



Sun Cluster 3.0 Error Messages Guide

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303-4900
U.S.A. 650-960-1300

Part Number 806-1426
November 2000, Revision A

Copyright Copyright 2000 Sun Microsystems, Inc. 901 San Antonio Road, Palo Alto, California 94303-4900 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

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 other countries, exclusively licensed through X/Open Company, Ltd. For Netscape Communicator™, the following notice applies: (c) Copyright 1995 Netscape Communications Corporation. All rights reserved.

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, and Solaris are trademarks, registered trademarks, or service marks 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 an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a).

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 2000 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, Californie 94303 Etats-Unis. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

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. La notice suivante est applicable à Netscape Communicator™: (c) Copyright 1995 Netscape Communications Corporation. Tous droits réservés.

Sun, Sun Microsystems, le logo Sun, AnswerBook2, docs.sun.com, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, 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.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REPENDRE A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



Contents

	Preface	5
1.	Introduction	9
2.	Message IDs 100000 - 199999	11
	Error Message List	11
3.	Message IDs 200000 - 299999	41
	Error Message List	41
4.	Message IDs 300000 - 399999	69
	Error Message List	69
5.	Message IDs 400000 - 499999	97
	Error Message List	97
6.	Message IDs 500000 - 599999	125
	Error Message List	125
7.	Message IDs 600000 - 699999	155
	Error Message List	155
8.	Message IDs 700000 - 799999	181
	Error Message List	181
9.	Message IDs 800000 - 899999	205
	Error Message List	205
10.	Message IDs 900000 - 999999	237

Preface

The *SunTM Cluster 3.0 Error Messages Guide* contains a list of error messages that might be seen on the console or in the syslog file while running Sun Cluster. For each error message, there is an explanation and a suggested solution.

This document is intended for experienced system administrators with extensive knowledge of Sun software and hardware.

The instructions in this book assume knowledge of the SolarisTM operating environment and expertise with the volume manager software used with Sun Cluster.

Using UNIX Commands

This document contains information on error messages generated by Sun Cluster. This document may not contain information on basic UNIX[®] commands and procedures such as shutting down the system, booting the system, and configuring devices.

See one or more of the following for this information:

- AnswerBook2TM online documentation for the Solaris software environment
- Other software documentation that you received with your system
- Solaris operating environment man pages

Typographic Conventions

Typeface or Symbol	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be superuser to do this.
	Command-line variable; replace with a real name or value	To delete a file, type <code>rm filename</code> .

Shell Prompts

Shell	Prompt
C shell	<i>machine_name%</i>
C shell superuser	<i>machine_name#</i>
Bourne shell and Korn shell	\$
Bourne shell and Korn shell superuser	#

Related Documentation

Application	Title	Part Number
Software Installation	<i>Sun Cluster 3.0 Installation Guide</i>	806-1419
Hardware Installation	<i>Sun Cluster 3.0 Hardware Guide</i>	806-1420
Data Services	<i>Sun Cluster 3.0 Data Services Installation and Configuration Guide</i>	806-1421
API Development	<i>Sun Cluster 3.0 Data Services Developers' Guide</i>	806-1422
Administration	<i>Sun Cluster 3.0 System Administration Guide</i>	806-1423
Concepts	<i>Sun Cluster 3.0 Concepts</i>	806-1424
Release Notes	<i>Sun Cluster 3.0 Release Notes</i>	806-1428

Sun Documentation on the Web

The docs.sun.comSM web site enables you to access Sun technical documentation on the Web. You can browse the docs.sun.com archive or search for a specific book title or subject at:

<http://docs.sun.com>

Getting Help

If you have problems installing or using Sun Cluster, contact your service provider and provide the following information:

- Your name and email address (if available)
- Your company name, address, and phone number

- The model and serial numbers of your systems
- The release number of the operating environment (for example, Solaris 7)
- The release number of Sun Cluster (for example, Sun Cluster 3.0)

Use the following commands to gather information on your system for your service provider:

Command	Function
<code>prtconf -v</code>	Displays the size of the system memory and reports information about peripheral devices
<code>psrinfo -v</code>	Displays information about processors
<code>showrev --p</code>	Reports which patches are installed
<code>prtdiag -v</code>	Displays system diagnostic information
<code>scinstall -pv</code>	Displays Sun Cluster release and package version information

Also have available the contents of the `/var/adm/messages` file.

Introduction

The chapters in this book provide a list of error messages that can appear on the consoles of cluster members while running Sun Cluster software. Each message includes the following information:

- Message ID

The message ID is an internally-generated ID that uniquely identifies the message.

- Description

The description is an expanded explanation of the error that was encountered including any background information that might aid you in determining what caused the error.

- Solution

The solution is the suggested action or steps that you should take to recover from any problems caused by the error.

The Message ID is a number ranging between 100000 and 999999. The chapters are divided into ranges of Message IDs. Within each chapter, the messages are ordered by Message ID. The best way to find a particular message is by searching on the Message ID.

Throughout the messages, you will see `printf(1)` formatting characters such as `%s` or `%d`. These characters will be replaced with a string or a decimal number in the displayed error message.

Message IDs 100000 - 199999

Error Message List

The following list is ordered by the message ID.

100088 :fatal: Got error <%d> trying to read CCR when making resource group <%s> managed; aborting node

Rgmd failed to read updated resource from the CCR on this node.

Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

100293 :dl_bind: kstr_msg failed %d error

Could not bind to the private interconnect.

Reboot of the node might fix the problem.

100396 :clxecd: unable to arm failfast.

clxecd problem could not enable one of the mechanisms which causes the node to be shutdown to prevent data corruption, when clxecd program dies.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

100555 :libsecurity: getnetconfigent error: %s

A client of the rpc.pmfd, rpc.fed or rgmd server was not able to initiate an rpc connection, because it could not get the network information. The pmfadm or

scha command exits with error. The rpc error is shown. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

102218 :couldn't initialize ORB, possibly because machine is booted in non-cluster mode

could not initialize ORB.

Please make sure the nodes are booted in cluster mode.

102340 :Prog <%=> step <%=>: authorization error.

An attempted program execution failed, apparently due to a security violation; this error should not occur. This failure is considered a program failure.

Correct the problem identified in the error message. If necessary, examine other syslog messages occurring at about the same time to see if the problem can be diagnosed. Save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance in diagnosing the problem.

104165 :Resource <%=> of Resource Group <%=> failed pingpong check on node <%=>.

Message logged for failed scha_control pingpong check methods on specific node.

No user action required.

104914 :CCR: Failed to set epoch on node %s errno = %d.

The CCR was unable to set the epoch number on the indicated node. The epoch was set by CCR to record the number of times a cluster has come up. This information is part of the CCR metadata.

There may be other related messages on the indicated node, which may help diagnose the problem, for example: If the root file system is full on the node, then free up some space by removing unnecessary files. If the root disk on the afflicted node has failed, then it needs to be replaced.

105337 :WARNING: thr_getspecific %d

The rgmd has encountered a failed call to thr_getspecific(3T). The error message indicates the reason for the failure. This error is non-fatal.

If the error message is not self-explanatory, contact your authorized Sun service provider for assistance in diagnosing the problem.

106181 :WARNING: lkcm_act: %d returned from udlm_rcv_message (the error was successfully masked from upper layers).

Unexpected error during a poll for dlm messages.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

108357 :lookup: unknown binding type <%d>

During a name server lookup an unknown binding type was encountered.

No action required. This is informational message.

108990 :CMM: Cluster members: %s.

This message identifies the nodes currently in the cluster.

This is an informational message, no user action is needed.

109102 :%s should be larger than %s.

The value of Thorough_Probe_Interval specified in scrgadm command or in CCR table was smaller than Cheap_Probe_Interval.

Reissue the scrgadm command with appropriate values as indicated.

109105 :(%s) setitimer failed: %d: %s (UNIX errno %d)

Call to setitimer() failed. The "setitimer" man page describes possible error codes. udlmctl will exit.

Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

110012 :lkcm_dreg failed to communicate to CMM ... will probably failfast: %s

Could not deregister udlm from ucmm. This node will probably failfast.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

110097 :Major number for driver (%s) does not match the one on other nodes.

The driver identified in this message does not have the same major number across cluster nodes, and devices owned by the driver are being used in global device services.

Look in the /etc/name_to_major file on each cluster node to see if the major number for the driver matches across the cluster. If a driver is missing from

the /etc/name_to_major file on some of the nodes, then most likely, the package the driver ships in was not installed successfully on all nodes. If this is the case, install that package on the nodes that don't have it. If the driver exists on all nodes but has different major numbers, see the documentation that shipped with this product for ways to correct this problem.

111527 :Method <%s> on resource <%s>: unknown command.

An internal logic error in the rgmd has prevented it from successfully executing a resource method.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

111697 :Failed to delete scalable service in group %s for IP %s Port %d%c%s: %s.

A call to the underlying scalable networking code failed.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

114036 :clxeccd: Error %d from putmsg

clxeccd program has encountered a failed putmsg(2) system call. The error message indicates the error number for the failure.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

114550 :Unable to create <%s>: %s.

The HA-NFS stop method attempted to create the specified file but failed.

Check the error message for the reason of failure and correct the situation. If unable to correct the situation, reboot the node.

115256 :file specified in USER_ENV %s doesn't exist

"User_env" property was set when configuring the resource. File specified in "User_env" property does not exist or is not readable. File should be specified with fully qualified path.

Specify existing file with fully qualified file name when creating resource. If resource is already created, please update resource property "User_env".

115461 :in libsecurity __rpc_get_local_uid failed

A server (rpc.pmfd, rpc.fed or rgmd) refused an rpc connection from a client because it failed the Unix authentication, because it is not making the rpc call over the loopback interface. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

115987 :execvp: %s

The rpc.pmfd server was not able to exec a new process, possibly due to bad arguments. The message contains the system error. The server does not perform the action requested by the client, and an error message is output to syslog.

Investigate that the file path to be executed exists. If all looks correct, save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

117498 :scha_resource_get error (%d) when reading extension property %s

Error occurred in API call scha_resource_get.

Check syslog messages for errors logged from other system modules. Stop and start fault monitor. If error persists then disable fault monitor and report the problem.

118046 :rebalance: no primary node could be found for resource group <%s>.

The rgmd is unable to bring the resource group online because all of its potential masters are down.

Repair and reboot broken nodes so they may rejoin the cluster; or use scrgadm(1M) to edit the Nodelist property of the resource group so that it includes nodes that are cluster members.

119120 :clconf: Key length is more than max supported length in clconf_ccr read

In reading configuration data through CCR, found the key length is more than max supported length.

Check the CCR configuraton information.

119649 :clcomm: Unregister of pathend state proxy failed

The system failed to unregister the pathend state proxy.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

120470 :(%s) t_sndudata: tli error: %s

Call to t_sndudata() failed. The "t_sndudata" man page describes possible error codes. udlmctl will exit.

Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

121858 :tag %s: not suspended, cannot resume

The user sent a resume command to the rpc.fed server for a tag that is not suspended. An error message is output to syslog.

Check the tag name.

123526 :Prog <%s> step <%s>: Execution failed: no such method tag.

An internal error has occurred in the rpc.fed daemon which prevents step execution. This is considered a step failure.

Examine other syslog messages occurring at about the same time to see if the problem can be identified. Save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance in diagnosing the problem. Re-try the edit operation.

124232 :clcomm: solaris xdoor fcntl failed: %s

A fcntl operation failed. The "fcntl" man page describes possible error codes.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

124810 :fe_method_full_name() failed for resource <%s>, resource group <%s>, method <%s>

Due to an internal error, the rgmd was unable to assemble the full method pathname. This is considered a method failure. Depending on which method was being invoked and the Failover_mode setting on the resource, this might cause the resource group to fail over or move to an error state.

Examine other syslog messages occurring at about the same time to see if the problem can be identified. Save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance in diagnosing the problem.

125159 :Load balancer setting distribution on %s:

The load balancer is setting the distribution for the specified service group.

This is an informational message, no user action is needed.

126142 :fatal: new_str strcpy: %s (UNIX error %d)

The rgmd failed to allocate memory, most likely because the system has run out of swap space. The rgmd will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

The problem is probably cured by rebooting. If the problem recurs, you might need to increase swap space by configuring additional swap devices. See `swap(1M)` for more information.

126318 :fatal: Unknown object type bound to %s

The low-level cluster machinery has encountered a fatal error. The rgmd will produce a core file and will cause the node to halt or reboot to avoid the possibility of data corruption.

Save a copy of the `/var/adm/messages` files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

126467 :HA: not implemented for userland

An invocation was made on an HA server object in user land. This is not currently supported.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

127182 :fatal: thr_create returned error: %s (UNIX error %d)

The rgmd failed in an attempt to create a thread. The rgmd will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

Fix the problem described by the UNIX error message. The problem may have already been corrected by the node reboot.

127411 :Error in reading /etc/mnttab: getmntent() returns <%d>

Failed to read `/etc/mnttab`.

Check with system administrator and make sure `/etc/mnttab` is properly defined.

127624 :must be superuser to start %s

Process `ucmmd` did not get started by superuser. `ucmmd` is going to exit now.

None. This is an internal error.

130822 :CMM: join_cluster: failed to register ORB callbacks with CMM.

The system can not continue when callback registration fails.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

131492 :pxvfs::mount(): global mounts are not enabled (need to run "clconfig -g" first)

A global mount command is attempted before the node has initialized the global file system name space. Typically this caused by trying to perform a global mount while the system is booted in single user mode.

If the system is not at run level 2 or 3, change to run level 2 or 3 using the init(1M) command. Otherwise, check message logs for errors during boot.

132032 :clexecd: strdup returned %d. Exiting.

clexecd program has encountered a failed strdup(3C) system call. The error message indicates the error number for the failure.

If the error number is 12 (ENOMEM), install more memory, increase swap space, or reduce peak memory consumption. If error number is something else, contact your authorized Sun service provider to determine whether a workaround or patch is available.

134417 :Global service <%s> of path <%s> is in maintainance.

Service is not supported by HA replica.

Resume the service by using scswitch(1m).

135918 :CMM: Quorum device %ld (%s) added; votecount = %d, bitmask of nodes with configured paths = 0x%llx.

The specified quorum device with the specified votecount and configured paths bitmask has been added to the cluster. The quorum subsystem treats a quorum device in maintenance state as being removed from the cluster, so this message will be logged when a quorum device is taken out of maintenance state as well as when it is actually added to the cluster.

This is an informational message, no user action is needed.

137294 :method_full_name: strdup failed

The rgmd server was not able to create the full name of the method, while trying to connect to the rpc.fed server, possibly due to low memory. An error message is output to syslog.

Investigate if the host is running out of memory. If not save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

137606 :clcomm: Pathend %p: disconnect_node not allowed

The system maintains state information about a path. The `disconnect_node` operation is not allowed in this state.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

138972 :could not set timeout: %s

A client was not able to make an rpc connection to a server (`rpc.pmfd`, `rpc.fed` or `rgmd`) because it could not set the rpc call timeout. The rpc error is shown. An error message is output to syslog.

Save the `/var/adm/messages` file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

139773 :clxecd: Error %d from strdup

`clxecd` program has encountered a failed `strdup(3C)` system call. The error message indicates the error number for the failure.

If the error number is 12 (`ENOMEM`), install more memory, increase swap space, or reduce peak memory consumption. If error number is something else, contact your authorized Sun service provider to determine whether a workaround or patch is available.

140225 :The request to relocate resource %s completed successfully.

The resource named was relocated to a different node.

This is an informational message, no user action is needed.

141062 :Failed to connect to host %s and port %d: %s.

An error occurred while fault monitor attempted to probe the health of the data service.

Wait for the fault monitor to correct this by doing restart or failover. For more error description, look at the syslog messages.

141236 :Failed to format stringarray for property %s from value %s.

The validate method for the scalable resource network configuration code was unable to convert the property information given to a usable format.

Verify the property information was properly set when configuring the resource.

141242 :HA: revoke not implemented for replica_handler

An attempt was made to use a feature that has not been implemented.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

141970 :in libsecurity caller has bad uid: get_local_uid=%d
authsys=%d desired uid=%d

A server (rpc.pmfd, rpc.fed or rgmd) refused an rpc connection from a client because it has the wrong uid. The actual and desired uids are shown. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

142779 :Unable to open failfast device

A server (rpc.pmfd or rpc.fed) was not able to establish a link to the failfast device, which ensures that the host aborts if the server dies. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

143622 :PNM: adapter %s is %s

A network adapter has been determined to be either "ok" or "faulty" by PNM, based on a network adapter fault detection algorithm.

For a network adapter determined to be "faulty", check that the physical connections between the adapter and its router are intact, including the adapter, any cables, hubs, and switches. Replace any broken component accordingly.

143694 :lkcm_act: caller is already registered

Message indicating that udlm is already registered with ucmm.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

144303 :fatal: uname: %s (UNIX error %d)

A uname(2) system call failed. The rgmd will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

145270 :Cannot determine if the server is secure: assuming non-secure.

While parsing the Netscape configuration file to determine if the Netscape server is running under secure or non-secure mode an error occurred. This error results in the Data Service assuming a non-secure Netscape server, and will probe the server as such.

