



# Logical Domains (LDoms) 1.0.2 Release Notes

---

Sun Microsystems, Inc.  
[www.sun.com](http://www.sun.com)

Part No. 820-3599-10  
February 2008, Revision A

Submit comments about this document at: <http://www.sun.com/hwdocs/feedback>

Copyright 2008 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.sun.com/patents> and one or more additional patents or pending patent applications in the U.S. and in other countries.

U.S. Government Rights - Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, Java, Solaris, Solaris Security Toolkit, JumpStart, OpenBoot, Sun Fire, Netra, Sun Blade, SunSolve, Sun BluePrints, Sun Management Center, Sun Explorer, SPARC, UltraSPARC, and SunVTS are service marks, trademarks, or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon architecture developed by Sun Microsystems, Inc.

The Adobe PostScript logo is a trademark of Adobe Systems, Incorporated.

Products covered by and information contained in this service manual are controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical biological weapons or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals lists is strictly prohibited.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

---

Copyright 2008 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, Californie 95054, États-Unis. Tous droits réservés.

Sun Microsystems, Inc. détient les droits de propriété intellectuels relatifs à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et ce sans limitation, ces droits de propriété intellectuelle peuvent inclure un ou plus des brevets américains listés à l'adresse <http://www.sun.com/patents> et un ou les brevets supplémentaires ou les applications de brevet en attente aux Etats - Unis et dans les autres pays.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, Java, Solaris, JumpStart, OpenBoot, Sun Fire, SunSolve, et SunVTS sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux États-Unis et dans d'autres pays.

Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

Le logo Adobe PostScript est une marque déposée de Adobe Systems, Incorporated.

Les produits qui font l'objet de ce manuel d'entretien et les informations qu'il contient sont regis par la législation américaine en matière de contrôle des exportations et peuvent être soumis au droit d'autres pays dans le domaine des exportations et importations. Les utilisations finales, ou utilisateurs finaux, pour des armes nucléaires, des missiles, des armes biologiques et chimiques ou du nucléaire maritime, directement ou indirectement, sont strictement interdites. Les exportations ou réexportations vers des pays sous embargo des États-Unis, ou vers des entités figurant sur les listes d'exclusion d'exportation américaines, y compris, mais de manière non exclusive, la liste de personnes qui font objet d'un ordre de ne pas participer, d'une façon directe ou indirecte, aux exportations des produits ou des services qui sont regis par la législation américaine en matière de contrôle des exportations et la liste de ressortissants spécifiquement désignés, sont rigoureusement interdites.

LA DOCUMENTATION EST FOURNIE "EN L'ÉTAT" ET TOUTES AUTRES CONDITIONS, DÉCLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES DANS LA LIMITE DE LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE À LA QUALITÉ MARCHANDE, À L'APTITUDE À UNE UTILISATION PARTICULIÈRE OU À L'ABSENCE DE CONTREFAÇON.



Adobe PostScript

# Contents

---

|  |          |
|--|----------|
| <b>Logical Domains (LDoms) 1.0.2 Release Notes</b>                                   | <b>1</b> |
| Changes for This Release   | 1        |
| Supported Platforms  | 2        |
| Required, Recommended, and Optional Software and Required Patches                    | 2        |
| Required and Recommended Software  | 3        |
| Domains Where Solaris OS Patches Are Required or Recommended To Be Installed         | 3        |
| Required Solaris 10 11/06 OS Patches   | 3        |
| Required Solaris 10 8/07 Patch   | 4        |
| Required System Firmware Patches   | 4        |
| Location of Patches  | 4        |
| Additional Recommended Software  | 5        |
| Optional Software  | 5        |
| Location of Documentation  | 5        |
| Supported Network Adapters   | 5        |
| ▼ To Determine If a Network Adapter Is GLDv3-Compliant                               | 6        |
| OpenBoot Firmware Not Available After Solaris OS Has Started If Domaining Is Enabled | 6        |
| Cards Not Supported  | 7        |
| Memory Size Requirements   | 7        |

|  |    |
|--|----|
| Logical Domain Channels (LDCs) and Logical Domains   | 8  |
| Software That Can Be Used With the Logical Domains Manager   | 10 |
| System Controller Software That Interacts With Logical Domains Software                            | 11 |
| General Notes and Issues   | 12 |
| System Controller and Service Processor Are Interchangeable Terms                                  | 12 |
| Limited Number of Logical Domain Configurations Can Be Stored on the SC                            | 12 |
| Rebooting the Control Domain When Guest Domains are Running  | 12 |
| Cleanly Shutting Down and Power Cycling a Logical Domains System                                   | 13 |
| ▼ To Power Off a System With Multiple Active Domains   | 13 |
| ▼ To Power Cycle the System  | 13 |
| Memory Size Requested Might Be Different From Memory Allocated                                     | 13 |
| Dynamic Reconfiguration of Virtual CPUs With Cryptographic Units                                   | 14 |
| Split PCI Regresses in FMA Functionality From Non-Logical Domains Systems                          | 14 |
| Logical Domain Variable Persistence  | 14 |
| Bugs Affecting Logical Domains 1.0.2 Software  | 16 |
| Format Oddities and Core Dump When Using ZVOL (Bug ID 6363316)                                     | 16 |
| Logical Domains Manager Can Erroneously Assign an Offline CPU to a Logical Domain (Bug ID 6431107) | 16 |
| Currently Cannot Install Guest Domains From DVDs (Bug ID 6434615)                                  | 17 |
| Some <code>format(1M)</code> Command Options Do Not Work With Virtual Disks (Bug ID 6437722)       | 17 |
| Virtual Disk Should Support Multihost Disk Control Operations (Bug ID 6437772)                     | 17 |
| Logical Domains Manager Should Validate Disk Paths and Network Devices (Bug ID 6447740)            | 17 |
| Network Devices  | 19 |
| One CPU Might Show 100% Utilization Under Heavy Network Loads (Bug ID 6492023)                     | 19 |

|   |    |
|---|----|
| Hang Can Occur With Guest OS in Simultaneous Operations (Bug ID 6497796)  | 19 |
| Panic Message on the Control Domain When Syncing a Guest Domain (Bug ID 6501168)  | 20 |
| The <code>iostat(1M)</code> Command Does Not Provide Meaningful I/O Statistics on Domain With Virtual Disks (Bug ID 6503157)              | 20 |
| Behavior of the <code>ldm stop-domain</code> Command Needs to Be Improved in Some Cases (Bug ID 6506494)                                  | 20 |
| Cannot Set Security Keys With Logical Domains Running (Bug ID 6510214)  | 21 |
| The <code>vntsd(1M)</code> Command Needs to Validate the <code>listen-to</code> IP Address (Bug ID 6512526)                               | 21 |
| Virtual Disk Server Should Export Volumes as Full Disks (Bug ID 6514091)  | 22 |
| The <code>add-vnet</code> Subcommand Allows a Virtual Network Device With the Same MAC Address as Another Logical Domain (Bug ID 6515615) | 23 |
| The <code>intrstat</code> Tool Is Not Supported for LDOMs VIO Interrupts (Bug ID 6543601)   | 23 |
| Certain <code>ldm</code> Subcommands Return Misleading Messages If One or More Arguments Are Missing (Bug IDs 6519049 and 6573220)        | 23 |
| Logical Domains Manager Forgets Variable Changes After a Power Cycle (Bug ID 6520041)   | 24 |
| Disks Managed by Veritas DMP Cannot Be Exported to Other Domains (Bug ID 6522993)   | 24 |
| CPU DR Is Disabled For All Logical Domains That Have MAU Cryptographic Units Bound (Bug ID 6525647)                                       | 25 |
| Disk Recovery Fails in a Service Domain When the Disk Device Is Actively Used as a Virtual Disk (Bug ID 6528156)                          | 25 |
| When Running Cluster Software, Selecting the <code>ok</code> Prompt Upon a Logical Domain Shutdown Can Cause a Panic (Bug ID 6528556)     | 26 |
| ▼ To Force the Primary Domain to Stop at the <code>ok</code> Prompt   | 26 |
| ▼ To Force All Other Domains to Stop at the <code>ok</code> Prompt  | 27 |
| ZFS Volumes Need to Have the Same Version of Solaris Software Running on the Service Domain and the Guest Domain (Bug ID 6528974)         | 28 |

- Page Retirement Does Not Persist in the Logical Domains Environment (Bug IDs 6531030 and 6531058) 28
- Aggregated Network Devices Are Not Supported by the Virtual Switch (Bug ID 6531266) 29
- Fault Manager Daemon (`fmd`) Does Not Correctly Recover From a Logical Domain Channel Reset (Bug ID 6533308) 29
- Using the `server-secure.driver` With an NIS Enabled System, Whether or Not LDoms Is Enabled (Bug ID 6533696) 30
- Network Performance Is Substantially Worse in a Logical Domain Guest Than in a Non-LDoms Configuration (Bug ID 6534438) 31
- Logical Domain Time-of-Day Changes Do Not Persist Across a Power Cycle of the Host (Bug ID 6536572) 32
- When the Logical Domains Manager is Running, OpenBoot PROM Variables Cannot be Modified by the `eeeprom(1M)` Command (Bug ID 6540368) 32
- Errors to Buses in a Split-PCI Configuration Might Not Get Logged (Bug ID 6542295) 32
- Emulex-based Fibre Channel Host Adapters Not Supported in Split-PCI Configuration on Sun Fire T1000 Servers (Bug ID 6544004) 33
- Starting and Stopping SunVTS Multiple Times Can Cause Host Console to Become Unusable (Bug ID 6549382) 33
- Infiniband PCI-X Card Lost if Primary Domain is Rebooted (Bug ID 6551861, 6552598, 6563713, 6589192 and 6625687) 33
- Do Not Use Control-C From the `prtdiag(1M)` Command (Bug ID 6552999 and 6630571) 34
- Do Not Specify a Virtual Switch Interface as a Network Device (Bug ID 6554177) 34
- Virtual Network Interface Does Not Handle Memory Allocation Failure Properly (Bug ID 6556778) 34
- Virtual Disks Created From Files Do Not Support EFI Labels (Bug ID 6558966) 35
- Occasionally During a Solaris OS Reboot, You Can Receive an Unsuccessful Read or Write Message (Bug ID 6560890) 35
- The `prtdiag` and `prtpicl` Utilities Do Not Work in Guest Domains (Bug ID 6563508) 35
- ▼ To Disable the PRI Pseudo Driver 36

