



Logical Domains (LDoms) 1.0 Release Notes

Sun Microsystems, Inc.
www.sun.com

Part No. 819-6429-11
May 2007, Revision B

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

Copyright 2007 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, JumpStart, OpenBoot, Sun Fire, SunSolve, 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 © 2007 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, Etats-Unis. Tous droits réservés.

Sun Microsystems, Inc. détient les droits de propriété intellectuelle 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 de 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 Etats-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 Etats-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 régi 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'ETAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE.



Adobe PostScript

Contents

Logical Domains (LDoms) 1.0 Release Notes 1

Changes for This Release 1

Supported Servers 1

Required and Recommended Software and Patches 2

Location of Documentation 2

Supported Network Adapters 3

▼ To Determine If a Network Adapter is GLDv3-Compliant 3

Graphics Card Support 4

Memory Size Requirements 4

Software That Can Be Used With the Logical Domains Manager 5

General Notes and Issues 6

The `SUNWldomu` and `SUNWldomr` Packages Are Missing From Metaclusters in the Solaris 10 11/06 OS 6

Some Features Are Not Available Currently on System Firmware 6.4.x for Netra T2000 Servers 6

Sun x8 Express 1/10G Ethernet Adapter (`nxge` driver) Not Supported 7

Rebooting the Control Domain When Guest Domains are Running or Bound is Not Supported 7

Cleanly Shutting Down and Power Cycling a Logical Domains System 7

▼ To Power Off a System With Multiple Active Domains 7

▼ To Power Cycle the System 7

Removing or Updating a Virtual Switch Can Cause Networking to Fail	8
Interrupts and Logical Domains	8
XML Format Is Undergoing Changes	10
The <code>ldm stop-domain</code> Command Can Time Out If the Domain Is Heavily Loaded	10
Memory Size Requested Might Be Different From Memory Allocated	11
Power Off the Host Before Resetting the SC When the Logical Domains Manager Is Running	11
Dynamic Reconfiguration of Virtual CPUs with Cryptographic Units	11
Split PCI Regresses in FMA Functionality from Non-Logical Domains Systems	12
Logical Domain Variable Persistence	12
The <code>eprom(1M)</code> Command Cannot Be Used to Reset EEPROM Values to Null In Logical Domains Systems	13
Logical Domains Manager Restart During <code>ldm</code> Command Execution	14
Bugs Affecting Logical Domains 1.0 Software	14
Logical Domains Manager Might Erroneously Assign an Offline CPU to a Logical Domain (Bug ID 6431107)	14
Guest Domains Currently Cannot Install the OS, LDoms, and Applications from DVDs (Bug ID 6434615)	15
Some <code>format(1M)</code> Command Options Do Not Work With Virtual Disks (Bug IDs 6437722 and 6531557)	15
Running Snoop on a Physical Interface Before the Solaris Virtual Switch Driver Attaches Can Cause a Panic (Bug ID 6473778)	15
Rebooting After a Break Hangs the Guest (Bug ID 6488115)	15
One CPU Might Show 100% Utilization Under Heavy Network Loads (Bug ID 6492023)	16
Warning Should Not Print When MD Generation Number Does Not Change (Bug ID 6495154)	16
Service Domain Containing the Virtual Switch Might Panic on Halting, Rebooting, or Net Installing (Bug ID 6496374)	16
Hang Can Occur With Guest OS in Simultaneous Operations (Bug ID 6497796)	17

SC Alert Messages Seen Each Time the SC is Reset (Bug ID 6499117)	17
Rebooting Multiple Guest Domains Continuously Can Cause OS to Hang (Bug ID 6501039)	17
Panic Message on the Control Domain When Syncing a Guest Domain (Bug ID 6501168)	18
Guest Domains Can Become Unresponsive During Repeated Reboots (Bug ID 6505472)	18
Behavior of the <code>ldm stop-domain</code> Command Needs To Be Improved in Some Cases (Bug ID 6506494)	18
Cannot Set Security Keys With Logical Domains Running (Bug ID 6510214)	19
The <code>vntsd(1M)</code> Command Needs to Validate the <code>listen-to</code> IP Address (Bug ID 6512526)	19
Logical Domain Can Panic or Lose Network Connectivity With Low Memory (Bug ID 6512604)	20
Virtual Disk Server Should Export ZFS Volumes as Full Disks (Bug ID 6514091)	21
The <code>add-vnet</code> Subcommand Allows a Virtual Network Device With the Same MAC Address as Another Logical Domain (Bug ID 6515615)	21
Restrictions on Using the I/O MMU Bypass Mode (Bug ID 6517338)	22
Need Logical Domains Manager Support for Contiguous Guest Memory Mappings (Bug ID 6517343)	22
System Hangs When Doing a Boot Net Installation With a ZFS File (Bug ID 6517957)	23
Certain <code>ldm</code> Subcommands Returns Misleading Message If One or More Arguments are Missing (Bug ID 6519049)	23
The <code>vntsd</code> Might Close Console Connections Unexpectedly (Bug ID 6520018)	24
Canceling a Delayed Reconfiguration Operation Results in a Misleading Message (Bug ID 6520530)	24
Service Domain Running <code>vcc</code> and <code>vntsd</code> Can Panic in Certain Situations (Bug ID 6521890)	25
Disks Managed by Veritas DMP Cannot Be Exported to Other Domains (Bug ID 6522993)	25

- Rarely, When a Guest Domain Boots, It Fails to Establish a Connection With the Virtual Switch (Bug IDs 6523891 and 6523926) 26
- Continuous Console Output From the Control Domain Causes Logical Domains Manager Requests to the SC to Fail (Bug ID 6524255) 27
- Service Domain Panics If It Fails to Map Pages for a Virtual Disk Backed by a File (Bug ID 6524333) 27
- OpenBoot Firmware Does Not Handle Logical Domain Channel Resets Properly (Bug ID 6524613) 28
- CPU Dynamic Reconfiguration Is Disabled for Logical Domains That Have Crypto Units Bound (Bug ID 6525647) 28
- Configuring Logical Domains With More than 32 Virtual I/O Devices or Services Can Result in a Hang During Boot (Bug ID 6526280) 28
- Booting Multiple Logical Domains in a Loop Can Cause an Unresponsive Domain (Bug ID 6526814) 29
- System MAC Address is Modified on I/O Domains When Adding a vnet Device (Bug ID 6526868) 29
- Logical Domains Manager May Stop With Invalid Logical Domain Name (Bug ID 6527206) 30
- Access to a Virtual Disk From a Guest Domain Can Hang If an I/O Operation Is Overlong (Bug ID 6527265) 30
- The `ldm start-domain -i xml-file` Command Causes the Logical Domains Manager to Stop and Restart (Bug ID 6527347) 31
- Memory Constraints Should Be Enforced and Reporting Improved (Bug ID 6527483) 31
- Attempt to Store Boot Command Variable During a Reboot Can Time Out (Bug ID 6527622) 31
- Disk Recovery Fails in a Service Domain When the Disk Device Is Actively Used as a Virtual Disk (Bug ID 6528156) 32
- Newly Configured Virtual Network Device Fails to Establish a Connection With the Virtual Switch (Bug ID 6528180) 33
- When Running Cluster Software, Selecting the `ok` Prompt Upon a Logical Domain Shutdown Can Cause Panic (Bug ID 6528556) 33
- ▼ To Force the Primary Domain to Stop at the `ok` Prompt 34
 - ▼ To Force All Other Domains to Stop at the `ok` Prompt 34

