.

Cannot Migrate a Domain Between a System That Has EFI GPT Disk Labels and a System That Does Not Have EFI GPT Disk Labels

System firmware versions 8.4, 9.1, and XCP2230 introduced support for EFI GPT disk labels. By default, virtual disks that are installed when running at least the Oracle Solaris 11.1 OS on those systems have an EFI GPT disk label. You cannot read this disk label on older versions of firmware (such as 9.0.*x*, 8.3, 7.*x*, or XCP2221). This situation precludes you from performing a live migration or a cold migration to a system that runs a system firmware version without EFI GPT support. Note that a cold migration also fails in this situation, which is different than the previous limitations.

To determine whether your virtual disk has an EFI GPT disk label, run the `devinfo -i` command on the raw device. The following examples show whether the virtual disk has an SMI VTOC or an EFI GPT disk label.

- **SMI VTOC disk label.** When your virtual disk has an SMI VTOC, you can perform a migration to firmware regardless of whether it supports EFI.

  This example indicates that the device has a VTOC label because the `devinfo -i` command reports device-specific information.

  ```
  # devinfo -i /dev/rdsk/c2d0s2
  /dev/rdsk/c2d0s2        0       0       73728   512     2
  ```

- **EFI GPT disk label.** When your virtual disk has an EFI GPT disk label, you can perform a migration only to firmware that has EFI support.

  This example indicates that the device has an EFI GPT disk label because the `devinfo -i` command reports an error.

```
# devinfo -i /dev/rdsk/c1d0s0
devinfo: /dev/rdsk/c1d0s0: This operation is not supported on EFI
labeled devices
```

SR-IOV Issues

SPARC M5 and SPARC M6: Virtual Function Add Operation Triggers a PCIe Fabric Error Panic When Attempting to Enable SR-IOV on the PCIe Bus That Hosts rKVMS Devices

Bug ID 21548747: If you use any SR-IOV devices on a SPARC M5 or SPARC M6 system on a PCIe bus that hosts rKVMS devices, attempting to enable SR-IOV after a dynamic bus add DBUS operation triggers an FMA error panic.

Workaround: Restrict the use of a PCIe bus that hosts rKVMS devices to the `primary` domain. Only use the static PCIe bus assignment method to remove and add this bus to the `primary` domain.

Fibre Channel Physical Function Configuration Changes Require Several Minutes to Complete

Bug ID 21299683: After you create or destroy virtual functions, it might take up to five minutes before you can attempt to create or destroy more virtual functions from the Fibre Channel physical function.

If you attempt to perform these operations before five minutes elapse, the operations fail with a message similar to the following:

```
The attempt to offline the pf /SYS/PCI-EM4/IOVFC.PF0 in domain
primary failed.
Error message from svc:/ldoms/agents in domain primary:
CMD_OFFLINE Failed. ERROR: devices or resources are busy.
```

Sometimes the states of the physical function, the virtual function, or both appear as `INV` in `ldm list-io` output. For example, in the following sample output, physical function 0 and its child virtual functions have an `INV` status.

```
primary# ldm list-io pci_5
NAME                                      TYPE   BUS      DOMAIN    STATUS
----                                      ----   ---      ------    ------
pci_5                                     BUS    pci_5    primary   IOV
```

```
/SYS/RCSA/PCIE5                          PCIE   pci_5    primary   OCC
/SYS/RCSA/PCIE6                          PCIE   pci_5    primary   OCC
/SYS/RCSA/PCIE5/IOVFC.PF0                PF     pci_5    primary   INV
/SYS/RCSA/PCIE5/IOVFC.PF1                PF     pci_5    primary
/SYS/RCSA/PCIE5/IOVFC.PF0.VF0            VF     pci_5             INV
/SYS/RCSA/PCIE5/IOVFC.PF0.VF1            VF     pci_5             INV
/SYS/RCSA/PCIE5/IOVFC.PF0.VF2            VF     pci_5             INV
/SYS/RCSA/PCIE5/IOVFC.PF0.VF3            VF     pci_5             INV
/SYS/RCSA/PCIE5/IOVFC.PF0.VF4            VF     pci_5             INV
/SYS/RCSA/PCIE5/IOVFC.PF0.VF5            VF     pci_5             INV
/SYS/RCSA/PCIE5/IOVFC.PF0.VF6            VF     pci_5             INV
```

Workaround: To clear the `INV` status on the physical and virtual functions, perform the following steps:

1. Wait five minutes.
2. Retry the IOV operation to create or destroy the virtual function on the Fibre Channel physical function.

You can perform all necessary IOV configurations for the Fibre Channel physical function in a single command by using the `ldm create-vf -n max` or `ldm destroy-vf -n max` command.

`ldm remove-io` Command Reports a Timeout and Fails to Remove the Last SR-IOV Virtual Function From an I/O Domain

Bug ID 20731016: When you use the `ldm remove-io` command to remove the last SR-IOV virtual function from an I/O domain, the command might report a timeout and fail to remove the virtual function.

Workaround: If this problem occurs, perform the following steps:

1. Verify that the `system/management/hwmgmtd` package is installed on the system.

   ```
   # pkg info system/management/hwmgmtd
   ```

2. Disable the `svc:/system/sp/management` service.

   ```
   # svcadm disable -st svc:/system/sp/management
   ```

3. Retry the `ldm remove-io` command.
4. When the SR-IOV virtual function is successfully removed, enable the `svc:/system/sp/management` service.

   ```
   # svcadm enable svc:/system/sp/management
   ```

Fibre Channel Physical Function Is Faulted by FMA and Disabled

Bug IDs 18168525 and 18156291: You must connect the Fibre Channel PCIe card to a Fibre Channel switch that supports NPIV and that is compatible with the PCIe card. If you do not use this configuration, using the `format` command, or creating or destroying a virtual function might cause the physical function to be faulted by FMA and disabled. If this fault occurs, the message is similar to the following:

```
SUNW-MSG-ID: PCIEX-8000-0A, TYPE: Fault, VER: 1, SEVERITY: Critical
EVENT-TIME: event-time
PLATFORM: platform-type
SOURCE: eft, REV: 1.16
EVENT-ID: event-ID
DESC: A problem was detected for a PCIEX device.
AUTO_RESPONSE: One or more device instances may be disabled
IMPACT: Loss of services provided by the device instances associated with
this fault
REC-ACTION: Use 'fmadm faulty' to provide a more detailed view of this event.
Please refer to the associated reference document at
http://support.oracle.com/msg/PCIEX-8000-0A for the latest service procedures
and policies regarding this diagnosis.
```

Workaround: If the card has been faulted by FMA, first check its connections and ensure that the card is not directly connected to storage. Then, perform the step that matches your configuration:

- **Card is directly connected to storage –** Correctly configure the Fibre Channel PCIe card by connecting it to a Fibre Channel switch that supports NPIV and is compatible with the PCIe card. Then, run the `fmadm repair` command to override the FMA diagnosis.
- **Card is not directly connected to storage –** Replace the card.

Warnings Appear on Console When Creating Fibre Channel Virtual Functions

Bug ID 17623156: When you create Fibre Channel virtual functions, you might see the following warnings:

```
WARNING: kmem_cache_destroy: 'px0_emlxs3_3_cache2'
  (3000383e030) not empty
WARNING: vmem_destroy('px0_emlxs3_3_vmem_top'):
  leaked 262144 identifiers
```

These messages do not affect the normal operation of the system and you can ignore them.

Workaround: None.

Fujitsu M10 Server Has Different SR-IOV Feature Limitations

On a Fujitsu M10 server you can assign PCIe endpoint devices and SR-IOV virtual functions from a particular PCIe bus to a maximum of 24 domains. The maximum is 15 domains for supported SPARC T-Series and SPARC M-Series platforms.

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.