|   |    |
|---|----|
| Under Certain Circumstances, Reverting to a Configuration Previously Saved on the SC Can Cause the Logical Domains Manager to Crash (Bug ID 6571091)            | 36 |
| Virtual Disk Service Should Support Unformatted Disks (Bug ID 6575050)  | 37 |
| Inconsistent Console Behavior When Not Using the Virtual Console (Bug ID 6581309)   | 37 |
| Guest Domain Can Lose Connection to the Virtual Switch When the Service Domain is Rebooted (Bug ID 6581720)   | 38 |
| SP <code>setdate</code> Command Changes the Solaris OS Date When Using Non-Default Logical Domain Configurations (Bug ID 6586305)                               | 39 |
| Do Not Use Non-Numeric Values for Port Numbers in <code>ldm</code> Commands (Bug ID 6590848)  | 39 |
| Virtual Disk Server Prints File Lookup Error During Service Domain Boot (Bug ID 6591399)  | 39 |
| Logical Domains Manager Does Not Retire Resources On Guest Domain After a Panic and Reboot (Bug ID 6591844)   | 40 |
| Logical Domains Manager Accepts Multiple Entries of I/O Constraints (Bug ID 6591905)  | 40 |
| Console Ports Dynamically Assigned When a Logical Domain Is Bound Become Hard Constraints If the Logical Domains Manager Subsequently Restarts (Bug ID 6596652) | 41 |
| XVR-200 Graphics Adapter On <code>pci@7c0</code> Leaf Causes PCI-E Panic and Termination (Bug ID 6598882)   | 41 |
| Guest Domain With More Than Four Virtual Networks on the Same Network Using DHCP Can Become Unresponsive (Bug ID 6603974)                                       | 42 |
| The <code>cmd</code> Core Dumps On a Hardened, Single Strand Control Domain (Bug ID 6604253)  | 42 |
| The <code>set-vcpu</code> Command Does Not Warn About MAUs Becoming Unreachable On Initial Reconfiguration (Bug ID 6609051)                                     | 42 |
| The <code>sun4v</code> Channel Nexus Generates Invalid Interrupt Cookies (Bug ID 6616313)   | 43 |
| Sun SPARC Enterprise T5120 Server Panics on Reboot (Bug ID 6620322)   | 43 |
| Attempting to <code>wanboot</code> a Logical Domain Using a Solaris 10 8/07 OS Installation DVD Causes Hang (Bug ID 6624950)                                    | 43 |

|   |    |
|---|----|
| LDoms Disk and Network Services Renamed After 1.0.1 to 1.0.2 Upgrade (Bug ID 6626770)                                       | 44 |
| Logical Domains Manager Gets Fatal Error On Startup After File System Full Error (Bug ID 6628063)                           | 44 |
| The <code>vldc_chpoll</code> Command Returns <code>ENOTACTIVE</code> After an SC Reset (Bug ID 6629230)                     | 45 |
| The <code>vxdump</code> Utility Does Not Correctly Export Disks (Bug ID 6637560)  | 45 |
| Configuring 17 Virtual Networks on a Virtual Switch Causes Guest Domains Boot Panic (Bug ID 6639934)                        | 45 |
| Addition of Virtual Disk or Network Devices Under Delayed Reconfiguration Can Fail (Bug ID 6646690)                         | 46 |
| Core Dump at <code>restore_bindings</code> (Bug ID 6648291)   | 46 |
| The <code>set-vcpu</code> Command Used in a Guest Domain During Delayed Reconfiguration Causes Termination (Bug ID 6651993) | 46 |
| The <code>ldm panic-domain</code> Command Returns Incorrect Error Message (Bug ID 6655083)                                  | 47 |
| LDoms Multidomain Functionality Does Not Support SNMP 1.5.4 on Some Systems (Bug ID 6655981)                                | 47 |
| Simultaneous Net Installation of Multiple Domains Fails When In A Common Console Group (Bug ID 6656033)                     | 47 |
| Multiple Memory Configuration Changes Under Delayed Reconfiguration Causes Hypervisor Termination (Bug ID 6657785)          | 48 |
| Hypervisor Does Not Generate <code>ereport</code> After Initial <code>pciexrc1</code> I/O Error (Bug ID 6657972)            | 48 |
| Guest Loses Virtual Disk Services When Using a Configuration Saved During Delayed Reconfiguration (Bug ID 6659100)          | 48 |
| Domains Stuck Initializing In OpenBoot Firmware After a <code>no-free-memory</code> Bind Failure (Bug ID 6661675)           | 49 |
| LDoms Issue That Applies Only to Solaris 10 11/06 OS  | 49 |
| Removing or Updating a Virtual Switch Can Cause Networking to Fail  | 50 |
| LDoms Bugs Fixed for Solaris 10 8/07 OS   | 50 |
| Networking  | 50 |
| Disk  | 51 |



|                            |    |
|----------------------------|----|
| General                    | 51 |
| Bugs Fixed for LDoms 1.0.2 | 51 |
| Fixed in Patch 127111-08   | 51 |



# Logical Domains (LDoms) 1.0.2 Release Notes

---

These release notes contain changes for this release, supported platforms, a matrix of required software and patches, and other pertinent information about this release, including bugs that affect Logical Domains 1.0.2 software.

---

## Changes for This Release

The major changes for this release of Logical Domains 1.0.2 software are to provide support for:

- Sun SPARC<sup>®</sup> Enterprise T5140 and T5240 Servers
- Up to 64 domains on Sun UltraSPARC<sup>®</sup> T2 processors
- Up to 128 domains on Sun UltraSPARC T2 Plus processors

---

# Supported Platforms

Logical Domains (LDoms) Manager 1.0.2 software is supported on the following platforms:

**TABLE 1** Supported Platforms

| <b>Name</b>                                  | <b>Reference</b>   |
|--|--|
| <b>Sun UltraSPARC T1-based Servers:</b>      |  |
| Sun Fire™ or SPARC Enterprise T1000 Server   | <i>Sun Fire or SPARC Enterprise T1000 Server Administration Guide</i>    |
| Sun Fire or SPARC Enterprise T2000 Server    | <i>Sun Fire or SPARC Enterprise T2000 Server Administration Guide</i>    |
| Netra™ T2000 Server                          | <i>Netra T2000 Server Administration Guide</i>                           |
| Netra CP3060 Blade                           | <i>Netra CP3060 Board Product Notes</i>                                  |
| Sun Blade™ T6300 Server Module               | <i>Sun Blade T6300 Server Module Administration Guide</i>                |
| <b>Sun UltraSPARC T2-based Servers:</b>      |  |
| Sun SPARC Enterprise T5120 and T5220 Servers | <i>Sun SPARC Enterprise T5120 and T5220 Servers Administration Guide</i> |
| Sun Blade T6320 Server Module                | <i>Sun Blade T6320 Server Module Product Notes</i>                       |
| Netra CP3260 Blade                           | <i>Netra CP3260 Board Product Notes</i>                                  |
| Netra T5220 Server                           | <i>Sun Netra T5220 Server Product Notes</i>                              |
| <b>Sun UltraSPARC T2 Plus-based Servers:</b> |  |
| Sun SPARC Enterprise T5140 and T5240 Servers | <i>Sun SPARC Enterprise T5140 and T5240 Servers Administration Guide</i> |

---

---

# Required, Recommended, and Optional Software and Required Patches

This section lists the required, recommended, and optional software for use with Logical Domains software.

# Required and Recommended Software

Following is a matrix of required and recommended software at a minimum for use with Logical Domains software.

**TABLE 2** Required Software Matrix

| Supported Servers                    | Logical Domains Manager | System Firmware     |
|--------------------------------------|-------------------------|---------------------|
| Sun UltraSPARC T1-based servers      | 1.0.2                   | 6.6.x recommended*  |
| Sun UltraSPARC T2-based servers      | 1.0.2                   | 7.1.x recommended\< |
| Sun UltraSPARC T2 Plus-based servers | 1.0.2                   | 7.1.x required      |

\* 6.5.x firmware minimum

\ 7.0.x firmware minimum

## Domains Where Solaris OS Patches Are Required or Recommended To Be Installed

Following are the domains on which Solaris OS patches are required or recommended to be installed.

**TABLE 3** Domains Where Solaris OS Patches Must Be Installed

| Solaris OS Version  | Patch ID  | Control     | I/O         | Service     | Guest       |
|---------------------|-----------|-------------|-------------|-------------|-------------|
| Solaris 10 OS 11/06 | 124921-02 | Required    | Required    | Required    | Required    |
|                     | 125043-01 | Required    | Required    | Required    | Required    |
|                     | 127111-08 | Recommended | Recommended | Recommended | Recommended |
| Solaris 10 OS 8/07  | 127111-08 | Required    | Recommended | Recommended | Recommended |

## Required Solaris 10 11/06 OS Patches

Following are the required patches for Solaris 10 11/06 OS for use with Logical Domains software:

- 124921-02 at a minimum, which contains updates to the Logical Domains 1.0.1 drivers and utilities. Logical Domains networking will be broken without this patch.
- 125043-01 at a minimum, which contains updates to the console (`qcn`) drivers. This patch depends on kernel update (KU) 118833-36, so if this is not already updated on your system, you must install it also.

## Required Solaris 10 8/07 Patch

Following is the required patch for Solaris 10 8/07 OS:

- 127111-08 at a minimum.

## Required System Firmware Patches

Following are the required system firmware patches at a minimum for use with Logical Domains 1.0.2 software on supported servers:

**TABLE 4** Required System Firmware Patches

| Patches   | Supported Servers                            |
|-----------|--|
| 136927-01 | Sun Fire and SPARC Enterprise T2000 Servers  |
| 136928-01 | Sun Fire and SPARC Enterprise T1000 Servers  |
| 136929-01 | Netra T2000 Server                           |
| 136930-01 | Netra CP3060 Blade                           |
| 136931-01 | Sun Blade T6300 Server Module                |
| 136932-01 | Sun SPARC Enterprise T5120 and T5220 Servers |
| 136933-01 | Sun Blade T6320 Server Module                |
| 136934-01 | Netra T5220 Server                           |
| 136935-01 | Netra CP3260 Blade                           |
| 136936-01 | Sun SPARC Enterprise T5140 and T5240 Servers |

## Location of Patches

You can find the required Solaris OS and system firmware patches at the SunSolve<sup>SM</sup> site:

<http://sunsolve.sun.com>

## Additional Recommended Software

**Solaris Security Toolkit 4.2 software** – This software can help you secure the Solaris OS in the control domain and other domains. Refer to the *Solaris Security Toolkit 4.2 Administration Guide* and *Solaris Security Toolkit 4.2 Reference Manual* for more information.

## Optional Software

- **Logical Domains (LDMs) Management Information Base (MIB) 1.0.1 software** – This software can help you enable third party applications to perform remote monitoring and a few control operations. Refer to the *Logical Domains (LDMs) MIB 1.0.1 Administration Guide* and *Release Notes* for more information.
- **Libvirt for LDMs 1.0.1 software** – This software provides virtual library (`libvirt`) interfaces for Logical Domains (LDMs) software so that virtualization customers can have consistent interfaces. The `libvirt` library (version 0.3.2) included in this software interacts with the Logical Domains Manager 1.0.1 software running on Solaris 10 Operating System (OS) to support Logical Domains virtualization technology. Refer to the *Libvirt for LDMs 1.0.1 Administration Guide* and *Release Notes* for more information.

---

## Location of Documentation

The *Logical Domains (LDMs) 1.0.2 Administration Guide* and *Logical Domains (LDMs) 1.0.2 Release Notes* can be found at:

<http://docs.sun.com/>

The *Beginners Guide to LDMs: Understanding and Deploying Logical Domains* can be found at the Sun BluePrints™ site.