Check the Netscape configuration file to make sure that it exists and that it contains information about whether the server is running as a secure server or not.

145770 :CMM: Monitoring disabled.

Transport path monitoring has been disabled in the cluster. It is enabled by default.

This is an informational message, no user action is needed.

145893 :CMM: Unable to read quorum information. Error = %d.

The specified error was encountered while trying to read the quorum information from the CCR. This is probably because the CCR tables were modified by hand, which is an unsupported operation. The node will panic.

Reboot the node in non-cluster (-x) mode and restore the CCR tables from the other nodes in the cluster or from backup. Reboot the node back in cluster mode. The problem should not reappear.

146238 :CMM: Halting to prevent split brain with node %ld.

Due to a connection failure with the specified node, the CMM is failing this node to prevent split brain partial connectivity.

Any interconnect failure should be resolved, and/or the failed node rebooted.

146961 :Signal %d terminated the child process.

An unexpected signal caused the termination of the program that checks the availability of name service.

Save a copy of the /var/adm/messages files on all nodes. If a core file was generated, submit the core to your service provider. Contact your authorized Sun service provider for assistance in diagnosing the problem.

148023 :method <%s> completed successfully for resource <%s>, resource group <%s>

RGM invoked a callback method for the named resource, as a result of a cluster reconfiguration, scha_control GIVEOVER, or scswitch. The method completed successfully.

This is an informational message, no user action is needed.

148393 :Unable to create thread. Exiting.\n

clxecd program has encountered a failed thr_create(2) system call. The error message indicates the error number for the failure.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

148465 :Prog <%s> step <%s>: RPC connection error.

An attempted program execution failed, due to an RPC connection problem. This failure is considered a program failure.

Examine other syslog messages occurring around the same time on the same node, to see if the cause of the problem can be identified. If the same error recurs, you might have to reboot the affected node.

148526 :fatal: Cannot get local nodename

An internal error has occurred. The rgmd will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

148902 :No node was specified as part of property %s for element %s. The property must be specified as %s=Weight%cNode,Weight%cNode,...

The property was specified incorrectly.

Set the property using the correct syntax.

149184 :clcomm: inbound_invo::signal:_state is 0x%x

The internal state describing the server side of a remote invocation is invalid when a signal arrives during processing of the remote invocation.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

150105 :This list element in System property %s has an invalid IP address (hostname): %s.

The system property that was named does not have a valid hostname or dotted-decimal IP address string.

Change the value of the property to use a valid hostname or dotted-decimal IP address string.

150535 :clcomm: Could not find %s(): %s

The function `get_libc_func` could not find the specified function for the reason specified. Refer to the man pages for `"dlsym"` and `"dlerror"` for more information.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

152159 :WARNING: lkcm_sync: udlm_send_reply failed, forcing reconfiguration

A reconfiguration will start.

None.

152478 :Monitor_retry_count or Monitor_retry_interval is not set.

The resource properties `Monitor_retry_count` or `Monitor_retry_interval` has not set. These properties control the restarts of the fault monitor.

Check whether the properties are set. If not, set these values using `scrgadm(1M)`.

152546 :ucm_callback for stop_trans generated exception %d

ucmm callback for stop transition failed.

Save a copy of the `/var/adm/messages` files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

153018 :WARNING: missing msg, expected: dont_care, %d, %d, but received: %d %d, %d. FORCING reconfiguration.

Unexpected message received by udlm. This will trigger an OPS reconfiguration.

None.

154317 :launch_validate: fe_method_full_name() failed for resource <%s>, resource group <%s>, method <%s>

Due to an internal error, the `rgmd` was unable to assemble the full method pathname for the `VALIDATE` method. This is considered a `VALIDATE` method failure. This in turn will cause the failure of a creation or update operation on a resource or resource group.

Examine other syslog messages occurring at about the same time to see if the problem can be identified. Re-try the creation or update operation. If the problem recurs, save a copy of the `/var/adm/messages` files on all nodes and contact your authorized Sun service provider for assistance.

155479 :ERROR: VALIDATE method timeout property of resource
<%s> is not an integer

The indicated resource's VALIDATE method timeout, as stored in the CCR, is not an integer value. This might indicate corruption of CCR data or rgmd in-memory state; the VALIDATE method invocation will fail. This in turn will cause the failure of a creation or update operation on a resource or resource group.

Use `scrgadm(1M) -pvv` to examine resource properties. If the VALIDATE method timeout or other property values appear corrupted, the CCR might have to be rebuilt. If values appear correct, this may indicate an internal error in the rgmd. Re-try the creation or update operation. If the problem recurs, save a copy of the `/var/adm/messages` files on all nodes and contact your authorized Sun service provider for assistance.

156527 :Unable to execute <%s>: <%s>.

Sun Cluster was unable to execute a command.

The problem could be caused by: 1) No more process table entries for a `fork()`
2) No available memory For the above two causes, the only option is to reboot the node. The problem might also be caused by: 3) The command that could not execute is not correctly installed For the above cause, the command might have the wrong path or file permissions. Correctly install the command.

157213 :CCR: The repository on the joining node %s could not be recovered, join aborted.

The indicated node failed to update its repository with the ones in current membership. And it will not be able to join the current membership.

There may be other related messages on the indicated node, which help diagnose the problem, for example: If the root disk failed, it needs to be replaced. If the root disk is full, remove some unnecessary files to free up some space.

158530 :CMM: Halting because this node is severely short of resident physical memory; availrmem = %ld pages,
tune.t_minarmem = %ld pages.

The local node does not have sufficient resident physical memory due to which it may declare other nodes down. To prevent this action, the local node is going to halt.

There may be other related messages that may indicate the cause for the node having reached the low memory state. Resolve the problem and reboot the node. If unable to resolve the problem, contact your authorized Sun service provider to determine whether a workaround or patch is available

158836 :Endpoint %s initialization error - errno = %d, failing associated pathend.

Communication with another node could not be established over the path.

Any interconnect failure should be resolved, and/or the failed node rebooted.

159059 :IP address (hostname) %s from %s at entry %d in list property %s does not belong to any network resource used by resource %s.

The hostname or dotted-decimal IP address string in the message does not resolve to an IP address equal to any resolved IP address from the named resource's Network_resources_used property. Any explicitly named hostname or dotted-decimal IP address string in the named list property must resolve to an IP address equal to a resolved IP address from Network_resources_used.

Either modify the hostname or dotted-decimal IP address string from the entry in the named property or modify Network_resources_used so that the entry resolves to an IP address equal to a resolved IP address from Network_resources_used.

159501 :host %s failed: %s

The rgm is not able to establish an rpc connection to the rpc.fed server on the host shown, and the error message is shown. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

159592 :clcomm: Cannot make high %d less than current total %d

An attempt was made to change the flow control policy parameter specifying the high number of server threads for a resource pool. The system does not allow the high number to be reduced below current total number of server threads.

No user action required.

160167 :Server successfully started.

Informational message. Oracle server has been successfully started by HA-Oracle.

None

160400 :fatal: fcntl(F_SETFD): %s (UNIX error %d)

This error should not occur. The rgmd will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

Save a copy of the `/var/adm/messages` files on all nodes, and of the `rgmd` core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

160619 :Could not enlarge buffer for DBMS log messages: %m

Fault monitor could not allocate memory for reading RDBMS log file. As a result of this error, fault monitor will not scan errors from log file. However it will continue fault monitoring.

Check if syetem is low on memory. If problem persists, please stop and start the fault monitor.

161275 :reservation fatal error(UNKNOWN) - Illegal command line option

The device fencing program has suffered an internal error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available. Copies of `/var/adm/messages` from all nodes should be provided for diagnosis. It may be possible to retry the failed operation, depending on the nature of the error. If the message specifies the "node_join" transition, then this node may be unable to access shared devices. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access shared devices. In either case, it may be possible to reacquire access to shared devices by executing "`/usr/cluster/lib/sc/run_reserve -c node_join`" on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group has failed to start on this node. If another node was available to host the device group, then it should have been started on that node. If desired, it may be possible to switch the device group to this node with the `scswitch` command. If no other node was available, then the device group will not have been started. The `scswitch` command may be used to retry the attempt to start the device group. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group has failed. The desired action may be retried.

161991 :Load balancer for group "%s" setting weight for node %s to %d

This message indicates that the user has set a new weight for a particular node from an old value.

This is an informational message, no user action is needed.

162419 :ERROR: launch_method: cannot get Failover_mode for resource <%s>, assuming NONE.

A method execution has failed or timed out. For some reason, the `rgmd` is unable to obtain the `Failover_mode` property of the resource. The `rgmd`

assumes a setting of NONE for this property, therefore avoiding the outcome of rebooting the node (for STOP method failure) or failing over the resource group (for START method failure). For these cases, the resource is placed into a STOP_FAILED or START_FAILED state, respectively.

Save a copy of the /var/adm/messages files on all nodes, and contact your authorized Sun service provider for assistance in diagnosing the problem.

163027 :CMM: Quorum device %s: owner set to node %ld.

The specified node has taken ownership of the specified quorum device.

This is an informational message, no user action is needed.

164168 :PNM: nafo%d: state transition from %s to %s on %s

A state transition has happened for a NAFO group. Transition to DOUBT happens when the active adapter is determined to be faulty by PNM. Transition to DOWN happens when all adapters in a NAFO group is determined to be faulty, resulting in a loss of network connectivity to a given subnet.

If a NAFO group transitions to DOWN state, check for error messages about adapters being faulty and take suggested user action:s accordingly. No user action: is needed for other state transitions.

164757 :reservation fatal error(%s) - realloc() error, errno %d

The device fencing program has been unable to allocate required memory.

Memory usage should be monitored on this node and steps taken to provide more available memory if problems persist. Once memory has been made available, the following steps may need to taken: If the message specifies the "node_join" transition, then this node may be unable to access shared devices. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access shared devices. In either case, access to shared devices can be reacquired by executing "/usr/cluster/lib/sc/run_reserve -c node_join" on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group has failed to start on this node. If another node was available to host the device group, then it should have been started on that node. The device group can be switched back to this node if desired by using the scswitch command. If no other node was available, then the device group will not have been started. The scswitch command may be used to start the device group. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group has failed. The desired action may be retried.

165512 :reservation error(%s) - my_map_to_did_device() error in other_node_status()

The device fencing program has suffered an internal error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available. Copies of /var/adm/messages from all nodes should be provided for diagnosis. It may be possible to retry the failed operation, depending on the nature of the error. If the message specifies the "node_join" transition, then this node may be unable to access shared devices. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access shared devices. In either case, it may be possible to reacquire access to shared devices by executing "/usr/cluster/lib/sc/run_reserve -c node_join" on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group has failed to start on this node. If another node was available to host the device group, then it should have been started on that node. If desired, it may be possible to switch the device group to this node with the scswitch command. If no other node was available, then the device group will not have been started. The scswitch command may be used to retry the attempt to start the device group. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group has failed. The desired action may be retried.

165527 :Oracle UDLM package is not properly installed. %s not found.

Oracle udlm package installation problem.

Make sure Oracle UDLM package is properly installed.

166362 :clexecd: Got back %d from I_RECVFD. Looks like parent is dead.

Parent process in the clexecd program is dead.

If the node is shutting down, ignore the message. If not, the node on which this message is seen, will shutdown to prevent to prevent data corruption. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

166560 :Maximum Primaries is %d. It should be 1.

Invalid value has set for Maximum Primaries. The value should be 1.

Reset this value using scrgadm(1M).

167108 :Starting Oracle server.

Informational message. Oracle server is being started by HA-Oracle.

None

167253 :Server stopped successfully.

Informational message. Oracle server successfully stopped.

None

168150 :INTERNAL ERROR CMM: Cannot bind quorum algorithm object to local name server.

There was an error while binding the quorum subsystem object to the local name server.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

168383 :Service not started

There was a problem detected in the initial startup of the service.

Attempt to start the service by hand to see if there are any apparent problems with the application. Correct these problems and attempt to start the data service again.

168630 :could not read cluster name

Could not get cluster name. Perhaps the system is not booted as part of the cluster.

Make sure the node is booted as part of a cluster.

168970 :sun_udlm_read_oracle_cfg: open failed: %s ... will use default values

Could not read parameter values from config file. Will use default values instead.

None.

169606 :Unable to create thread. Exiting.

clxecd program has encountered a failed thr_create(2) system call. The error message indicates the error number for the failure.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

169608 :INTERNAL ERROR: scha_control_action: invalid action
<%d>

The scha_control function has encountered an internal logic error. This will cause scha_control to fail with a SCHA_ERR_INTERNAL error, thereby preventing a resource-initiated failover.

Please save a copy of the /var/adm/messages files on all nodes, and report the problem to your authorized Sun service provider.

169765 :Configuration file not found.

Internal error. Configuration file for online_check not found.

Please report this problem.

171786 :listener %s is not running. Attempting restart.

Listen monitor has detected failure of listener. Monitor will attempt to restart the listener.

None

171878 :in libsecurity setnetconfig failed when initializing the client: %s - %s

A client was not able to make an rpc connection to a server (rpc.pmfd, rpc.fed or rgmd) because it could not establish a rpc connection for the network specified. The rpc error and the system error are shown. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

172566 :Stopping oracle server using shutdown abort

Informational message. Oracle server will be stopped using "shutdown abort" command.

Examine "Stop_timeout" property of the resource and increase "Stop_timeout" if you don't wish to use "shutdown abort" for stopping Oracle server.

173733 :Failed to retrieve the resource type property %s for %s: %s.

The query for a property failed. The reason for the failure is given in the message.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

174751 :Failed to retrieve the process monitor facility tag.

Failed to create the tag that has used to register with the process monitor facility.

Check the syslog messages that occurred just before this message. In case of internal error, save the /var/adm/messages file and contact authorized Sun service provider.

174928 :ERROR: process_resource: resource <%s> is offline pending boot, but no BOOT method is registered

A non-fatal internal error has occurred in the rgmd state machine.

Since this problem might indicate an internal logic error in the rgmd, please save a copy of the /var/adm/messages files on all nodes, the output of an scstat -g command, and the output of a scrgadm -pvv command. Report the problem to your authorized Sun service provider.

175370 :svc_restore_priority: Could not restore original scheduling parameters: %s

The server was not able to restore the original scheduling mode. The system error message is shown. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

175553 :clconf: Your configuration file is incorrect! The type of property %s is not found

Could not find the type of property in the configuration file.

Check the configuration file.

176860 :Error: Unable to update scha_control timestamp file <%s> for resource <%s>

The rgmd failed in a call to utime(2) on the local node. This may prevent the anti-"pingpong" feature from working, which may permit a resource group to fail over repeatedly between two or more nodes. The failure of the utime call might indicate a more serious problem on the node.

Examine other syslog messages occurring around the same time on the same node, to see if the source of the problem can be identified.

177070 :Got back %d in revents of the control fd. Exiting.

clxecd program has encountered an error.

The clxecd program will exit and the node will be halted or rebooted to prevent data corruption. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

177252 :reservation warning(%s) - MHI0CGRP_INRESV error will retry in %d seconds

The device fencing program has encountered errors while trying to access a device. The failed operation will be retried

This is an informational message, no user action is needed.

177899 :t_bind (open_cmd_port) failed

Call to `t_bind()` failed. The "`t_bind`" man page describes possible error codes. `ucmmd` will exit and the node will abort.

Save the `/var/adm/messages` file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

179364 :CCR: Invalid CCR metadata.

The CCR could not find valid metadata on all nodes of the cluster.

Boot the cluster in `-x` mode to restore the cluster repository on all the nodes in the cluster from backup. The cluster repository is located at `/etc/cluster/ccr/`.

181193 :Cannot access file <%s>, err = <%s>

The `rgmd` has failed in an attempt to `stat(2)` a file used for the anti-"pingpong" feature. This may prevent the anti-pingpong feature from working, which may permit a resource group to fail over repeatedly between two or more nodes. The failure to access the file might indicate a more serious problem on the node.

Examine other syslog messages occurring around the same time on the same node, to see if the source of the problem can be identified.

183071 :Cannot Execute %s: %s.

Failure in executing the command.

Check the syslog message for the command description. Check whether the system is low in memory or the process table is full and take appropriate action. Make sure that the executable exists.

183799 :clconf: CSR not initialized

While executing task in `clconf` and modifying the state of proxy, found component CSR not initialized.

Check the CSR component in the configuration file.

184139 :scvxvmlg warning - found no match for %s, removing it

The program responsible for maintaining the VxVM device namespace has discovered inconsistencies between the VxVM device namespace on this node and the VxVM configuration information stored in the cluster device configuration system. If configuration changes were made recently, then this message should reflect one of the configuration changes. If no changes were made recently or if this message does not correctly reflect a change that has been made, the VxVM device namespace on this node may be in an inconsistent state. VxVM volumes may be inaccessible from this node.

If this message correctly reflects a configuration change to VxVM diskgroups then no action is required. If the change this message reflects is not correct,

then the information stored in the device configuration system for each VxVM diskgroup should be examined for correctness. If the information in the device configuration system is accurate, then executing `"/usr/cluster/lib/dcs/scvxvmlg"` on this node should restore the device namespace. If the information stored in the device configuration system is not accurate, it must be updated by executing `"/usr/cluster/bin/scconf -c -D name=diskgroup_name"` for each VxVM diskgroup with inconsistent information.