Oracle Solaris OS Bugs Affecting the Oracle VM Server for SPARC 3.3 Software

The following Oracle Solaris OS bugs have been fixed in the fully qualified Oracle Solaris OS releases. These bugs might still be present in Oracle Solaris 10 OS versions. To avoid these problems, ensure that you run one of the Oracle Solaris OS versions that is associated with the bug ID.

To obtain details about the bugs in this table, review the bug reports.

TABLE 1–1 Resolved Oracle Solaris OS Bugs

Bug ID	Bug Description	Resolved in Oracle Solaris OS Versions
15707426	Logical Domains agent service does not come online if the system log service does not come online	Oracle Solaris 11 Oracle Solaris 10 1/13 with at lea ID 147147-26
15701258	Virtual CPU timeout failures during DR	Oracle Solaris 11 Oracle Solaris 10 1/13 with at lea ID 147147-26
15560811	Oracle Solaris 11: zones configured with an automatic network interface might fail to start	Oracle Solaris 11
15422900	Guest domain with too many virtual networks on the same network using DHCP can become unresponsive	Oracle Solaris 11

Bugs Affecting the Oracle VM Server for SPARC 3.3 Software

IOV Information Update Might Take Four Minutes After Running the cfgadm configure or cfgadm unconfigure Command

Bug ID 21953704: The `ldm list-io` command might not show the most up-to-date IOV information immediately after running a `cfgadm` command. You might have to wait as long as four minutes for the updated information to be available.

Workaround: None.

ovmtcreate Generates Incorrect OVF File If the Locale Is Not C

Bug ID 21780045: The `ovmtcreate` utility generates a NULL string for the `Version` information in the OVF file if the locale is not the C locale (non-English locale environment).

The values for the `Version` and `FullVersion` properties are null as shown by the XML lines that appear in bold in this example:

```
<ovf:VirtualSystem ovf:id="templates">
        <ovf:Info>Oracle VM Template</ovf:Info>
        <ovf:ProductSection ovf:class="com.oracle.ovmt">
                <ovf:Info>Oracle VM Template</ovf:Info>
                <ovf:Product>Oracle VM Template</ovf:Product>
                <ovf:Version></ovf:Version>
                <ovf:FullVersion></ovf:FullVersion>
```

When the `ovmtdeploy` utility uses the templates that you created by using the `ovmtcreate` utility in the non-C locale environment, a Java exception occurs because the templates include the NULL strings.

```
# /opt/ovmtutils/bin/ovmtdeploy -d guest10 -o /export/home/ovm \
/export/home/templates.ova

Oracle Virtual Machine for SPARC Deployment Utility
ovmtdeploy Version
Copyright (c) 2014, 2015, Oracle and/or its affiliates. All rights reserved.

STAGE 1 - EXAMINING SYSTEM AND ENVIRONMENT
------------------------------------------
Checking user privilege
Performing platform & prerequisite checks
Checking for required services
Named resourced available

2 - ANALYZING ARCHIVE & RESOURCE REQUIREMENTS
-------------------------------------------------
Checking .ova format and contents
Validating archive configuration
Exception in thread "main" java.lang.NullPointerException
        at ovfparse.OvfParse.getTagValue(OvfParse.java:233)
        at ovfparse.VmProduct.<init>(VmProduct.java:33)
        at ovfparse.VmSys.<init>(VmSys.java:72)
        at ovfparse.OvfParse.parseOVFByDOM(OvfParse.java:371)
        at ovfparse.OvfParse.<init>(OvfParse.java:56)
        at ovmtdeploy.Ovmtdeploy.exec(Ovmtdeploy.java:1841)
        at ovmtdeploy.Ovmtdeploy.main(Ovmtdeploy.java:1946)
```

Workaround: Perform the following steps:

1. Edit the OVF file to add the version numbers to the contents of the `Version` and `FullVersion` properties.
2. Re-archive the template `ova` by using the `gtar` command.

   For example:

   ```
   # /usr/bin/gtar -cf templates.ova templates.ovf templates.mf System.img.gz
   ```

3. Run the `ovmtdeploy` utility with `-k` option to skip checksum verification.

ldm add-vsan Fails After Replacing the PCIe Card

Bug ID 21674282: When you replace a PCIe card in the same slot, using the `ldm add-vsan` command that specifies an alias for the physical SCSI HBA device (`/SYS`) might fail.

Workaround: Do not specify the device name alias. Instead, specify the full device path name (`/pci`) for the `ldm add-vsan` command.

ovmtcreate Fails If the Service Domain Has Multiple Virtual Disk Servers

Bug ID 21635033: When a service domain has more than one virtual disk server (`vds`), running the `ovmtcreate` utility for a guest domain might fail because the utility checks only the first `vds` instance in the service domain.

For example, running the `ovmtcreate` utility for the `gdom3` domain fails if the virtual disk is configured as follows:

- The `primary` domain has four virtual disk servers (`vds`)
- The virtual disk server device that corresponds to the virtual disk on the `gdom3` domain is associated with `vds3`

In the following sample output, the lines in bold show that `vds0` is the first virtual disk server and virtual disk server device for the `gdom3` virtual disk is not `vds0`.

```
primary# ldm list -l -p -o disk
VERSION 1.15

DOMAIN|name=primary|
VDS|name=vds0|nclients=1
|vol=vol0|opts=|dev=/export/home/ovm/gdom0.img|mpgroup=
VDS|name=vds1|nclients=1
|vol=vol0|opts=|dev=/export/home/ovm/gdom1.img|mpgroup=
VDS|name=vds2|nclients=1
|vol=vol0|opts=|dev=/export/home/ovm/gdom2.img|mpgroup=
VDS|name=cdrom|nclients=3
|vol=1|opts=|dev=/export/home/ovm/sol-113_1.iso|mpgroup=
|vol=2|opts=|dev=/export/home/ovm/sol-113_2.iso|mpgroup=
|vol=3|opts=|dev=/export/home/ovm/sol-113_3.iso|mpgroup=
|vol=4|opts=|dev=/export/home/ovm/sol-113_4.iso|mpgroup=
VDS|name=vds3|nclients=1
|vol=disk0|opts=|dev=/export/home/ovm/gdom3.img|mpgroup=
DOMAIN|name=gdom0|
VDISK|name=vdisk0|vol=vol0@vds0|timeout=|dev=disk@0|server=primary|mpgroup=|id=0
VDISK|name=cdrom|vol=1@cdrom|timeout=|dev=disk@1|server=primary|mpgroup=|id=1
DOMAIN|name=gdom1|
VDISK|name=vdisk0|vol=vol0@vds1|timeout=|dev=disk@0|server=primary|mpgroup=|id=0
VDISK|name=cdrom|vol=2@cdrom|timeout=|dev=disk@1|server=primary|mpgroup=|id=1
DOMAIN|name=gdom2|
VDISK|name=vdisk0|vol=vol0@vds2|timeout=|dev=disk@0|server=primary|mpgroup=|id=0
```

```
VDISK|name=cdrom|vol=3@cdrom|timeout=|dev=disk@1|server=primary|mpgroup=|id=1
DOMAIN|name=gdom3|
VDISK|name=vdisk0|vol=disk0@vds3|timeout=|dev=disk@0|server=primary|mpgroup=|id=0
```

The following `ldm list` command shows the `gdom3` domain status:

```
primary# ldm list
NAME         STATE      FLAGS   CONS    VCPU  MEMORY   UTIL  NORM  UPTIME
primary      active     -n-cv-  UART    32    46848M   0.3%  0.3%  1d 51m
gdom0        active     -n----  5000    24    24G      0.0%  0.0%  1d 35m
gdom1        active     -n----  5001    24    24G      0.0%  0.0%  8d 18h 21m
gdom2        active     -n----  5002    24    24G      0.0%  0.0%  8d 17h 43m
gdom3        bound      ------  5003    24    24G
```

The following command shows the error you receive when running the `ovmtcreate` command for the `gdom3` domain:

```
# /opt/ovmtutils/bin/ovmtcreate -d gdom3 -o /export/home/ovmt
STAGE 1 - EXAMINING SYSTEM AND ENVIRONMENT
------------------------------------------
Performing platform & prerequisite checks
Checking user permissions
Checking for required packages
Checking for required services
Checking directory permissions

STAGE 2 - ANALYZING DOMAIN
--------------------------
Retrieving and processing attributes
Checking domain state
Getting domain resource settings
Discovering network topology
Discovering disk topology
ERROR: VDS Device  does not exist or not readable
```

Workaround: Ensure that the service domain has only one virtual disk server before you run the `ovmtcreate` utility.