<http://www.sun.com/blueprints/0207/820-0832.html>

---

## Supported Network Adapters

In a logical domains environment, the virtual switch service running in a service domain can directly interact with GLDv3-compliant network adapters. Though non-GLDv3 compliant network adapters can be used in these systems, the virtual

switch cannot interface with them directly. Refer to “Configuring Virtual Switch and Service Domain for NAT and Routing” in the *Logical Domains (LDoms) 1.0.2 Administration Guide* for information about how to use non-GLDv3 compliant network adapters.

## ▼ To Determine If a Network Adapter Is GLDv3-Compliant

1. Use the Solaris OS `dladm(1M)` command, where, for example, `bge0` is the network device name.

```
# dladm show-link bge0
bge0                type: non-vlan    mtu: 1500        device: bge0
```

2. Look at `type`: in the output:

- GLDv3-compliant drivers will have a type of `non-vlan` or `vlan`.
- Non-GLDv3-compliant drivers will have a type of `legacy`.

---

## OpenBoot Firmware Not Available After Solaris OS Has Started If Domaining Is Enabled

---

**Note** – Domaining is always enabled on all supported platforms, except Sun UltraSPARC T1-based platforms.

---

Domaining is enabled once a logical domains configuration created by the Logical Domains Manager is instantiated. If domaining is enabled, the OpenBoot™ firmware is not available after the Solaris OS has started, because it is removed from memory.

To reach the `ok` prompt from the Solaris OS, you must halt the domain. You can use the Solaris OS `halt(1M)` command to halt the domain. For more information, refer to “Result of a Solaris OS `halt(1M)` Command” in the *Logical Domains (LDoms) 1.0.2 Administration Guide*.



---

# Cards Not Supported

The following cards are *not* supported for this LDomS 1.0.2 software release:

- Sun XVR-200 Graphics Accelerator
- Sun XVR-300 Graphics Accelerator
- Sun Dual Port 4x IB Host Channel Adapter PCI-X Card
- Dual Port 4x PCI Express Infiniband Host Channel Adapter - Low Profile



---

**Caution** – If these unsupported configurations are used with LDomS 1.0.2, stop and unbind all logical domains before the control domain is rebooted. Failure to do so can result in a system crash causing the loss of all the logical domains that are active in the system.

---

The following bug IDs are filed to provide the support for the currently unsupported cards: 6552598, 6563713, 6589192, and 6598882.

---

# Memory Size Requirements

Logical Domains software does not impose a memory size limitation when creating a domain. The memory size requirement is a characteristic of the guest operating system. Some Logical Domains functionality might not work if the amount of memory present is less than the recommended size. For recommended and minimum size memory requirements, refer to the installation guide for the operating system you are using. The default size for a swap area is 512 megabytes. Refer to “System Requirements and Recommendations” in the *Solaris 10 Installation Guide: Planning for Installation and Upgrade*.

The OpenBoot™ PROM has a minimum size restriction for a domain. Currently, that restriction is 12 megabytes. If you have a domain less than that size, the Logical Domains Manager will automatically boost the size of the domain to 12 megabytes. Refer to the release notes for your system firmware for information about memory size requirements.

## Booting a Large Number of Domains

As sun4v systems with greater thread counts are released, you can have more domains per system than previous releases:

- Up to 64 in Sun SPARC Enterprise T5120 and T5220 servers

- up to 128 in Sun SPARC Enterprise T51420 and T5240 servers

If unallocated virtual CPUs are available, assign them to the service domain to help process the virtual I/O requests. Allocate 4 to 8 virtual CPUs to the service domain when creating more than 32 domains.

Since maximum domain configurations have only a single CPU in the service domain, do not put unnecessary stress on the single CPU when configuring and using the domain.

The virtual switch (`vsw`) services should be spread over all the network adapters available in the machine. For example, if booting 128 domains on a Sun SPARC Enterprise T5240 server, create 4 `vsw` services, each serving 32 virtual net (`vnet`) instances. Do not have more than 32 `vnet` instances per `vsw` service, because having more than 60 `vnet` instance tied to a single `vsw` causes hard hangs in the service domain.

To run the maximum configurations, a machine needs 64 gigabytes of memory (and up to 128 gigabytes in the Sun SPARC Enterprise T5240 server, if possible) so that the guest domains contain an adequate amount of memory. The guest domains require a minimum of 512 megabytes of memory but can benefit from having more than that, depending on the workload running in the domain and the configuration of the domain (number of virtual devices in the domain). Memory and swap space usage increases in a guest domain when the `vsw` services used by the domain provides services to many virtual networks (in multiple domains). This is due to the peer-to-peer links between all the `vnet` connected to the `vsw`.

The service domain benefits from having extra memory. Four gigabytes is the recommended minimum when running more than 64 domains. Start domains serially rather than all at once. Start domains in groups of 10 or less and wait for them to boot before starting the next batch. The same advice applies to installing domains.

## Logical Domain Channels (LDCs) and Logical Domains

There is a limit to the number of LDCs available in any logical domain. For Sun UltraSPARC T1-based platforms, that limit is 256; for all other platforms, the limit is 512. Practically speaking, this only becomes an issue on the control domain, because the control domain has at least part, if not all, of the I/O subsystem allocated to it, and because of the potentially large number of LDCs created for both virtual I/O data communications and the Logical Domains Manager control of the other logical domains.

---

**Note** – The examples in this section are what happens on Sun UltraSPARC T1-based platforms. However, the behavior is the same if you go over the limit on other supported platforms.

---

If you try to add a service, or bind a domain, so that the number of LDC channels exceeds the limit on the control domain, the operation fails with an error message similar to the following:

13 additional LDCs are required on guest primary to meet this request, but only 9 LDCs are available

The following guidelines can help prevent creating a configuration that could overflow the LDC capabilities of the control domain:

1. The control domain allocates 12 LDCs for various communication purposes with the hypervisor, Fault Management Architecture (FMA), and the system controller (SC), independent of the number of other logical domains configured.
2. The control domain allocates one LDC to every logical domain, including itself, for control traffic.
3. Each virtual I/O service on the control domain consumes one LDC for every connected client of that service.

For example, consider a control domain and 8 additional logical domains. Each logical domain needs at a minimum:

- Virtual network
- Virtual disk
- Virtual console

Applying the above guidelines yields the following results (numbers in parentheses correspond to the preceding guideline number from which the value was derived):

$12(1) + 9(2) + 8 \times 3(3) = 45$  LDCs in total.

Now consider the case where there are 32 domains instead of 8, and each domain includes 3 virtual disks, 3 virtual networks, and a virtual console. Now the equation becomes:

$12 + 33 + 32 \times 7 = 269$  LDCs in total.

Depending upon the capabilities of your platform the Logical Domain Manager will either accept or reject the configurations.

---

# Software That Can Be Used With the Logical Domains Manager

This section details the software that is compatible with and can be used with the Logical Domains software in the control domain.

- **SunVTS™ 6.4** functionality is available in the control domain and guest domains on LDoms 1.0.2–enabled systems.

**Sun VTS 6.3** functionality is available for all hardware configured in the control domain on Sun Fire and SPARC Enterprise T1000 servers and Sun Fire and SPARC Enterprise T2000 servers with LDoms 1.0 software enabled. If you attempt to execute in a guest domain, SunVTS 6.3 software exits after printing a message.

SunVTS is Sun's Validation Test Suite, which provides a comprehensive diagnostic tool that tests and validates Sun hardware by verifying the connectivity and proper functioning of most hardware controllers and devices on Sun servers. For more information about SunVTS, refer to the *SunVTS User's Guide* for your version of SunVTS.

- **Sun™ Management Center 4.0 Version 3 Add-On Software** can be used only on the control domain with the Logical Domains Manager software enabled. Sun Management Center is an open, extensible system monitoring and management solution that uses Java™ and a variant of the Simple Network Management Protocol (SNMP) to provide integrated and comprehensive enterprise-wide management of Sun products and their subsystem, component, and peripheral devices. Support for hardware monitoring within the Sun Management Center environment is achieved through the use of appropriate hardware server module add-on software, which presents hardware configuration and fault reporting information to the Sun Management Center management server and console. Refer to the *Sun Management Center 4.0 Version 3 Add-On Software Release Notes: For Sun Fire, SunBlade, Netra, and SunUltra Systems for more information about using Sun Management Center 4.0 Version 3 on the supported servers*.
- **Sun™ Explorer 5.7 Data Collector can be used with the Logical Domains Manager 1.0.2** software enabled on the control domain. Sun Explorer is a diagnostic data collection tool. The tool comprises shell scripts and a few binary executables. Refer to the *Sun Explorer User's Guide* for more information about using the Sun Explorer Data Collector.
- **Solaris™ Cluster** software can be used only on an I/O domain, because it works only with the physical hardware, not the virtualized hardware. Refer to Sun Cluster documentation for more information about the Sun Cluster software.

---

# System Controller Software That Interacts With Logical Domains Software

The following system controller (SC) software interacts with the Logical Domains 1.0.2 software:

- **Sun Integrated Lights Out Manager (ILOM) 2.0** firmware is the system management firmware you can use to monitor, manage, and configure Sun UltraSPARC T2-based server platforms. ILOM is preinstalled on these platforms and can be used on the control domain on LDom-supported servers with the Logical Domains Manager 1.0.2 software enabled. Refer to the *Sun Integrated Lights Out Manager 2.0 User's Guide* for features and tasks that are common to Sun rackmounted servers or blade servers that support ILOM. Other user documents present ILOM features and tasks that are specific to the server platform you are using. You can find the ILOM platform-specific information within the documentation set that accompanies your system.
- **Advanced Lights Out Manager (ALOM) Chip Multithreading (CMT) Version 1.3** software can be used on the control domain on UltraSPARC® T1-based servers with the Logical Domains Manager 1.0.2 software enabled. Refer to "Using LDom With ALOM CMT" in the *Logical Domains (LDoms) 1.0.2 Administration Guide*. The ALOM system controller enables you to remotely manage and administer your supported CMT servers. ALOM enables you to monitor and control your server either over a network or by using a dedicated serial port for connection to a terminal or terminal server. ALOM provides a command-line interface that you can use to remotely administer geographically distributed or physically inaccessible machines. For more information about using ALOM CMT Version 1.3 software, refer to the *Advanced Lights Out Management (ALOM) CMT v1.3 Guide*.
- **Netra Data Plane Software Suite 1.1** is a complete board software package solution. The software provides an optimized rapid development and runtime environment on top of multistrand partitioning firmware for Sun CMT platforms. The Logical Domains Manager contains some `ldm` subcommands (`add-vdpcs`, `rm-vdpcs`, `add-vdpcc`, and `rm-vdpcc`) for use with this software. Refer to the *Netra Data Plane Software Suite 1.1* documentation for more information about this software.

---

## General Notes and Issues

This section contains general notes and issues concerning the Logical Domains 1.0.2 software.