185089 :CCR: Updating table %s failed to startup on node %s.

The operation to update the indicated table failed to start on the indicated node.

There may be other related messages on the nodes where the failure occurred, which may help diagnose the problem. If the root disk failed, it needs to be replaced. If the indicated table was deleted by accident, boot the offending node(s) in -x mode to restore the indicated table from other nodes in the cluster. The CCR tables are located at `/etc/cluster/ccr/`. If the root disk is full, remove some unnecessary files to free up some space.

185465 :No action on DBMS Error %s: %ld

Database server returned error. Fault monitor does not take any action on this error.

No action required.

185720 :lkdb_parm: lib initialization failed

initializing a library to get the static lock manager parameters failed.

This is an internal error. Save the contents of `/var/adm/messages`, `/var/cluster/ucmm/ucmm_reconf.log` and `/var/cluster/ucmm/dlm*/logs/*` from all the nodes and contact your Sun service representative.

185839 :IP address (hostname) and Port pairs %s%c%d and %s%c%d in property %s, at entries %d and %d, effectively duplicate each other. The port numbers are the same and the resolved IP addresses are the same.

The two list entries at the named locations in the named property have port numbers that are identical, and also have IP address (hostname) strings that resolve to the same underlying IP address. An IP address (hostname) string and port entry should only appear once in the property.

Specify the property with only one occurrence of the IP address (hostname) string and port entry.

185974 :Default Oracle paramter file %s does not exist

Oracle Parameter file has not been specified. Default parameter file indicated in the message does not exist.

Please make sure that parameter file exists at the location indicated in message or specify "Parameter_file" property for the resource.

186306 :Conversion of hostnames failed for %s.

The hostname or IP address given could not be converted to an integer.

Add the hostname to the /etc/inet/hosts file. Verify the settings in the /etc/nsswitch.conf file include "files" for host lookup.

186484 :PENDING_METHODS: bad resource state <%s> (%d) for resource <%s>

The rgmd state machine has discovered a resource in an unexpected state on the local node. This should not occur and may indicate an internal logic error in the rgmd.

Look for other syslog error messages on the same node. Save a copy of the /var/adm/messages files on all nodes, and report the problem to your authorized Sun service provider.

186524 :reservation error(%s) - do_scsi2_release() error for disk %s

The device fencing program has encountered errors while trying to access a device. All retry attempts have failed.

The action which failed is a scsi-2 ioctl. These can fail if there are scsi-3 keys on the disk. To remove invalid scsi-3 keys from a device, use "scdidadm -R" to repair the disk (see scdidadm man page for details). If there were no scsi-3 keys present on the device, then this error is indicative of a hardware problem, which should be resolved as soon as possible. Once the problem has been resolved, the following actions may be necessary: If the message specifies the "node_join" transition, then this node may be unable to access the specified device. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access the device. In either case, access can be reacquired by executing "/usr/cluster/lib/sc/run_reserve -c node_join" on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group may have failed to start on this node. If the device group was started on another node, it may be moved to this node with the scswitch command. If the device group was not started, it may be started with the scswitch command. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group may have failed. If so, the desired action may be retried.

186612 :_cladm CL_GET_CLUSTER_NAME failed; perhaps system is not booted as part of cluster

Could not get cluster name. Perhaps the system is not booted as part of the cluster.

Make sure the node is booted as part of a cluster.

187307 :invalid debug_level: "%s"

Invalid debug_level argument passed to udlmctl. udlmctl will not startup.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

191225 :clcomm: Created %d threads, wanted %d for pool %d

The system creates server threads to support requests from other nodes in the cluster. The system could not create the desired minimum number of server threads. However, the system did succeed in creating at least 1 server thread. The system will have further opportunities to create more server threads. The system cannot create server threads when there is inadequate memory. This message indicates either inadequate memory or an incorrect configuration.

There are multiple possible root causes. If the system administrator specified the value of "maxusers", try reducing the value of "maxusers". This reduces memory usage and results in the creation of fewer server threads. If the system administrator specified the value of "cl_comm:min_threads_default_pool" in "/etc/system", try reducing this value. This directly reduces the number of server threads. Alternatively, do not specify this value. The system can automatically select an appropriate number of server threads. Another alternative is to install more memory. If the system administrator did not modify either "maxusers" or "min_threads_default_pool", then the system should have selected an appropriate number of server threads. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

191270 :IP address (hostname) string %s in property %s, entry %d does not resolve to an IP address that belongs to one of the resources named in property %s.

The IP address or hostname named does not belong to one of the network resources designated for use by this resource

Either select a different IP address to use that is in one of the network resources used by this resource or create a network resource that contains the named IP address and designate that resource as one of the network resources used by this resource.

191409 :scvxvmlg warning - chown(%s) failed

The program responsible for maintaining the VxVM namespace was unable to access the global device namespace. If configuration changes were recently

made to VxVM diskgroups or volumes, this node may be unaware of those changes. Recently created volumes may be inaccessible from this node.

Verify that the /global/.devices/node@N (N = this node's node number) is mounted globally and is accessible. If no configuration changes have been recently made to VxVM diskgroups or volumes and all volumes continue to be accessible from this node, then no further action is required. If changes have been made, the device namespace on this node can be updated to reflect those changes by executing "/usr/cluster/lib/dcs/scvxvmlg". If the problem persists, contact your authorized Sun service provider to determine whether a workaround or patch is available.

191492 :CCR: CCR unable to read root file system.

The CCR failed to read repository due to root file system failure on this node.

The root file system needs to be replaced on the offending node.

191506 :ERROR: enabled resource <%s> in resource group <%s> depends on disabled resource <%s>

An enabled resource was found to depend on a disabled resource. This should not occur and may indicate an internal logic error in the rgmd.

Look for other syslog error messages on the same node. Save a copy of the /var/adm/messages files on all nodes, and report the problem to your authorized Sun service provider.

191772 :Failed to configure the networking components for scalable resource %s for method %s.

The processing that is required for scalable services did not complete successfully.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

191957 :The property %s does not have a legal value.

The property named does not have a legal value.

Assign the property a legal value.

192183 :freeze_adjust_timeouts: call to rpc.fed failed, tag <%s> err <%d> result <%d>

The rgmd failed in its attempt to suspend timeouts on an executing method during temporary unavailability of a global device group. This could cause the resource method to time-out. Depending on which method was being invoked and the Failover_mode setting on the resource, this might cause the resource group to fail over or move to an error state.

No action is required if the resource method execution succeeds. If the problem recurs, rebooting this node might cure it. Save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance in diagnosing the problem.

192518 :Cannot access start script %s: %s

The start script is not accessible and executable. This may be due to the script not existing or the permissions not being set properly.

Make sure the script exists, is in the proper directory, and has read and execute permissions set appropriately.

192619 :reservation error(%s) - Unable to open device %s

The device fencing program has encountered errors while trying to access a device. All retry attempts have failed.

This may be indicative of a hardware problem, which should be resolved as soon as possible. Once the problem has been resolved, the following actions may be necessary: If the message specifies the "node_join" transition, then this node may be unable to access the specified device. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access the device. In either case, access can be reacquired by executing "/usr/cluster/lib/sc/run_reserve -c node_join" on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group may have failed to start on this node. If the device group was started on another node, it may be moved to this node with the scswitch command. If the device group was not started, it may be started with the scswitch command. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group may have failed. If so, the desired action may be retried.

193137 :Service group "%s" deleted

The service group by that name is no longer known by the scalable services framework.

This is an informational message, no user action is needed.

193263 :Service is online.

While attempting to check the health of the data service, probe detected that the resource status is fine and it is online.

This is informational message. No user action is needed.

193933 :CMM: Votecount changed from %d to %d for node %s.

The specified node's votecount has been changed as indicated.

This is an informational message, no user action is needed.

194512 :Failed to stop HA-NFS system fault monitor.

Process monitor facility has failed to stop the HA-NFS system fault monitor.

Use pmfadm(1M) with -s option to stop the HA-NFS system fault monitor with tag name "cluster.nfs.daemons". If the error still persists, then reboot the node.

194810 :clcomm: thread_create failed for resource_thread

The system could not create the needed thread, because there is inadequate memory.

There are two possible solutions. Install more memory. Alternatively, reduce memory usage. Since this happens during system startup, application memory usage is normally not a factor.

195286 :CMM: Placing reservation on quorum device %s failed with error %d.

The specified error was encountered while trying to place a reservation on the specified quorum device, hence this node can not take ownership of this quorum device.

There may be other related messages on this and other nodes connected to this quorum device that may indicate the cause of this problem. Refer to the quorum disk repair section of the administration guide for resolving this problem.

195867 :clexecd: Unexpected eventmask %x in revents of the control fd.

clexecd program has encountered an error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

196233 :INTERNAL ERROR: launch_method: method tag <%s> not found in method invocation list for resource group <%s>

An internal error has occurred. The rgmd will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

197307 :Resource contains invalid hostnames.

The hostnames that has to be made available by this logical host resource are invalid.

It is advised to keep the hostnames in /etc/inet/hosts file and enable "files" for host lookup in nsswitch.conf file. Any of the following situations might have occurred. 1) If hosts are not in /etc/inet/hosts file then make sure the nameserver is reachable and has host name entries specified. 2) Invalid hostnames might have been specified while creating the logical host resource. If this is the case, use the scrgadm command to respecify the hostnames for this logical host resource.

197456 :CCR: Fatal error: Node will be killed.

Some fatal error occurred on this node during the synchronization of cluster repository. This node will be killed to allow the synchronization to continue.

Look for other messages on this node that indicated the fatal error occurred on this node. For example, if the root disk on the afflicted node has failed, then it needs to be replaced.

197997 :clexecd: dup2 of stdin returned with errno %d while exec"ing (%s). Exiting.

clexecd program has encountered a failed dup2(2) system call. The error message indicates the error number for the failure.

The clexecd program will exit and the node will be halted or rebooted to prevent data corruption. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

198216 :t_bind cannot bind to requested address

Call to t_bind() failed. The "t_bind" man page describes possible error codes. ucmmmd will exit and the node will abort.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

198284 :Failed to start fault monitor.

The fault monitor for this data service was not started. There may be prior messages in syslog indicating specific problems.

The user should correct the problems specified in prior syslog messages. This problem may occur when the cluster is under load and Sun Cluster cannot start the application within the timeout period specified. You may consider increasing the Monitor_Start_timeout property. Try switching the resource group to another node using scswitch (1M).

198542 :No network resources found for resource.

No network resources were found for the resource.

Declare network resources used by the resource explicitly using the property `Network_resources_used`. For the resource name and resource group name, check the syslog tag.

198851 :fatal: Got error <%d> trying to read CCR when disabling resource <%s>; aborting node

Rgmd failed to read updated resource from the CCR on this node.

Save a copy of the `/var/adm/messages` files on all nodes, and of the `rgmd` core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

199467 :clcomm::ObjectHandler::_unreferenced called

This operation should never be executed.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

Message IDs 200000 - 299999

Error Message List

The following list is ordered by the message ID.

201878 :clconf: Key length is more than max supported length in clconf_file_io

In reading configuration data through CCR FILE interface, found the data length is more than max supported length.

Check the CCR configuraton information.

203680 :fatal: Unable to bind to nameserver

The low-level cluster machinery has encountered a fatal error. The rgmd will produce a core file and will cause the node to halt or reboot to avoid the possibility of data corruption.

Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

203739 :Resource %s uses network resource %s in resource group %s, but the property %s for resource group %s does not include resource group %s. This dependency must be set.

For all network resources used by a scalable resource, a dependency on the resource group containing the network resource should be created for the resource group of the scalable resource.

Use the `scrgadm(1M)` command to update the `RG_dependencies` property of the scalable resource's resource group to include the resource groups of all network resources that the scalable resource uses.

204163 :clcomm: error in copyin for state_balancer

The system failed a copy operation supporting statistics reporting.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

204584 :clxecd: Going down on signal %d.

clxecd program got a signal indicated in the error message.

clxecd program will exit and node will be halted or rebooted to prevent data corruption. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

205445 :check_and_start(): Out of memory

System runs out of memory in function `check_and_start()`

Install more memory, increase swap space, or reduce peak memory consumption.

205873 :Permissions incorrect for %s. s bit not set.

Permissions of `$ORACLE_HOME/bin/oracle` are expected to be `"-rwsr-s-x"` (set-group-ID and set-user-ID set). These permissions are set at the time of Oracle installation. Fault monitor will not function correctly without these permissions.

Check file permissions. Check Oracle installation. Relink Oracle, if necessary.

206501 :CMM: Monitoring re-enabled.

Transport path monitoring has been enabled back in the cluster, after being disabled.

This is an informational message, no user action is needed.

206947 :ON_PENDING_MON_DISABLED: bad resource state <%s> (%d)
for resource <%s>

The `rgmd` state machine has discovered a resource in an unexpected state on the local node. This should not occur and may indicate an internal logic error in the `rgmd`.

Look for other syslog error messages on the same node. Save a copy of the `/var/adm/messages` files on all nodes, and report the problem to your authorized Sun service provider.

207481 :getlocalhostname() failed for resource <%s>, resource
group <%s>, method <%s>

The rgmd was unable to obtain the name of the local host, causing a method invocation to fail. Depending on which method is being invoked and the Failover_mode setting on the resource, this might cause the resource group to fail over or move to an error state.

Examine other syslog messages occurring at about the same time to see if the problem can be identified. Save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance in diagnosing the problem.

208216 :ERROR: resource group <%s> has RG_dependency on non-existent resource group <%s>

A non-existent resource group is listed in the RG_dependencies of the indicated resource group. This should not occur and may indicate an internal logic error in the rgmd.

Look for other syslog error messages on the same node. Save a copy of the /var/adm/messages files on all nodes, and report the problem to your authorized Sun service provider.

208596 :clcomm: Path %s being initiated

A communication link is being established with another node.

No action required.

208701 :%s error status ignored in step %s

Ignoring the error status from step execution since this does not affect outcome of the step.

None.

209274 :path_check_start(): Out of memory

Run out of memory in function path_check_start().

Install more memory, increase swap space, or reduce peak memory consumption.

210725 :Warning: While trying to lookup host %s, the length of the returned address (%d) was longer than expected (%d). The address will be truncated.

The value of the resolved address for the named host was longer than expected.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

211198 :Completed successfully.

Data service method completed successfully.

No action required.

212337 :(%s) scan of seqnum failed on "%s", ret = %d

Could not get the sequence number from the udln message received.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

212667 :PNM: could not start due to lock %s

An attempt was made to start multiple instances of the PNM daemon pnmd(1M), or pnmd(1M) has problem acquiring a lock on the named file.

Check if another instance of pnmd is already running. If not, remove the named lock file and start pnmd using the /etc/init.d/pnm script.

213112 :latch_intention(): IDL exception when communicating to node %d

An inter-node communication failed, probably because a node died.

No action is required; the rgmd should recover automatically.

215538 :Not all hostnames brought online.

Failed to bring all the hostnames online. Only some of the ip addresses are online.

Use ifconfig command to make sure that the ip addresses are available. Check for any error message before this error message for a more precise reason for this error. Use scswitch command to move the resource group to a different node. If problem persists, reboot.

216087 :rebalance: resource group <%s> is being switched updated or failed back, cannot assign new primaries

The indicated resource group has lost a master due to a node death. However, the RGM is unable to switch the resource group to a new master because the resource group is currently in the process of being modified by an operator action, or is currently in the process of "failing back" onto a node that recently joined the cluster.

Use scstat(1M) -g to determine the current mastery of the resource group. If necessary, use scswitch(1M) -z to switch the resource group online on desired nodes.

216244 :CCR: Table %s has invalid checksum field. Reported: %s, actual: %s.

The indicated table has an invalid checksum that does not match the table contents. This causes the consistency check on the indicated table to fail.

Boot the offending node in -x mode to restore the indicated table from backup or other nodes in the cluster. The CCR tables are located at /etc/cluster/ccr/.

216379 :Stopping fault monitor using pmfadm tag %s

Informational message. Fault monitor will be stopped using Process Monitoring Facility (PMF), with the tag indicated in message.

None

216774 :WARNING: update_state:udlm_send_reply failed

A warning for udlm state update and results in udlm abort.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

217093 :Call failed: %s

A client was not able to make an rpc connection to a server (rpc.pmfd, rpc.fed or rgmd) to execute the action shown. The rpc error message is shown. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

218227 :Error accessing policy string

This message appears when the customer is initializing or changing a scalable services load balancer, by starting or updating a service. The Load_Balancing_String is missing.

Add a Load_Balancing_String parameter when creating the resource group.

220849 :CCR: Create table %s failed.

The CCR failed to create the indicated table.

The failure can happen due to many reasons, for some of which no user action is required because the CCR client in that case will handle the failure. The cases for which user action is required depends on other messages from CCR on the node, and include: If it failed because the cluster lost quorum, reboot the cluster. If the root file system is full on the node, then free up some space by removing unnecessary files. If the root disk on the afflicted node has failed, then it needs to be replaced. If the cluster repository is corrupted as indicated by other CCR messages, then boot the offending node(s) in -x mode to restore the cluster repository backup. The cluster repository is located at /etc/cluster/ccr/.