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 M10 servers only.

Software running on Oracle SPARC systems 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 system 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 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.

DRM Is Unresponsive When a Domain Has Only a Few Virtual CPUs

Bug ID 21561834: If the number of virtual CPUs in a domain drops below four, DRM might fail to add virtual CPUs to the domain even when utilization significantly exceeds the upper utilization level. If the `util-upper` property value is greater than the default value of 70, DRM might fail to add virtual CPUs even if the domain has more than four virtual CPUs.

Workaround: Set the DRM policy's `elastic-margin` property value to at least 15.

```
primary# ldm set-policy elastic-margin=15 name=policy-name  domain-name
```

If the `util-upper` property value is greater than 70, set the DRM policy's `elastic-margin` property value to at least 20.

```
primary# ldm set-policy elastic-margin=20 name=policy-name  domain-name
```

Note – Ensure that the `elastic-margin` property value is lower than the `util-upper` property value.

Fujitsu M10 servers: Running ldm set-socket on an Active Domain Might Cause Unstable Operation of the Logical Domains Manager

Bug ID 21527087: In rare cases, using the `ldm set-socket` command to specify sockets for a running domain might cause the following unexpected behavior:

- The Logical Domains Manager might crash
- The `ldm set-socket` command completes but not all of the domain's CPUs and memory are remapped to the specified sockets

However, if the physical partition (PPAR) has more than 12 sockets, do not use the `ldm set-socket --restored-degraded` and `ldm set-socket socket_id=`*id* commands while the domain is running. If you do run these commands on a running domain, the `ldmd` state might become corrupted.

Workaround: Stop the domain before executing an `ldm set-socket` command.

It is always safe to clear an active domain's socket constraints by using the `ldm set-socket` command to set the `socket_id` property to a NULL value.

Random device busy or ldm remove-io Command Failures While Removing One or More PCIe Buses

Bug ID 21510615: Sometimes, you might get a persistent `device busy` or `ldm remove-io` failures while removing one or more PCIe buses.

Workaround: Check the `gdm` service, disable manually (or check and kill Xorg) and retry the `ldm remove-io` operation.

```
# svcs | grep gdm
# svcadm disable -st svc:/application/graphical-login/gdm:default
```

Or:

```
# ps -ef | grep Xorg
# pkill Xorg
```

Fujitsu M10 servers: Inconsistent Socket Constraints Might Cause the Logical Domains Manager to Crash During CPU Removal

Bug ID 21367043: In rare circumstances, socket constraints might become out of synchronization with the bound CPU and memory resources of a domain. The `ldm rm-vcpu`, `ldm set-vcpu`, `ldm rm-core`, and `ldm set-core` commands might cause the Logical Domains Manager to crash with the following error message in the `ldmd` SMF log:

```
fatal error: xcalloc(0,4) : one of number or size is <= 0 at line 1183
of affinity_core.c
```

Workaround: Clear the domain's socket constraints by using the following commands:

```
primary# ldm list-socket domain-name
primary# ldm set-socket socket_id= domain-name
```

`ldmpower` Causes a Segmentation Fault of ldmd

Bug ID 21369897: While administering a guest domain, running the `ldmpower` command causes a segmentation fault of the `ldmd` daemon.

Workaround: Do not execute the `ldmpower` command while performing addition or removal operations on a guest domain.

Fatal Error in the PCIe Fabric Causes a Root Domain to Panic

Bug IDs 21352084, 21861284, and 21861327: In rare circumstances, a root domain might panic if it receives an I/O error and starts to analyze the error while an I/O domain is reset.

The panic message is similar to the following:

```
panic[cpu15]/thread=2a1017d3c20:
Fatal error has occured in: PCIe fabric.(0x2)(0x245)
```

The ereports are dumped to the console at the time of the panic. The ereports show that some status register values, including the `pcie_ue_status` value, are all `FF`s. After the panic, the root domain reboots itself and recovers.

Workaround: None.

Slow I/O on Virtual SCSI HBA Guest Domain When One of the Service Domains is Down With a Virtual SCSI HBA Timeout Set

Bug ID 21321166: I/O throughput is sometimes slower when using a virtual SCSI HBA MPxIO path to an offline service domain.

Workaround: Disable the path to the offline service domain by using the `mpathadm disable` *path* command until the service domain is returned to service.

Fujitsu M10 servers: ldm shrink-socket Command Removes Additional Memory If the Memory Block Is Unaligned

Bug ID 21299404: If you use the `ldm shrink-socket` command to perform a memory DR operation and one of the domain's memory blocks is not 256-Mbyte aligned, the command might remove an additional 256 Mbytes of memory from the active domain. If the domain's memory is fragmented, the `ldmd` daemon might attempt to further remove additional memory.

Workaround: None.

ldm list-group Shows the Same Memory and I/O In Both the /SYS/MB and Other Resource Groups