ZFS Volumes Need to Have the Same Version of Solaris Software Running on the Service Domain and the Guest Domain (Bug ID 6528974)	35
Rebooting the primary Domain While Other Active or Bound Domains Are Present and Rebooting an I/O Domain Are Not Supported (Bug ID 6529426)	35
Virtual Disk Server Does Not Close Underlying Physical Device or File Correctly (Bug ID 6530040)	36
Fault Management Daemon in the primary Domain Appears to Hang If You Do Not Enable the Logical Domains Manager Daemon (Bug ID 6530948)	37
Page Retirement Is Not Persistent in the Logical Domains Environment (Bug IDs 6531030 and 6531058)	37
Aggregated Network Devices Are Not Supported by the Virtual Switch (Bug ID 6531266)	37
Loopback Devices Cannot Be Used as Virtual Disks (Bug ID 6532144)	38
Serial Port Cannot Be Used or Opened (Bug ID 6532334)	38
Second PCI Path Does Not Persist Across primary Domain Reboots (Bug ID 6532604)	38
Using the <code>server-secure.driver</code> With an NIS-enabled System, LDoms or No LDoms (Bug ID 6533696)	39
▼ To Reset Your System	40
Cannot Forcibly Stop a Domain If the Domain Has Any Bound PCI Buses (Bug ID 6536420)	41
Logical Domain Time-of-Day Changes Do Not Persist Across a Power Cycle of the Host (Bug ID 6536572)	41
XML Output of the <code>list-constraints</code> Subcommand Is Missing the <code>IO:</code> Configuration Value (Bug ID 6537156)	41
MAC Addresses Are Not Listed in the Correct Format in XML Input and Output Files (Bug ID 6537172)	42
SC Command <code>bootmode bootscript</code> Does Not Work in an LDoms Environment (Bug ID 6540125)	42
OpenBoot <code>power-off</code> Command Behavior With Logical Domains 1.0 Software (Bug ID 6540632)	43
When the <code>ldm add-vnet</code> Command Fails, the Error Message Is Not Clear (Bug ID 6541158)	43

- Machine Description Too Large for Logical Domains Manager's Memory Allocator (Bug ID 6541171) 44
- Do Not Use an `ldm set-vnet` Command to Change a MAC Address on an Active Domain (Bug ID 6541284) 44
- Occasionally, the Logical Domains Manager Does Not Send a Warning That a Disk Is in Use (Bug ID 6541323) 45
- Errors to Buses in a Split-PCI Configuration Might Not Get Logged (Bug ID 6542295) 45
- Forced Stop of a Logical Domain Can Result in a Panic on the Next Boot (Bug ID 6543418) 46
- Adding a Virtual Disk to a Bound Guest Does Not Take Effect (Bug ID 6543735) 46
- Avoid Using the OpenBoot `watch-net-all` Command on a Domain That Contains a Virtual Network (Bug ID 6543748) 46
- During `wanboot` or `waninstall`, Miniroot Download Time Can Increase Significantly (Bug ID 6543749) 47
- Emulex-based Fibre Channel Host Adapters Not Supported in Split-PCI Configuration on Sun Fire T1000 Servers (Bug ID 6544004) 47
- Logical Domain With No Virtual I/O Devices Might Not Respond (Bug ID 6544197) 47
- ▼ To Create a Fictitious Disk Device 48
- Logical Domains Manager Incorrectly Assumes All `vcc` Devices Use the Same Port Range If the Logical Domains Manager Database Does Not Exist (Bug ID 6544772) 49
- Upgrading the Solaris OS Could Cause Problems With Virtual Networking (Bug ID 6544866) 50
- Adding Non-Existent Disk Device to Single-CPU Domain Causes Hang (Bug ID 6544946) 50
- Virtual Disk Can Lose a Label When Rebinding a Domain (Bug ID 6544963) 50
- ▼ To Check If a Label Is Valid 52
- Watchdog Time-out Can Be Triggered by Heavy Network Loads (Bug ID 6545470) 53
- Logical Domains Manager Not Propagating `factory-default` TOD-offset (Bug ID 6546491) 53

Virtual Switch Requires e1000g Driver Update Following Solaris 10 Upgrade
(Bug ID 6547846) 54

Starting and Stopping SunVTS Multiple Times Can Cause Host Console to
Become Unusable (Bug ID 6549382) 54

Documentation Errata 55

The vntsd(1M) Man Page Is Missing Information About Using the Double
Tilde (~~) (Bug ID 6513007) 55

Logical Domains (LDoms) 1.0 Release Notes

These release notes contain the following information about this release of the Logical Domains 1.0 software:

- [“Changes for This Release” on page 1](#)
- [“Supported Servers” on page 1](#)
- [“Required and Recommended Software and Patches” on page 2](#)
- [“Location of Documentation” on page 2](#)
- [“Supported Network Adapters” on page 3](#)
- [“Software That Can Be Used With the Logical Domains Manager” on page 5](#)
- [“Graphics Card Support” on page 4](#)
- [“Memory Size Requirements” on page 4](#)
- [“General Notes and Issues” on page 6](#)
- [“Bugs Affecting Logical Domains 1.0 Software” on page 14](#)
- [“Documentation Errata” on page 55](#)

Changes for This Release

This release of Logical Domains 1.0 software is specifically to add support for the Sun Blade™ T6300 Server Module.

Supported Servers

Logical Domains (LDoms) Manager 1.0 software is supported on the following servers:

- Sun Fire™ or SPARC® Enterprise T1000 Server – for more information, refer to the *Sun Fire or SPARC Enterprise T1000 Server Administration Guide*.
- Sun Fire or SPARC Enterprise T2000 Server – for more information, refer to the *Sun Fire or SPARC Enterprise T1000 Server Administration Guide*.
- Netra™ T2000 Server – for more information, refer to the *Netra T2000 Server Administration Guide*.
- Netra CP3060 Blade – for more information, refer to the *Netra CP3060 Board Product Notes*.

- Sun Blade™ T6300 Server Module – for more information, refer to the *Sun Blade T6300 Server Module Administration Guide*.

Required and Recommended Software and Patches

The following software is required or recommended. The minimum versions required for software and patches are listed.

- *(Required)* **Solaris 10 11/06 Operating System** – Refer to the Solaris 10 Collection for more information.
- **Solaris 10 11/06 OS patches:**
 - *(Required)* **124921-02**, which contains updates to the Logical Domains 1.0 drivers and utilities. Logical Domains networking will be broken without this patch.
 - *(Required)* **125043-01**, which contains updates to the console (qcn) drivers. This patch depends on **KU 118833-36**, so if this is not already updated on your system, you must do so.
 - *(Recommended)* **125800-02**, which is a fix for “[Fault Management Daemon in the primary Domain Appears to Hang If You Do Not Enable the Logical Domains Manager Daemon \(Bug ID 6530948\)](#)” on page 37, which only occurs if the Logical Domains Manager is not installed.

You can find the patches at the SunSolveSM site:

<http://sunsolve.sun.com>

- *(Required)* **System firmware version 6.4.x** – Refer to your server’s administration guide for more information.
- *(Required)* **Logical Domains Manager 1.0 software**
- *(Recommended)* **Solaris Security Toolkit 4.2 software** – Refer to the *Solaris Security Toolkit 4.2 Administration Guide* and *Solaris Security Toolkit 4.2 Reference Manual* for more information.

Location of Documentation

The *Logical Domains (LDoms) 1.0 Administration Guide* and *Logical Domains (LDoms) 1.0 Release Notes* can be found at the following site:

http://www.sun.com/products-n-solutions/hardware/docs/Software/enterprise_computing/systems_management/ldoms/ldoms1_0/index.html

The *Beginners Guide to LDomS: Understanding and Deploying Logical Domains* can be found at the following 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 Administration Guide* for information about how to use non-GLDv3 compliant network adapters.

Currently, the following adapters with their corresponding drivers are supported by the virtual switch:

- Intel PRO/1000 Gigabit family of network interface controllers (e1000g driver)
- Gigabit Ethernet driver for Broadcom BCM57xx (bge driver)

▼ 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`.

Graphics Card Support

The following graphics card can be used with LDom software on the Sun Fire and SPARC Enterprise T2000 servers:

- X3777, XVR-200 Graphics Accelerator
- The graphics card is not virtualized, so it will be shown as a regular physical device owned by the I/O domain. The graphics card provides the graphic capability for that particular I/O domain.

Following are the specifics:

- If all I/O devices are owned by one domain, the graphics card can be installed in any PCI-E slot (0, 1, or 2).
- If the I/O devices are split between two domains, the graphics card must be installed into either PCI-E slot 1 or 2, both of which share the same PCI-E switch with USB interface, so that the graphics card, USB mouse, and keyboard are in the same domain and can be used.

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. For the Solaris 10 11/06 OS, 512 MB is the recommended size of memory to install or upgrade, and 128 MB is the minimum size. The default size for a swap area is 512 MB. For the Solaris 10 11/06 OS, refer to "System Requirements and Recommendations" in the Solaris 10 11/06 Installation Guide: Planning for Installation and Upgrade.

The OpenBoot™ PROM has a minimum size restriction for a domain. For system firmware 6.4.x, that restriction is 12 MB. If you have a domain less than that size, the Logical Domains Manager will automatically boost the size of the domain to 12 MB. Refer to the release notes for your system firmware for information about memory size requirements.