### System Controller and Service Processor Are Interchangeable Terms

For discussions in Logical Domains documentation, the terms system controller (SC) and service processor (SP) are interchangeable.

### Limited Number of Logical Domain Configurations Can Be Stored on the SC

Currently, there is a limit of 8 configurations for logical domains that can be saved on the system controller using the `ldm add-config` command, not including the `factory-default` configuration.

### Rebooting the Control Domain When Guest Domains are Running

When you reboot the control domain when guest domains are running, you will encounter the following bugs:

- [“Guest Domain Can Lose Connection to the Virtual Switch When the Service Domain is Rebooted \(Bug ID 6581720\)”](#) on page 38
- [“Virtual Disk Server Prints File Lookup Error During Service Domain Boot \(Bug ID 6591399\)”](#) on page 39

# Cleanly Shutting Down and Power Cycling a Logical Domains System

If you have made any configuration changes since last saving a configuration to the SC, before you attempt to power off or power cycle a Logical Domains system, make sure you save the latest configuration that you want to keep.

## ▼ To Power Off a System With Multiple Active Domains

1. **Shut down and unbind all the non-I/O domains.**
2. **Shut down and unbind any active I/O domains.**
3. **Halt the primary domain.**

Because no other domains are bound, the firmware automatically powers off the system.

## ▼ To Power Cycle the System

1. **Shut down and unbind all the non-I/O domains.**
2. **Shut down and unbind any active I/O domains.**
3. **Reboot the primary domain.**

Because no other domains are bound, the firmware automatically power cycles the system before rebooting it. When the system restarts, it boots into the Logical Domains configuration last saved or explicitly set.

# Memory Size Requested Might Be Different From Memory Allocated

Under certain circumstances, the Logical Domains (LDom)s Manager rounds up the requested memory allocation to either the next largest 8-kilobyte or 4-megabyte multiple. This can be seen in the following example output of the `ldm list-domain -l` command, where the constraint value is smaller than the actual allocated size:

|         |              |             |       |
|---------|--------------|-------------|-------|
| Memory: |              |             |       |
|         | Constraints: | 1965 M      |       |
|         | raddr        | paddr5      | size  |
|         | 0x1000000    | 0x291000000 | 1968M |

## Dynamic Reconfiguration of Virtual CPUs With Cryptographic Units

Currently, there is an issue related to dynamic reconfiguration (DR) of virtual CPUs if a logical domain contains one or more cryptographic (mau) units:

- DR of virtual CPUs is completely disabled on all active logical domains that contain any cryptographic units (Bug ID 6525647).

## Split PCI Regresses in FMA Functionality From Non-Logical Domains Systems

Currently, Fault Management Architecture (FMA) diagnosis of I/O devices in a Logical Domains environment might not work correctly. The problems are:

- Input/output (I/O) device faults diagnosed in a non-control domain are not logged on the control domain. These faults are only visible in the logical domain that owns the I/O device.
- I/O device faults diagnosed in a non-control domain are not forwarded to the system controller. As a result, these faults are not logged on the SC and there are no fault actions on the SC, such as lighting of light-emitting diodes (LEDs) or updating the dynamic field-replaceable unit identifiers (DFRUIDs).
- Errors associated with a root complex that is not owned by the control domain are not diagnosed properly. These errors can cause faults to be generated against the diagnosis engine (DE) itself.

## Logical Domain Variable Persistence

With domaining enabled, variable updates persist across a reboot, but not across a power cycle, unless the variable updates are either initiated from OpenBoot firmware on the control domain, or followed by saving the configuration to the SC.

In this context, it's important to note that a reboot of the control domain could initiate a power cycle of the system:

- When the control domain reboots, if there are no bound guest domains, and no delayed reconfiguration in progress, the SC powercycles the system.
- When the control domain reboots, if there are guest domains bound or active (or the control domain is in the middle of a delayed reconfiguration), the SC does not powercycle the system.

LDom variables for a domain can be specified using any of the following methods:



- At the OpenBoot prompt
- Using the Solaris OS `eeeprom(1M)` command
- Using the Logical Domains Manager CLI (`ldm`)
- Modifying, in a limited fashion, from the system controller (SC) using the `bootmode` command; that is, only certain variables, and only when in the `factory-default` configuration.

The goal is that, variable updates made using any of these methods always persist across reboots of the domain, and always reflect in any subsequent logical domain configurations saved to the SC.

In Logical Domains 1.0.2 software, there are a few cases where variable updates do not persist as expected:

- With domaining enabled (the default in all cases except the UltraSPARC T1000 & T2000 systems running in `factory-default` configuration), all methods of updating a variable (OpenBoot firmware, `eeeprom` command, `ldm` subcommand) persist across reboots of that domain, but not across a power cycle of the system, unless a subsequent logical domain configuration is saved to the SC. In addition, in the control domain, updates made using OpenBoot firmware persist across a power cycle of the system; that is, even without subsequently saving a new logical domain configuration to the SC.
- When domaining is not enabled, variable updates specified through the Solaris OS `eeeprom(1M)` command persist across a reboot of the primary domain into the same `factory-default` configuration, but do not persist into a configuration saved to the SC. Conversely, in this scenario, variable updates specified using the Logical Domains Manager do not persist across reboots, but are reflected in a configuration saved to the SC.

So, when domaining is not enabled, if you want a variable update to persist across a reboot into the same `factory-default` configuration, use the `eeeprom` command. If you want it saved as part of a new logical domains configuration saved to the SC, use the appropriate Logical Domains Manager command.

- In all cases, when reverting to the `factory-default` configuration from a configuration generated by the Logical Domains Manager, all LDoms variables start with their default values.

The following bug IDs have been filed to resolve these issues: 6520041, 6540368, and 6540937.

---

## Bugs Affecting Logical Domains 1.0.2 Software

This section summarizes the bugs that you might encounter when using this version of the software. The bug descriptions are in numerical order by bug ID. If a recovery procedure and a workaround are available, they are specified.

### Format Oddities and Core Dump When Using ZVOL (Bug ID 6363316)

Format oddities and a core dump occur when using the ZFS Volume Emulation Driver (ZVOL) and when the Logical Domains environment has virtual disks with an Extensible Firmware Interface (EFI) label. Selecting such disks with the `format(1M)` command cause a core dump.

### Logical Domains Manager Can Erroneously Assign an Offline CPU to a Logical Domain (Bug ID 6431107)

When the Fault Management Architecture (FMA) places a CPU offline, it records that information, so that when the machine is rebooted the CPU remains offline. The offline designation persists in a non-Logical Domains environment.

However, in a Logical Domains environment, this persistence is not always maintained for CPUs in guest domains. The Logical Domains Manager does not currently record data on fault events sent to it. This means that a CPU in a guest domain that has been marked as faulty, or one that was not allocated to a logical domain at the time the fault event is replayed, can subsequently be allocated to another logical domain with the result that it is put back online.

## Currently Cannot Install Guest Domains From DVDs (Bug ID 6434615)

The Solaris 10 OS virtual disk drivers (vdc and vds) currently do not support the CDIO(7I) ioctls that are needed to install guest domains from DVDs. Therefore, it is not possible at this time to install a guest domain from a DVD. However, a guest domain can access a CD/DVD to install applications. If the CD/DVD device is added to the guest domain, and the guest is booted from another virtual disk, the CD/DVD can be mounted in the guest domain after the boot operation.

## Some format(1M) Command Options Do Not Work With Virtual Disks (Bug ID 6437722)

Refer to “Operating the Solaris OS With Logical Domains” in Chapter 5 of the *Logical Domains (LDoms) 1.0.2 Administration Guide* for specific information.

## Virtual Disk Should Support Multihost Disk Control Operations (Bug ID 6437772)

The Solaris OS virtual disk drivers (vdc and vds) currently do not support multihost disk control operations (MHI(7I) ioctls).

## Logical Domains Manager Should Validate Disk Paths and Network Devices (Bug ID 6447740)

If a disk device listed in a guest domain’s configuration is either non-existent, already opened by another process, or otherwise unusable, the disk cannot be used by the virtual disk server (vds) but the Logical Domains Manager does not emit any warning or error when the domain is bound or started.

When the guest tries to boot, messages similar to the following are printed on the guest’s console:

```
WARNING: /virtual-devices@100/channel-devices@200/disk@0: Timeout
connecting to virtual disk server... retrying
```

In addition, if a network interface specified using the `net-dev=` parameter does not exist or is otherwise unusable, the virtual switch is unable to communicate outside the physical machine, but the Logical Domains Manager does not emit any warning or error when the domain is bound or started.

*Recovery:*

In the case of an **errant virtual disk service device or volume**, perform the following steps:

1. Stop the domain owning the virtual disk bound to the errant device or volume.
2. Issue the `ldm rm-vdsdev` command to remove the errant virtual disk service device.
3. Issue the `ldm add-vdsdev` command to correct the physical path to the volume.
4. Restart the domain owning the virtual disk.

In the case of an **errant `net-dev=` property specified for a virtual switch**, perform the following steps:

1. Issue the `ldm set-vsw` command with the corrected `net-dev=` property.
2. Reboot the domain hosting the virtual switch in question.

## Network Devices

If a disk device listed in a guest domain's configuration is being used by software other than the Logical Domains Manager (for example, if it is mounted in the service domain), the disk cannot be used by the virtual disk server (`vdss`), but the Logical Domains Manager does not emit a warning that it is in use when the domain is bound or started.

When the guest domain tries to boot, a message similar to the following is printed on the guest's console:

```
WARNING: /virtual-devices@100/channel-devices@200/disk@0: Timeout
connecting to virtual disk server... retrying
```

*Recovery:* Unbind the guest domain, and unmount the disk device to make it available. Then bind the guest domain, and boot the domain.

## One CPU Might Show 100% Utilization Under Heavy Network Loads (Bug ID 6492023)

Under heavy network loads, a single CPU in the service domain might show 100% utilization dealing with the network traffic. (This would show in the `sys` column from `mpstat`.)

*Workaround:* Attach at least two, and preferably four, CPUs to the service domain containing the virtual switch to ensure that the system remains responsive under a heavy load, or reduce the load on the system.

## Hang Can Occur With Guest OS in Simultaneous Operations (Bug ID 6497796)

Under rare circumstances, when an `ldom` variable, such as `boot-device`, is being updated from within a guest domain by using the `eeeprom(1M)` command at the same time that the Logical Domains Manager is being used to add or remove virtual CPUs from the same domain, the guest OS can hang.

*Workaround:* Ensure that these two operations are not performed simultaneously.

*Recovery:* Use the `ldm stop-domain` and `ldm start-domain` commands to stop and start the guest OS.

## Panic Message on the Control Domain When Syncing a Guest Domain (Bug ID 6501168)

If too many guest domains are performing I/O to a control or I/O domain, and if that domain is in the middle of panicking, the interrupt request pool of 64 entries overflows and the system cannot save a crash dump. The panic message is as follows:

```
intr_req pool empty
```

*Workaround:* None.