Bug ID 21283102: The `ldm list-rsrc-group` command might show the same memory and I/O resource information under both the `/SYS/MB` (motherboard) and other resource groups. For example:

```
primary# ldm list-group
NAME                                    CORE  MEMORY   IO
/SYS/PM0                                32    64G      4
/SYS/PM1                                32    256G     4
/SYS/PM2                                32    128G     4
```

```
/SYS/PM3                                  32    128G     4
/SYS/MB                                   0     576G     16

primary# ldm list-group -a -l
NAME                                      CORE  MEMORY   IO
/SYS/PM0                                  32    64G      4

CORE
    CID                                                 BOUND
    0, 1                                                primary
    2, 3, 4, 5, 6, 7, 8, 9
    10, 11, 12, 13, 14, 15, 16, 17
    18, 19, 20, 21, 22, 23, 24, 25
    26, 27, 28, 29, 30, 31

MEMORY
    PA               SIZE             BOUND
    0x0              57M              _sys_
    0x3900000        32M              _sys_
    0x5900000        94M              _sys_
    0xb700000        393M             _sys_
    0x24000000       192M             _sys_
    0x30000000       31488M
    0x7e0000000      64M              _sys_
    0x7e4000000      64M              _sys_
    0x7e8000000      384M             _sys_
    0x80000000000    32G

IO
    DEVICE           PSEUDONYM        BOUND
    pci@300          pci_0            primary
    pci@340          pci_1            primary
    pci@380          pci_2            primary
    pci@3c0          pci_3            primary
--------------------------------------------------------------------------------

NAME                                      CORE  MEMORY   IO
/SYS/PM1                                  32    256G     4

CORE
    CID                                                 BOUND
    32, 33, 34, 35, 36, 37, 38, 39
    40, 41, 42, 43, 44, 45, 46, 47
    48, 49, 50, 51, 52, 53, 54, 55
    56, 57, 58, 59, 60, 61, 62, 63

MEMORY
    PA               SIZE             BOUND
    0x100000000000   768M
    0x100030000000   24G              primary
    0x100630000000   105728M
    0x180000000000   128G

IO
    DEVICE           PSEUDONYM        BOUND
    pci@400          pci_4            primary
    pci@440          pci_5            primary
    pci@480          pci_6            primary
    pci@4c0          pci_7            primary
```

```
--------------------------------------------------------------------------------

NAME                                    CORE  MEMORY   IO
/SYS/PM2                                32    128G     4

CORE
    CID                                                 BOUND
    64, 65, 66, 67, 68, 69, 70, 71
    72, 73, 74, 75, 76, 77, 78, 79
    80, 81, 82, 83, 84, 85, 86, 87
    88, 89, 90, 91, 92, 93, 94, 95

MEMORY
    PA               SIZE             BOUND
    0x200000000000   64G
    0x280000000000   64G

IO
    DEVICE           PSEUDONYM        BOUND
    pci@500          pci_8            primary
    pci@540          pci_9            primary
    pci@580          pci_10           primary
    pci@5c0          pci_11           primary
--------------------------------------------------------------------------------

NAME                                    CORE  MEMORY   IO
/SYS/PM3                                32    128G     4

CORE
    CID                                                 BOUND
    96, 97, 98, 99, 100, 101, 102, 103
    104, 105, 106, 107, 108, 109, 110, 111
    112, 113, 114, 115, 116, 117, 118, 119
    120, 121, 122, 123, 124, 125, 126, 127

MEMORY
    PA               SIZE             BOUND
    0x300000000000   64G
    0x380000000000   64G

IO
    DEVICE           PSEUDONYM        BOUND
    pci@600          pci_12           primary
    pci@640          pci_13           primary
    pci@680          pci_14           primary
    pci@6c0          pci_15           primary
--------------------------------------------------------------------------------

NAME                                    CORE  MEMORY   IO
/SYS/MB                                 0     576G     16

MEMORY
    PA               SIZE             BOUND
    0x0              57M              _sys_
    0x3900000        32M              _sys_
    0x5900000        94M              _sys_
    0xb700000        393M             _sys_
    0x24000000       192M             _sys_
    0x30000000       31488M
```

```
    0x7e0000000      64M             _sys_
    0x7e4000000      64M             _sys_
    0x7e8000000      384M            _sys_
    0x80000000000    32G
    0x100000000000   768M
    0x100030000000   24G             primary
    0x100630000000   105728M
    0x180000000000   128G
    0x200000000000   64G
    0x280000000000   64G
    0x300000000000   64G
    0x380000000000   64G

IO
    DEVICE           PSEUDONYM       BOUND
    pci@300          pci_0           primary
    pci@340          pci_1           primary
    pci@380          pci_2           primary
    pci@3c0          pci_3           primary
    pci@400          pci_4           primary
    pci@440          pci_5           primary
    pci@480          pci_6           primary
    pci@4c0          pci_7           primary
    pci@500          pci_8           primary
    pci@540          pci_9           primary
    pci@580          pci_10          primary
    pci@5c0          pci_11          primary
    pci@600          pci_12          primary
    pci@640          pci_13          primary
    pci@680          pci_14          primary
    pci@6c0          pci_15          primary
```

Workaround: See the detailed information for memory and I/O in the following columns to determine whether the same resource information is shown:

- **Memory:** `PA`, `SIZE` and `BOUND`
- **I/O:** `DEVICE`, `PSEUDONYM` and `BOUND`

Virtual SCSI HBA Does Not See Dynamic LUN Changes Without a Reboot

Bug ID 21188211: If LUNs are added to or removed from a virtual SAN after a virtual SCSI HBA is configured, the `ldm rescan-vhba` command sometimes does not show the new LUN view.

Workaround: Remove the virtual SCSI HBA and then re-add it. Check to see whether the LUNs are seen. If the removal and re-add operations are unsuccessful, you must reboot the guest domain.

Logical Domains Manager Should Not Rely on Polling to Obtain Configuration Status From the DIO Agent

Bug ID 21114622: When you execute the `ldm create-vf` or `ldm destroy-vf` command, the associated physical function driver is detached and re-attached, which can take a significant but

unquantifiable amount of time. The amount of time depends the number of virtual functions that are involved and the complexity of the target hardware device.

Running the `ldm list-io` command might show that the physical function (and its child virtual functions) have the `INV` (invalid) status.

Currently, the Logical Domains Manager polls the agent for a period of time and then stops polling. If the polling period is too short, the device might show the `INV` status indefinitely.

Note – The fix for bug 20772410 should reduce the occurrence of this problem.

Workaround: From the root domain that owns the physical function device, restart the `ldoms/agents` service.

```
primary# svcadm restart ldoms/agents
```

Run this command if the `INV` status persists for at least six minutes after issuing the `ldm create-vf` or `ldm destroy-vf` command.

vhba Should Support SCSI HBAs When MPxIO Is Enabled in the Service Domain

Bug ID 20951004: `vhba` should support SCSI HBAs when MPxIO is enabled in the service domain.

Workaround: Disable MPxIO for all the initiator ports on the service domain by running the following command:

```
# stmsboot -d
```

Suppress FRU-Monitor Alerts When PCI Bus Is Re-Assigned From the primary Domain to a Guest Domain

Bug ID 20882700: When a PCIe device (or an SR-IOV virtual function) is removed from or added to a domain, the Oracle Solaris 11.3 `fmd` fault management daemon reports the event in exactly the same way as if an FRU had been physically removed or added.

You might see console messages and messages in the `/var/adm/messages` file similar to the following:

-

```
SUNW-MSG-ID: FMD-8000-A0, TYPE: Alert, VER: 1, SEVERITY: Minor
EVENT-TIME: Tue May 19 18:39:41 PDT 2015
PLATFORM: unknown, CSN: unknown, HOSTNAME: starbuck
SOURCE: software-diagnosis, REV: 0.1
EVENT-ID: 5077e6c3-6a15-457e-a55b-cb72ea5f9728
DESC: FRU has been added to the system.
```

```
AUTO-RESPONSE: FMD topology will be updated.
IMPACT: System impact depends on the type of FRU.
REC-ACTION: Use fmadm faulty to provide a more detailed view of this event.
Please refer to the associated reference document at
http://support.oracle.com/msg/FMD-8000-A0 for the latest service procedures
and policies regarding this diagnosis.
```

-

```
# fmadm faulty
--------------- ------------------------------------  ----------- ---------
TIME            EVENT-ID                              MSG-ID      SEVERITY

--------------- ------------------------------------  ----------- ---------
Apr 14 10:04:00 2d981602-975c-4861-9f26-e37360eca697  FMD-8000-CV Minor

Problem Status    : open
Diag Engine       : software-diagnosis / 0.1
System
    Manufacturer  : Oracle Corporation
    Name          : SPARC T7-2
    Part_Number   : T7_2
    Serial_Number : T7_2
    Host_ID       : 86582a8c

----------------------------------------
Suspect 1 of 1 :
   Problem class : alert.oracle.solaris.fmd.fru-monitor.fru-remove
   Certainty   : 100%

   FRU
     Status          : active/not present
     Location        : "/SYS/MB/PCIE1"
     Manufacturer    : unknown
     Name            : unknown
     Part_Number     : unknown
     Revision        : unknown
     Serial_Number   : unknown
     Chassis
        Manufacturer  : Oracle-Corporation
        Name          : SPARC-T7-2
        Part_Number   : T7_2
        Serial_Number : T7_2
   Resource
     Status          : active/not present

Description : FRU '/SYS/MB/PCIE1' has been removed from the system.

Response    : FMD topology will be updated.

Impact      : System impact depends on the type of FRU.

Action      : Use 'fmadm faulty' to provide a more detailed view of this event.
              Please refer to the associated reference document at
              http://support.oracle.com/msg/FMD-8000-CV for the latest service
              procedures and policies regarding this diagnosis.
```

Workaround: You can ignore these alerts as long as they were generated by explicit administrator actions to add or remove an I/O device from a domain.

mpathadm Shows Incorrect Path State Output for a Virtual SCSI HBA When a Fibre Channel Cable Is Pulled

Bug ID 20876502: Pulling the SAN cable from a service domain that is part of a virtual SCSI HBA MPxIO guest domain configuration causes the Path State column of the `mpathadm` output to show incorrect values. In addition, pulling the cable leads to I/O operation failures in the guest domain.