222512 :fatal: could not create death_ff

The daemon indicated in the message tag (rgmd or ucmmmd) was unable to create a failfast device. The failfast device kills the node if the daemon process dies either due to hitting a fatal bug or due to being killed inadvertently by an operator. This is a requirement to avoid the possibility of data corruption. The daemon will produce a core file and will cause the node to halt or reboot.

Save a copy of the /var/adm/messages files on all nodes, and of the core file generated by the daemon. Contact your authorized Sun service provider for assistance in diagnosing the problem.

223145 :gethostbyname failed for (%s)

Failed to get information about a host. The "gethostbyname" man page describes possible reasons.

Make sure entries in /etc/hosts, /etc/nsswitch.conf and /etc/netconfig are correct to get information about this host.

223458 :INTERNAL ERROR CMM: quorum_algorithm_init called already.

This is an internal error during node initialization, and the system can not continue.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

224718 :Failed to create scalable service in group %s for IP %s Port %d%c%s: %s.

A call to the underlying scalable networking code failed. This call may fail because the IP, Port, and Protocol combination listed in the message conflicts with the configuration of an existing scalable resource. A conflict can occur if the same combination exists in a scalable resource that is already configured on the cluster. A combination may also conflict if there is a resource that uses Load_balancing_policy LB_STICKY_WILD with the same IP address as a different resource that also uses LB_STICKY_WILD.

Try using a different IP, Port, and Protocol combination. Otherwise, save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

224783 :clcomm: Path %s has been deleted

A communication link is being removed with another node. The interconnect may have failed or the remote node may be down.

Any interconnect failure should be resolved, and/or the failed node rebooted.

225882 :Internal: Unknown command type (%d)

An internal error has occurred in the rgmd while trying to connect to the rpc.fed server.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

226914 :scswitch: internal error: bad nodename %s in nodelist of resource group %s

The indicated resource group's Nodelist property, as stored in the CCR, contains an invalid nodename. This might indicate corruption of CCR data or rgmd in-memory state. The scswitch command will fail.

Use scstat(1M) -g and scrgadm(1M) -pvv to examine resource group properties. If the values appear corrupted, the CCR might have to be rebuilt. If values appear correct, this may indicate an internal error in the rgmd. Contact your authorized Sun service provider for assistance in diagnosing and correcting the problem.

227214 :Error: duplicate method <%s> launched on resource <%s> in resource group <%s>

Due to an internal error, the rgmd state machine has attempted to launch two different methods on the same resource on the same node, simultaneously. The rgmd will reject the second attempt and treat it as a method failure.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

228461 :CMM: Issuing a SCSI2 Release failed on quorum device %s with error %d.

This node encountered the specified error while issuing a SCSI2 Release operation on the specified quorum device. The quorum code will either retry this operation or will ignore this quorum device.

There may be other related messages that may provide more information regarding the cause of this problem. SCSI2 operations fail with an error code of EACCES if SCSI3 keys are present on the device. Scrub the SCSI3 keys off of the quorum device.

228999 :reservation error(%s) - Ignoring device %s because it requires scsi-3 support

The device fencing program has discovered a device which is connected to more than 2 nodes.

Connecting shared devices to more than 2 nodes is not yet supported. The connectivity of the device should be reduced to 2 nodes.

231770 :ns: Could not initialize ORB: %d

could not initialize ORB.

Please make sure the nodes are booted in cluster mode.

231991 :WARNING: lkcm_dreg: udln_send_reply failed

Could not deregister udln with ucmm.

None.

232201 :Invalid port number returned.

Invalid port number was retrieved for the Port_list property of the resource.

Any of the following situations may occur. Different user action is required for these different scenarios. 1) If a new resource has created or updated, check whether it has valid port number. If port number is not valid, provide valid port number using scrgadm(1M) command. 2) Check the syslog messages that have occurred just before this message. If it is "Out of memory" problem, then correct it. 3) For all other cases, treat it as an Internal error. Contact your authorized Sun service provider.

232501 :Validation failed. ORACLE_HOME/bin/svrmgrl not found
ORACLE_HOME=%s

Oracle binaries (svrmgrl) not found in ORACLE_HOME/bin directory.
ORACLE_HOME specified for the resource is indicated in the message.
HA-Oracle will not be able to manage resource if ORACLE_HOME is incorrect.

Specify correct ORACLE_HOME when creating resource. If resource is already created, please update resource property "ORACLE_HOME".

232565 :Scalable services enabled.

This means that the scalable services framework is set up in the cluster. Specifically, is is printed out for the node that has joined the cluster and for which services have been downloaded. Once the services have been downloaded, those services are ready to participate as scalable services.

This is an informational message, no user action is needed.

232920 :-d must be followed by a hex bitmask

Incorrect arguments used while setting up sun specific startup parameters to the Oracle unix dlm.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

233017 :Successfully stopped %s.

The resource was successfully stopped by Sun Cluster.

No user action is required.

233053 :SharedAddress offline.

The status of the sharedaddress resource is offline.

This is informational message. No user action required.

233327 :Switchover (%s) error: failed to mount FS (%d)

The file system specified in the message could not be hosted on the node the message came from.

Check /var/adm/messages to make sure there were no device errors. If not, contact your authorized Sun service provider to determine whether a workaround or patch is available.

233956 :Error in reading message in child process: %m

Error occurred when reading message in fault monitor child process. Child process will be stopped and restarted.

If error persists, then disable the fault monitor and resport the problem.

233961 :scvxvmlg error - symlink(%s, %s) failed

The program responsible for maintaining the VxVM namespace was unable to access the global device namespace. If configuration changes were recently made to VxVM diskgroups or volumes, this node may be unaware of those changes. Recently created volumes may be inaccessible from this node.

Verify that the /global/.devices/node@N (N = this node's node number) is mounted globally and is accessible. If no configuration changes have been recently made to VxVM diskgroups or volumes and all volumes continue to be accessible from this node, then no further action is required. If changes have been made, the device namespace on this node can be updated to reflect those changes by executing "/usr/cluster/lib/dcs/scvxvmlg". If the problem persists, contact your authorized Sun service provider to determine whether a workaround or patch is available.

234438 :INTERNAL ERROR: Invalid resource property type <%d> on resource <%s>; aborting node

An attempted creation or update of a resource has failed because of invalid resource type data. This may indicate CCR data corruption or an internal logic error in the rgmd. The rgmd will produce a core file and will force the node to halt or reboot.

Use `scrgadm(1M) -pvv` to examine resource properties. If the resource or resource type properties appear to be corrupted, the CCR might have to be rebuilt. If values appear correct, this may indicate an internal error in the `rgmd`. Re-try the creation or update operation. If the problem recurs, save a copy of the `/var/adm/messages` files on all nodes and contact your authorized Sun service provider for assistance.

234463 :INTERNAL ERROR: process_resource: resource group <%s> is pending_mon_disable but contains resource <%s> in STOP_FAILED state

During a resource monitor disable (`scswitch -M -n`), the `rgmd` has discovered a resource in `STOP_FAILED` state. This may indicate an internal logic error in the `rgmd`, since updates are not permitted on the resource group until the `STOP_FAILED` error condition is cleared.

Look for other syslog error messages on the same node. Save a copy of the `/var/adm/messages` files on all nodes, and report the problem to your authorized Sun service provider.

236733 :lookup of oracle dba gid failed.

Could not find group id for dba. `udlm` will not startup.

Make sure `/etc/nswitch.conf` and `/etc/group` files are valid and have correct information to get the group id of dba.

237149 :clcomm: Path %s being constructed

A communication link is being established with another node.

No action required.

237724 :Failed to retrieve hostname: %s.

The call back method has failed to determine the hostname. Now the callback methods will be executed in `/var/core` directory.

No user action is needed. For detailed error message, look at the syslog message.

239415 :Failed to retrieve the cluster handle: %s.

Access to the object named failed. The reason for the failure is given in the message.

Save a copy of the `/var/adm/messages` files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

239735 :Couldn't parse policy string %s

This message appears when the customer is initializing or changing a scalable services load balancer, by starting or updating a service. The Load_Balancing_String is invalid.

Check the Load_Balancing_String value specified when creating the resource group and make sure that a valid value is used.

240107 :resource %s state on node %s change to R_ONLINE

This is a notification from the rgmd that a resource's state has changed. This may be used by system monitoring tools.

This is an informational message, no user action is needed.

240376 :No protocol was given as part of property %s for element %s. The property must be specified as %s=PortNumber%cProtocol,PortNumber%cProtocol,...

The property named does not have a legal value.

Assign the property a legal value.

240388 :Prog <%s> step <%s>: timed out.

A step has exceeded its configured timeout and was killed by ucmmmd. This in turn will cause a reconfiguration of OPS.

Other syslog messages occurring just before this one might indicate the reason for the failure. After correcting the problem that caused the step to fail, the operator may retry reconfiguration of OPS.

241147 :Invalid value %s for property %s.

An invalid value was supplied for the property.

Supply "conf" or "boot" as the value for DNS_mode property.

241441 :clexecd: ioctl(I_RECVFD) returned %d. Returning %d to clexecd.

clexecd program has encountered a failed ioctl(2) system call. The error message indicates the error number for the failure.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

241948 :Failed to retrieve resource <%s> extension property <%s>: %s.

An internal error occurred in the rgmd while checking a resource property.

Look for other syslog error messages on the same node. Save a copy of the /var/adm/messages files on all nodes, and report the problem to your authorized Sun service provider.

242214 :clxecd: fork1 returned %d. Returning %d to clxecd.

clxecd program has encountered a failed fork1(2) system call. The error message indicates the error number for the failure.

If the error number is 12 (ENOMEM), install more memory, increase swap space, or reduce peak memory consumption. If error number is something else, contact your authorized Sun service provider to determine whether a workaround or patch is available.

243444 :CMM: Issuing a SCSI2 Tkown failed for quorum device with error %d.

This node encountered the specified error while issuing a SCSI2 Tkown operation on a quorum device. This will cause the node to conclude that it has been unsuccessful in preempting keys from the quorum device, and therefore the partition to which it belongs has been preempted. If a cluster gets divided into two or more disjoint subclusters, exactly one of these must survive as the operational cluster. The surviving cluster forces the other subclusters to abort by grabbing enough votes to grant it majority quorum. This is referred to as preemption of the losing subclusters.

There will be other related messages that will identify the quorum device for which this error has occurred. If the error encountered is EACCES, then the SCSI2 command could have failed due to the presence of SCSI3 keys on the quorum device. Scrub the SCSI3 keys off of it, and reboot the preempted nodes.

243639 :Scalable service instance [%s,%s,%d] deregistered on node %s.

The specified scalable service had been deregistered on the specified node. Now, the gif node cannot redirect packets for the specified service to this node.

This is an informational message, no user action is needed.

243965 :udlm_ack_msg: udp is null!

Can not acknowledge a message received from udlmctl because the address to acknowledge to is null.

None.

243996 :Failed to retrieve resource <%s> extention property <%s>

Can not get extention property.

Look for other syslog error messages on the same node. Save a copy of the /var/adm/messages files on all nodes, and report the problem to your authorized Sun service provider.

244116 :clcomm: socreate on routing socket failed with error = %d

The system prepares IP communications across the private interconnect. A socket create operation on the routing socket failed.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

245186 :reservation warning(%s) - MHIOCGRP_PREEMPTANDABORT error will retry in %d seconds

The device fencing program has encountered errors while trying to access a device. The failed operation will be retried

This is an informational message, no user action is needed.

247682 :recv_message: cm_reconfigure: %s

udlm received a message to reconfigure.

None. OPS is going to reconfigure.

247868 :in libsecurity: file %s not readable or bad content

The rpc.pmfd, rpc.fed or rgmd server was not able to read an rpcbind information cache file, or the file's contents are corrupted. The affected component should continue to function by calling rpcbind directly.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

248031 :scvxvmlg warning - %s does not exist, creating it

The program responsible for maintaining the VxVM device namespace has discovered inconsistencies between the VxVM device namespace on this node and the VxVM configuration information stored in the cluster device configuration system. If configuration changes were made recently, then this message should reflect one of the configuration changes. If no changes were made recently or if this message does not correctly reflect a change that has been made, the VxVM device namespace on this node may be in an inconsistent state. VxVM volumes may be inaccessible from this node.

If this message correctly reflects a configuration change to VxVM diskgroups then no action is required. If the change this message reflects is not correct, then the information stored in the device configuration system for each VxVM diskgroup should be examined for correctness. If the information in the device

configuration system is accurate, then executing `"/usr/cluster/lib/dcs/scvxvmlg"` on this node should restore the device namespace. If the information stored in the device configuration system is not accurate, it must be updated by executing `"/usr/cluster/bin/sconf -c -D name=diskgroup_name"` for each VxVM diskgroup with inconsistent information.

249804 :INTERNAL ERROR CMM: Failure creating sender thread.

An instance of the userland CMM encountered an internal initialization error. This is caused by inadequate memory on the system.

Add more memory to the system. If that does not resolve the problem, contact your authorized Sun service provider to determine whether a workaround or patch is available.

249934 :Method <%s> failed to execute on resource <%s> in resource group <%s>, error: <%d>

A resource method failed to execute, due to a system error number identified in the message. The indicated error number appears not to match any of the known `errno` values described in `intro(2)`. This is considered a method failure. Depending on which method is being invoked and the `Failover_mode` setting on the resource, this might cause the resource group to fail over or move to an error state, or it might cause an attempted edit of a resource group or its resources to fail.

Other syslog messages occurring at about the same time might provide evidence of the source of the problem. If not, save a copy of the `/var/adm/messages` files on all nodes, and (if the `rgmd` did crash) a copy of the `rgmd` core file, and contact your authorized Sun service provider for assistance.

250133 :Failed to open the device %s: %s.

This is an internal error. System failed to perform the specified operation.

For specific error information check the syslog message. Provide the following information to your authorized Sun service provider to diagnose the problem. 1) Saved copy of `/var/adm/messages` file 2) Output of `"ls -l /dev/sad"` command 3) Output of `"modinfo | grep sad"` command.

250709 :CMM: Initialization for quorum device %s failed with error EACCES. Will retry.

This node is not able to access the specified quorum device because the node is still fenced off. A retry will be attempted.

This is an informational message, no user action is needed.

250800 :clconf: Not found clexecd on node %d for %d seconds. Giving up!

Could not find clexecd to execute the program on a node. Indicated giving up after retries.

No action required. This is informational message.

251552 :Failed to validate configuration.

The data service is not properly configured.

Look at the prior syslog messages for specific problems and take corrective action.

254131 :resource group %s removed.

This is a notification from the rgmd that the operator has deleted a resource group. This may be used by system monitoring tools.

This is an informational message, no user action is needed.

254692 :scswitch: internal error: bad state <%s> (<%d>) for resource group <%s>

While attempting to execute an operator-requested switch of the primaries of a resource group, the rgmd has discovered the indicated resource group to be in an invalid state. The switch action will fail.

This may indicate an internal error or bug in the rgmd. Contact your authorized Sun service provider for assistance in diagnosing and correcting the problem.

254744 :CMM: Open failed for quorum device %s. Unable to scrub device.

The open operation failed for the specified quorum device while it was being added into the cluster. The add of this quorum device will fail.

The quorum device has failed or the path to this device may be broken. Refer to the disk repair section of the administration guide for resolving this problem. Retry adding the quorum device after the problem has been resolved.

255115 :Retrying to retrieve the resource type information.

An update to cluster configuration occurred while resource type properties were being retrieved

Ignore the message.

255929 :in libsecurityauthsys_create_default failed

A client was not able to make an rpc connection to a server (rpc.pmfd, rpc.fed or rgmd) because it failed the authentication process. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

258357 :Method <%s> failed to execute on resource <%s> in resource group <%s>, error: <%s>

A resource method failed to execute, due to a system error described in the message. For an explanation of the error message, consult intro(2). This is considered a method failure. Depending on which method was being invoked and the Failover_mode setting on the resource, this might cause the resource group to fail over or move to an error state, or it might cause an attempted edit of a resource group or its resources to fail.

If the error message is not self-explanatory, other syslog messages occurring at about the same time might provide evidence of the source of the problem. If not, save a copy of the /var/adm/messages files on all nodes, and (if the rgmd did crash) a copy of the rgmd core file, and contact your authorized Sun service provider for assistance.

258909 :clxeccd: sigfillset returned %d. Exiting.

clxeccd program has encountered a failed sigfillset(3C) system call. The error message indicates the error number for the failure.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

259455 :in fe_set_env_vars malloc failed

The rgmd server was not able to allocate memory for the environment name, while trying to connect to the rpc.fed server, possibly due to low memory. An error message is output to syslog.

Investigate if the host is running out of memory. If not save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

259810 :reservation error(%s) - do_scsi3_reserve() error for disk %s

The device fencing program has encountered errors while trying to access a device. All retry attempts have failed.

For the user action required by this message, see the user action for message 192619.