## The `iostat(1M)` Command Does Not Provide Meaningful I/O Statistics on Domain With Virtual Disks (Bug ID 6503157)

The `iostat(1M)` command does not return any meaningful information when run on a domain with virtual disks. This is because the LDom's `vDisk` client driver (`vdc`) does not measure I/O activity nor save any info to `kstats` which could be read by the `iostat` command.

*Workaround:* Gather the I/O statistics on the service domain exporting the virtual disks.

## Behavior of the `ldm stop-domain` Command Needs to Be Improved in Some Cases (Bug ID 6506494)

There are some cases where the behavior of the `ldm stop-domain` command is confusing.

If the Solaris OS is halted on the domain; for example, by using the `halt(1M)` command; and the domain is at the prompt `"r)eboot, o)k prompt, h)alt?,"` the `ldm stop-domain` command fails with the following error message:

```
LDom <domain name> stop notification failed
```

*Workaround:* Force a stop by using the `ldm stop-domain` command with the `-f` option.

```
# ldm stop-domain -f ldom
```

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

```
LDom <domain name> stop notification failed
```

*Recovery:* If you restart the domain from the `kmdb` prompt, the stop notification is handled, and the domain does stop.

## Cannot Set Security Keys With Logical Domains Running (Bug ID 6510214)

In a Logical Domains environment, there is no support for setting or deleting wide-area network (WAN) boot keys from within the Solaris OS 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 only valid for a single use.

## The `vntsd(1M)` Command Needs to Validate the `listen-to` IP Address (Bug ID 6512526)

The Solaris 10 OS `vntsd(1M)` command does not validate the `listen_addr` property in the `vntsd` command's Service Management Facility (SMF) manifest. If the `listen_addr` property is invalid, `vntsd` fails to bind the IP address and exits.

*Recovery:*

1. Update the SMF `listen_addr` property with the correct IP address.

2. Refresh vntsd.

```
# svcadm refresh vntsd
```

3. Restart vntsd.

```
# svcadm restart vntsd
```

## Virtual Disk Server Should Export Volumes as Full Disks (Bug ID 6514091)

When a ZFS, SVM, or VxVM volume is exported as a virtual disk to another domain, then the other domain sees that virtual disk as a disk with a single slice (s0), and the disk cannot be partitioned. As a consequence, such a disk is not usable by the Solaris installer, and you cannot install Solaris on the disk.

For example, `/dev/zvol/dsk/tank/zvol` is a ZFS volume that is exported as a virtual disk from the `primary` domain to `domain1` using these commands:

```
# ldm add-vdsdev /dev/zvol/dsk/tank/zvol disk_zvol@primary-vds0
# ldm add-vdisk vdisk0 disk_zvol@primary_vds0 domain1
```

The `domain1` sees only one device for that disk (for example, `c0d0s0`), and there is no other slice for that disk; for example, no device `c0d0s1`, `c0d0s2`, `c0d0s3`....

*Workaround:* You can create a file and export that file as a virtual disk. This example creates a file on a ZFS system:

```
# mkfile 30g /tank/test/zfile
# ldm add-vdsdev /tank/test/zfile disk_zfile@primary-vds0
# ldm add-vdisk vdisk0 disk_zfile@primary-vds0 domain1
```

---

**Note** – When exporting a ZFS, SVM, or VxVM volume as a virtual disk, be aware that you will have to change your configuration once this bug is fixed. The instructions for changing the configuration will be provided.

---



## The `add-vnet` Subcommand Allows a Virtual Network Device With the Same MAC Address as Another Logical Domain (Bug ID 6515615)

When creating logical domains with virtual switches and virtual network devices, the Logical Domains Manager does not prevent you from creating these devices with the same given MAC address. This can become a problem if the logical domains with virtual switches and virtual networks that have conflicting MAC addresses are in a bound state simultaneously.

*Workaround:* Ensure that you do not bind logical domains whose `vsw` and `vnet` MAC addresses might conflict with another `vsw` or `vnet` MAC address.

## The `intrstat` Tool Is Not Supported for LDOMs VIO Interrupts (Bug ID 6543601)

The `intrstat` tool is not supported for LDOMs virtual I/O (VIO) interrupts. Without the `intrstat` tool, you cannot monitor interrupts targeted at virtual devices; that is, virtual disk client and server, virtual switch, virtual network device and virtual console. This does not impact normal operations.

## Certain `ldm` Subcommands Return Misleading Messages If One or More Arguments Are Missing (Bug IDs 6519049 and 6573220)

Misleading error messages are returned from certain `ldm` subcommands that take two or more required arguments, if one or more of those required arguments is missing.

For example, if the `add-vsw` subcommand is missing the `vswitch-name` or `ldom` argument, you receive an error message similar to the following:

```
# ldm add-vsw net-dev=e1000g0 primary
Illegal name for service: net-dev=e1000g0
```

For another example, if the `add-vnet` command is missing the *vswitch-name* of the virtual switch service with which to connect, you receive an error message similar to the following:

```
# ldm add-vnet mac-addr=08:00:20:ab:32:40 vnet1 ldg1
Illegal name for VNET interface: mac-addr=08:00:20:ab:32:40
```

As another example, if you fail to add a logical domain name at the end of an `ldm add-vcc` command, you receive an error message saying that the `port-range=` property must be specified.

*Recovery:* Refer to the *Logical Domains (LDoms) Manager 1.0.1 Man Page Guide* or the `ldm` man page for the required arguments of the `ldm` subcommands, and retry the commands with the correct arguments.

## Logical Domains Manager Forgets Variable Changes After a Power Cycle (Bug ID 6520041)

This issue is summarized in [“Logical Domain Variable Persistence”](#) on page 14.

## Disks Managed by Veritas DMP Cannot Be Exported to Other Domains (Bug ID 6522993)

In a service domain, disks that are managed by Veritas Dynamic Multipathing (DMP) cannot be exported as virtual disks to other domains. If a disk that is managed by Veritas DMP is added to a virtual disk server (`vds`) and then added as a virtual disk to a guest domain, the domain is unable to access and use that virtual disk. In such a case, the service domain reports the following errors in the `/var/adm/messages` file after binding the guest domain:

```
vd_setup_vd(): ldi_open_by_name(/dev/dsk/c4t12d0s2) = errno 16
vds_add_vd(): Failed to add vdisk ID 0
```

*Recovery:* If Veritas Volume Manager (VxVM) is installed on your system, you can either disable Veritas DMP for the disks you want to use as virtual disks or disable the exclusive open done by the `vds` driver.

You can disable the exclusive open done by the `vds` driver by setting the kernel global variable `vd_open_flags` to `“0x3”`.

You can disable the exclusive open on the running system with the following command:

```
# echo 'vd_open_flags/W 0x3' | mdb ?kw
```

You also need to add the change in `/etc/system` to make it persistent across reboots:

```
# set vds:vd_open_flags = 0x3
```

## CPU DR Is Disabled For All Logical Domains That Have MAU Cryptographic Units Bound (Bug ID 6525647)

Due to problems with the Solaris Crypto Framework and its handling of CPU dynamic reconfiguration (DR) events that affect MAU cryptographic units, CPU DR is disabled for all logical domains that have any crypto units bound to it.

*Workaround:* To be able to use CPU DR on the control domain, all the crypto units must be removed from it while the system is running in the `factory-default` configuration, before saving a new configuration to the SC. To perform CPU DR on all other domains, stop the domain first so it is in the bound state.

## Disk Recovery Fails in a Service Domain When the Disk Device Is Actively Used as a Virtual Disk (Bug ID 6528156)

The virtual disk server opens the physical disk exported as a virtual disk device at the time of the bind operation. In certain cases, a recovery operation on the physical disk following a disk failure may not be possible if the guest domain is bound.

For instance, when a RAID or a mirror Solaris™ Volume Manager (SVM) volume is used as a virtual disk by another domain, and if there is a failure on one of the components of the SVM volume, then the recovery of the SVM volume using the `metareplace` command or using a hot spare does not start. The `metastat` command shows the volume as resynchronizing, but there is no progress in the synchronization.

Similarly, when a Fibre Channel Arbitrated Loop (FC\_AL) device is used as a virtual disk, you must use the Solaris OS `luxadm(1M)` command with a loop initialization primitive sequence (`forceclip` subcommand) to reinitialize the physical disk after unbinding the guest.

---

**Note** – Recovery mechanisms may fail in a similar manner for other devices, if the mechanism requires that the device being recovered is not actively in use.

---

*Recovery:* To complete the recovery or SVM resynchronization, stop and unbind the domain using the SVM volume as a virtual disk. Then resynchronize the SVM volume using the `metasync` command.

## When Running Cluster Software, Selecting the `ok` Prompt Upon a Logical Domain Shutdown Can Cause a Panic (Bug ID 6528556)

If Solaris™ Cluster software is in use with Logical Domains software, and the cluster is shut down, the console of each logical domain in the cluster displays the following prompt:

```
r) reboot, o) k prompt, h) alt?
```

If the `ok` prompt (`o` option) is selected, the system can panic.

*Workarounds:*

- Select `halt` (`h` option) at the prompt on the logical domain console to avoid the panic.
- To force the logical domain to stop at the `ok` prompt, even if the `OpenBoot auto-boot?` variable is set to `true`, follow one of the two following procedures.

### ▼ To Force the Primary Domain to Stop at the `ok` Prompt

Use this procedure *only* for the `primary` domain.

1. Issue the following ALOM command to reset the domain:

```
sc> poweron
```

The OpenBoot banner is displayed on the console:

```
Sun Fire T200, No Keyboard  
Copyright 2007 Sun Microsystems, Inc. All rights reserved.  
OpenBoot 4.26.0, 4096 MB memory available, Serial #68100096.  
Ethernet address 0:14:4f:f:20:0, Host ID: 840f2000.
```

2. Issue the following ALOM command to send a break to the domain immediately after the OpenBoot banner displays:

```
sc> break -y
```

The logical domain immediately drops to the ok prompt.

## ▼ To Force All Other Domains to Stop at the ok Prompt

Use this procedure for all logical domains *except* the primary domain.

1. Issue the following command from the control domain to disable the `auto-boot?` variable for the logical domain:

```
# ldm set-var auto-boot?=false domain-name
```

2. Issue the following command from the control domain to reset the logical domain:

```
# ldm start-domain domain-name
```

The logical domain stops at the ok prompt.

3. Issue the following OpenBoot command to restore the value of the `auto-boot?` variable:

```
ok setenv auto-boot? true
```

## ZFS Volumes Need to Have the Same Version of Solaris Software Running on the Service Domain and the Guest Domain (Bug ID 6528974)

If a guest domain is running the Solaris 10 OS and using a virtual disk built from a ZFS volume provided by a service domain running the Solaris™ Express or OpenSolaris™ programs, then the guest domain might not be able to access that virtual disk.

The same problem can occur with a guest domain running the Solaris Express or OpenSolaris programs using a virtual disk built from a ZFS volume provided by a service domain running Solaris 10 OS.

*Workaround:* Ensure that the guest domain and the service domain are running the same version of Solaris software (Solaris 10 OS, Solaris Express, or OpenSolaris).