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™ 6.3** functionality can be used on all hardware configured in the control domain on Sun Fire T1000 and Sun Fire T2000 Servers that have Logical Domains Manager 1.0 software enabled. 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 6.3 User’s Guide*.
- **Sun™ Management Center 3.6 Version 6 Add-on Software** can be used only on the control domain with the Logical Domains Manager 1.0 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. For more information about using Sun Management Center on the supported servers, refer to the *Sun Management Center 3.6 Version 6 Add-On Software Release Notes: For Sun Fire, SunBlade, Netra, and SunUltra Systems*.
- **Advanced Lights Out Manager (ALOM) Chip Multithreading (CMT) Version 1.3** software can be used on the control domain on supported servers with the Logical Domains Manager 1.0 software enabled. Refer to “Using LDOMs With ALOM CMT” in the *Logical Domains (LDoms) 1.0 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*.
- **Sun™ Explorer 5.7 Data Collector** can be used with the Logical Domains Manager 1.0 software enabled on the control domain. Sun Explorer is a diagnostic data collection tool. The tool comprises shell scripts and a few binary executables. For more information about using the Sun Explorer Data Collector, refer to the *Sun Explorer User’s Guide*.

General Notes and Issues

This section contains general notes and issues concerning the Logical Domains 1.0 software.

The SUNWldomu and SUNWldomr Packages Are Missing From Metaclusters in the Solaris 10 11/06 OS

The SUNWldomu package is missing from the SUNWCreq metacluster, and both the SUNWldomu and SUNWldomr packages are missing from the SUNWCrnet metacluster (Bug ID 6484072).

If either of those metaclusters are installed on a machine to be used with Logical Domains software, you must install the packages manually.

Some Features Are Not Available Currently on System Firmware 6.4.x for Netra T2000 Servers

Currently, System Firmware 6.4.x does not support the following features on Netra T2000 Servers:

- `scadm(1M)` command, which administers the system controller (SC)
- sun4u-compatible Platform Information and Control Library (PICL)
- sun4u-compatible `prtdiag(1M)` command, which displays system diagnostic information
- Configuration of 4 hard disk drives (HDDs)

The following bug IDs are filed to add this support:

- Add Support for Netra T2000 4-HDD Configuration (6531680 and 6532118)
- Create a New Service Channel to Support PICL and `scadm` (6535555)
- Enable RMC protocol in ALOM-CMT to support `scadm` on Netra T2000 (6536617)

Sun x8 Express 1/10G Ethernet Adapter (nxge driver) Not Supported

This Logical Domains 1.0 software release does not support the Sun x8 Express 1/10G Ethernet Adapter (nxge driver).

Rebooting the Control Domain When Guest Domains are Running or Bound is Not Supported

Rebooting the control domain when guest domains are running or bound is not supported, because doing so could cause the control domain to hang.

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.

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 might cause networking to fail because of Bug ID 6486145.

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.

Interrupts and Logical Domains

Each thread of the UltraSPARC® T1 processor has a limited capacity for handling multiple, simultaneous, outstanding interrupts. Overall, only 256 interrupts can be outstanding for each thread allocated to a logical domain. The chief producers of interrupts are the I/O subsystem and Logical Domain Channels (LDCs), which are the main interdomain communication mechanism.

When configuring logical domains, avoid creating more interrupt producers than 256 times the number of threads assigned to the domain. Practically speaking, this only becomes an issue on the control domain, because of the fact that it 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 Logical Domains Manager control of the other logical domains.

The following guidelines can help you prevent creation of a configuration that could overflow the interrupt capabilities of the control domain:

1. There are $(256 * \# \text{ threads in domain})$ available slots for outstanding interrupts.
2. Each I/O bridge consumes 64 of these potential interrupts.

3. Each LDC consumes two interrupts.
4. The control domain allocates 12 LDCs for various communication purposes with the Hypervisor, Fault Management Architecture (FMA), and the system controller (SC).
5. The control domain allocates one LDC to every logical domain, including itself, for control traffic
6. 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 with four threads that includes both I/O bridges, and 8 additional configured logical domains. Each logical domain needs at a minimum:

- Virtual network
- Virtual disk
- Virtual console

Applying the above guidelines yields the following statistics. The numbers on the statistics lines correlate with the numbers on the guidelines.

1. There are $256 * 4 = 1024$ available interrupts.
2. The two I/O bridges consume 128 interrupts, leaving 896 for LDCs.
- 3 & 4. There are 24 interrupts consumed for standard control domain services, leaving 872.
- 3 & 5. There are 20 interrupts ((8 domains + 2 LDCs for loopback control channel) * 2 interrupts/LDC) consumed for control traffic to the domains, leaving 852.
- 3 & 6. There are 48 interrupts (3 services * 2 interrupts/service * 8 domains) consumed for virtual I/O services, leaving 804.

This configuration will not have problems with potential interrupt overload.

Now consider the case where there are 16 domains instead of 8, and the control domain is reduced to a single thread. The equation for the number of spare interrupt slots in this scenario is:

$$256 - 128 - 24 - 36 - 96 = -28$$

This configuration has the potential for interrupt overload and should be avoided.

XML Format Is Undergoing Changes

The XML format produced by the `-x` option to the `ldm list-constraints` subcommand and consumed by the `-i` option of several `ldm` subcommands is undergoing changes.

If you use the format that is currently being produced for generating scripts or manual editing, you will have to make changes later for your scripts or manual editing.

The `ldm stop-domain` Command Can Time Out If the Domain Is Heavily Loaded

An `ldm stop-domain` command can time out before the domain completes shutting down. When this happens, an error similar to the following is returned by the Logical Domains Manager:

```
LdDom ldg8 stop notification failed
```

However, the domain could still be processing the shutdown request. Use the `ldm list-domain` command to verify the status of the domain. For example:

```
# ldm list-domain ldg8
Name      State  Flags  Cons  VCPU  Memory  Util  Uptime
ldg8     active s----  5000   22   3328M  0.3%  1d 14h 31m
```

The preceding list shows the domain as active, but the `s` flag indicates that the domain is in the process of stopping. This should be a transitory state.

The following example shows the domain has now stopped:

```
# ldm list-domain ldg8
Name      State  Flags  Cons  VCPU  Memory  Util  Uptime
ldg8     bound  -----  5000   22   3328M
```

Memory Size Requested Might Be Different From Memory Allocated

Under certain circumstances, the Logical Domains (LDMs) Manager rounds up the requested memory allocation to either the next largest 8 KB or 4 MB 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	

Power Off the Host Before Resetting the SC When the Logical Domains Manager Is Running

Attempting to reset the system controller while the Logical Domains Manager is running can result in undefined behavior.

Workaround: When running the Logical Domains Manager, power-off the host completely before resetting the system controller.

Dynamic Reconfiguration of Virtual CPUs with Cryptographic Units

Currently, there are several issues 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).
- Virtual CPUs cannot be removed from a domain that is bound but not started if doing so would remove all links to any of the cryptographic units that are currently bound to that domain (Bug ID 6540665).
- The error message when the second restriction (Bug ID 6540665) is encountered is somewhat ambiguous (Bug ID 6540671):

Cannot remove 3 VCPUs because MAU bindings exist.

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 which 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

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 in all cases, variable updates made using any of these methods always persistent across reboots of the domain, and always reflect in any subsequent logical domain configurations saved to the SC.

In Logical Domains 1.0 software, there are a few cases where variable updates do not persist:

- When running in a `factory-default` configuration, 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.

When running the `factory-default` configuration, if you want a variable update to persist across a reboot into the same `factory-default` configuration, use the `eeprom` 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.

- Once domaining has been enabled (that is, the machine is running in a configuration generated by the Logical Domains Manager, not the `factory-default` configuration), all methods of updating a variable (OpenBoot firmware, `eeprom` 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 reverting to the `factory-default` configuration from a configuration generated by the Logical Domains Manager, all LDOMs variables start with their default values.

The `eeprom(1M)` Command Cannot Be Used to Reset EEPROM Values to Null In Logical Domains Systems

The `eeprom(1M)` command cannot be used to reset EEPROM values to null in Logical Domains systems. The following example shows what happens if you attempt this:

```
primary# eeprom boot-file=
eeprom: OPROMSETOPT: Invalid argument
boot-file: invalid property.
```

The same command works correctly on non-Logical Domains systems as shown in this example:

```
# eeprom boot-file=
# eeprom boot-file
boot-file: data not available.
```

Logical Domains Manager Restart During `ldm` Command Execution

If the Logical Domains Manager stops and then restarts during execution of any Logical Domains Manager `ldm` command, the program returns the following error message:

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

Recovery: This message usually indicates the command did not successfully complete. Verify that is the case, then reissue the command if appropriate.