261123 :resource group %s state change to managed.

This is a notification from the rgmd that a resource group's state has changed. This may be used by system monitoring tools.

This is an informational message, no user action is needed.

262295 :Failback bailing out because resource group <%s>is being updated or switched

The rgmd was unable to failback the specified resource group to a more preferred node because the resource group was already in the process of being updated or switched.

This is an informational message, no user action is needed.

262898 :Name service not available.

The monitor_check method detected that name service is not responsive.

Check if name service is configured correctly. Try some commands to query name services, such as ping and nslookup, and correct the problem. If the error still persists, then reboot the node.

263258 :CCR: More than one copy of table %s has the same version but different checksums. Using the table from node %s.

The CCR detects that two valid copies of the indicated table have the same version but different contents. The copy on the indicated node will be used by the CCR.

This is an informational message, no user action is needed.

263606 :unpack_rg_seq: rname_to_r error <%s>

Due to an internal error, the rgmd was unable to find the specified resource data in memory.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

265925 :CMM: Cluster lost operational quorum; aborting.

Not enough nodes are operational to maintain a majority quorum, causing the cluster to fail to avoid a potential split brain.

The nodes should be rebooted.

266059 :security_svc_reg failed.

The rpc.pmfd server was not able to initialize authentication and rpc initialization. This happens while the server is starting up, at boot time. The server does not come up, and an error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

266834 :CMM: Our partition has been preempted.

The cluster partition to which this node belongs has been preempted by another partition during a reconfiguration. The preempted partition will abort. If a cluster gets divided into two or more disjoint subclusters, exactly one of these must survive as the operational cluster. The surviving cluster forces the other subclusters to abort by grabbing enough votes to grant it majority quorum. This is referred to as preemption of the losing subclusters.

There may be other related messages that may indicate why quorum was lost. Determine why quorum was lost on this node partition, resolve the problem and reboot the nodes in this partition.

267558 :Error when reading property %s.

Unable to read property value using API. Property name is indicated in message. Syslog messages may give more information on errors in other modules.

Check syslog messages. Please report this problem.

267589 :launch_fed_prog: call to rpc.fed failed for program
<%s>, step <%s>

Launching of fed program failed due to a failure of ucmmmd to communicate with the rpc.fed daemon. If the rpc.fed process died, this might lead to a subsequent reboot of the node.

Examine other syslog messages occurring at about the same time to see if the problem can be identified and if it recurs. Save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance.

267673 :Validation failed. ORACLE binaries not found
ORACLE_HOME=%s

Oracle binaries not found under ORACLE_HOME. ORACLE_HOME specified for the resource is indicated in the message. HA-Oracle will not be able to manage Oracle if ORACLE_HOME is incorrect.

Specify correct ORACLE_HOME when creating resource. If resource is already created, please update resource property "ORACLE_HOME".

267724 :stat of file system %s failed.

HA-NFS fault monitor reports a probe failure on a specified file system.

Make sure the specified path exists.

269240 :clconf: write_ccr routine shouldn't be called from
kernel

Routine write_ccr that writes a clconf tree out to CCR should not be called from kernel.

No action required. This is informational message.

269902 :reservation fatal error(%s) - Unable to find gdev property

A required rawdisk device group property is missing.

Executing `"/usr/cluster/bin/scgdevs -L"` on this node should generate the required property. If this successfully creates the required property, it should be possible to retry the failed operation. If the message specifies the "node_join" transition, then this node may be unable to access shared devices. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access shared devices. In either case, it may be possible to reacquire access to shared devices by executing `"/usr/cluster/lib/sc/run_reserve -c node_join"` on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group has failed to start on this node. If another node was available to host the device group, then it should have been started on that node. If desired, it may be possible to switch the device group to this node with the `scswitch` command. If no other node was available, then the device group will not have been started. The `scswitch` command may be used to retry the attempt to start the device group. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group has failed. The desired action may be retried. If the problem persists, contact your authorized Sun service provider to determine whether a workaround or patch is available.

270043 :reservation warning(%s) - MHIOCENFAILFAST error will retry in %d seconds

The device fencing program has encountered errors while trying to access a device. The failed operation will be retried

This is an informational message, no user action is needed.

272238 :reservation warning(%s) - MHIOCGRP_RESERVE error will retry in %d seconds

The device fencing program has encountered errors while trying to access a device. The failed operation will be retried

This is an informational message, no user action is needed.

272732 :scvxmlg warning - chmod(%s) failed

The program responsible for maintaining the VxVM namespace was unable to access the global device namespace. If configuration changes were recently made to VxVM diskgroups or volumes, this node may be unaware of those changes. Recently created volumes may be inaccessible from this node.

Verify that the `/global/.devices/node@N` (N = this node's node number) is mounted globally and is accessible. If no configuration changes have been

recently made to VxVM diskgroups or volumes and all volumes continue to be accessible from this node, then no further action is required. If changes have been made, the device namespace on this node can be updated to reflect those changes by executing `"/usr/cluster/lib/dcs/scvxvmlg"`. If the problem persists, contact your authorized Sun service provider to determine whether a workaround or patch is available.

273018 :INTERNAL ERROR CMM: Failure starting CMM.

An instance of the userland CMM encountered an internal initialization error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

273354 :CMM: Node %s (nodeid = %d) is dead.

The specified node has died. It is guaranteed to be no longer running and it is safe to take over services from the dead node.

The cause of the node failure should be resolved and the node should be rebooted if node failure is unexpected.

273638 :The entry %s and entry %s in property %s have the same port number: %d.

The two entries in the list property duplicate port number.

Remove one of the entries or change its port number.

274421 :Port %d%c%s is listed twice in property %s, at entries %d and %d.

The port number in the message was listed twice in the named property, at the list entry locations given in the message. A port number should only appear once in the property.

Specify the property with only one occurrence of the port number.

274605 :Server is online.

Informational message. Oracle server is online.

None

274887 :clcomm: solaris xdoor: rejected invo: door_return returned, errno = %d

An unusual but harmless event occurred. System operations continue unaffected.

No user action is required.

274901 :Invalid protocol %s given as part of property %s.

The property named does not have a legal value.

Assign the property a legal value.

276380 : "pmfadm -k": Error signaling <%s>: %s

An error occurred while rpc.pmfad attempted to send a signal to one of the processes of the given tag. The reason for the failure is also given. The signal was sent as a result of a "pmfadm -k" command.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

276672 :reservation error(%s) - did_get_did_path() error

The device fencing program has suffered an internal error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available. Copies of /var/adm/messages from all nodes should be provided for diagnosis. It may be possible to retry the failed operation, depending on the nature of the error. If the message specifies the "node_join" transition, then this node may be unable to access shared devices. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access shared devices. In either case, it may be possible to reacquire access to shared devices by executing "/usr/cluster/lib/sc/run_reserve -c node_join" on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group has failed to start on this node. If another node was available to host the device group, then it should have been started on that node. If desired, it may be possible to switch the device group to this node with the scswitch command. If no other node was available, then the device group will not have been started. The scswitch command may be used to retry the attempt to start the device group. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group has failed. The desired action may be retried.

277995 :(%s) msg of wrong version %d, expected %d

Expected to receive a message of a different version. udlmctl will fail.

Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

279084 :CMM: node reconfiguration #%lld completed.

The cluster membership monitor has processed a change in node or quorum status.

This is an informational message, no user action is needed.

279152 :listener %s probe successful.

Informational message. Listener monitor successfully completed first probe.

None

279309 :Failfast: Invalid failfast mode %s specified. Returning default mode PANIC.

An invalid value was supplied for the failfast mode. The software will use the default PANIC mode instead.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

280108 :clcomm: unable to rebind %s to name server

The name server would not rebind this entity.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

280256 :clnt_tp_create failed: %s

A client was not able to make an rpc connection to a server (rpc.pmfd, rpc.fed or rgmd) because it could not create the rpc handle. The rpc error is shown. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

281386 :dl_attach: DL_OK_ACK rtnd prim %u

Wrong primitive returned to the DL_ATTACH_REQ.

Reboot the node. If the problem persists, check the documentation for the private interconnect.

281428 :Failed to retrieve the resource group handle: %s.

An API operation on the resource group has failed.

For the resource group name, check the syslog tag. For more details, check the syslog messages from other components. If the error persists, reboot the node.

281680 :fatal: couldn't initialize ORB, possibly because machine is booted in non-cluster mode

The rgmd was unable to initialize its interface to the low-level cluster machinery. This might occur because the operator has attempted to start the rgmd on a node that is booted in non-cluster mode. The rgmd will produce a core file, and in some cases it might cause the node to halt or reboot to avoid data corruption.

If the node is in non-cluster mode, boot it into cluster mode before attempting to start the rgmd. If the node is already in cluster mode, save a copy of the /var/

adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

281819 :%s exited with error %s in step %s

A ucmm step execution failed in the indicated step.

None. See /var/adm/messages for previous errors and report this problem if it occurs again during the next reconfiguration.

282406 :fork1 returned %d. Exiting.

clxecd program has encountered a failed fork1(2) system call. The error message indicates the error number for the failure.

If the error number is 12 (ENOMEM), install more memory, increase swap space, or reduce peak memory consumption. If error number is something else, contact your authorized Sun service provider to determine whether a workaround or patch is available.

282508 :INTERNAL ERROR: r_state_at_least: state <%s> (%d)

A non-fatal internal error has occurred in the rgmd state machine.

Since this problem might indicate an internal logic error in the rgmd, please save a copy of the /var/adm/messages files on all nodes, the output of an scstat -g command, and the output of a scrgadm -pvv command. Report the problem to your authorized Sun service provider.

282828 :reservation warning(%s) - MHIocrelease error will retry in %d seconds

The device fencing program has encountered errors while trying to access a device. The failed operation will be retried

This is an informational message, no user action is needed.

283262 :HA: rm_state_machine::service_suicide() not yet implemented

Unimplemented feature was activated.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

284006 :reservation fatal error(UNKNOWN) - Out of memory

The device fencing program has been unable to allocate required memory.

Memory usage should be monitored on this node and steps taken to provide more available memory if problems persist. Once memory has been made available, the following steps may need to be taken: If the message specifies the

"node_join" transition, then this node may be unable to access shared devices. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access shared devices. In either case, access to shared devices can be reacquired by executing `"/usr/cluster/lib/sc/run_reserve -c node_join"` on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group has failed to start on this node. If another node was available to host the device group, then it should have been started on that node. The device group can be switched back to this node if desired by using the `scswitch` command. If no other node was available, then the device group will not have been started. The `scswitch` command may be used to start the device group. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group has failed. The desired action may be retried.

284644 :Warning: node %d has a weight assigned to it for property %s, but node %d is not in the %s for resource %s.

A node has a weight assigned but the resource can never be active on that node, therefore it doesn't make sense to assign that node a weight.

This is an informational message, no user action is needed. Optionally, the weight that is assigned to the node can be omitted.

286722 :scvxvmlg error - remove(%s) failed

The program responsible for maintaining the VxVM namespace was unable to access the global device namespace. If configuration changes were recently made to VxVM diskgroups or volumes, this node may be unaware of those changes. Recently created volumes may be inaccessible from this node.

Verify that the `/global/.devices/node@N` (N = this node's node number) is mounted globally and is accessible. If no configuration changes have been recently made to VxVM diskgroups or volumes and all volumes continue to be accessible from this node, then no further action is required. If changes have been made, the device namespace on this node can be updated to reflect those changes by executing `"/usr/cluster/lib/dcs/scvxvmlg"`. If the problem persists, contact your authorized Sun service provider to determine whether a workaround or patch is available.

286807 :clnt_tp_create_timed of program %s failed %s.

HA-NFS fault monitor was not able to make an rpc connection to an nfs server.

Save the `/var/adm/messages` file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

289194 :Can't perform failover: Failover mode set to NONE.

Cannot perform failover of the data service. Failover mode is set to NONE.

This is informational message. If failover is desired, then set the Failover_mode value to SOFT or HARD using scrgadm(1M).

289503 :Unable to re-compute NFS resource list.

The list of HA-NFS resources online on the node has gotten corrupted.

Make sure there is space available in /tmp. If the error is showing up despite that, reboot the node.

290735 :Conversion of hostnames failed.

Data service is unable to convert the specified hostname into an IP address.

Check the syslog messages that occurred just before, to check whether there is any internal error. If there is, then contact your authorized Sun service provider. Otherwise, if the logical host and shared address entries are specified in the /etc/inet/hosts file, check these entries are correct. If this is not the reason then check the health of the name server.

291986 :dl_bind ack bad len %d

Sanity check. The message length in the acknowledgment to the bind request is different from what was expected. We are trying to open a fast path to the private transport adapters.

Reboot of the node might fix the problem.

292013 :clcomm: UioBuf: uio was too fragmented - %d

The system attempted to use a uio that had more than DEF_IOV_MAX fragments.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

295666 :clcomm: setrlimit(RLIMIT_NOFILE): %s

During Galileo initialization within this user process, the setrlimit call failed with the specified error.

Read the man page for setrlimit for a more detailed description of the error.

295838 :Listener %s started.

Informational message. HA-Oracle successfully started Oracle listener.

None

297061 :clcomm: can't get new reference

An attempt was made to obtain a new reference on a revoked handler.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

297139 :CCR: More than one data server has override flag set for the table %s. Using the table from node %s.

297178 :Error opening procfs control file <%s> for tag <%s>: %s

The rpc.pmfd server was not able to open a procfs control file, and the system error is shown. procfs control files are required in order to monitor user processes.

Investigate if the machine is running out of memory. If this is not the case, save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

297325 :The node portion of %s at position %d in property %s is not a valid node identifier or node name.

An invalid node was specified for the named property. The position index, which starts at 0 for the first element in the list, indicates which element in the property list was invalid.

Specify a valid node instead.

297536 :Could not host device service %s because this node is being shut down

An attempt was made to start a device group on this node while the node was being shutdown.

If the node was not being shutdown during this time, or if the problem persists, please contact your authorized Sun service provider to determine whether a workaround or patch is available.

297867 :(%s) t_bind: tli error: %s

Call to t_bind() failed. The "t_bind" man page describes possible error codes. udlmctl will exit.

Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

298911 :setrlimit: %s

The rpc.pmfd server was not able to set the limit of files open. The message contains the system error. This happens while the server is starting up, at boot time. The server does not come up, and an error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

299417 :in libsecurity strong Unix authorization failed

A server (rgmd) refused an rpc connection from a client because it failed the Unix authentication. This happens if a caller program using scha public api, either in its C form or its CLI form, is not running as root or is not making the rpc call over the loopback interface. An error message is output to syslog.

Check that the calling program using the scha public api is running as root and is calling over the loopback interface. If both are correct, save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

Message IDs 300000 - 399999

Error Message List

The following list is ordered by the message ID.

300397 :resource %s property changed.

This is a notification from the rgmd that a resource's property has been edited by the cluster administrator. This may be used by system monitoring tools.

This is an informational message, no user action is needed.

301573 :clcomm: error in copyin for cl_change_flow_settings

The system failed a copy operation supporting a flow control state change.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

301603 :fatal: cannot create any threads to handle switchback

The rgmd was unable to create a sufficient number of threads upon starting up. This is a fatal error. The rgmd will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

Make sure that the hardware configuration meets documented minimum requirements. Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

301635 :clexecd: close returned %d. Exiting.

clexecd program has encountered a failed thr_create(3THR) system call. The error message indicates the error number for the failure.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

302670 :udlm_setup_port: fcntl: %d

A server was not able to execute fcntl(). udlm exits and the node aborts and panics.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

303231 :mount_client_impl::remove_client() failed attempted RM change_repl_prov_status() to remove client, spec %s, name %s

The system was unable to remove a PXFS replica on the node that this message was seen.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

304365 :clcomm: Could not create any threads for pool %d

The system creates server threads to support requests from other nodes in the cluster. The system could not create any server threads during system startup. This is caused by a lack of memory.

There are two solutions. Install more memory. Alternatively, take steps to reduce memory usage. Since the creation of server threads takes place during system startup, application memory usage is normally not a factor.

305298 :cm_callback_impl abort_trans: exiting

ucmm callback for abort transition failed.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

307195 :clcomm: error in copyin for cl_read_flow_settings

The system failed a copy operation supporting flow control state reporting.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

307531 :open: %s

While starting up, the rgmd daemon was not able to open /dev/console. The message contains the system error. This will prevent the rgmd from starting on this node.

Examine other syslog messages occurring at about the same time to see if the problem can be identified. Save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance in diagnosing and correcting the problem.

308612 :reservation error(%s) - my_map_to_did_device() error in is_scsi3_disk()

The device fencing program has suffered an internal error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available. Copies of /var/adm/messages from all nodes should be provided for diagnosis. It may be possible to retry the failed operation, depending on the nature of the error. If the message specifies the "node_join" transition, then this node may be unable to access shared devices. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access shared devices. In either case, it may be possible to reacquire access to shared devices by executing "/usr/cluster/lib/sc/run_reserve -c node_join" on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group has failed to start on this node. If another node was available to host the device group, then it should have been started on that node. If desired, it may be possible to switch the device group to this node with the scswitch command. If no other node was available, then the device group will not have been started. The scswitch command may be used to retry the attempt to start the device group. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group has failed. The desired action may be retried.