## Page Retirement Does Not Persist in the Logical Domains Environment (Bug IDs 6531030 and 6531058)

When a memory page of a guest domain is diagnosed as faulty, the Logical Domains Manager retires the page in the logical domain. If the logical domain is stopped and restarted again, the page is no longer in a retired state.

The `fmadm faulty -a` command shows whether the page from either the control or guest domain is faulty, but the page is not actually retired. This means the faulty page can continue to generate memory errors.

*Workaround:* Use the following command in the control domain to restart the Fault Manager daemon, `fmd(1M)`:

```
primary# svcadm restart fmd
```

## Aggregated Network Devices Are Not Supported by the Virtual Switch (Bug ID 6531266)

Currently, the virtual switch (*vsw*) does not support the use of aggregated network interfaces. If a virtual switch instance is told to use an aggregated device (*aggr15* in this example), then a warning message similar to the following appears on the console during boot:

```
WARNING: mac_open aggr15 failed
```

*Recovery:* Configure the virtual switch to use a supported GLDv3-compliant network interface, and then reboot the domain.

## Fault Manager Daemon (*fmd*) Does Not Correctly Recover From a Logical Domain Channel Reset (Bug ID 6533308)

If you reset the system controller while the host is powered on, subsequent error reports and faults are not delivered to the host.

*Recovery:* Use one of the following methods to recover.

- Restart the *fmd*(1M):

```
# svcadm disable fmd
# svcadm enable fmd
```

- Reboot.
- Reload the Encoding Table Management (ETM) module:

```
# fmadm unload etm
# fmadm load /usr/platform/sun4v/lib/fm/fmd/plugins/etm.so
```

## Using the `server-secure.driver` With an NIS Enabled System, Whether or Not LDom Is Enabled (Bug ID 6533696)

On a system configured to use the Network Information Services (NIS) or NIS+ name service, if the Solaris Security Toolkit software is applied with the `server-secure.driver`, NIS or NIS+ fails to contact external servers. A symptom of this problem is that the `ypwhich(1)` command, which returns the name of the NIS or NIS+ server or map master, fails with a message similar to the following:

```
Domain atlas some.atlas.name.com not bound on nis-server-1.c
```

The recommended Solaris Security Toolkit driver to use with the Logical Domains Manager is `ldm_control-secure.driver`, and NIS and NIS+ work with this recommended driver.

If you are using NIS as your name server, you cannot use the Solaris Security Toolkit profile `server-secure.driver`, because you may encounter Solaris OS Bug ID 6557663, IP Filter causes panic when using `ipnat.conf`. However, the default Solaris Security Toolkit driver, `ldm_control-secure.driver`, is compatible with NIS.

### To Reset Your System

Log in to the system console from the system controller, and if necessary, switch to the ALOM mode by typing:

```
# #.
```

Power off the system by typing the following command in ALOM mode:

```
sc> poweroff
```

Power on the system.

```
sc> poweron
```

Switch to the console mode at the ok prompt:

```
sc> console
```



Power on the system.

```
ok boot -s
```

Edit the file `/etc/shadow`, and change the first line of the shadow file that has the root entry to:

```
root::6445::::
```

Log in to the system and do one of the following:

- Add file `/etc/ipf/ipnat.conf`
- Undo the Solaris Security Toolkit, and apply another driver

```
# /opt/SUNWjass/bin/jass-execute -ui  
# /opt/SUNWjass/bin/jass-execute -a ldm_control-secure.drivert
```

## Network Performance Is Substantially Worse in a Logical Domain Guest Than in a Non-LDoms Configuration (Bug ID 6534438)

The virtual networking infrastructure adds additional overhead to communications from a logical domain. All packets are sent through a virtual network device, which, in turn, passes the packets to the virtual switch. The virtual switch then sends the packets out through the physical device. The lower performance is seen due to the inherent overheads of the stack.

*Workarounds:* Do one of the following depending on your server:

- On Sun UltraSPARC T1-based servers, such as the Sun Fire T1000 and T2000, and Sun UltraSPARC T2+ based servers such as the Sun SPARC Enterprise T5140 and T5240, assign a physical network card to the logical domain using a split-PCI configuration. For more information, refer to “Configuring Split PCI Express Bus to Use Multiple Logical Domains” in the *Logical Domains (LDoms) 1.0.2 Administration Guide*.
- On Sun Ultra SPARC T2-based servers, such as the Sun SPARC Enterprise T5120 and T5220 servers, assign a Network Interface Unit (NIU) to the logical domain.

## Logical Domain Time-of-Day Changes Do Not Persist Across a Power Cycle of the Host (Bug ID 6536572)

If the time or date on a logical domain is modified, for example using the `ntpdate` command, the change persists across reboots of the domain but not across a power cycle of the host.

*Workaround:* For time changes to persist, save the configuration with the time change to the SC and boot from that configuration.

## When the Logical Domains Manager is Running, OpenBoot PROM Variables Cannot be Modified by the `eeprom(1M)` Command (Bug ID 6540368)

This issue is summarized in [“Logical Domain Variable Persistence”](#) on page 14.

## Errors to Buses in a Split-PCI Configuration Might Not Get Logged (Bug ID 6542295)

During operations in a split-PCI configuration, if a bus is unassigned to a domain or is assigned to a domain but not running the Solaris OS, any error in that bus or any other bus may not get logged. Consider the following example:

In a split-PCI configuration, Bus A is not assigned to any domain, and Bus B is assigned to the primary domain. In this case, any error that occurs on Bus B might not be logged. (The situation occurs only during a short time period.) The problem resolves when the unassigned Bus A is assigned to a domain and is running the Solaris OS, but by then some error messages might be lost.

*Workaround:* When using a split-PCI configuration, quickly verify that all buses are assigned to domains and are running the Solaris OS.

## Emulex-based Fibre Channel Host Adapters Not Supported in Split-PCI Configuration on Sun Fire T1000 Servers (Bug ID 6544004)

The following message appears at the `ok` prompt if an attempt is made to boot a guest domain that contains Emulex-based Fibre Channel host adapters (Sun Part Number 375-3397):

```
ok> FATAL:system is not bootable, boot command is disabled
```

These adapters are not supported in a split-PCI configuration on Sun Fire T1000 servers.

## Starting and Stopping SunVTS Multiple Times Can Cause Host Console to Become Unusable (Bug ID 6549382)

If SunVTS™ is started and stopped multiple times, it is possible that switching from the SC console to the host console, using the console SC command can result in either of the following messages being repeatedly emitted on the console:

```
Enter #. to return to ALOM.  
Warning: Console connection forced into read-only mode
```

*Recovery:* Reset the SC using the `resetsc` command.

## Infiniband PCI-X Card Lost if Primary Domain is Rebooted (Bug ID 6551861, 6552598, 6563713, 6589192 and 6625687)

The following Infiniband cards are not supported with LDomS 1.0.1 and 1.0.2:

- Sun Dual Port 4x IB Host Channel Adapter PCI-X
- Dual port 4x PCI Express Infiniband Host Channel Adapter - Low Profile.

*Workaround:* If one of these unsupported configuration is used with LDoms all logical domains be stopped and unbound, before the primary/control domain rebooted. Failure to do so may result in the device becoming unusable and the system will not recognize the card.

## Do Not Use Control-C From the `prtdiag(1M)` Command (Bug ID 6552999 and 6630571)

Normally, when the verbose (-v) option is specified to the `prtdiag(1M)` command in the control domain, additional environmental status information is displayed. If the output of this information is interrupted by issuing a Control-C, the PICL daemon, `picld(1M)`, can enter a state which prevents it from supplying the environmental status information to the `prtdiag` command from that point on, and the additional environmental data is no longer displayed.

*Workaround:* Restart the `picld(1M)` SMF service in the control domain using the following command:

```
# svcadm restart picld
```

## Do Not Specify a Virtual Switch Interface as a Network Device (Bug ID 6554177)

Do not specify a virtual switch (`vsw`) interface as the network device for a virtual switch configuration. That is, do not specify a virtual switch interface as the `net-dev` property for the `ldm add-vswitch` or `ldm set-vswitch` commands.

## Virtual Network Interface Does Not Handle Memory Allocation Failure Properly (Bug ID 6556778)

If you attempt to plumb 12 virtual networks on a guest domain, the 12th virtual network hangs the guest domain if the memory is 512M or lower.

*Workaround:* Provide the guest domain with at least 1 gigabyte of memory or plumb fewer virtual networks.

## Virtual Disks Created From Files Do Not Support EFI Labels (Bug ID 6558966)

If a virtual disk is backed by a file then this virtual disk can not be labeled with an EFI label and can not directly be added in a ZFS pool.

*Workaround:* The disk has to be labeled with a VTOC label using the `format(1M)` command. The disk can be added to a ZFS pool by creating a VTOC label with a slice covering the entire disk (for example, slice 0) and adding that slice to the ZFS pool instead of adding the entire disk. For example, use `zpool create xyzpool c0d1s0` instead of `zpool create xyzpool c0d1`.

## Occasionally During a Solaris OS Reboot, You Can Receive an Unsuccessful Read or Write Message (Bug ID 6560890)

Occasionally during a Solaris OS boot, a console message from the Domain Services (ds) module reports that reading or writing from a logical domain channel was unsuccessful. The reason code (131) indicates that the channel has been reset. Below are examples of the console message:

```
NOTICE: ds@1: ldc_read returned 131_  
WARNING: ds@0: send_msg: ldc_write failed (131):
```

*Recovery:* None. These console messages do not affect the normal operation of the system and can be ignored.

## The `prtdiag` and `prtpicl` Utilities Do Not Work in Guest Domains (Bug ID 6563508)

The `prtpicl(1M)` and `prtdiag(1M)` utilities do not work in a guest domain. Each utility produces the following error message, and neither utility displays any other information:

```
picl_initialize failed: Daemon not responding
```

In these situations, the PICL daemon, `picld(1M)`, is in a hung state.

*Workarounds:* Use one of the following workarounds:

- Run the `prtdiag(1M)` and `prtpicl(1M)` utilities only on the control domain.
- Disable the PRI pseudo driver using the following procedure.

## ▼ To Disable the PRI Pseudo Driver

### 1. Stop the PICL service:

```
# svcadm disable picl
```

### 2. Disable the PRI pseudo device driver:

```
# rem_drv ds_pri
```

### 3. Re-start the PICL service:

```
# svcadm enable picl
```

## Under Certain Circumstances, Reverting to a Configuration Previously Saved on the SC Can Cause the Logical Domains Manager to Crash (Bug ID 6571091)

After reverting to a logical domain configuration previously saved using the `ldm add-config` command, the Logical Domains Manager might crash with the following error message:

```
"0L != clientp->published_name".
```

*Workaround:* When creating virtual I/O clients and services do not use the canonical names which the Logical Domains Manager applies when there is no match in the constraints database. These names are:

| Device             | Canonical Name Format        |
|--------------------|------------------------------|
| <code>vdisk</code> | <code>vdiskNN</code>         |
| <code>vnet</code>  | <code>vnetNN</code>          |
| <code>vsw</code>   | <code>ldom-name-vswNN</code> |

| Device | Canonical Name Format        |
|--------|------------------------------|
| vcc    | <i>ldom-name-vccNN</i>       |
| vds    | <i>dom-name-vdsNN</i>        |
| vdsdev | <i>ldom-name-vdsNN-volVV</i> |

*NN* and *VV* refer to monotonically increasing instance numbers.