Bugs Affecting Logical Domains 1.0 Software

This section summarizes the bugs that you may 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.

Logical Domains Manager Might 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.

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

Guest Domains Currently Cannot Install the OS, LDoms, and Applications from DVDs (Bug ID 6434615)

The Solaris 10 OS virtual disk drivers (`vdc` and `vdsc`) currently do not support the `CDIO(7I) ioctl`s that are needed to install from DVDs.

Some `format(1M)` Command Options Do Not Work With Virtual Disks (Bug IDs 6437722 and 6531557)

Refer to “Operating the Solaris OS With Logical Domains” in Chapter 5 of the *Logical Domains (LDoms) 1.0 Administration Guide* for specific information.

Running Snoop on a Physical Interface Before the Solaris Virtual Switch Driver Attaches Can Cause a Panic (Bug ID 6473778)

Running a snoop session on a physical interface (for example, the `e1000g0` which the virtual switch has been instructed to use) before the driver has attached, and then canceling the snoop session after the driver has attached, can cause the system to panic with a recursive `rw_enter` panic.

This is an issue only if the virtual switch has been explicitly unloaded previously by use of the `rem_drv(1M)` command.

Recovery: Reboot the domain containing the virtual switch.

Workaround: Run snoop only after the virtual switch has attached.

Rebooting After a Break Hangs the Guest (Bug ID 6488115)

If you send a break on the guest console using the `~#` option of the virtual network terminal server daemon, `vntsd(1M)`, and type `r` for reboot, the guest hangs with no response from the domain.

Recovery: Stop and start the guest domain from the control domain using the `ldm stop-domain` and `ldm start-domain` commands.

Workaround: Issue a reboot at the command line from within the guest domain.

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

Under heavy network loads, one CPU might show 100% utilization dealing with the network traffic.

Workaround: Attach several CPUs to the domain containing the virtual switch to ensure that the system remains responsive under a heavy load.

Warning Should Not Print When MD Generation Number Does Not Change (Bug ID 6495154)

Messages like the following might be seen on the console every time the Logical Domains Manager starts:

```
unix: WARNING: machine_descrip_update: new MD has the same
generation (1359) as the old MD
```

Workaround: Ignore the warning.

Service Domain Containing the Virtual Switch Might Panic on Halting, Rebooting, or Net Installing (Bug ID 6496374)

On halting, rebooting, or net installing a guest domain, the service domain containing the virtual switch might panic with the following message.

```
turnstile_block: unowned mutex
```

Once you establish the connection between the virtual switch and the virtual networks in the guest domains, this bug is not an issue.

Recovery: Reboot the service domain containing the virtual switch.

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

Under rare circumstances, when an ldm 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 `ldm stop` and `ldm start` - Stop and start the guest OS.

SC Alert Messages Seen Each Time the SC is Reset (Bug ID 6499117)

The following messages might be seen on the ALOM-CMT console or in output of the ALOM-CMT `showlogs` command when the system controller is reset:

```
SC Alert: Can't connect to port 2800.  
SC Alert: Can't connect to port 2900.
```

Workaround: Ignore these messages as they have no effect on the system.

Rebooting Multiple Guest Domains Continuously Can Cause OS to Hang (Bug ID 6501039)

Under rare circumstances, if a guest domain is rebooted at a time when it is experiencing high interrupt activity, the OS might hang.

Workaround: None.

Recovery: Use `ldm stop` and `ldm start` - 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 will overflow and the system will not be able to save a crash dump. The panic message is as follows:

```
intr_req pool empty
```

Guest Domains Can Become Unresponsive During Repeated Reboots (Bug ID 6505472)

Following repeated reboots, a domain's user interface can become unresponsive. This happens because of failure to re-establish connection with the service domain. (The system should still respond to network activity, such as the `ping(1M)` command from a remote system and should also respond to a system abort sequence.)

Recovery:

1. **Stop the unresponsive domain using the `ldm stop-domain` command.**
2. **Restart the domain using the `ldm start-domain` command.**

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 `ldom 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 currently 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
```

Logical Domain Can Panic or Lose Network Connectivity With Low Memory (Bug ID 6512604)

Under certain conditions, domains configured with less than one gigabyte of memory and a large number of virtual network devices can either hang or panic with the following stack trace:

```
panic[cpu0]/thread=2a10011fcc0:
recursive_mutex_enter, lp=300058a4ad8 owner=2a10011fcc0
thread=2a10011fcc0

vpanic()
mutex_vector_enter+0x2e0()
vgen_sendmsg+0x3c()
vgen_send_version_negotiate+0x78()
vgen_handshake+0x84()
vgen_reset_hphase+0x178()
vgen_handshake_reset+0x20()
vgen_handshake_retry+4()
vgen_hwatchdog+0x24()
callout_execute+0x98()
taskq_thread+0x1a4()
thread_start+4()
```

Workaround: Reconfigure the logical domain with either more memory or fewer virtual network devices.

Recovery: Apply the workaround, and reboot the domain.

Virtual Disk Server Should Export ZFS 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 which 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` only sees 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, and 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.

Restrictions on Using the I/O MMU Bypass Mode (Bug ID 6517338)

There is currently an option to turn the bypass mode on when adding a I/O bus. (Refer to the description of the `ldm add-io` command in Appendix A of the *Logical Domains (LDoms) 1.0 Administration Guide* or the `ldm` man page for more information.) If you want to change the bypass property of a bus that is already allocated, do one of the following:

- Do an `ldm remove-io` command followed by an `ldm add-io` command *only* while the Logical Domains Manager is running in configuration mode; that is, after booting the `factory-default` configuration.

Note – Attempting to do an `ldm remove-io` command immediately followed by an `ldm add-io` command, so that the `bypass` mode property changes, causes the Logical Domains Manager to terminate and any delayed reconfiguration in process is cancelled.

- If possible, do an `ldm remove-io` command, reboot the domain, then do the `ldm add-io` command.

Note – It might not be possible to reboot the domain after removing an I/O bus, depending on whether that I/O bus contains critical disk or network devices needed to boot.

Need Logical Domains Manager Support for Contiguous Guest Memory Mappings (Bug ID 6517343)

If any domains require contiguous memory mappings, be sure to bind those domains before any other domains get bound and then unbound. Otherwise, memory fragmentation could result in discontinuous memory being assigned to the subject domains.

Note – The Bearer Plane domain in the Netra™ Data Plane Software (NDPS) environment could have this problem. Otherwise, any other sun4v-compliant OS should not have this problem.

System Hangs When Doing a Boot Net Installation With a ZFS File (Bug ID 6517957)

A service domain providing access to a virtual disk whose back end is a Zettabyte File System (ZFS) file may hang when this virtual disk is used by another domain.

Workaround: To prevent this problem, configure the service domain with a sufficient amount of memory. Four gigabytes of memory is shown to be sufficient for most cases.

Certain `ldm` Subcommands Returns Misleading Message If One or More Arguments are Missing (Bug ID 6519049)

A misleading error message is 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, an error message like the following is returned:

```
# 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:

```
# 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
```

Recovery: Refer to Appendix A of the *Logical Domains (LDDoms) 1.0 Administration Guide* or the `ldm` man page for the required arguments of the `ldm` subcommands.

The `vntsd` Might Close Console Connections Unexpectedly (Bug ID 6520018)

Under rare circumstances, when a domain is added or removed, without an intervening attempt to connect to the console through `telnet`, the virtual network terminal server daemon, `vntsd(1M)`, fails to clean up the console state correctly. A subsequent attempt to use this `vntsd` TCP port for a domain will result in the connection being terminated unexpectedly.

Workaround: Restart `vntsd(1M)`:

```
# svcadm enable vntsd
```

Canceling a Delayed Reconfiguration Operation Results in a Misleading Message (Bug ID 6520530)

As part of the normal processing of canceling a delayed reconfiguration operation using the `ldm remove-reconf` command, the Logical Domains Manager exits. The Logical Domains Manager relies on the Service Management Facility (SMF) to restart it. This causes the following message to be returned:

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

Recovery: Verify that the Logical Domains Manager did exit as part of the normal operation of the cancel operation, by looking in the Logical Domains Manager daemon (`ldmd`) log in `/var/svc/log/ldoms-ldmd:default.log` for this message:

```
warning: Exiting to re-sync after cancelling a delayed reconfig  
operation
```

Workaround: Ignore the message.