308800 :ERROR: rebalance: <%s> is pending_methods on node <%d>

An internal error has occurred in the locking logic of the rgmd. This error should not occur. It may prevent the rgmd from bringing the indicated resource group online.

Look for other syslog error messages on the same node. Save a copy of the /var/adm/messages files on all nodes, and report the problem to your authorized Sun service provider.

310953 :clnt_control of program %s failed %s.

HA-NFS fault monitor failed to reset the retry timeout for retransmitting the rpc request.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

311463 :Failover attempt failed.

The failover attempt of the resource is rejected or encountered an error.

For more detailed error message, check the syslog messages. Check whether the Pingpong_interval has appropriate value. If not, adjust it using scrgadm(1M). Otherwise, use scswitch to switch the resource group to a healthy node.

311808 :Can not open /etc/mnttab: %s

Error in open /etc/mnttab, the error message is followed.

Check with system administrator and make sure /etc/mnttab is properly defined.

312053 :Cannot execute %s: %s.

Failure in executing the command.

Check the syslog message for the command description. Check whether the system is low in memory or the process table is full and take appropriate action. Make sure that the executable exists.

313867 :Unknown step: %s

Request to run an unknown udln step.

None.

314314 :prog <%s> step <%s> terminated due to receipt of signal <%d>

ucmmd step terminated due to receipt of a signal.

Examine other syslog messages occurring at about the same time to see if the problem can be identified and if it recurs. Save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance.

314356 :resource %s enabled.

This is a notification from the rgmd that the operator has enabled a resource. This may be used by system monitoring tools.

This is an informational message, no user action is needed.

315446 :node id <%d> is out of range

The low-level cluster machinery has encountered an error.

Look for other syslog messages occurring just before or after this one on the same node; they may provide a description of the error.

316344 :Adapter %s is not a valid NAFO name on this node.

Validation of the adapter information has failed. The specified NAFO group does not exist on this node.

Create appropriate NAFO group on this node or recreate the logical host with correct NAFO.

319047 :CMM: Issuing a SCSI2 Trown failed for quorum device %s with error %d.

This node encountered the specified error while issuing a SCSI2 Trown operation on the indicated quorum device. This will cause the node to conclude that it has been unsuccessful in preempting keys from the quorum device, and therefore the partition to which it belongs has been preempted. If a cluster gets divided into two or more disjoint subclusters, exactly one of these must survive as the operational cluster. The surviving cluster forces the other subclusters to abort by grabbing enough votes to grant it majority quorum. This is referred to as preemption of the losing subclusters.

If the error encountered is EACCES, then the SCSI2 command could have failed due to the presence of SCSI3 keys on the quorum device. Scrub the SCSI3 keys off of it, and reboot the preempted nodes.

319048 :CCR: Cluster has lost quorum while updating table %s, it is possibly in an inconsistent state - ABORTING.

The cluster lost quorum while the indicated table was being changed, leading to potential inconsistent copies on the nodes.

Check if the indicated table are consistent on all the nodes in the cluster, if not, boot the cluster in -x mode to restore the indicated table from backup. The CCR tables are located at /etc/cluster/ccr/.

319375 :clxecd: wait_for_signals got NULL.

clxecd problem encountered an error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

321245 :resource <%s> is disabled but not offline

While attempting to execute an operator-requested enable or disable of a resource, the rgmd has found the indicated resource to have its Onoff_switch property set to DISABLED, yet the resource is not offline. This suggests corruption of the RGM's internal data and will cause the enable or disable action to fail.

This may indicate an internal error or bug in the rgmd. Contact your authorized Sun service provider for assistance in diagnosing and correcting the problem.

321667 :clcomm: cl_comm: not booted in cluster mode.

Attempted to load the cl_comm module when the node was not booted as part of a cluster.

Users should not explicitly load this module.

322675 :Some NFS system daemons are not running.

HA-NFS fault monitor checks the health of statd, lockd, mountd and nfsd daemons on the node. It detected that one or more of these are not currently running.

No action. The monitor would restart these. If it doesn't, reboot the node.

322797 :Error registering provider "%s" with the framework.

The device configuration system on this node has suffered an internal error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

322862 :clcomm: error in copyin for cl_read_threads_min

The system failed a copy operation supporting flow control state reporting.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

322879 :clcomm: Invalid copyargs: node %d pid %d

The system does not support copy operations between the kernel and a user process when the specified node is not the local node.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

322908 :CMM: Failed to join the cluster: error = %d.

The local node was unsuccessful in joining the cluster.

There may be other related messages on this node that may indicate the cause of this failure. Resolve the problem and reboot the node.

323498 :libsecurity: NULL RPC to program %ld failed will retry %s

A client of the rpc.pmfd, rpc.fed or rgmd server was not able to initiate an rpc connection, because it could not execute a test rpc call, and the program will retry to establish the connection. The message shows the specific rpc error. The program number is shown. To find out what program corresponds to this number, use the rpcinfo command. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

323966 :NAFO group %s has status %s so no action will be taken.

The status of the NAFO group has become stable.

This is an informational message, no user action is needed.

324478 :(%s): Error %d from read

An error was encountered in the clexecd program while reading the data from the worker process.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

325322 :clcomm: error in copyin for state_resource_pool

The system failed a copy operation supporting statistics reporting.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

326043 :reservation fatal error(%s) - release_resv_lock()
returned exception

The device fencing program has suffered an internal error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available. Copies of /var/adm/messages from all nodes should be provided for diagnosis. It may be possible to retry the failed operation, depending on the nature of the error. If the message specifies the "node_join" transition, then this node may be unable to access shared devices. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access shared devices. In either case, it may be possible to reacquire access to shared devices by executing "/usr/cluster/lib/sc/run_reserve -c node_join" on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group has failed to start on this node. If another node was available to host the device group, then it should have been started on that node. If desired, it may be possible to switch the device group to this node with the scswitch command. If no other node was available, then the device group will not have been started. The scswitch command may be used to retry the attempt to start the device group. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group has failed. The desired action may be retried.

327057 :SharedAddress stopped.

The stop method is completed and the resource is stopped.

This is informational message. No user action required.

329286 :(%s) instead of UDLM_ACK got a %d

Did not receive an acknowledgement from udln as was expected.

None.

329429 :reservation fatal error(%s) - host_name not specified

The device fencing program has suffered an internal error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available. Copies of /var/adm/messages from all nodes should be provided for diagnosis. It may be possible to retry the failed operation, depending on the nature of the error. If the message specifies the "node_join" transition, then this node may be unable to access shared devices. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access shared devices. In either case, it may be possible to reacquire access to shared devices by executing "/usr/cluster/lib/sc/run_reserve -c node_join" on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group has failed to start on this node. If another node was available to host the device group, then it should have been started on that node. If desired, it may be possible to switch the device group to this node with the scswitch command. If no other node was available, then the device group will not have been started. The scswitch command may be used to retry the attempt to start the device group. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group has failed. The desired action may be retried.

329496 :unlatch_intention(): IDL exception when communicating to node %d

An inter-node communication failed, probably because a node died.

No action is required; the rgmd should recover automatically.

329778 :clconf: Data length is more than max supported length in clconf_ccr read

In reading configuration data through CCR, found the data length is more than max supported length.

Check the CCR configuraton information.

329847 :Warning: node %d has a weight of 0 assigned to it for property %s.

The named node has a weight of 0 assigned to it. A weight of 0 means that no new client connections will be distributed to that node.

Consider assigning the named node a non-zero weight.

330063 :error in vop open %x

Opening a private interconnect interface failed.

Reboot of the node might fix the problem.

330526 :CMM: Number of steps specified in registering callback = %d; should be <= %d.

The number of steps specified during registering a CMM callback exceeds the allowable maximum. This is an internal error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

331221 :CMM: Max detection delay specified is %ld which is larger than the max allowed %ld.

The maximum of the node down detection delays is larger than the allowable maximum. The maximum allowed will be used as the actual maximum in this case.

This is an informational message, no user action is needed.

333069 :Failed to retrieve nodeid for %s.

The nodeid for the given name could not be determined.

Make sure that the name given is a valid node identifier or node name.

334697 :Failed to retrieve the cluster property %s: %s.

The query for a property failed. The reason for the failure is given in the message.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

334992 :clutil: Adding deferred task after threadpool shutdown id %s

During shutdown this operation is not allowed.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

335206 :Failed to get host names from the resource.

Retrieving the IP addresses from the network resources from this resource group has failed.

Internal error or API call failure might be the reasons. Check the error messages that occurred just before this message. If there is internal error, contact your authorized Sun service provider. For API call failure, check the syslog messages from other components. For the resource name and resource group name, check the syslog tag.

335591 :Failed to retrieve the resource group property %s: %s.

An API operation has failed while retrieving the resource group property. Low memory or API call failure might be the reasons.

In case of low memory, the problem will probably be cured by rebooting. If the problem reoccurs, you might need to increase swap space by configuring additional swap devices. Otherwise, if it is API call failure, check the syslog messages from other components. For resource group name and the property name, check the current syslog message.

336860 :read %d for %snum_ports

Could not get information about the number of ports udlm uses from config file udlm.conf.

Check to make sure udlm.conf file exist and has entry for udlm.num_ports. If everything looks normal and the problem persists, contact your Sun service representative.

337008 :rgm_comm_impl::_unreferenced() called unexpectedly

The low-level cluster machinery has encountered a fatal error. The rgmd will produce a core file and will cause the node to halt or reboot to avoid the possibility of data corruption.

Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

337166 :Error setting environment variable %s.

An error occurred while setting the environment variable LD_LIBRARY_PATH. This is required by the fault monitor for the nsldap data service. The fault monitor appends the ldap server root path, including the lib directory, to the LD_LIBRARY_PATH environment variable

Check that there is a lib directory under the server root of the nsldap data service which pertains to this resource. If this directory has been removed, then it must be replaced by reinstalling Netscape Directory Server, or whatever other means are appropriate.

337212 :resource type %s removed.

This is a notification from the rgmd that a resource type has been deleted. This may be used by system monitoring tools.

This is an informational message, no user action is needed.

338839 :clexecd: Could not create thread. Error: %d. Sleeping for %d seconds and retrying.

clxecd program has encountered a failed thr_create() system call. The error message indicates the error number for the failure. It will retry the system call after specified time.

If the message is seen repeatedly, contact your authorized Sun service provider to determine whether a workaround or patch is available.

339424 :Could not host device service %s because this node is being removed from the list of eligible nodes for this service

A switchover/failover was attempted to a node that was being removed from the list of nodes that could host this device service.

This is an informational message, no user action is needed.

339521 :CCR: Lost quorum while starting to update table %s.

The cluster lost quorum when CCR started to update the indicated table.

Reboot the cluster.

339657 :Issuing a restart request.

This is informational message. We are about to call API function to request for restart. In case of failure, follow the syslog messages after this message.

No user action is needed.

339954 :fatal: cannot create any threads to launch callback methods

The rgmd was unable to create a thread upon starting up. This is a fatal error. The rgmd will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

Make sure that the hardware configuration meets documented minimum requirements. Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

340287 :idl_set_timestamp(): IDL Exception

The rgmd has encountered an error that prevents the scha_control function from successfully setting a ping-pong time stamp, presumably because a node died. This does not prevent the attempted failover from succeeding, but in the worst case, might prevent the anti-"pingpong" feature from working, which may permit a resource group to fail over repeatedly between two or more nodes.

Examine syslog output on the node that rebooted, to determine the cause of node death. The syslog output might indicate further remedial actions.

342336 :clcomm: Pathend %p: path_down not allowed in state %d

The system maintains state information about a path. A path_down operation is not allowed in this state.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

343307 :Could not open file %s: %s.

System has failed to open the specified file.

Check whether the permissions are valid. This might be the result of a lack of system resources. Check whether the system is low in memory and take appropriate action. For specific error information, check the syslog message.

344470 :Unable to get status for NAFO group %s.

The specified NAFO group is not in functional state. Logical host resource can't be started without a functional NAFO.

LogicalHostname resource will not be brought online on this node. Check the messages(pnmd errors) that encountered just before this message for any NAFO or adapter problem. Correct the problem and rerun the scrgadm.

345342 :Failed to connect to %s port %d.

The data service fault monitor probe was trying to connect to the host and port specified and failed. There may be a prior message in syslog with further information.

Make sure that the port configuration for the data service matches the port configuration for the underlying application.

346036 :libsecurity: unexpected getnetconfigent error

A client of the rpc.pmfd, rpc.fed or rgmd server was not able to initiate an rpc connection, because it could not get the network information. The pmfadm or scha command exits with error. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

347091 :resource type %s added.

This is a notification from the rgmd that the operator has created a new resource type. This may be used by system monitoring tools.

This is an informational message, no user action is needed.

347344 :Did not find a valid port number to match field <%s> in configuration file <%s>: %s.

A failure occurred extracting a port number for the field within the configuration file. The field exists and the file exists and is accessible. The value for the field

may not exist or may not be an integer greater than zero. An error in environment may have occurred, indicated by a non-zero errno value at the end of the message.

Check to see if the value for the field in the configuration file exists and is an integer greater than zero. If there is an error in the field value, fix the value and retry the operation.

348240 :clxecd: putmsg returned %d.

clxecd program has encountered a failed putmsg(2) system call. The error message indicates the error number for the failure.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

353557 :Filesystem (%s) is locked and cannot be frozen

The file system has been locked with the _FIOLFS ioctl. It is necessary to perform an unlock _FIOLFS ioctl. The growfs(1M) or lockfs(1M) command may be responsible for this lock.

An _FIOLFS LOCKFS_ULOCK ioctl is required to unlock the file system.

355950 :HA: unknown invocation result status %d

An invocation completed with an invalid result status.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

356795 :CMM: Reconfiguration step %d was forced to return.

One of the CMM reconfiguration step transitions failed, probably due to a problem on a remote node. A reconfiguration is forced assuming that the CMM will resolve the problem.

This is an informational message, no user action is needed.

357263 :munmap: %s

The rpc.pmfd server was not able to delete shared memory for a semaphore, possibly due to low memory, and the system error is shown. This is part of the cleanup after a client call, so the operation might have gone through. An error message is output to syslog.

Investigate if the machine is running out of memory. If this is not the case, save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

357915 :Error: Unable to stat directory <%s> for scha_control timestamp file

The rgmd failed in a call to stat(2) on the local node. This may prevent the anti-"pingpong" feature from working, which may permit a resource group to fail over repeatedly between two or more nodes. The failure of the stat call might indicate a more serious problem on the node.

Examine other syslog messages occurring around the same time on the same node, to see if the source of the problem can be identified.

358211 :monitor_check: the failover requested by scha_control for resource <%s>, resource group <%s> was not completed because of error: %s

A scha_control(1HA,3HA) GIVEOVER attempt failed, due to the error listed.

Examine other syslog messages on all cluster members that occurred about the same time as this message, to see if the problem that caused the error can be identified and repaired.

359648 :Service has failed.

Probe is detected a failure in the data service. The data service cannot be restarted on the same node, since there are frequent failures. Probe is setting resource status as failed.

Wait for the fault monitor to failover the data service. Check the syslog messages and configuration of the data service.

361048 :ERROR: rgm_run_state() returned non-zero while running boot methods

The rgmd state machine has encountered an error on this node.

Look for preceding syslog messages on the same node, which may provide a description of the error.

361489 :in libsecurity __rpc_negotiate_uid failed for transport %s

A server (rpc.pmfd, rpc.fed or rgmd) was not able to start because it could not establish a rpc connection for the network. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

362463 :clcomm: Endpoint %p: path_down not allowed in state %d

The system maintains information about the state of an Endpoint. The path_down operation is not allowed in this state.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

362519 :dl_attach: DL_ERROR_ACK bad PPA

Could not attach to the physical device. We are trying to open a fast path to the private transport adapters.

Reboot of the node might fix the problem.

362657 :Error when sending response from child process: %m

Error occurred when sending message from fault monitor child process. Child process will be stopped and restarted.

If error persists, then disable the fault monitor and resport the problem.

363505 :check_for_ccrdata failed strdup for (%s)

Call to strdup failed. The "strdup" man page describes possible reasons.

Install more memory, increase swap space or reduce peak memory consumption.

363972 :reservation message(%s) - Waiting for reservation lock

Locking is used by the device fencing program to ensure correct behavior when different nodes see different cluster memberships. This node is waiting for an instance of the device fencing program on another node to complete.

The lock should eventually be granted. If node failures are involved, the lock will not be granted until node deaths are assured, which may take a few minutes. If the lock is eventually granted, no user action is required. If the lock is not granted, your authorized Sun service provider should be contacted to help diagnose the problem.

364188 :Validation failed. Listener_name not set

"Listener_name" property of the resource is not set. HA-Oracle will not be able to manage Oracle listener if Listener name is not set.

Specify correct "Listener_name" when creating resource. If resource is already created, please update resource property.

364510 :The specified Oracle dba group id (%s) does not exist

Group id of oracle dba does not exist.

Make sure /etc/nswitch.conf and /etc/group files are valid and have correct information to get the group id of dba.