## Virtual Disk Service Should Support Unformatted Disks (Bug ID 6575050)

A physical disk that is unformatted or that does not have a valid disk label, either a Volume Table of Contents (VTOC) or an Extensible Firmware Interface (EFI) label, cannot be exported as a virtual disk to another domain.

Trying to export such a disk as a virtual disk fails when you attempt to bind the domain to which the disk is exported. A message similar to this one is issued and stored in the messages file of the service domain exporting the disk:

```
vd_setup_vd(): vd_read_vtoc returned errno 22 for /dev/dsk/c1t44d0s2
vds_add_vd(): Failed to add vdisk ID 1
```

To export a physical disk that is unformatted or that does not have a valid disk label, use the `format(1M)` command first in the service domain to write a valid disk label (VTOC or EFI) onto the disk to be exported.

## Inconsistent Console Behavior When Not Using the Virtual Console (Bug ID 6581309)

Console behavior on the control domain is inconsistent when a graphics device and keyboard are specified for console use. This occurs when the OpenBoot variables `input-device` and `output-device` are set to anything other than the default value of `virtual-console`.

If the control domain is set this way, some console messages are sent to the graphics console and others are sent to the virtual console. This results in incomplete information on either console. In addition, when the system is halted, or a break is sent to the console, control is passed to the virtual console which requires keyboard input over the virtual console. As a result, the graphics console appears to hang.

*Workaround:* To avoid this problem, use only the virtual console. From the OpenBoot ok prompt, ensure that the default value of `virtual-console` is set for both the `input-device` and variables.

*Recovery:* Once the graphics console appears hung, do the following:

1. Connect to the virtual console from the system processor to provide the required input.
2. Press the carriage return on the virtual console keyboard once to see the output on the virtual console.

If these solutions do not work for your configuration or if you have further questions, contact Sun Services.

## Guest Domain Can Lose Connection to the Virtual Switch When the Service Domain is Rebooted (Bug ID 6581720)

Under certain conditions, after a service domain is rebooted while a guest domain is running, the virtual network (`vnet`) device on the guest fails to establish a connection with the virtual switch on the service domain. As a result, the guest domain cannot send and receive network packets.

*Workarounds:* Use one of the following workarounds on the domain with the virtual network:

- Unplumb and replumb the `vnet` interface. You can do this if the domain with `vnet` cannot be rebooted. For example:

```
# ifconfig vnet0 down
# ifconfig vnet0 unplumb
# ifconfig vnet0 plumb
# ifconfig vnet0 ip netmask mask broadcast + up
```

- Add the following lines to the `/etc/system` file on the domain with `vnet` and reboot the domain:

```
set vnet:vgen_hwd_interval = 5000
set vnet:vgen_max_hretries = 6
```



## SP `setdate` Command Changes the Solaris OS Date When Using Non-Default Logical Domain Configurations (Bug ID 6586305)

If you use the service processor (SP) `setdate` command after you configure non-default logical domains and save them to the SP, the date on non-default logical domains changes.

*Workaround:* Configure the SP date using the `setdate` command *before* you configure the logical domains and save them on the SP.

*Recovery:* If you use the SP `setdate` command after you save the non-default logical domains configurations on the SP, you would need to boot each non-default logical domains to the Solaris OS and correct the date. Refer to the `date(1)` or `ntpdate(1M)` commands in the Solaris 10 OS Reference Manual Collection for more information about correcting the date.

## Do Not Use Non-Numeric Values for Port Numbers in `ldm` Commands (Bug ID 6590848)

The current behavior for the port number argument to the `ldm set-vcons` command, as well as the port range arguments to the `ldm {add,set}-vcc` commands, is to ignore anything starting with a non-numeric value. For example, if the value `0.051` is passed in as the port number for a virtual console, rather than returning an error, the value is interpreted as `0`, which tells the Logical Domains Manager to use automatic port allocation.

*Workaround:* Do not use non-numeric values in port numbers for any `ldm` commands.

## Virtual Disk Server Prints File Lookup Error During Service Domain Boot (Bug ID 6591399)

When a service domain is rebooted while some guest domains are bound, you can see messages similar to these from the virtual disk server:

```
vd_setup_file(): Cannot lookup file (/export/disk_image_s10u4_b12.1) errno=2
vd_setup_vd(): Cannot use device/file (/export/disk_image_s10u4_b12.1) errno=2
```

These messages indicate that the specified file or device is to be exported to a guest domain, but that this file or device is not ready to be exported yet.

*Workaround:* These messages are usually harmless and should stop once the service domain has completed its boot sequence. If similar messages are printed after the service domain is fully booted, you may want to check whether the specified file or device is accessible from the service domain.

## Logical Domains Manager Does Not Retire Resources On Guest Domain After a Panic and Reboot (Bug ID 6591844)

If a CPU or memory fault occurs, the affected domain might panic and reboot. If FMA attempts to retire the faulted component while the domain is rebooting, the Logical Domains Manager is not able to communicate with the domain, and the retire fails. In this case, the `fmadm faulty` command lists the resource as degraded.

*Recovery:* Wait for the domain to complete rebooting and then force FMA to replay the fault event by restarting `fmd` on the control domain using this command:

```
# svcadm restart fmd
```

## Logical Domains Manager Accepts Multiple Entries of I/O Constraints (Bug ID 6591905)

It is possible to erroneously add duplicate I/O constraints when configuring a logical domain.

*Workaround:*

1. Show the duplicate entries by using this command:

```
# ldm list -l
```

2. Remove the duplicate I/O entries using this command:

```
# rm-io bus ldom
```

## Console Ports Dynamically Assigned When a Logical Domain Is Bound Become Hard Constraints If the Logical Domains Manager Subsequently Restarts (Bug ID 6596652)

When logical domains are being configured with no specific console port specified for any logical domain, any Logical Domains Manager (which may happen automatically as part of a delayed reconfiguration or LDoms manager exit), may change the LDoms Manager console port configuration state from what the user has originally entered. This may result in the following error message when attempting to bind a logical domain:

```
Unable to bind client vcons0
```

Workaround: Check the actual configuration state for the guest which failed to bind using the command:

```
# ldm ls-constraints
```

The output should show that the port constraint in the console matches with one of the bound guests. Use the `ldm destroy` command to completely remove the guest. Create the guest from scratch without any constraint set on the console, or use another console port not currently assigned to any bound guest.

## XVR-200 Graphics Adapter On `pci@7c0` Leaf Causes PCI-E Panic and Termination (Bug ID 6598882)

If a XVR-200 graphics adapter is installed in a PCI-Express slot on the `pci@7c0` leaf of a Sun Fire or SPARC Enterprise T2000 server, rebooting the domain can cause a panic and an hypervisor termination.

The XVR-200 card is not supported with the LDoms 1.0.2 release.

## Guest Domain With More Than Four Virtual Networks on the Same Network Using DHCP Can Become Unresponsive (Bug ID 6603974)

If you configure more than four virtual networks (vnets) in a guest domain on the same network using the Dynamic Host Protocol (DHCP), the guest domain can eventually become unresponsive while running network traffic.

*Recovery:* Issue an `ldm stop-domain ldom` command followed by an `ldm start-domain ldom` command on the guest domain (*ldom*) in question.

*Workaround:* Avoid such configurations.

## The fmd Core Dumps On a Hardened, Single Strand Control Domain (Bug ID 6604253)

If you run the Solaris 10 11/06 OS and you harden drivers on the primary domain which is configured with only one strand, rebooting the primary domain or restarting the `fmd` can result in an `fmd` core dump. The `fmd` dumps core while it cleans up its resources, and this does not affect the FMA diagnosis.

*Workaround:* Add a few more strands into the primary domain. For example,

```
# ldm add-vcpu 3 primary
```

## The set-vcpu Command Does Not Warn About MAUs Becoming Unreachable On Initial Reconfiguration (Bug ID 6609051)

When removing CPUs from a domain that is in delayed reconfiguration mode, if all the CPUs which are bound to that domain and on the same core are removed, and the Modular Arithmetic Unit (MAU) on that core had also been bound to the same domain, that MAU becomes orphaned. It is no longer reachable by the domain to which it is bound, nor is it made available to any other domain which has CPUs bound to the same core. In addition, there is no warning or error returned at the time the MAU is orphaned.

*Workaround:* Remove sufficient MAUs from the domain prior to removing CPUs, so that removing CPUs does not result in MAUs becoming unreachable.

- On UltraSPARC T1-based systems, there is one MAU for every four CPU strands

- On UltraSPARC T2-based systems, there is one MAU for every eight CPU strands  
To find out which MAUs are bound to the domain, type:

```
# ldm ls -l ldom
```

To remove MAUs from a domain, type:

```
# ldm rm-mau number ldom
```

## The sun4v Channel Nexus Generates Invalid Interrupt Cookies (Bug ID 6616313)

The sun4v channel nexus generates interrupt cookies that are put in the first word for the devmondo generated for each channel by the hypervisor. The devhandle for the channel nexus is 0x200. If the generated devinos are less than 0x1ff (511), the cookie is valid. If the generated devinos are greater than 0x1ff (511), the cookie is invalid.

## Sun SPARC Enterprise T5120 Server Panics on Reboot (Bug ID 6620322)

The kernel acquires a lock with incorrect CPU owner information before entering into the PROM service routines. This incorrect information can cause a panic.

## Attempting to wanboot a Logical Domain Using a Solaris 10 8/07 OS Installation DVD Causes Hang (Bug ID 6624950)

Wanbooting a logical domain using a miniroot created from a Solaris 10 8/07 OS installation DVD hangs during boot of the miniroot.

## LDoms Disk and Network Services Renamed After 1.0.1 to 1.0.2 Upgrade (Bug ID 6626770)

If the LDoms Manager database is not preserved across an upgrade or fresh installation of the Solaris OS on the primary domain, then any devalias for virtual disks and networks in guests which reference non-canonical device names no longer point to valid device names when the guest next enters OpenBoot firmware, on a reboot for example.

This can cause problems if these devalias are used in the guest's OpenBoot parameters. If the `boot-file` was set to a disk devalias, for example, it may no longer be valid and booting fails.

*Recovery:* To recover and allow the domain to boot if you have not saved the LDoms Manager constraint database, change any affected devalias values to refer to the device by either its new name or the full path name of the device.

*Workaround:* Follow the recommendations on installing or upgrading Solaris on the primary domain in the *Logical Domains (LDoms) 1.0.2 Administration Guide*. Specifically, save and restore the LDoms Manager constraint database at `/var/opt/SUNWldm/ldom-db.xml` across an OS upgrade.