Workaround: Plug in the SAN cable and run the `ldm rescan-vhba` command for all the virtual SCSI HBAs to the service domain that has the cable attached. After performing this workaround, the guest domain should resume performing I/O operations.

device busy Error When Attempting to Remove a PCIe Bus That Hosts an SES-Enabled Storage Device

Bug ID 20774477: If you use SES-enabled storage devices, you might see a `device busy` error when you attempt to remove a PCIe bus that hosts these devices. To determine whether you are using this type of storage device, search for the `ses` or `enclosure` string in the `ldm list-io -l` output for the PCIe bus.

Workaround: Perform one of the following workarounds to remove the PCIe bus:

- **Dynamically remove the PCIe bus.**
  1. Disable the FMD service.

     ```
     primary# svcadm disable -st svc:/system/fmd
     ```
  2. Remove the PCIe bus.

     ```
     primary# ldm remove-io bus
     ```
  3. Re-enable the FMD service.

     ```
     primary# svcadm enable svc:/system/fmd
     ```
- **Statically remove the PCIe bus.**
  1. Place the root domain that has the PCIe bus in a delayed reconfiguration.

     ```
     primary# ldm start-reconf root-domain
     ```
  2. Remove the PCIe bus.

     ```
     primary# ldm remove-io bus
     ```
  3. Perform a reboot from the root domain console.

     ```
     root-domain# reboot
     ```

rcm_daemon Might Emit a Message on the Console During an ldm remove-io Operation

Bug ID 20619894: If the `system/management/hwmgmtd` package is not installed, a dynamic bus remove operation causes the `rcm_daemon` to print the following message on the console:

```
rcm_daemon[839]: rcm script ORCL,pcie_rc_rcm.pl: svcs: Pattern 'sp/management'
doesn't match any instances
```

Workaround: You can safely ignore this message.

Dynamic Bus Removal Should be Aware of Virtual SAN Devices Prior to Removing a PCIe Bus

Bug ID 20532270: Be aware of any direct I/O or dynamic bus removal operations that attempt to remove the physical SCSI HBA from the virtual SAN's control.

If you perform an `ldm remove-io` operation on a PCIe resource that is referenced by a virtual SAN device, that device is unusable if it has never been referenced by an `ldm add-vhba` command. If the `ldm remove-io` operation occurs after you run the `ldm add-vhba` command, the `vsan` module prevents the PCIe resource from being removed.

Workaround: Delete the virtual SAN.

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.

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.

Panic When Using the ldm rm-io *virtual-function* Command to MPxIO That Contains a Virtual SCSI HBA

Bug ID 20046234: When a virtual SCSI HBA and a Fibre Channel SR-IOV device can view the same LUNs in a guest domain when MPxIO is enabled, a panic might occur. The panic occurs if the Fibre Channel SR-IOV card is removed from the guest domain and then re-added.

Workaround: Do not configure a guest domain with Fibre Channel SR-IOV and a virtual SCSI HBA when both have MPxIO enabled.

ixgbevf Nodes on an I/O Domain Are Reported as Disabled By the ipadm Command and Non-Existent By the ifconfig Command

Bug ID 20004281: When a `primary` domain is power cycled, `ixgbevf` nodes on the I/O domain might be reported as disabled by the `ipadm` command, and as nonexistent by the `ifconfig` command.

Workaround: Re-enable the IP interfaces:

```
# svcadm restart network/physical:default
```

HGXE Interfaces Are Unusable If Assigned by Using Direct I/O to an I/O Domain

Bug ID 19943809: The `hxge` driver cannot use interfaces inside an I/O domain when the card is assigned by using the direct I/O feature.

The following warning is issued to the system log file:

```
WARNING: hxge0 : <== hxge_setup_mutexes: failed 0x1
```

Workaround: Add the following line to the `/etc/system` and reboot:

```
set px:px_force_intx_support=1
```

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.

Oracle VM Server for SPARC No Longer Keeps Track of Freed MAC Addresses

Bug ID 19078763: Oracle VM Server for SPARC no longer keeps track of freed MAC addresses. MAC addresses are now allocated by randomly selecting an address and then confirming that address is not used by any logical domains on the local network.

Sun Storage 16 Gb Fibre Channel Universal HBA Firmware Does Not Support Bandwidth Controls

Bug ID 18083904: The firmware for Sun Storage 16 Gb Fibre Channel Universal HBA, Emulex cards does not support setting bandwidth controls. The HBA firmware ignores any value that you specify for the `bw-percent` property.

Workaround: None.

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.

Issues Might Arise When FMA Detects Faulty Memory

Bug ID 17576087: Performing a power cycle of the system to a saved configuration might not restore the memory after the faulty memory has been replaced.

Workaround: After you replace the faulty memory, perform a power cycle of the system to the `factory-default` configuration. Then, perform a power cycle of the system to the configuration that you want to use.

DLMP Does Not Work in a Guest Domain on Top of a Virtual Network Device or an SR-IOV Virtual Function

You cannot configure a DLMP aggregation on an SR-IOV NIC virtual function or a virtual network device in a guest domain.

Cannot Install the Oracle Solaris 11.1 OS Using an EFI GPT Disk Label on Single-Slice Virtual Disk

Bug ID 17422973: The installation of the Oracle Solaris 11.1 OS on a single-slice disk might fail with the following error on a SPARC T4 server that runs at least system firmware version 8.4.0, a SPARC T5, SPARC M5, or SPARC M6 server that runs at least system firmware version 9.1.0, and a Fujitsu M10 server that runs at least XCP version 2230:

```
cannot label 'c1d0': try using fdisk(1M) and then provide a specific slice
Unable to build pool from specified devices: invalid vdev configuration
```

Workaround: Relabel the disk with an SMI label.

Guest Domain Panics at lgrp_lineage_add(mutex_enter: bad mutex, lp=10351178)

Bug ID 17020950: After migrating an active domain from a SPARC T4 platform to a SPARC T5, SPARC M5, or SPARC M6 platform that was bound using firmware version 8.3, performing a memory dynamic reconfiguration might result in a guest domain panic.

Workaround: Before you perform the migration, update the SPARC T4 system with version 8.4 of the system firmware. Then, rebind the domain.

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 3.3 Administration Guide*
- "How to Remove an InfiniBand Virtual Function From a Root Domain" in *Oracle VM Server for SPARC 3.3 Administration Guide*

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
   ```

Subdevices Under a PCIe Device Revert to an Unassigned Name

Bug ID 16299053: After disabling a PCIe device, you might experience unexpected behavior. The subdevices that are under the disabled PCIe device revert to the non-assigned names while the PCIe device is still owned by the domain.

Workaround: If you decide to disable a PCIe slot on the ILOM, ensure that the PCIe slot is not assigned to a domain by means of the direct I/O (DIO) feature. That is, first ensure that the PCIe slot is assigned to the corresponding root domain before disabling the slot on the ILOM.

If you disable the PCIe slot on the ILOM while the PCIe slot is assigned to a domain with DIO, stop that domain and reassign the device to the root domain for the correct behavior.

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
```

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
0000100001540be30 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
```

`ixgbevf` Device in SR-IOV Domains Might Become Disabled When Rebooting the primary Domain

Bug ID 16224353: After rebooting the primary domain, `ixgbevf` instances in `primary` domain might not work.

Workaround: None.

SPARC M5-32 and SPARC M6-32: LSI-SAS Controller Is Incorrectly Exported With SR-IOV

Bug ID 16071170: On a SPARC M5-32 or a SPARC M6-32 system, the internal SAS controllers are exported as SR-IOV-enabled controllers even though these cards do not support SR-IOV.

The Oracle VM Server for SPARC log shows the following messages when attempting to create the physical function on these cards:

```
Dec 11 04:27:54 warning: Dropping pf
pci@d00/pci@1/pci@0/pci@0/pci@4/LSI,sas@0: no IOV capable driver
Dec 11 04:27:54 warning: Dropping pf
pci@d80/pci@1/pci@0/pci@c/pci@0/pci@4/LSI,sas@0: no IOV capable driver
Dec 11 04:27:54 warning: Dropping pf
pci@c00/pci@1/pci@0/pci@c/pci@0/pci@4/LSI,sas@0: no IOV capable driver
Dec 11 04:27:54 warning: Dropping pf
pci@e00/pci@1/pci@0/pci@0/pci@0/pci@4/LSI,sas@0: no IOV capable driver
```

The system has four LSI SAS controller ports, each in one IOU of the SPARC M5-32 and SPARC M6-32 assembly. This error is reported for each port.

Workaround: You can ignore these messages. These messages indicate only that the LSI-SAS controller devices on the system are capable of SR-IOV but no SR-IOV support is available for this hardware.

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

Bug ID 16068376: On a T5-8 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.

No Error Message When a Memory DR Add is Partially Successful

Bug ID 15812823: In low free-memory situations, not all memory blocks can be used as part of a memory DR operation due to size. However, these memory blocks are included in the amount of free memory. This situation might lead to a smaller amount of memory being added to the domain than expected. No error message is shown if this situation occurs.

Workaround: None.

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.

To ensure that the system remains in a state in which none of the previous actions have taken place, see Using the ldm init-system Command to Restore Domains on Which Physical I/O Changes Have Been Made (`https://support.oracle.com/epmos/faces/DocumentDisplay?id=1575852.1`).

Control Domain Requires the Lowest Core in the System

Bug ID 15778392: The control domain requires the lowest core in the system. So, if core ID 0 is the lowest core, it cannot be shared with any other domain if you want to apply the whole-core constraint to the control domain.

For example, if the lowest core in the system is core ID 0, the control domain should look similar to the following output:

```
# ldm ls -o cpu primary
NAME
primary

VCPU
VID    PID    CID    UTIL STRAND
0      0      0      0.4%   100%
1      1      0      0.2%   100%
2      2      0      0.1%   100%
3      3      0      0.2%   100%
4      4      0      0.3%   100%
5      5      0      0.2%   100%
6      6      0      0.1%   100%
7      7      0      0.1%   100%
```

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 systems, 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.us.oracle.com/oip/faces/secure/km/DocumentDisplay.jspx?id=1325454.1`).

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

ldm stop Command Might Fail If Issued Immediately After an ldm start Command

Bug ID 15759601: If you issue an `ldm stop` command immediately after an `ldm start` command, the `ldm stop` command might fail with the following error:

```
LDom domain-name stop notification failed
```

Workaround: Reissue the `ldm stop` command.

System Panics When Rebooting a primary Domain That Has a Very Large Number of Virtual Functions Assigned

Bug ID 15750727: A system might panic when you reboot a `primary` domain that has a very large number of virtual functions assigned to it.

Workaround: Perform one of the following workarounds:

- Decrease the virtual function number to reduce the number of failed virtual functions. This change might keep the chip responsive.
- Create more Interrupt Resource Management (IRM) pools for the `ixgbe` virtual function because only one IRM pool is created by default for all the `ixgbe` virtual functions on the system.

Partial Core primary Fails to Permit Whole-Core DR Transitions

Bug ID 15748348: When the `primary` domain shares the lowest physical core (usually `0`) with another domain, attempts to set the whole-core constraint for the `primary` domain fail.

Workaround: Perform the following steps:

1. Determine the lowest bound core that is shared by the domains.

   ```
   # ldm list -o cpu
   ```

2. Unbind all the CPU threads of the lowest core from all domains other than the `primary` domain.

   As a result, CPU threads of the lowest core are not shared and are free for binding to the `primary` domain.

3. Set the whole-core constraint by doing one of the following:
   - Bind the CPU threads to the `primary` domain, and set the whole-core constraint by using the `ldm set-vcpu -c` command.
   - Use the `ldm set-core` command to bind the CPU threads and set the whole-core constraint in a single step.

Cannot Use Oracle Solaris Hot-Plug Operations to Hot-Remove a PCIe Endpoint Device

Bug ID 15721872: You cannot use Oracle Solaris hot-plug operations to hot-remove a PCIe endpoint device after that device is removed from the `primary` domain by using the `ldm rm-io` command. For information about replacing or removing a PCIe endpoint device, see "Making PCIe Hardware Changes" in *Oracle VM Server for SPARC 3.3 Administration Guide*.

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.

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

Bug ID 15668368: A SPARC T3-1 system 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 system, 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
```

Guest Domains Running Oracle Solaris 10: Memory DR Removal Operations With Multiple Plumbed NIU nxge Instances Can Hang Indefinitely and Never Complete

Bug ID 15667770: When multiple NIU `nxge` instances are plumbed on a domain, the `ldm rm-mem` and `ldm set-mem` commands, which are used to remove memory from the domain, might never complete. To determine whether the problem has occurred during a memory removal operation, monitor the progress of the operation with the `ldm list -o status` command. You might have encountered this problem if the progress percentage remains constant for several minutes.

Workaround: Cancel the `ldm rm-mem` or `ldm set-mem` command, and check whether a sufficient amount of memory was removed. If not, a subsequent memory removal command to remove a smaller amount of memory might complete successfully.

If the problem has occurred on the `primary` domain, do the following:

1. Start a delayed reconfiguration operation on the primary domain.

   ```
   # ldm start-reconf primary
   ```
2. Assign the desired amount of memory to the domain.
3. Reboot the `primary` domain.

If the problem occurred on another domain, stop the domain before adjusting the amount of memory that is assigned to the domain.

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.

Dynamic Reconfiguration of MTU Values of Virtual Network Devices Sometimes Fails

Bug ID 15631119: If you modify the maximum transmission unit (MTU) of a virtual network device on the control domain, a delayed reconfiguration operation is triggered. If you subsequently cancel the delayed reconfiguration, the MTU value for the device is not restored to the original value.

Recovery: Rerun the `ldm set-vnet` command to set the MTU to the original value. Resetting the MTU value puts the control domain into delayed reconfiguration mode, which you need to cancel. The resulting MTU value is now the original, correct MTU value.

```
# ldm set-vnet mtu=orig-value vnet1 primary
# ldm cancel-op reconf primary
```

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.

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.

PCI Express Dual 10-Gigabit Ethernet Fiber Card Shows Four Subdevices in ldm list-io -l Output

Bug ID 15597025: When you run the `ldm ls-io -l` command on a system that has a PCI Express Dual 10-Gigabit Ethernet Fiber card (X1027A-Z) installed, the output might show the following:

```
primary# ldm ls-io -l
...
pci@500/pci@0/pci@c PCIE5 OCC primary
```

```
network@0
network@0,1
ethernet
ethernet
```

The output shows four subdevices even though this Ethernet card has only two ports. This anomaly occurs because this card has four PCI functions. Two of these functions are disabled internally and appear as `ethernet` in the `ldm ls-io -l` output.

Workaround: You can ignore the `ethernet` entries in the `ldm ls-io -l` output.

ldm Commands Are Slow to Respond When Several Domains Are Booting

Bug ID 15572184: An `ldm` command might be slow to respond when several domains are booting. If you issue an `ldm` command at this stage, the command might appear to hang. Note that the `ldm` command will return after performing the expected task. After the command returns, the system should respond normally to `ldm` commands.

Workaround: Avoid booting many domains simultaneously. However, if you must boot several domains at once, refrain from issuing further `ldm` commands until the system returns to normal. For instance, wait for about two minutes on Sun SPARC Enterprise T5140 and T5240 servers and for about four minutes on the Sun SPARC Enterprise T5440 server or Sun Netra T5440 server.

Oracle Solaris 11: Zones Configured With an Automatic Network Interface Might Fail to Start

Bug ID 15560811: In Oracle Solaris 11, zones that are configured with an automatic network interface (`anet`) might fail to start in a domain that has Logical Domains virtual network devices only.