Service Domain Running vcc and vntsd Can Panic in Certain Situations (Bug ID 6521890)

If an `ldm stop-domain -f` command is used, a subsequent `ldm unbind-domain` command can result in a panic of the service domain running the virtual console concentrator (`vcc`) and the virtual network terminal server daemon (`vntsd`), as it fails to properly close connections as a result of the unbind operation and sends the following message:

```
recursive_mutex_enter in ldc_set_cb_mode
```

Following is a typical stack trace:

```
panic[cpu3]/thread=300096f1620: recursive_mutex_enter, lp=
3000c4f0548 owner=300096f1620 thread=300096f1620

    unix:mutex_vector_enter+188()
    ldc:ldc_set_cb_mode+1c()
    vcc:i_vcc_ldc_fini+dc()
    vcc:i_vcc_close_port+20()
    vcc:vcc_close+f8()
    specfs:spec_close+18c()
    genunix:fop_close+1c()
    genunix:closef+4c()
    genunix:closeandsetf+3a8()
    genunix:close+8()
```

Recovery: Reboot the service domain.

Workaround: Before issuing the `ldm stop-domain -f` command, ensure that all console `telnet` connections to the guest have been closed.

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

In a service domain, disks which are managed by Veritas Dynamic Multipathing (DMP) cannot be exported as virtual disks to other domains. If a disk which is managed by Veritas DMP is added to a virtual disk server (`vd`s) 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, disable Veritas DMP for the disks you want to use as virtual disks.

Rarely, When a Guest Domain Boots, It Fails to Establish a Connection With the Virtual Switch (Bug IDs 6523891 and 6523926)

Under rare conditions, when a guest boots, it fails to establish a connection with the virtual switch (vsw). As a consequence, it is not able to send network packets to the domain containing the virtual switch and, hence, to the outside world.

For example, if the guest is using DHCP to obtain its IP address then you might see a similar message to the following on the guest console:

```
Failed to configure IPv4 DHCP interface(s): vnet0
```

Communications with other guest domains are unaffected. This is only a potential problem when booting a guest domain. Once the guest domain has booted the operating system and is sending network packets, then this is no longer an issue.

Workaround: Stop and unbind the guest domain, and then re-bind and restart it. If this fails to resolve the problem, then stop, unbind the guest domain, and remove and re-add the virtual switch driver before rebinding and restarting the guest domain.

Continuous Console Output From the Control Domain Causes Logical Domains Manager Requests to the SC to Fail (Bug ID 6524255)

If a command is entered on the console that causes a continuous stream of characters to be output, and on another terminal the `ldm add-config` command is entered to store an logical domain configuration on the SC, the Logical Domains Manager might time out and the command would fail. The following message is returned by the Logical Domains Manager when this happens:

```
Error: Operation failed because of an error communicating with the
system controller
```

Recovery: Try the `ldm add-config` command again after the console output completes or is interrupted.

Workaround: The problem only occurs if you are actually connected to the console from the SC, which is accomplished using the `ALOM console` command. Disconnect from the console, using the documented escape sequence of `#.` to return to the SC console, and then the problem is no longer present.

Service Domain Panics If It Fails to Map Pages for a Virtual Disk Backed by a File (Bug ID 6524333)

A service domain can panic if a guest domain is using a virtual disk backed by a file, and the service domain has an I/O error while accessing the file. The service domain crashes only if the following conditions are present:

- A file was created on the device.
- This file was exported as a virtual disk to another domain.
- The other domain is using that file.
- The device where the file is stored has an I/O error while the file is been accessed by the guest domain.

The I/O error is usually caused by an hardware problem; for example, a disk error or storage access problem.

Recovery: Fix the I/O problem in the service domain; for example, change the faulty disk.

OpenBoot Firmware Does Not Handle Logical Domain Channel Resets Properly (Bug ID 6524613)

If virtual disk in a guest domain is not properly configured, then an attempt to boot can take about three minutes to time out. During this time-out period, the system might look hung as no progress indicator is displayed in the console.

CPU Dynamic Reconfiguration Is Disabled for Logical Domains That Have Crypto Units Bound (Bug ID 6525647)

The way that the Solaris Crypto Framework handles CPU dynamic reconfiguration (DR) events that affect crypto units, CPU DR is disabled for all logical domains that have any crypto units bound.

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

Configuring Logical Domains With More than 32 Virtual I/O Devices or Services Can Result in a Hang During Boot (Bug ID 6526280)

When a domain is configured with more than 32 virtual I/O devices or services, some of the devices fail to get configured correctly. When this problem is encountered, and the impacted device is a virtual disk server, it can manifest as a hang during boot.

This problem is not encountered if there are no more than 32 virtual I/O devices or services in a domain. The number of devices in a logical domain can be determined from the device tree using the following command:

```
# prtconf -c /devices/virtual-devices@100/channel-devices@200
```

Recovery: Stop and unbind the domain with more than 32 virtual I/O devices or services. Reconfigure the domain to contain less than 32 virtual I/O devices or services, and bind and start the domain.

Booting Multiple Logical Domains in a Loop Can Cause an Unresponsive Domain (Bug ID 6526814)

When booting multiple logical domains in parallel, one of the logical domains may become completely unresponsive early in the boot process. Attempts to send the `break` command to the logical domain will have no effect, and processing in the domain will not make any forward progress. The logical domain enters this state because it is incorrectly waiting for a message from the Logical Domains Manager that will never arrive.

Recovery: You can force the Logical Domains Manager to send the unresponsive logical domain the type of message it is expecting. Use the `ldm set-var` command to set an Logical Domains variable for the logical domain. The actual value of the variable is not important; for example, you could use the following:

```
# ldm set-var a=b ldom
```

Once the logical domain processes the message generated by the `ldm set-var` command, the boot should continue normally.

System MAC Address is Modified on I/O Domains When Adding a vnet Device (Bug ID 6526868)

When adding the first virtual network (`vnet`) device to a logical domain, the MAC address of the domain, as contained in the system banner, changes. As a further consequence, the host ID of the domain also changes.

Recovery: To perform operations like JumpStart, make sure to specify the MAC address of the interface over which the netboot will occur, and not the system MAC address. In addition, any software that is dependent on the host ID should be configured with the final, resultant host ID after the domain has been completely configured.

Logical Domains Manager May Stop With Invalid Logical Domain Name (Bug ID 6527206)

If the `ldm stop-domain` command has a valid logical domain name followed by an invalid logical domain name, the command fails with the following error message:

```
LDom "<invalid LDom name>" was not found
```

A subsequent command, such as `ldm stop-domain` can cause the Logical Domains Manager to stop, and the command returns the following error message:

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

Recovery: The Logical Domains Manager restarts automatically, and you can retry the command.

Access to a Virtual Disk From a Guest Domain Can Hang If an I/O Operation Is Overlong (Bug ID 6527265)

When a read or write I/O operation on a virtual disk takes a long time to complete, and an `ioctl` is issued on the same virtual disk, it can trigger a virtual disk hang with the following console message:

```
NOTICE: [2] disk access failed.
```

This problem occurs when a disk managed by Solaris I/O Multipathing software is exported as a virtual disk and a storage or path failure forces the multipathing software to switch the disk access to use another path. Under these circumstances a pending read or write I/O operation can take a long time to complete. At the same time, if a `format(1M)` or `prtvtoc(1M)` command is issued on that virtual disk, a hang might result.

Recovery: If the access to the virtual disk hangs, reboot the domain using that virtual disk.

Workaround: Avoid doing `ioctl` operations on disks that have active read or write I/O operations. For example, do not use the `format(1M)` or `prtvtoc(1M)` command on a virtual disk when the disk is mounted and read/write operations are being executed on the disk.

The `ldm start-domain -i xml-file` Command Causes the Logical Domains Manager to Stop and Restart (Bug ID 6527347)

Issuing the `ldm start-domain` with the `-i xml-file` option causes the Logical Domains Manager to stop, resulting in the following error message:

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

The Logical Domains Manager does restart automatically, and the domain is created and started successfully.

Memory Constraints Should Be Enforced and Reporting Improved (Bug ID 6527483)

If the amount of memory allocated to a logical domain does not meet the minimum size as required by the OpenBoot PROM (currently set to 12MB, but could change), the Logical Domains Manager silently increases the allocation to meet the minimum.