## Logical Domains Manager Gets Fatal Error On Startup After File System Full Error (Bug ID 6628063)

If the file system on which the LDoms database file (`/var/opt/SUNWldm/ldom-db.xml`) resides on the control domain becomes full, the LDoms Manager may fail to properly update the database after a configuration state change. This can be detected by looking for the following warning in the LDoms Manager Service Management Facility (SMF) log file (`/var/svc/log/ldoms-ldmd:default.log`):

```
warning: Database could not be saved to file
```

When this occurs in some cases the LDoms Manager may fail to come up when restarted and an error of the following form is seen in LDoms Manager SMF log file,

```
fatal error: Server physio minor 4 not available
```

*Recovery:* Delete the LDOMs database and restart the LDOMs Manager. If non-canonical device names have been used and you have not saved the LDOMs Manager constraint database, change any affected devalias values to refer to the device by either its new name or the full path name of the device.

*Workaround:* Ensure that the file system on which the LDOMs database is stored does not become full. For example, one could make `/var/opt/SUNWldm` a separate file system.

## The `vldc_chpoll` Command Returns ENOTACTIVE After an SC Reset (Bug ID 6629230)

The `scadm` command that uses the logical domain channel (LDC) connection can hang following a SC reset.

*Workaround:* Reboot the host to reestablish connection with the SC, or all applications are forced to close and reopen the channel.

## The `vxdmp` Utility Does Not Correctly Export Disks (Bug ID 6637560)

If a physical disk is exported as a virtual disk through the Veritas Dynamic Multipathing (DMP) framework (that is, using `/dev/vx/dmp/cXdXtXs2`), the physical disk is not exported correctly, and it appears as a single slice disk in the guest domain.

*Workaround:* The physical disk should be exported without using the Veritas DMP framework. The disk should be exported using `/dev/dsk/cXdXtXs2` instead of `/dev/vx/dmp/cXdXtXs2`.

## Configuring 17 Virtual Networks on a Virtual Switch Causes Guest Domains Boot Panic (Bug ID 6639934)

Guest domains panic during boot after adding the 17th virtual network (`vnet`) to a virtual switch (`vsw`) service.

*Workaround:* Do not configure more than 15 virtual networks on a virtual switch.

## Addition of Virtual Disk or Network Devices Under Delayed Reconfiguration Can Fail (Bug ID 6646690)

If virtual devices are added to an active domain, and virtual devices are removed from that domain before that domain reboots, then the added devices do not function once the domain reboots.

*Recovery:* Remove and then add the non-functional virtual devices, making sure that all the remove requests precede all the add requests, and then reboot the domain.

*Workaround:* On an active domain, do not add and remove any virtual devices without an intervening reboot of the domain.

## Core Dump at `restore_bindings` (Bug ID 6648291)

In certain sparse memory configurations, an attempt to either add more memory than is available, or to create more than 32 memory blocks in a bound or active domain (the maximum supported limit) can result in the LDom Manager terminating. When this happens, the following message is returned on the failed request:

```
Receive failed: logical domain manager not responding
```

SMF then restarts the LDom Manager, and the system is fully functional once it is restarted.

## The `set-vcpu` Command Used in a Guest Domain During Delayed Reconfiguration Causes Termination (Bug ID 6651993)

During a single delayed reconfiguration operation, do not attempt to add CPUs to a domain if any were previously removed during the same delayed reconfiguration. Either cancel the existing delayed reconfiguration, if possible, or commit it by rebooting the target domain, and then add the CPUs.



Failure to heed this restriction, under certain circumstances, can lead to the hypervisor returning a parse error to the LDom Manager, resulting in the following error message on the add attempt:

```
Receive failed: logical domain manager not responding
```

The LDom Manager restarts, and the add request is ignored.

## The `ldm panic-domain` Command Returns Incorrect Error Message (Bug ID 6655083)

If the hypervisor rejects an `ldm panic-domain` request (because the domain is already resetting for example), the error message returned by the LDom Manager is misleading:

```
Invalid LDom ldg23
```

This implies the domain is invalid which is not the case.

## LDoms Multidomain Functionality Does Not Support SNMP 1.5.4 on Some Systems (Bug ID 6655981)

LDoms multidomain functionality does not support SNMP 1.5.4 on Sun SPARC Enterprise T5140 and Sun SPARC Enterprise T5240 systems. Only a single global domain is supported.

## Simultaneous Net Installation of Multiple Domains Fails When In A Common Console Group (Bug ID 6656033)

Simultaneous net installation of multiple guest domains fails on Sun SPARC Enterprise T5140 and Sun SPARC Enterprise T5240 systems with a common console group.

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

## Multiple Memory Configuration Changes Under Delayed Reconfiguration Causes Hypervisor Termination (Bug ID 6657785)

Performing multiple `add-mem`, `set-mem`, and `rm-mem` operations as part of a single delayed reconfiguration operation can cause a hypervisor termination when the domain is restarted, shutting down the entire system.

*Workaround:* Avoid performing multiple `add-mem`, `set-mem`, or `rm-mem` operations on a domain as part of a single delayed reconfiguration operation.

## Hypervisor Does Not Generate `ereport` After Initial `pciexrc1` I/O Error (Bug ID 6657972)

The first error always reports correctly. The problem is if all CPUs are not configured, the hypervisor does not generate an `ereport` on the SP on the second and subsequent errors. The system has to power cycle to get back to normal.

## Guest Loses Virtual Disk Services When Using a Configuration Saved During Delayed Reconfiguration (Bug ID 6659100)

After a delayed reconfiguration on a guest domain and a subsequent power cycle, the guest fails to boot with the following message:

```
Boot device: /virtual-devices@100/channel-devices@200/disk@0 File
and args:
WARNING: /virtual-devices@100/channel-devices@200/disk@0: Timeout
connecting to virtual disk server... retrying
```

This happens when a configuration is saved to the SP while a delayed reconfiguration is pending.

*Workarounds:* Either do not save the configuration to SP after the delayed reconfiguration is completed and the guest rebooted, or run the following on the primary domain after the guest is rebooted after a delayed reconfiguration:

```
# ldm stop ldom
# ldm unbind ldom
# ldm bind ldom
# ldm start ldom
```

The guest should now boot successfully.

## Domains Stuck Initializing In OpenBoot Firmware After a no-free-memory Bind Failure (Bug ID 6661675)

If a `bind-domain` operation, or a request to increase the allocated memory of a bound or active domain, fails due to lack of available memory, it is possible for the next successful such operation on that domain to assign an incorrect real address (RA), resulting in the domain hanging on boot.

*Recovery:* Perform two subsequent successful requests. The first successful one runs the risk of triggering this condition; the second one works properly. Each of these requests needs to be either a `bind-domain` operation, or an attempt to add memory to an already bound or active domain. For example, adding memory to an inactive domain does not meet the requirement of a successful request for the purposes of recovering from this bug.

*Workaround:* Avoid the trigger condition by not attempting to allocate more memory than is available on the system.

---

## LDoms Issue That Applies Only to Solaris 10 11/06 OS

The following LDoms issue applies only if you have Solaris 10 11/06 OS running on your system.

# Removing or Updating a Virtual Switch Can Cause Networking to Fail

Once the virtual switch driver (`vswitch`) has attached, either as part of the normal Solaris OS boot sequence, or as a result of an explicit Solaris OS `add_drv(1M)` command, removing or updating the driver can cause networking to fail.

*Workaround:* Once `vswitch` has attached, do not remove the driver using the Solaris OS `rem_drv(1M)` command or update the driver using the Solaris OS `update_drv(1M)` command.

*Recovery:* If you do remove the driver using the `rem_drv` command and then attempt to reattach it using the `add_drv` command, you must reboot after the `add_drv` command completes to ensure the networking restarts correctly. Similarly, you must also reboot after an `update_drv` command completes to ensure the networking does not fail.

---

## LDoms Bugs Fixed for Solaris 10 8/07 OS

The following LDoms bugs were fixed for the Solaris 10 8/07 OS:

### Networking

- 6405380 LDoms `vswitch` needs to be modified to support network interfaces
- 6418780 `vswitch` needs to be able to process updates to its MD node
- 6447559 `vswitch` should take advantage of multiple unicast address support
- 6474949 `vswitch` panics if `mac_open` of the underlying network device fails
- 6492423 `vswitch` multi-ring code hangs when queue thread not started
- 6492705 `vsw` warning messages should identify device instance number
- 6512604 `handshake` `untimeout()` race condition in `vnet`
- 6517019 `vgen_multicst` does not handle `kmem_zalloc` failure
- 6496374 `vsw`: “`turnstile_block: unowned mutex`” panic on a diskless-clients test bed
- 6514591 `vsw`: fix for 6496374 causes `softhang`
- 6523926 `handshake` restart can fail following reboot under certain conditions
- 6523891 `vsw` needs to update lane state correctly for RDX pkts
- 6556036 `vswitch` panics when trying to boot over `vnet` interface

## Disk

- 6520626 Assertion panic in vdc following primary domain reboot
- 6527265 Hard hang in guest ldom on issuing the format command
- 6534269 vdc incorrectly allocs mem handle for synchronous DKIOCFUSHWRITECACHE calls
- 6547651 fix for 6524333 badly impact performance when writing to a vdisk
- 6524333 Service domain panics if it fails to map pages for a disk on file
- 6530040 vds does not close underlying physical device or file properly

## General

- 6488115 reboot from guest via break hangs
- 6495154 mdeg should not print a warning when the MD generation number does not change
- 6520018 vntsd gets confused and immediately closes newly established console connections
- 6505472 RC1 build: guest ldg(s) soffhang during repeat boot
- 6521890 recursive mutex\_enter in ldc\_set\_cb\_mode
- 6528180 link state change is not handled under certain conditions in ldc
- 6526280 Guest with 64 vdisk devices hangs during boot
- 6528758 'ds\_cap\_send: invalid handle' message during LDom boot

---

## Bugs Fixed for LDom 1.0.2

Sun recommends the latest patch be installed. The following LDom bugs were fixed for the LDom 1.0.2 software release:

### Fixed in Patch 127111-08

- 6578761 System hangs in ds\_cap\_fini() and ds\_cap\_init()
- 6593231 Domain Services logging facility must manage memory better
- 6616313 cnex incorrectly generates interrupt cookies
- 6630945 vntsd runs out of file descriptor with very large domain counts
- 6501039 rebooting multiple guests continuously causes a reboot thread to hang

- 6527622 Attempt to store boot command variable during a reboot can time out
- 6589682 IO-DOMAIN-RESET (Ontario-AA): kern\_postprom panic on tavor-pci configuration (reboot)
- 6605716 halting the system should not override auto-boot? on the next poweron
- 6519849 vnet hot lock in vnet\_m\_tx affecting performance
- 6530331 vsw when plumbed and in prog mode should write its mac address into HW
- 6531557 format(1m) does not work with virtual disks
- 6536262 vds occasionally sends out-of-order responses
- 6544946 Adding non existent disk device to single cpu domain causes hang
- 6566086 vdc needs an I/O timeout
- 6573657 vds type-conversion bug prevents raw disk accesses from working
- 6575216 Guests may lose access to disk services (VDS) if IO domain is rebooted
- 6578918 disk image should have a device id