366225 :Listener %s stopped successfully

Informational message. HA-Oracle successfully stopped Oracle listener.

None

367617 :reservation fatal error(%s) - Invalid file format "%s"

The device fencing program has suffered an internal error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available. Copies of /var/adm/messages from all nodes should be provided for diagnosis. It may be possible to retry the failed operation, depending on the nature of the error. If the message specifies the "node_join" transition, then this node may be unable to access shared devices. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access shared devices. In either case, it may be possible to reacquire access to shared devices by executing "/usr/cluster/lib/sc/run_reserve -c node_join" on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group has failed to start on this node. If another node was available to host the device group, then it should have been started on that node. If desired, it may be possible to switch the device group to this node with the scswitch command. If no other node was available, then the device group will not have been started. The scswitch command may be used to retry the attempt to start the device group. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group has failed. The desired action may be retried.

367864 :svc_init failed.

The rpc.pmfd server was not able to initialize server operation. This happens while the server is starting up, at boot time. The server does not come up, and an error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

368363 :Failed to retrieve the current primary node.

Cannot retrieve the current primary node for the given resource group.

Check the syslog messages that occurred just before this message, to see whether there is any internal error has occurred. If it is, contact your authorized Sun service provider. Otherwise, Check if the resource group is in STOP_FAILED state. If it is, then clear the state and bring the resource group online.

368819 :t_rcvudata in recv_request: %s

Call to t_rcvudata() failed. The "t_rcvudata" man page describes possible error codes. udln will exit and the node will abort.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

369460 :udln_send_reply %s: udp is null!

Can not communicate with udlmctl because the address to send to is null.

None. udlm will handle this error.

370949 :created %d threads to launch resource callback methods;
desired number = %d

The rgmd was unable to create the desired number of threads upon starting up. This is not a fatal error, but it may cause RGM reconfigurations to take longer because it will limit the number of tasks that the rgmd can perform concurrently.

Make sure that the hardware configuration meets documented minimum requirements. Examine other syslog messages on the same node to see if the cause of the problem can be determined.

371297 :%s: Invalid command line option. Use -S for secure mode

rpc.sccheckd should always be invoked in secure mode. If this message shows up, someone has modified configuration files that affects server startup.

Reinstall cluster packages or contact your service provider.

371369 :CCR: CCR data server on node %s unreachable while
updating table %s.

While the TM was updating the indicated table in the cluster, the specified node went down and has become unreachable.

The specified node needs to be rebooted.

372880 :CMM: Quorum device %ld (gdevname %s) can not be acquired
by the current cluster members. This quorum device is held by
node%s %s.

This node does not have its reservation key on the specified quorum device, which has been reserved by the specified node or nodes that the local node can not communicate with. This indicates that in the last incarnation of the cluster, the other nodes were members whereas the local node was not, indicating that the CCR on the local node may be out-of-date. In order to ensure that this node has the latest cluster configuration information, it must be able to communicate with at least one other node that was a member of the previous cluster incarnation. These nodes holding the specified quorum device may either be down or there may be up but the interconnect between them and this node may be broken.

If the nodes holding the specified quorum devices are up, then fix the interconnect between them and this node so that communication between them is restored. If the nodes are indeed down, boot one of them.

372887 :HA: repl_mgr: exception occurred while invoking RMA

An unrecoverable failure occurred in the HA framework.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

373148 :The port portion of %s at position %d in property %s is not a valid port.

The property named does not have a legal value. The position index, which starts at 0 for the first element in the list, indicates which element in the property list was invalid.

Assign the property a legal value.

373816 :clcomm: copyinstr: max string length %d too long

The system attempted to copy a string from user space to the kernel. The maximum string length exceeds length limit.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

374006 :prog <%s> failed on step <%s> retcode <%d>

ucmmd step failed on a step.

Examine other syslog messages occurring at about the same time to see if the problem can be identified and if it recurs. Save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance.

374377 :NAFO Failure.

The NAFO group hosting the LogicalHostname has failed.

The LogicalHostname resource would be failed over to a different node. If that fails, check the system logs for other messages. Also, correct the networking problem on the node so that the NAFO group in question is healthy again.

374738 :dl_bind: DL_BIND_ACK bad sap %u

SAP in acknowledgment to bind request is different from the SAP in the request. We are trying to open a fast path to the private transport adapters.

Reboot of the node might fix the problem.

377245 :request addr > max "%s & %s"

Error from udlm on an address request. Udlm exits and the nodes aborts and panics.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

377347 :CMM: Node %s (nodeid = %ld) is up; new incarnation number = %ld.

The specified node has come up and joined the cluster. A node is assigned a unique incarnation number each time it boots up.

This is an informational message, no user action is needed.

378427 :prog <%s> step <%s> terminated due to receipt of signal

ucmmd step terminated due to receipt of a signal.

Examine other syslog messages occurring at about the same time to see if the problem can be identified and if it recurs. Save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance.

378807 :clexecd: %s: sigfillset returned %d. Exiting.

clexecd program has encountered a failed sigfillset(3C) system call. The error message indicates the error number for the failure.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

378872 :%s operation failed: %s.

Specified system operation could not complete successfully.

This is as an internal error. Contact your authorized Sun service provider with the following information. 1) Saved copy of /var/adm/messages file. 2) Output of "ifconfig -a" command.

379450 :reservation fatal error(%s) - fenced_node not specified

The device fencing program has suffered an internal error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available. Copies of /var/adm/messages from all nodes should be provided for diagnosis. It may be possible to retry the failed operation, depending on the nature of the error. If the message specifies the "node_join" transition, then this node may be unable to access shared devices. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access shared devices. In either case, it may be possible to reacquire access to shared devices by executing "/usr/cluster/lib/sc/run_reserve -c node_join" on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group has failed to start on this node. If another node was available to host the device group, then it should have been started on that node. If desired, it may be possible to switch the device group to this node with the scswitch command. If no other node was available, then the device group will not have been started. The scswitch command may be used to retry the attempt to start the device group. If the failure occurred during

the "primary_to_secondary" transition, then the shutdown or switchover of a device group has failed. The desired action may be retried.

380064 :switchback attempt failed on resource group <%s> with error <%s>

The rgmd was unable to failback the specified resource group to a more preferred node. The additional error information in the message explains why.

Examine other syslog messages occurring around the same time on the same node. These messages may indicate further action.

380365 :(%s) t_rcvudata, res %d, flag %d: tli error: %s

Call to t_rcvudata() failed. The "t_sndudata" man page describes possible error codes. udlmctl will exit.

Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

380897 :rebalance: WARNING: resource group <%s> is <%s> on node <%d>, resetting to OFFLINE.

The resource group has been found to be in the indicated state and is being reset to OFFLINE. This message is a warning only and should not adversely affect the operation of the RGM.

Since this problem might indicate an internal logic error in the rgmd, please save a copy of the /var/adm/messages files on all nodes, the output of an scstat -g command, and the output of a scrghadm -pvv command. Report the problem to your authorized Sun service provider.

381244 :in libsecurity mkdir of %s failed: %s

The rpc.pmfd, rpc.fed or rgmd server was not able to create a directory to contain "cache" files for rpcbind information. The affected component should still be able to function by directly calling rpcbind.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

381765 :sema_post: %s

The rpc.pmfd server was not able to act on a semaphore. The message contains the system error. The server does not perform the action requested by the client, and an error message is output to syslog.

Investigate if the machine is running out of memory. If this is not the case, save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

382169 :Share path name %s not absolute.

A path specified in the dfstab file does not begin with "/"

Only absolute path names can be shared with HA-NFS.

382252 :Share path %s: file system %s is not mounted.

The specified file system, which contains the share path specified, is not currently mounted.

Correct the situation with the file system so that it gets mounted.

382995 :ioctl in negotiate_uid failed

Call to ioctl() failed. The "ioctl" man page describes possible error codes. ucmmd will exit and the node will abort.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

383706 :NULL value returned for the resource property %s.

NULL value was specified for the extension property of the resource.

For the property name check the syslog message. Any of the following situations might have occurred. Different user action is needed for these different scenarios.

1) If a new resource is created or updated, check whether the value of the extension property is empty. If it is, provide valid value using scrgadm(1M). 2)

For all other cases, treat it as an Internal error. Contact your authorized Sun service provider.

384549 :CCR: Could not backup the CCR table %s errno = %d.

The indicated error occurred while backing up indicated CCR table on this node. The errno value indicates the nature of the problem. errno values are defined in the file /usr/include/sys/errno.h. An errno value of 28(ENOSPC) indicates that the root file system on the indicated node is full. Other values of errno can be returned when the root disk has failed(EIO) or some of the CCR tables have been deleted outside the control of the cluster software(ENOENT).

There may be other related messages on this node, which may help diagnose the problem, for example: If the root file system is full on the node, then free up some space by removing unnecessary files. If the root disk on the afflicted node has failed, then it needs to be replaced. If the indicated CCR table was accidentally deleted, then boot this node in -x mode to restore the indicated CCR table from other nodes in the cluster or backup. The CCR tables are located at /etc/cluster/ccr/.

384621 :RDBMS probe successful

This message indicates that Fault monitor has successfully probed the RDBMS server

No action required. This is informational message.

385407 :t_alloc (open_cmd_port) failed with errno%d

Call to t_alloc() failed. The "t_alloc" man page describes possible error codes. ucmmmd will exit and the node will abort.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

385902 :pmf_search_children: Error signaling <%s>: %s

An error occurred while rpc.pmfd attempted to send a signal to one of the processes of the given tag. The reason for the failure is also given. The signal was sent to the process as a result of some event external to rpc.pmfd. rpc.pmfd "intercepted" the signal, and is trying to pass the signal on to the monitored process.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

386024 :ERROR: rebalance: duplicate nodeid <%d> in Nodelist of resource group <%s>; continuing

The same nodename appears twice in the Nodelist of the given resource group. Although non-fatal, this should not occur and may indicate an internal logic error in the rgmd.

Use scrgadm -pv to check the Nodelist of the affected resource group. Save a copy of the /var/adm/messages files on all nodes, and report the problem to your authorized Sun service provider.

386072 :chdir: %s

The rpc.pmfd server was not able to change directory. The message contains the system error. The server does not perform the action requested by the client, and an error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

386908 :Resource is already stopped.

An attempt was made to stop a resource that has already been stopped.

Using the ps command check to make sure that all processes for the Data Service have been stopped. Check syslog for any possible errors which may have occurred just before this message. If everything appears to be correct, then no action is required.

387003 :CCR: CCR metadata not found.

The CCR is unable to locate its metadata.

Boot the offending node in -x mode to restore the indicated table from backup or other nodes in the cluster. The CCR tables are located at /etc/cluster/ccr/.

387150 :scvxvmlg warning - found no matching volume for device node %s, removing it

The program responsible for maintaining the VxVM device namespace has discovered inconsistencies between the VxVM device namespace on this node and the VxVM configuration information stored in the cluster device configuration system. If configuration changes were made recently, then this message should reflect one of the configuration changes. If no changes were made recently or if this message does not correctly reflect a change that has been made, the VxVM device namespace on this node may be in an inconsistent state. VxVM volumes may be inaccessible from this node.

If this message correctly reflects a configuration change to VxVM diskgroups then no action is required. If the change this message reflects is not correct, then the information stored in the device configuration system for each VxVM diskgroup should be examined for correctness. If the information in the device configuration system is accurate, then executing "/usr/cluster/lib/dcs/scvxvmlg" on this node should restore the device namespace. If the information stored in the device configuration system is not accurate, it must be updated by executing "/usr/cluster/bin/scconf -c -D name=diskgroup_name" for each VxVM diskgroup with inconsistent information.

387232 :resource %s monitor enabled.

This is a notification from the rgmd that the operator has enabled monitoring on a resource. This may be used by system monitoring tools.

This is an informational message, no user action is needed.

387288 :clcomm: Path %s online

A communication link has been established with another node.

No action required.

389221 :could not open configuration file: %s

The specified configuration file could not be opened.

Check if the configuration file exists and has correct permissions. If the problem persists, contact your Sun Service representative.

389231 :clcomm: inbound_invo::cancel:_state is 0x%x

The internal state describing the server side of a remote invocation is invalid when a cancel message arrives during processing of the remote invocation.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

389516 :NULL value returned for the extension property %s.

NULL value was specified for the extension property of the resource.

For the property name check the syslog message. Any of the following situations might have occurred. Different user action is needed for these different scenarios.

1) If a new resource is created or updated, check whether the value of the extension property is empty. If it is, provide valid value using scrgadm(1M). 2) For all other cases, treat it as an Internal error. Contact your authorized Sun service provider.

389901 :ext_props(): Out of memory

System runs out of memory in function ext_props().

Install more memory, increase swap space, or reduce peak memory consumption.

390691 :NFS daemon down

HA-NFS fault monitor detected that an nfs daemon died and will automatically restart it later.

No action required.

391738 :(%s) bad poll revent: %x (hex)

Call to poll() failed. The "poll" man page describes possible error codes. udlmctl will exit.

Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

392782 :Failed to retrieve the property %s for %s: %s.

API operation has failed in retrieving the cluster property.

For property name, check the syslog message. For more details about API call failure, check the syslog messages from other components.

393385 :Service daemon not running.

Process group has died and the data service's daemon is not running. Updating the resource status.

Wait for the fault monitor to restart or failover the data service. Check the configuration of the data service.

393960 :sigaction failed in set_signal_handler

The ucmmmd has failed to initialize signal handlers by a call to sigaction(2). The error message indicates the reason for the failure. The ucmmmd will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

Save a copy of the /var/adm/messages files on all nodes and of the ucmmmd core. Contact your authorized Sun service provider for assistance in diagnosing the problem.

394584 :PNM: could not start due to configuration error in %s

An error is encountered when processing the named configuration file. The PNM daemon is thus not fully functional and adapter monitoring and failover are not enabled.

Use the pnmset(1M) command to reset the PNM configuration on the node. The pnmmd daemon will be reinitialized automatically.

395353 :Failed to check whether the resource is a network address resource.

While retrieving the IP addresses from the network resources, the attempt to check whether the resource is a network resource or not has failed.

Internal error or API call failure might be the reasons. Check the error messages that occurred just before this message. If there is internal error, contact your authorized Sun service provider. For API call failure, check the syslog messages from other components.

396134 :Register callback with NAFO %s failed: Error %d.

LogicalHostname resource was unable to register with NAFO for status updates.

Most likely it is result of lack of system resources. Check for memory availability on the node. Reboot the node if problem persists.

397020 :unix DLM abort failed

Failed to abort unix dlm. This is an error that can be ignored.

None.

397340 :Monitor initialization error. Unable to open resource: %s Group: %s: error %d

Error occurred in monitor initialization. Monitor is unable to get resource information using API calls.

Check syslog messages for errors logged from other system modules. Stop and start fault monitor. If error persists then disable fault monitor and report the problem.

398345 :Error %d setting policy %d %s

This message appears when the customer is initializing or changing a scalable services load balancer, by starting or updating a service. An internal error happened while trying to change the load balancing policy.

This is an internal error and it could happen if another RGM are operation were happening at the same time. The user action is to try it again. If it happens when another RMG update is not happening, contact your Sun Service provider for help.

398878 :reservation fatal error(%s) -
dcs_get_service_parameters() error, returned %d

The device fencing program has suffered an internal error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available. Copies of /var/adm/messages from all nodes should be provided for diagnosis. It may be possible to retry the failed operation, depending on the nature of the error. If the message specifies the "node_join" transition, then this node may be unable to access shared devices. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access shared devices. In either case, it may be possible to reacquire access to shared devices by executing "/usr/cluster/lib/sc/run_reserve -c node_join" on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group has failed to start on this node. If another node was available to host the device group, then it should have been started on that node. If desired, it may be possible to switch the device group to this node with the scswitch command. If no other node was available, then the device group will not have been started. The scswitch command may be used to retry the attempt to start the device group. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group has failed. The desired action may be retried.

399216 :clexecd: Got an unexpected signal %d in process %s
(pid=%d, ppid=%d)

clexecd program got a signal indicated in the error message.

clexecd program will exit and node will be halted or rebooted to prevent data corruption. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

399266 :Cluster goes into pingpong booting because of failure of
method <%s> on resource <%s>. RGM is not aborting this node.

A stop method is failed and Failover_mode is set to HARD, but the RGM has detected this resource group falling into pingpong behavior and will not abort the node on which the resource's stop method failed. This is most likely due to the failure of both resource's start and stop methods.

Save a copy of /var/adm/messages, check for both failed start and stop methods of the failing resource, and make sure to have the failure corrected. Refer to the procedure for clearing the ERROR_STOP_FAILED condition on a resource group in the Sun Cluster Administration Guide to restart resource group.

399753 :CCR: CCR data server failed to register with CCR transaction manager.

The CCR data server on this node failed to join the cluster, and can only serve readonly requests.

There may be other related CCR messages on this and other nodes in the cluster, which may help diagnose the problem. It may be necessary to reboot this node or the entire cluster.

Message IDs 400000 - 499999

Error Message List

The following list is ordered by the message ID.