Attempt to Store Boot Command Variable During a Reboot Can Time Out (Bug ID 6527622)

When the Solaris OS `reboot(1M)` command is issued to reboot a guest OS, the following messages can appear on the guest console:

```
WARNING: promif_ldom_setprop: ds response timeout  
WARNING: unable to store boot command for use on reboot
```

The reboot proceeds as usual, but all arguments passed to the OpenBoot PROM boot command; that is, arguments which appear after the `--` delimiter of the Solaris OS `reboot(1M)` command, are ignored by the boot code. The same warnings can occur even if no arguments are passed to the `reboot` command, because the system always attempts to store a default boot command.

Recovery: Once this occurs, there is no recovery.

Workaround: To prevent it from happening on future boots, you can do one of the following:

- Use the Solaris OS `eeeprom(1M)` command to set boot-related variables such as `boot-file`.
- Use the `ldm set-variable` command to set boot-related variables for the guest domain.

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 (`forceлип` 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.

Newly Configured Virtual Network Device Fails to Establish a Connection With the Virtual Switch (Bug ID 6528180)

When a new virtual device is added to a logical domain, it can fail to establish a connection with the virtual switch device. This results in loss of network connectivity to and from the logical domain. When this error is encountered, on inspection, it will reveal that the `/dev/vnetN` symbolic link for the virtual network instance is missing.

If present, and not in error, the link points to a corresponding `/devices` entry as shown here:

```
/dev/vnetN -> ../devices/virtual-devices@100/channel-devices@200/  
network@1:vnetN
```

Recovery: Do one of the following:

- Perform a reconfiguration boot of the logical domain, whenever a vnet device is added to the logical domain.
- If the logical domain is already booted, unconfigure the vnet interface, run the `devfsadm(1M)` command, and then reconfigure the interface.

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

If Sun™ 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)eboot, 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 procedures described below.

▼ 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 Solaris™ Express or OpenSolaris™, then the guest domain might not be able to access that virtual disk.

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

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

Rebooting the primary Domain While Other Active or Bound Domains Are Present and Rebooting an I/O Domain Are Not Supported (Bug ID 6529426)

Rebooting the primary domain while there are other active or bound domains, as well as rebooting an I/O domain, is not supported in this release of Logical Domains firmware. If, when in a configuration with multiple active or bound domains, you inadvertently reboot either the primary domain or any I/O domain, or if any such domain panics, the I/O devices owned by that domain are left in an undefined state. In addition, there is a possibility that the following error messages appear during the reboot:

```
WARNING: Unable to connect to Domain Service providers
WARNING: Unable to update ldom variable
WARNING: /pci@7c0/pci@0/pci@1/pci@0/isa@2: Could not open asr
packages
```

Recovery: If this occurs (regardless of whether the above errors are seen), the only recovery is to perform a clean power off and power on of the system. See [“Cleanly Shutting Down and Power Cycling a Logical Domains System”](#) on page 7 for this procedure.

Workaround: If the primary or other I/O domain needs to be reset, you must power cycle the system. See [“Cleanly Shutting Down and Power Cycling a Logical Domains System”](#) on page 7 for this procedure.

Virtual Disk Server Does Not Close Underlying Physical Device or File Correctly (Bug ID 6530040)

On a service domain, a file or a device which has been used as a virtual disk might still appear to be in use by the virtual disk server (vds), although it is not in use and no guest domain is running and bound.

In such a case, the `ldm list` command shows all domains (except the service domain) as inactive; for example:

```
# ldm list
Name      State    Flags  Cons  VCPU  Memory  Util  Uptime
primary   active  -t-cv  SP    4     1G      1.8%  3d 3h 28m
domain1   inactive -----
```

In contrast, the `fuser` command shows that devices are still in use by the vds driver; for example:

```
# fuser /dev/dsk/c2t42d3s2
/dev/dsk/c2t42d3s2: [vds, dev_path=/virtual-devices@100/
channel-devices@200/virtual-disk-server@0]
```

Recovery: Reboot the system. See [“Cleanly Shutting Down and Power Cycling a Logical Domains System”](#) on page 7 for more information.

Fault Management Daemon in the `primary` Domain Appears to Hang If You Do Not Enable the Logical Domains Manager Daemon (Bug ID 6530948)

If you upgrade the Logical Domains firmware, and you do not enable the Logical Domains Manager daemon, `ldmd`, the Solaris 10 OS Fault Management daemon `fmd(1M)`, on the `primary` domain appears to hang. The Fault Management daemon attempts to communicate with the Logical Domains Manager daemon, which does not exist or is down. Since each request has a 20-minute time-out, it looks like a hang. If your machine is FMA-clean, you might not experience the problem.

Page Retirement Is Not Persistent in the Logical Domains Environment (Bug IDs 6531030 and 6531058)

When a memory page of the `primary` domain is diagnosed as faulty, the Logical Domains Manager retires the page. The `fmd` command fails to obtain the page status and does not replay the page fault.

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

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.

Loopback Devices Cannot Be Used as Virtual Disks (Bug ID 6532144)

Exporting loopback (`lofi`) devices as virtual disks is not supported by Logical Domains 1.0 software. Exporting loopback devices can result in unexpected behavior including a system panic. The virtual disk server should be configured to export a disk image file directly.

Serial Port Cannot Be Used or Opened (Bug ID 6532334)

If you see an error message similar to the following during boot, your serial port will be unusable:

```
WARNING: /pci@7c0/pci@0/pci@1/pci@0/isa@2: Could not open asr
package.
```

There will be no corresponding device in either the OpenBoot device tree or the Solaris device tree, and no serial device drivers will be attached.

Recovery: Reset or power cycle the system.

Second PCI Path Does Not Persist Across primary Domain Reboots (Bug ID 6532604)

The Logical Domains Manager does not persist physical I/O constraints in its constraint database. As a result, if the Logical Domains Manager restarts, then logical domains in the inactive state have lost any previously specified physical I/O constraints.

Recovery: Re-add the constraint.

Using the `server-secure.driver` With an NIS-enabled System, LDomS or No LDomS (Bug ID 6533696)

On a system configured to use the Network Information Services (NIS) or NIS+ name service, if the Solaris™ Security Toolkit 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.
```

This is true whether the Solaris Security Toolkit is applied indirectly through the `ldm-install` script menu options or applied directly using this command:

```
# /opt/SUNWjass/bin/jass-execute -d server-secure.driver
```

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.

As an alternative, NIS and NIS+ do work with the `server-secure.driver` by following the steps at the end of this file:

```
/opt/SUNWjass/Files/etc/ipf/ipf.conf-server
```

Here are those steps quoted from the end of the file:

```
"If you are using NIS as your name service, you will need to allow name resolution to pass through your firewall. This is not possible with only ipf.conf, since NIS is an rpc service without a fixed port. Instead, use the proxy in ipnat to redirect rpc traffic, with rules like
```

```
map eri0 0/0 -> 192.1.1.1/32 proxy port 111 rpcbu/udp
```

```
in file /etc/ipf/ipnat.conf (replace "eri0" with your network adapter instance and "192.1.1.1" with the adapter's IP address, both from "ifconfig -a")."
```

If the `server-secure.driver` is used on a system configured to use NIS or NIS+, you might fail to log in because either there are no local user accounts or the superuser account requires a password change which cannot be made because NIS or NIS+ fails. If this occurs, you must reset your system, and you do lose your logical domains configuration.

▼ To Reset Your System

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

```
# #.
```

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

```
sc> poweroff
```

3. Power on the system.

```
sc> poweron
```

4. Switch to the console mode at the `ok` prompt:

```
sc> console
```

5. Boot the system to single user mode:

```
ok boot -s
```

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

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

7. You can now 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 -u  
# /opt/SUNWjass/bin/jass-execute -a ldm_control-secure.driver
```

Cannot Forcibly Stop a Domain If the Domain Has Any Bound PCI Buses (Bug ID 6536420)

The `ldm stop-domain -f` command is disabled if the domain has any PCI-Express I/O buses bound to it. This is currently the case with all the platforms supported by Logical Domains 1.0 software. In this case, an error message of the following form is returned:

```
LdDom ldg1 stop notification failed
```

Due to this restriction, if an I/O domain is unresponsive to console or network input, and is unable to process a domain service shutdown request from the Logical Domains Manager, then there is no way to perform an isolated stop of that domain.

Recovery: Shut down all the other domains, and power cycle the server.

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.