- **Workaround 1:** Assign one or more physical network devices to the guest domain. Use PCIe bus assignment, the Direct I/O (DIO), or the SR-IOV feature to assign a physical NIC to the domain.
- **Workaround 2:** If the zones configuration requirement is to have interzone communication only within the domain, create an `etherstub` device. Use the `etherstub` device as the "lower link" in the zones configuration so that virtual NICs are created on the `etherstub` device.
- **Workaround 3:** Use exclusive link assignment to assign a Logical Domains virtual network device to a zone. Assign virtual network devices, as needed, to the domain. You might also choose to disable inter-vnet links to be able to create a large number of virtual network devices.

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
```

Sometimes, Executing the uadmin 1 0 Command From a Logical Domains System Does Not Return the System to the OK Prompt

Bug ID 15511551: Sometimes, executing the `uadmin 1 0` command from the command line of a Logical Domains system does not leave the system at the `ok` prompt after the subsequent reset. This incorrect behavior is seen only when the Logical Domains variable `auto-reboot?` is set to `true`. If `auto-reboot?` is set to `false`, the expected behavior occurs.

Workaround: Use this command instead:

```
uadmin 2 0
```

Or, always run with `auto-reboot?` set to `false`.

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.

OpenBoot PROM Variables Cannot be Modified by the eeprom Command When the Logical Domains Manager Is Running

Bug ID 15387338: This issue is summarized in "Logical Domains Variable Persistence" in *Oracle VM Server for SPARC 3.3 Administration Guide* and affects only the control domain.

Cannot Set Security Keys With Logical Domains Running

Bug ID 15370442: The Logical Domains environment does not support setting or deleting wide-area network (WAN) boot keys from within the Oracle Solaris OS by using the `ickey(1M)` command. All `ickey` operations fail with the following error:

```
ickey: setkey: ioctl: I/O error
```

In addition, WAN boot keys that are set using OpenBoot firmware in logical domains other than the control domain are not remembered across reboots of the domain. In these domains, the keys set from the OpenBoot firmware are valid only for a single use.

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
```

Resolved Issues

The following enhancement requests and bugs have been fixed for the Oracle VM Server for SPARC 3.3 software release:

21446603	Fix suspended reconfiguration with mblock splitting
21439137	Unbound constrained domains vanish post ldmd crash on reconf cancellation
21438526	PCIe slot status remains EMP after successful hotplug enable of the connector
21434726	Ldom migration fails with Ldoms 3.3 Integration 15 and target ldmd cores
21427171	Disable unsupported 'dr-pdom' domain service messages
21427137	Disable move-core and move-memory commands
21419747	Problem in PAPL/M10
21367340	LDoms: auto changing physical-bindings constraint
21366945	Recovery Mode messages should indicate recovery is ongoing
21350128	Fix for 21240346 breaks DTrace probes in ldmd
21342706	PCIe slots stuck in UNK status following DBUS add operation
21328933	vhba timeout value does not persist across guest domain migration
21313658	Fujitsu socket remapping requires latest firmware
21306352	Board DR causes ldmd abort after ldmad changes in S11.2 SRU9
21296317	ldomMIB.so gets issues with athena_skip_cpu_memory and fj_force_flag symbols

21258580	Migration failed with DRM policy enabled
21258063	Memory operation block by "A memory blacklist operation is being processed"
21251095	Migrated guest domain panics on resumption on target machine
21249867	Migration results in hung ldmd on the target machine
21248066	mem_bind_real() reports "Cannot move memory" warning during a failed migration
21240346	Domains with Kernel Zones should emit meaningful output for migration failure
21239682	Add ovmtutils to ldoms-incorporation
21234461	FC SRIOV: "set-io port-wwn=<> node-wwn=<> VF" succeeding when not in delay-reconf
21209454	rgrp_get_split_mblk_cap_info: referenced symbol not found
21203360	CPU autoreplacement breaks CMI multi-instance
21202779	Domains left suspended if suspend fails but HV is reconfigured
21196845	Live Service on M10 removes memory other than specified resource group
21195903	Possible migration failures and memory corruption
21193823	Assertion failure on rm-memory -g as do_mem_dr() called while not on sequencer
21191079	The LDoms smf service on the newly deployed sparc system goes to maintenance
21189304	Assertion failure when CMI nop attempted for active domain
21181793	ovmtconfig fails to set the auto-boot property
21174268	Recovery with Fujitsu sockets may calculate resources incorrectly
21174261	init-system -f should ignore Fujitsu sockets
21174256	Recovery should ignore degraded sockets saved in bootset
21174253	Degraded Fujitsu sockets not reset to current bound resources
21174250	M10-specific code breaks recovery on non-M10 platforms
21173437	Board DR move failures if domains do not support mblock splitting
21173038	Board DR move failure with bound domain causes SEGV
21165224	add-domain -i fails with socket not a known resource
21164944	Blacklisting and board evacuation can run at the same time as DRM
21151265	/SYS/MB is rejected as invalid to resource group move and remove CLIs

21150868	ovmtcreate retains domain name in XML even when alternate specified
21150415	M10 specific code breaks remap CLIs on non-M10 platforms
21138802	Migration does not take advantage of mblock splitting support
21134659	set-vcpu of domain with physical-bindings tries to set incorrect number of CPUs
21129227	Disable evacuate-io/restore-io code for LDoms 3.3
21123013	Live migration pre-suspend API group check should be removed
21117019	Properties files should include other settings such as allowing NIS and LDAP
21116899	ldm list-group only shows partial memory on T5-2
21116432	ovmtutils: remove the standalone ovfparse script from /opt/ovmtutils/bin
21116138	Incorrect total memory size in ldm list-group -a -l command for /SYS/MB
21109812	ovmtprop -h should include banner and copyright message
21109793	Solaris boilerplate should not contain default values
21109771	ovmtconfig properties files should not contain default values
21093569	ovmtutils: ovmtlibrary should not set a limit on property description
21092364	Incorrect format strings in ldomcli/policy.c and rgrp_mem.c, nvlist lint errors
21092307	No error messages for XMPP commands in Recovery Mode
21084406	Out of sequence 'dr-pdom' progress messages sent to GM
21082758	Missing HOSTID in "ldm ls-constraints -x <guest>"
21081832	Logic error in SDIO NPRD recovery
21077022	"PRI component node has no topo-hc-name" warning in ldmd service log
21072214	Policy active does not set to vcpu-max when domain is loaded and has >vcpu-max
21070377	init-system reports Malformed XML Document
21060914	Illegal recovery_mode_boot_timeout values should be handled gracefully
21057594	Better policy checks between PM and SP policy needed
20996572	'core -g' commands should not run at the same time as DRM
20995171	Heartbeat thread parameters may be clobbered
20991993	Guest Live migration from T5-8 to T5-2 fails: API group 0x1 v1.5 is not supported
20976838	Assertion failed: ildp->rgrp_mem.ildom_sel_mblks == 0, file rgrp_mem.c:1746

20975146	Need a way to force full reconfig
20974426	Disable changes to autosave made by 18746688
20970759	ldomTable output shows 'Wrong Types' from LDoms 3.3 Integration 05
20960651	Wholecore DRM continues in off state after migration from Ldoms 3.2 to 3.3
20957412	Assert failure in platsvc.c:domain_suspend_data_handler()
20938662	ovmtcreate should not rely on ldoms package verification, but presence only
20938619	Suspend attempts wait when a prior suspend is in progress
20937100	ldm list-group should improve how the SPARC T7 motherboard is handled
20927881	Whole core DRM policy goes to OFF state after live migration
20927189	set-core command fails if existing number CPUs is not multiple of core-size
20925389	Fatal in pri_get_cmis() if all strands in chip degraded
20905098	Assertion fail during migration on OVM3.3
20898641	Update move-core to be consistent with move-memory
20891044	A panic/reboot during suspend leaves a guest 'suspending'
20890579	Remove ovmtutils from ldoms-incorporation
20888305	set-core in delayed reconfig allows exceeding max-cores limit - LDoms 3.3_b05
20886229	'move-mem _sys_' should only be allowed on certain platforms
20885337	Only attempt memory or core remap with compatible firmware
20885286	Add debug command to allocate RKEY table memory
20884770	Don't suspend a guest with physical IO if it is not supported
20881247	Addition of VCPUs does not happen when weighted mean util > Util upper value
20880305	Restore unlimited attack/decay
20878254	Support Live Service on Fujitsu M10
20868341	Min binding size for real memory should be 4MB not 256MB
20868330	Binding real address should allow 4MB granularity instead of 256MB
20868325	bind_real_status() returns incorrect offset if a mblock needs to be split
20868319	PA range selected by best_suitable_mblks() may not be optimal
20862729	Solaris expects nmblk=1 not 0 for requesting a memory map