XML Output of the `list-constraints` Subcommand Is Missing the IO: Configuration Value (Bug ID 6537156)

The `ldm list-constraints -x ldom` command fails to include physical I/O information in its XML output.

Workaround: Currently, the only way to add a physical I/O device is by using the `ldm add-io` or `ldm set-io` commands.

MAC Addresses Are Not Listed in the Correct Format in XML Input and Output Files (Bug ID 6537172)

When using the `ldm list-constraints` command with the `-x` option, all virtual network (`vnet`) or virtual switch (`vsw`) devices that had their MAC addresses manually specified will have those addresses incorrectly formatted in the XML output. Any subsequent attempt to create devices using this same XML syntax will fail to create the specified `vnet` or `vsw` devices.

Recovery: Verify that the MAC address in the XML file is in colon format (`xx:xx:xx:xx:xx:xx`) format and not in hexadecimal format (`0xxxxxxxxxxxxx`). Edit the XML file as necessary to correct the MAC address.

SC Command `bootmode bootscript` Does Not Work in an LDom Environment (Bug ID 6540125)

The following system controller (SC) command only works when running on the `factory-default` system configuration:

```
bootmode bootscript="script"
```

Where *script* is an OpenBoot command string that is run during OpenBoot firmware initialization.

When using a system configuration created by the Logical Domains Manager, the above SC command has no effect.

OpenBoot `power-off` Command Behavior With Logical Domains 1.0 Software (Bug ID 6540632)

The following table shows the expected behavior for the OpenBoot `power-off` command with Logical Domains 1.0 software.

TABLE 1 Behavior of OpenBoot `power-off` Command With LDOMs Software

Domain	Domaining Enabled?	Behavior of OpenBoot <code>power-off</code> Command
Control	Disabled	Domain powers off, and stays off until the <code>power-on</code> command is executed on the SC.
Control	Enabled, no guest running (all in stop/bound state)	Domain powers off, and stays off until the <code>power-on</code> command is executed on the SC.
Control	Enabled, multiple guests in bound or active state	Domain resets, and starts again.
Guest	N/A	Domain stops with no console activities. This is the equivalent of stopping a guest using the <code>ldm stop-domain</code> command. The guest remains in a bound state.

When the `ldm add-vnet` Command Fails, the Error Message Is Not Clear (Bug ID 6541158)

When the system shows you the following error message, it means that either your network is not set up correctly on that system, or the Logical Domains Manager daemon (`ldmd`) tried to open up a socket prior to your network fully running:

```
Automatic MAC allocation not available
```

Recovery: Check to ensure that your network is up and running. If so, stop and restart `ldmd` by using the following Solaris 10 OS Service Management Facility (SMF) commands:

```
# svcadm disable ldmd
# svcadm enable ldmd
```

Machine Description Too Large for Logical Domains Manager's Memory Allocator (Bug ID 6541171)

When attempting to bind a new domain or reconfigure an existing bound or active domain, you may encounter a failure that manifests itself with the error:

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

This can occur if the machine description (MD) the Logical Domains Manager builds describing the current configuration of the system turns out to be too large for the memory space allocated. When this occurs, the Logical Domains Manager terminates rather than sending an MD to the hypervisor that cannot be instantiated. Also, the Logical Domains Manager log will contain a message of this form:

```
fatal error: Configuration required a MD of size 0x814c0 bytes  
which is larger than the maximum supported size of 0x78000 bytes
```

Recovery: Scale back on either the number of virtual I/O devices or the number of domains configured in the system.

Do Not Use an `ldm set-vnet` Command to Change a MAC Address on an Active Domain (Bug ID 6541284)

Do not attempt to issue an `ldm set-vnet` command to modify the MAC address of the virtual network (`vnet`) device on an active logical domain. The command appears to succeed, but the change is not fully enacted until the domain reboots, and until then, networking over that device could fail to work.

Workaround: Stop the domain, perform an `ldm set-vnet` command, and start the domain.

Occasionally, the Logical Domains Manager Does Not Send a Warning That a Disk Is in Use (Bug ID 6541323)

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 (vds), 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.

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 might not get logged. Consider the following example:

In a split-PCI configuration the primary domain contains Bus B, and Bus A is not assigned to any 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 may be lost.

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

Forced Stop of a Logical Domain Can Result in a Panic on the Next Boot (Bug ID 6543418)

When using the force (-f) option to the `ldm stop-domain` command, the domain could panic with one of the following two signatures the next time the domain is started:

- `panic[cpu0]/thread=180e000: cpu1 failed to start (2)`
- `panic[cpu0]/thread=180e000: bad kernel MMU miss at TL 2`

Recovery: The condition that causes the panic is not persistent. After the panic, the logical domain can be restarted normally.

Workaround: Whenever possible, do not use the force (-f) option to the `ldm stop-domain` command. If it is absolutely necessary to use the force option, unbind the logical domain before restarting it.

Adding a Virtual Disk to a Bound Guest Does Not Take Effect (Bug ID 6543735)

The `ldm add-vdisk` command issued while an logical domain is in the bound state does not take effect, even though the `ldm` command appears to succeed, and the virtual disk (`vdisk`) appears in a subsequent `ldm list-domain` command.

Recovery: To make the added virtual disk visible to the system, issue the `ldm unbind-domain ldom` command followed by the `ldm bind-domain ldom` command, where `ldom` is the domain to which the `vdisk` was added

Workaround: Issue the `ldm add-vdisk` command only to a domain that is in the inactive state.

Avoid Using the OpenBoot `watch-net-all` Command on a Domain That Contains a Virtual Network (Bug ID 6543748)

Running the OpenBoot command `watch-net-all` on any domain that contains a virtual network (`vnet`) can result in one of the following errors:

- `ERROR: Last Trap: Memory Address not Aligned`
- `ERROR: Last Trap: Fast Data Access MMU Miss`

Recovery: Reset the domain before continuing. If the domain contains physical I/O devices, power-cycle the domain before continuing.

Workaround: Avoid using the command `watch-net-all` on domains that contain virtual networks.

During `wanboot` or `waninstall`, Miniroot Download Time Can Increase Significantly (Bug ID 6543749)

During `wanboot` or `waninstall`, the time it takes to download the miniroot can increase significantly when booting from a virtual network (`vnet`) device. Early tests showed miniroot download to be 5 to 6 times slower on a guest domain.

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 # 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.

Logical Domain With No Virtual I/O Devices Might Not Respond (Bug ID 6544197)

A logical domain that is not the `primary` domain and has no virtual I/O devices might fail to respond to configuration change messages from the Logical Domains Manager. A domain in this state also might not support such services as dynamic reconfiguration (DR) on a CPU or domain shutdown requests. This situation, in most cases, is encountered only by an I/O domain, because only a domain with physical I/O devices (that is, an IO domain) is likely to be configured with no virtual I/O devices.

A domain in a bound or running state encounters this situation if the domain is subsequently unbound and rebound, and there was an intervening Logical Domains Manager restart. The following conditions induce an Logical Domains Manager restart:

- Explicit administrator action; for example, using the following SMF command:

```
# svcadm restart svc:/ldoms/ldmd
```

- Logical Domains Manager process failure
- Control domain reboot
- System power cycle

One way to tell if the domain is in this degraded state is by determining if the Logical Domains Manager has assigned logical domain channel (LDC) 0 to the domain's console. You can obtain this information by issuing the following command:

```
# ldm list-bindings | grep Vcons  
Vcons: [via LDC:0]
```

If the output shows that the console has been allocated LDC 0 (as shown in the preceding example), the situation has been triggered.

Recovery: Once in this state, the only recovery is to destroy and then recreate the domain. When recreating the domain, also be sure to apply the workaround described below to prevent any future recurrence.

Workaround: Always make sure each domain includes at least one virtual I/O client before it is first bound. If you have no need for a virtual I/O client, you can create a fictitious disk device as shown in the following procedure.

▼ To Create a Fictitious Disk Device

Note – If your configuration already includes a virtual disk service, skip Step 1 and go to Step 2.

1. **Use the Logical Domains Manager to create a virtual disk service on one of your logical domains.**

For example:

```
# ldm add-vds primary-vds0 primary
```

2. On the domain hosting the virtual disk server, create a disk file to export.

For example:

```
# mkfile 10m /path/to/fictitious-disk-file
```

3. Export the file as a virtual disk.

For example:

```
# ldm add-vdsdev /path/to/fake-disk-file fake-volume@primary-vds0
```

4. Create the fictitious virtual disk device on the domain lacking any virtual I/O client.

For example:

```
# ldm add-vdisk fake-disk fake-volume@primary-vds0 no-vio-domain
```

Logical Domains Manager Incorrectly Assumes All vcc Devices Use the Same Port Range If the Logical Domains Manager Database Does Not Exist (Bug ID 6544772)

If the Logical Domains Manager database is deleted or otherwise lost, and the configuration includes virtual consoles that are bound to TCP ports outside the default range of 5000-5100, the Logical Domains Manager refuses to start.

Recovery: If the Logical Domains Manager is still running when the database is lost, perform any reconfiguration operation (for example, create a fictitious LDom variable) to cause a new database file to be created. If the Logical Domains Manager is not running, you must revert to the `factory-default` configuration and re-create the previous operating configuration to re-synchronize the Logical Domains Manager database with the configuration.

Workaround: To prevent this problem from occurring on the loss of the Logical Domains Manager database, follow both of these restrictions:

1. Do not configure any virtual console concentrator (vcc) device to allocate TCP ports outside the 5000-5100 range.
2. Do not configure more than one virtual console concentrator service per logical domain.

Upgrading the Solaris OS Could Cause Problems With Virtual Networking (Bug ID 6544866)

If you upgrade the Solaris OS image to a later version on any of the logical domains of an Logical Domains 1.0 system, problems with virtual networking might result. For example, if you execute an `ldm add-config` command with active domains after already booting with the active domains configured, the next time the system is restarted, those active guest domains will not have network support.

Recovery: Return to the last good logical domain configuration on the SC, unbind all domains, and rebind them. Then, use the `ldm add-config` command to set the new configuration.

Note – If you have a virtual network installed on the primary domain, this recovery procedure does not work. In this situation, you must remove and re-install the virtual network.

Workaround: None. However, always check for and install the latest Logical Domains Manager patches before upgrading the Solaris OS on a domain.

Adding Non-Existent Disk Device to Single-CPU Domain Causes Hang (Bug ID 6544946)

When a guest domain is configured to a virtual disk that is backed by a non-existent storage device, the domain can hang either during reconfiguration boot or running the `devfsadm(1M)` command. The error is encountered because the virtual disk driver fails to detach properly following a attach failure.

Workaround: Add more than one CPU to the domain

Recovery: Unconfigure or replace the non-existent disk device with a valid disk device and reboot the domain.

Virtual Disk Can Lose a Label When Rebinding a Domain (Bug ID 6544963)

In some cases, when a file is used as a virtual disk, the label of that virtual disk can be lost when rebinding a domain (`ldm bind-domain` command) using that file (or a copy of that file) as a virtual disk.

Workaround: To prevent this problem, run the file checksum (fcksum) script on any file that has been used as a virtual disk and for which the disk label or disk partitioning has been changed. The fcksum script follows:

CODE EXAMPLE 1 File Checksum (fcksum) Script

```
#!/bin/ksh

file=$1

if [ -z "$file" ]; then
    echo "usage: fcksum <file>"
    echo
    exit 1

fi

if [ ! -f "$file" ]; then
    echo "usage: fcksum <file>"
    echo
    echo "<file> should be a regular file"
    echo
    exit 1

fi

backup=label.$file.`date +%y%m%d_%H%M%S`
echo "Backing up original label in $backup"
dd if=$file of=$backup count=1 2>/dev/null

(
cat <<EOM
*0x1fe>/s/c
*0x1b8>/s/d
##(<c&0x8000>t
<c&0x7fff>c
<d^0x8000>d
.,<t="Changing checksum"
0x1fe,<t/w <c
.,<t="Changing dummy field"
0x1b8,<t/w <d
.,<t="Label checksum has been updated"
.,#<t="Label checksum is okay"
EOM
) | mdb -w $file
```

Note – It is easier to cut and paste this script from an HTML file than a PDF file. Both formats of these release notes are available at the web site specified in [“Location of Documentation” on page 2.](#)

The `fcksum` script checks to see whether the label will be correctly validated during the next `ldm bind-domain` command, and if not, the script will change the label and its checksum so that it can be correctly validated.

Run the script right after the domain using the virtual disk is unbound (`ldm unbind-domain` command) for the first time.

For example, if the *file-name* file is used by a domain as a virtual disk, and if the Solaris system is being installed onto that virtual disk, then run the script after the first `ldm unbind-domain` command on that domain.

Note – Run the `fcksum` script before any `ldm bind-domain` command on a domain; otherwise, you can lose the disk label.

▼ To Check If a Label Is Valid

- **Run the `fcksum` script on any file that has been used as a virtual disk and for which the disk label or disk partitioning has been changed.**

```
$ ./fcksum file-name
```

The script first backs up the existing label of the file in a file named `label.file.name.day_time`. Then one of the following occurs:

- If the script updates the label, you receive this output:

```
Backing up original label in label.file.070314_201917
Changing checksum
0x1fe:          0xe456 =          0x6456
Changing dummy field
0x1b8:          0      =          0x8000
Label checksum has been updated
```

- If the script does not update the label, you receive this output:

```
Backing up original label in label.file.070314_201005
Label checksum is okay
```

Watchdog Time-out Can Be Triggered by Heavy Network Loads (Bug ID 6545470)

A system in which the virtual switch has been configured to use the bge network interface can trigger the watchdog time-out under heavy network load conditions. This often happens when the CPU count in guest domains running network intensive workloads is significantly larger than the number of CPUs in the service domain.

Even though watchdog time-outs do not cause a system to reset, the system does become progressively more non-responsive. A message similar to the following may also appear on the console:

```
APR 19 17:05:47 ERROR: Watchdog timeout ignored because user is
running on a Logical Domains Configuration
```

If the watchdog message is seen, or you want to run network intensive loads in the guest domain, apply the following workaround. However, note that doing so may result in a slight degradation of network performance under certain loads.

Workaround: Set the following in the `/etc/system` file, and reboot the service domain.

```
set vsw_chain_len=20
```

Recovery: Apply the workaround, and power cycle the system.

Logical Domains Manager Not Propagating factory-default TOD-offset (Bug ID 6546491)

When the primary domain is reconfigured from its `factory-default` settings, the domain's time might be incorrectly reset when the domain is rebooted into the new configuration. This change can cause various problems, including delayed or improper Fault Management Architecture (FMA) fault diagnosis.

Workaround: Before changing the factory configuration, explicitly set the date before the new configuration is stored on the SC. The following Solaris command, which sets the date to its current value, is sufficient to avoid this situation:

```
% date `date +%H%M.%S`
```

Virtual Switch Requires e1000g Driver Update Following Solaris 10 Upgrade (Bug ID 6547846)

Sun Fire T1000 and T2000 servers with Intel PCI-Express network interfaces, and installed with early Solaris 10 OS releases, might have been configured to use the `ipge` Ethernet driver. This driver was a temporary support mechanism for these interfaces in the early releases of the Solaris 10 OS and has now been superseded by the Sun standard GLDv3-compliant `e1000g` driver. Additionally, in a LDoms environment the virtual switch requires that it be configured to use the GLDv3-compliant `e1000g` driver.

Systems freshly installed with the Solaris 10 11/06 OS at a minimum are automatically configured to exclusively use the `e1000g` driver. However, when these systems are upgraded from an earlier Solaris 10 OS release to use the Solaris 10 11/06 OS, you must manually convert these systems to use the Sun standard `e1000g` network driver.

Refer to SunSolve Doc ID 102502 for more information on enabling the systems to use `e1000g` drivers.

If you do not replace the `ipge` driver with the `e1000g` driver and configure the virtual switch to use the `e1000g` device, you will lose network connectivity to guest domains. A warning message similar to the following might appear on the console of the service domain.

```
WARNING: mac_open ipge0 failed
```

Recovery: Update the system to use the `e1000g` driver. Configure the virtual switch to use the `e1000g` driver.

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.

Documentation Errata

This section of the release notes contains errors in the Logical Domains 1.0 documentation.

The `vntsd(1M)` Man Page Is Missing Information About Using the Double Tilde (`~~`) (Bug ID 6513007)

The Solaris 10 11/06 OS `vntsd(1M)` man page is missing information about the use of the double tilde (`~~`). A tilde (`~`) appearing as the first character of a line is an escape signal that directs the console to perform a special console command. When connected to the console using `telnet` from within another `telnet` session, use the tilde-tilde (`~~`) sequence to output a tilde to the domain's console.