20862661	Logical Domains Manager should retry failed mem allocation for multiple mblocks on the same LGPG
20854392	Support for Fujitsu socket constraints
20854350	Enhance Board DR memory processing
20854296	Fujitsu CPU autoreplacement should use CPU remapping
20849084	Unable to clear the vnet maxbw using XML interface
20848996	LDoms 3.1: vnet maxbw is sent in Mbps instead of bps
20826013	Recovery Mode: boot timeout handling needs improvement
20822568	set-vnet does not detect dup mac addrs between auto generated alt-mac-addrs
20811827	timed_read() in sptok_init_virtual_channel() doesn't timeout
20811467	LDoms deployment on zfs filesystem is taking too long to install
20799692	OVMT utilities: ovmtlibrary -c list -i 1 -o causes sqlite3 error: syntax error
20796786	MIB missing ldomVdsdevTable, ldomVnetTable, ldomVdiskTable from LDoms 3.2 b13
20772407	NPRD: ldm looses track of PFs, reporting not-found on ldm ls-io -l <PF> cmds
20765597	OVMT utilities: ovmtutils build cleanup
20765043	OVMT utilities: add the ovmtutils build generated files to .hgignore
20764954	OVMT utilities: remove xindent python code which is not written by Oracle
20752170	Assertion failed: ldp->rgrp_mem.ildom_sel_mblks != 0 in rgrp_mem_remove()
20742183	Build updates hg controlled file after 18714584
20738497	OVMT utilities: add fortify rules for ovmtutils java code
20704924	Mblocks should not be split if it is not supported by a guest
20704724	Defer retained memory API group check
20704064	A DR memory query hangs when sent to a hung domain
20694418	list-dependencies does not handle mpgroups correctly
20692005	mem_allocate_real assert hit during migrate.py
20690717	migrate.py cannot move memory because there is no suitable free memory
20688980	Ldom 3.1.1.2/M5-32/ldmd coredumps while doing a live migration
20687207	Releasing a resource that is not held should return success
20680100	Assertion failed in function debug_chk_mem_align_page()

20648386	Held resources do not persist across a restart
20630625	disable_self() in ldmd_include.sh SMF helper function can drop S10 support
20604908	get_ldom_by_rc_name() does not always return the correct root domain
20586961	Problem with mgmtldmgr/ldommgr-mgmt
20586857	ldmd dumps core at complete_pending_evac_memory subsequent to core fault
20581427	Failed ldm move-mem may pass after add-mem
20570207	Recovery Mode hangs primary domain if Elastic PM
20540894	VHBA: ldm ls -o san,hba not implemented properly
20537580	Parfait 1.6 detects some null pointer dereferences and memory leaks
20532612	vHBA Logical Domains Manager integration
20514062	Remove support for Logical Domains Manager on Oracle Solaris 10 primary domain
20508772	Domains still suspended after failure in BDR sequence
20506186	Failed suspended reconfig can lead to inconsistent HV MD
20496285	Restore partial move failed to preserve PA LGPG
20487202	Enable Recovery Mode by default
20475338	Move memory failed to handle free memory = 0
20458698	3.2 fix for 19513561 is incomplete
20458388	ovmtconfig should supply timestamps in debug mode
20457740	Logical Domains Manager should support enabling trusted mode for VNETs
20432421	Fujitsu socket commands behave unexpectedly in delayed reconfiguration
20427316	vpci port destroyed after remove-io from one px
20426593	ldm list-rsrc-group incorrectly displays PCIE bus for CMIOU10 under CMIOU1
20416951	Disable generation of LDoms audit records
20399591	Migration bind should not split mblocks
20395425	Update LDM_RELEASE_NUM in Makefile.master
20373163	Assertion failed: 'tgt_num > 0' in core_remap_restore()
20372917	Logical Domains Manager should check DS channel state before suspending guests
20372132	cpu-arch=(migration-class1|generic) should not include DAX in the MDs

20365522	SPARC T7 domain cannot migrate back after it is migrated to another system
20340446	ovmtcreate fails to read correct number of CPU assigned to guest domain dax1
20321459	Recovery fails when service domain missing virtual disk back end
20305752	Add memory using mblock option fails after memory evacuation
20305649	SPARC M7 After SP reset, Solaris does not reflect policy change in ILOM
20288350	mdprint.py should be able to print the PRI and save PRI/guest MD's
20257979	PF/VF stuck in INV state when in delayed reconf & ldomsmanager is restarted
20245881	ldmd dumps core at mem_unconfigure_range() when system memory is faulted
20245294	Impacted domain memory removal does not use MEMDR_MINSIZE
20244972	All mblocks in a large page must be remapped together
20232511	Unable to migrate vcpu resource state T41 to T42 ldoms3.1 / 8.7.0 build 23
20215998	move-memory should pass a board structure to subroutines
20215712	dr_mem get-map needs to handle timeout case
20207846	Memory DR add should break large requests at 16G RA boundary
20207824	Fatal error when res ids run out due to memory fragmentation
20199059	Logical Domains Manager frees an arg pointer twice in add memory failure path
20187335	Crash in SSL_write() returning from a resource group operation
20187197	CPUs are not resumed after power cap under the limit
20157499	Memory DR silently adds less memory than requested
20124236	ldmd should log HVMDs, error info after failing to reconfigure the HV
20095724	"channel-endpoint" is missing "back" arc to "channel-endpoints"
20047766	Warning: Attempt to reset hvmem@0xb700000:0x18900000 to 0x40000b700000:0x1890000
20035182	Enable move-core and move-memory commands for Live Service
20034433	"channel %s has been restored" message is misleading
20027321	mv-mem does not allow _sys_ to be moved
20011636	TestFailure Not all resource pmstates were lowered from the previous state of 1
20005608	list-group output should sort resource groups numerically
20003817	move-mem with -s option needs to be revisited

19976062	Panic, assertion failed, file: cpu_idle.c, line: 737
19926799	move-mem should only do one suspend/resume
19914686	Missing memory info from ldm list-group -a -l on T7-2 or T5-2
19905790	DR DS requests can timeout when stalled by a CLI command
19898359	Logical Domains Manager error sending SET_HEARTBEAT message: Invalid argument
19887895	Broken link /var/opt/SUNWldm after installation from Unified Archive
19883744	move-memory error handling needs to be cleaned up
19880888	Logical Domains Manager unable to bring resources to initial power states on SPARC_T5-2
19723743	Dynamic device specific property support
19679858	ldm ls-constraints lose the configured hypervisor mpgroup function
19430230	Non-existent resources should be ignored during evacuation
19365876	IO Resource Group information is missing
18714584	Support for ovmt_utilities
18402532	Support SPARC M7 and SPARC T7 coprocessor interrupt bindings
18168555	Evacuate and restore support for PCIe buses and SR-IOV cards
18043073	set-core returns 0, but the domain is not whole-core constrained
17834232	DRM support for whole-core constrained domains
17563489	mem_add_in_sequencer() duplicates code in mem_add()
16713362	Recovery Mode should support PCI slot removal in non-primary root domains
15811513	ldm set-vcpu -c syntax should be deprecated
15694181	Remove obsolete code and deprecate max-ipc