400592 :UNIX DLM is asking for a reconfiguration to recover from a communication error. This message is acceptable during a reconfiguration already in progress.

The cluster will reconfigure.

None.

401115 :t_rcvudata (recv_request) failed

Call to t_rcvudata() failed. The "t_rcvudata" man page describes possible error codes. ucmmmd will exit and the node will abort.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

401573 :INTERNAL ERROR: START method not registered for resource <%s>

A non-fatal internal error has occurred in the rgmd state machine.

Since this problem might indicate an internal logic error in the rgmd, please save a copy of the /var/adm/messages files on all nodes, the output of an scstat -g command, and the output of a scrgadm -pvv command. Report the problem to your authorized Sun service provider.

402289 :t_bind: %s

Call to `t_bind()` failed. The "`t_bind`" man page describes possible error codes. `udlm` will exit and the node will abort.

This is an internal error. Save the contents of `/var/adm/messages`, `/var/cluster/ucmm/ucmm_reconf.log` and `/var/cluster/ucmm/dlm/*/logs/*` from all the nodes and contact your Sun service representative.

402992 :Failfast: Destroying failfast unit %s while armed.

The specified failfast unit was destroyed while it was still armed.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

404309 :in libsecurity cred flavor is not AUTH_SYS

A server (`rpc.pmfd`, `rpc.fed` or `rgmd`) refused an `rpc` connection from a client because the authorization is not of `UNIX` type. An error message is output to `syslog`.

Save the `/var/adm/messages` file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

404866 :method_full_name: malloc failed

The `rgmd` server was not able to create the full name of the method, while trying to connect to the `rpc.fed` server, probably due to low memory. An error message is output to `syslog`.

Investigate if the host is running out of memory. If not save the `/var/adm/messages` file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

405508 :clcomm: Adapter %s has been deleted

A network adapter has been removed.

No action required.

405552 :Unable to contact fault monitor, restarting service.

The process monitoring facility tried to send a message to the fault monitor noting that the data service application died. It was unable to do so.

Since some part (daemon) of the application has failed, it would be restarted. If fault monitor is not yet started, wait for it to be started by Sun Cluster framework. If fault monitor has been disabled, enable it using `scswitch`.

406522 :resource group %s state on node %s change to
RG_PENDING_OFF_STOP_FAILED

This is a notification from the `rgmd` that a resource group has had a `STOP` method failure or timeout on one of its resources. This may be used by system

monitoring tools. The resource group will move to the `ERROR_STOP_FAILED` state on the given node.

Refer to the procedure for clearing the `ERROR_STOP_FAILED` condition on a resource group in the Sun Cluster Administration Guide.

406610 :st_ff_arm failed: %s

The `rpc.pmfd` server was not able to initialize the failfast mechanism. This happens while the server is starting up, at boot time. The server does not come up, and an error message is output to `syslog`. The message contains the system error.

Save the `/var/adm/messages` file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

406635 :fatal: joiners_run_boot_methods: exiting early because of unexpected exception

The low-level cluster machinery has encountered a fatal error. The `rgmd` will produce a core file and will cause the node to halt or reboot to avoid the possibility of data corruption.

Save a copy of the `/var/adm/messages` files on all nodes, and of the `rgmd` core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

408214 :Failed to create scalable service group %s: %s.

A call to the underlying scalable networking code failed.

Save a copy of the `/var/adm/messages` files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

408282 :clcomm: RT or TS classes not configured

The system requires either real time or time sharing thread scheduling classes for use in user processes. Neither class is available.

Configure Solaris to support either real time or time sharing or both thread scheduling classes for user processes.

408742 :svc_setschedprio: Could not save current scheduling parameters: %s

The server was not able to save the original scheduling mode. The system error message is shown. An error message is output to `syslog`.

Save the `/var/adm/messages` file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

409443 :fatal: unexpected exception in rgm_init_pres_state

This node encountered an unexpected error while communicating with other cluster nodes during a cluster reconfiguration. The rgmd will produce a core file and will cause the node to halt or reboot.

Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

410860 :lkcm_act: cm_reconfigure failed: %s

ucmm reconfiguration failed. This could also point to a problem with the interconnect components.

None if the next reconfiguration succeeds. If not, save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

411254 :Cannot change the nafo adapter on the local host.

A different nafo name for property NetIfList is specified in scrgadm command. The nafo name on local node is set at resource creation time. Users may only update the value of property NetIfList for adding a nafo group on a new node.

Rerun the scrgadm command with proper value of property NetIfList.

411369 :Not found clexecd on node %d for %d seconds. Giving up!

Could not find clexecd to execute the program on a node. Indicated giving up after retries.

This is an informational message, no user action is needed.

412106 :Internal Error. Unable to get fault monitor name

This is an internal error. Could not determine fault monitor program name.

Please report this problem.

412366 :setsid failed: %s

Failed to run the "setsid" command. The "setsid" man page describes possible error codes.

None. ucmmmd will exit.

412533 :clcomm: validate_policy: invalid relationship moderate %d low %d pool %d

The system checks the proposed flow control policy parameters at system startup and when processing a change request. The moderate server thread level cannot be less than the low server thread level.

No user action required.

412558 :inet addr %s length %d = %s

Information about hosts.

None.

413513 :INTERNAL ERROR Failfast: ff_impl_shouldnt_happen.

An internal error has occurred in the failfast software.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

413569 :CCR: Invalid CCR table : %s.

CCR could not find a valid version of the indicated table on the nodes in the cluster.

There may be other related messages on the nodes where the failure occurred. They may help diagnose the problem. If the indicated table is unreadable due to disk failure, the root disk on that node needs to be replaced. If the table file is corrupted or missing, boot the cluster in -x mode to restore the indicated table from backup. The CCR tables are located at /etc/cluster/ccr/.

414680 :fatal: register_president: Don't have reference to myself

The low-level cluster machinery has encountered a fatal error. The rgmd will produce a core file and will cause the node to halt or reboot to avoid the possibility of data corruption.

Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

415842 :fatal: scswitch_onoff: invalid opcode <%d>

While attempting to execute an operator-requested enable or disable of a resource, the rgmd has encountered an internal error. This error should not occur. The rgmd will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

416483 :Failed to retrieve the resource information.

A Sun Cluster data service is unable to retrieve the resource information. Low memory or API call failure might be the reasons.

In case of low memory, the problem will probably be cured by rebooting. If the problem reoccurs, you might need to increase swap space by configuring additional swap devices. Otherwise, if it is API call failure, check the syslog messages from other components.

417144 :Must be root to start %s

The program or daemon has been started by someone not in superuser mode.

Login as root and run the program. If it is a daemon, it may be incorrectly installed. Reinstall cluster packages or contact your service provider.

419220 :%s restore operation failed.

In the process of creating a shared address resource the system was attempting to reconfigure the ip addresses on the system. The specified operation failed.

Use ifconfig command to make sure that all the ip addresses are present. If not, remove the shared address resource and run scrgadm command to recreate it. If problem persists, reboot.

419529 :INTERNAL ERROR CMM: Failure registering callbacks.

An instance of the userland CMM encountered an internal initialization error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

419972 :clcomm: Adapter %s is faulted

A network adapter has encountered a fault.

Any interconnect failure should be resolved, and/or a failed node rebooted.

420763 :Switchover (%s) error (%d) after failure to become secondary

The file system specified in the message could not be hosted on the node the message came from.

Check /var/adm/messages to make sure there were no device errors. If not, contact your authorized Sun service provider to determine whether a workaround or patch is available.

422214 :CMM: Votecount changed from %d to %d for quorum device %ld (%s).

The votecount for the specified quorum device has been changed as indicated.

This is an informational message, no user action is needed.

422541 :Failed to register with PDTserver

This means that we have lost communication with PDT server. Scalable services will not work any more. Probably, the nodes which are configured to be the primaries and secondaries for the PDT server are down.

Need to restart any of the nodes which are configured to be the primary or secondary for the PDT server.

423538 :WARNING: UDLM_PROCEED was picked up by a lkcm_act, returning LKCM_NOOP

An internal warning during udlm state update.

None.

423958 :resource group %s state change to unmanaged.

This is a notification from the rgmd that a resource group's state has changed. This may be used by system monitoring tools.

This is an informational message, no user action is needed.

424061 :Validation failed. ORACLE_HOME %s does not exist

Directory specified as ORACLE_HOME does not exist. ORACLE_HOME property is specified when creating Oracle_server and Oracle_listener resources.

Specify correct ORACLE_HOME when creating resource. If resource is already created, please update resource property "ORACLE_HOME".

424095 :scvxvmlg fatal error - %s does not exist

The program responsible for maintaining the VxVM namespace was unable to access the global device namespace. If configuration changes were recently made to VxVM diskgroups or volumes, this node may be unaware of those changes. Recently created volumes may be inaccessible from this node.

Verify that the /global/.devices/node@N (N = this node's node number) is mounted globally and is accessible. If no configuration changes have been recently made to VxVM diskgroups or volumes and all volumes continue to be accessible from this node, then no further action is required. If changes have been made, the device namespace on this node can be updated to reflect those changes by executing "/usr/cluster/lib/dcs/scvxvmlg". If the problem persists, contact your authorized Sun service provider to determine whether a workaround or patch is available.

424774 :Resource group <%s> requires operator attention due to STOP failure

This is a notification from the rgmd that a resource group has had a STOP method failure or timeout on one of its resources. The resource group is in ERROR_STOP_FAILED state. This may cause another operation such as

scswitch(1M), scrgadm(1M), or scha_control(1HA,3HA) to fail with a SCHA_ERR_STOPFAILED error.

Refer to the procedure for clearing the ERROR_STOP_FAILED condition on a resource group in the Sun Cluster Administration Guide.

424816 :Unable to set automatic MT mode.

The rpc.pmfd server was not able to set the multi-threaded operation mode. This happens while the server is starting up, at boot time. The server does not come up, and an error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

425053 :CCR: Can't access table %s while updating it on node %s
errno = %d.

The indicated error occurred while updating the the indicated table on the indicated node. The errno value indicates the nature of the problem. errno values are defined in the file /usr/include/sys/errno.h. An errno value of 28 (ENOSPC) indicates that the root file system on the node is full. Other values of errno can be returned when the root disk has failed (EIO) or some of the CCR tables have been deleted outside the control of the cluster software (ENOENT).

There may be other related messages on the node where the failure occurred. These may help diagnose the problem. If the root file system is full on the node, then free up some space by removing unnecessary files. If the indicated table was accidentally deleted, then boot the offending node in -x mode to restore the indicated table from other nodes in the cluster. The CCR tables are located at /etc/cluster/ccr/. If the root disk on the afflicted node has failed, then it needs to be replaced.

425551 :getnetconfignt (open_cmd_port) failed

Call to getnetconfignt failed and ucmmmd could not get network information. ucmmmd will exit and the node will abort.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

426221 :CMM: Reservation key changed from %s to %s for node %s
(id = %d).

The reservation key for the specified node was changed. This can only happen due to the CCR infrastructure being changed by hand, which is not a supported operation. The system can not continue, and the node will panic.

Boot the node in non-cluster (-x) mode, recover a good copy of the file /etc/cluster/ccr/infrastructure from one of the cluster nodes or from backup, and then boot this node back in cluster mode. If all nodes in the cluster exhibit this

problem, then boot them all in non-cluster mode, make sure that the infrastructure files are the same on all of them, and boot them back in cluster mode. The problem should not happen again.

426678 :rgmd died

An inter-node communication failed because the rgmd died on another node. To avoid data corruption, the failfast mechanism will cause that node to halt or reboot.

No action is required. The cluster will reconfigure automatically. Examine syslog output on the rebooted node to determine the cause of node death. The syslog output might indicate further remedial actions.

429663 :Node %s not in list of configured nodes

The specified scalable service could not be started on this node because the node is not in the list of configured nodes for this particular service.

If the specified service needs to be started on this node, use scrgadm to add the node to the list of configured nodes for this service and then restart the service.

429907 :clxecd: waitpid returned %d. Returning %d to clxecd.

clxecd program has encountered a failed waitpid(2) system call. The error message indicates the error number for the failure.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

430445 :Monitor initialization error. Incorrect arguments

Error occurred in monitor initialization. Arguments passed to the monitor by callback methods were incorrect.

This is an internal error. Disable the monitor and report the problem.

432473 :reservation fatal error(%s) - joining_node not specified

The device fencing program has suffered an internal error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available. Copies of /var/adm/messages from all nodes should be provided for diagnosis. It may be possible to retry the failed operation, depending on the nature of the error. If the message specifies the "node_join" transition, then this node may be unable to access shared devices. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access shared devices. In either case, it may be possible to reacquire access to shared devices by executing "/usr/cluster/lib/sc/run_reserve -c node_join" on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group has failed to start on this

node. If another node was available to host the device group, then it should have been started on that node. If desired, it may be possible to switch the device group to this node with the `scswitch` command. If no other node was available, then the device group will not have been started. The `scswitch` command may be used to retry the attempt to start the device group. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group has failed. The desired action may be retried.

432987 :Failed to retrieve nodeid.

Data service is failed to retrieve the host information.

If the logical host and shared address entries are specified in the `/etc/inet/hosts` file, check these entries are correct. If this is not the reason then check the health of the name server. For more error information, check the syslog messages.

433438 :Setup error. SUPPORT_FILE %s does not exist

This is an internal error. Support file is used by HA-Oracle to determine the fault monitor information.

Please report this problem.

433501 :fatal: priocntl: %s (UNIX error %d)

The daemon indicated in the message tag (`rgmd` or `ucmmd`) has encountered a failed system call to `priocntl(2)`. The error message indicates the reason for the failure. The daemon will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

If the error message is not self-explanatory, save a copy of the `/var/adm/` messages files on all nodes, and of the core file generated by the daemon. Contact your authorized Sun service provider for assistance in diagnosing the problem.

433895 :INTERNAL ERROR: Invalid resource property tunable flag <%d> for property <%s>; aborting node

An internal error occurred in the `rgmd` while checking whether a resource property could be modified. The `rgmd` will produce a core file and will force the node to halt or reboot.

Look for other syslog error messages on the same node. Save a copy of the `/var/adm/messages` files on all nodes, and report the problem to your authorized Sun service provider.

434480 :CCR: CCR data server not found.

The CCR data server could not be found in the local name server.

Reboot the node. Also contact your authorized Sun service provider to determine whether a workaround or patch is available.

435521 :Warning: node %d does not have a weight assigned to it for property %s, but node %d is in the %s for resource %s. A weight of %d will be used for node %d.

The named node does not have a weight assigned to it, but it is a potential master of the resource.

No user action is required if the default weight is acceptable. Otherwise, use `scrgadm(1M)` to set the `Load_balancing_weights` property to include the node that does not have an explicit weight set for it.

437100 :Validation failed. Invalid command line parameter %s %s

Unable to process parameters passed to the call back method. This is an internal error.

Please report this problem.

437236 :dl_bind: DLPI error %u

DLPI protocol error. We cannot bind to the physical device. We are trying to open a fast path to the private transport adapters.

Reboot of the node might fix the problem.

437975 :The property %s cannot be updated because it affects the scalable resource %s.

The property named is not allowed to be changed after the resource has been created.

If the property must be changed, then the resource should be removed and re-added with the new value of the property.

438420 :Interface %s is plumbed but is not suitable for global networking.

The specified adapter may be either point to point adapter or loopback adapter which is not suitable for global networking.

Reconfigure the appropriate NAFO group to exclude this adapter.

438700 :Some ip addresses might still be on loopback.

Some of the ip addresses managed by the specified SharedAddress resource were not removed from the loopback interface.

Use the `ifconfig` command to make sure that the ip addresses being managed by the SharedAddress implementation are present either on the loopback interface or on a physical adapter. If they are present on both, use `ifconfig` to delete them from the loopback interface. Then use `scswitch` to move the resource group

containing the SharedAddresses to another node to make sure that the resource group can be switched over successfully.

438866 :sysinfo in getlocalhostname failed

sysinfo call did not succeed. The "sysinfo" man page describes possible error codes.

This is an internal error. Please report this problem.

439099 :HA: hxdoor %d.%d does not exist on secondary

An HA framework hxdoor is missing.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

440406 :Cannot check online status. Server processes are running.

HA-Oracle could not check online status of Oracle server. Oracle server processes are running but Oracle server may or may not be online yet.

Examine "Connetc_string" property of the resource. Make sure that user id and password specified in connect string are correct and permissions are granted to user for connecting to the server.

440530 :Started the fault monitor.

The fault monitor for this data service was started successfully.

No action needed.

440792 :Warning: some resources in resource group <%s> failed to start

The indicated resource group is pending online. One or more resources' START methods failed to execute successfully. Because the resources' Failover_mode is set to NONE, the resource group is moving to the ONLINE state in spite of the resource start failures.

This is a warning message, no user action is needed. The operator may choose to issue an scswitch(1M) command to try switching the affected resource group to another node or to try restarting it on the same node.

441826 : "pmfadm -a" Action failed for <%s>

The given tag has exceeded the allowed number of retry attempts (given by the "pmfadm -n" option) and the action (given by the "pmfadm -a" option) was initiated by rpc.pmfd. The action failed (i.e., returned non-zero), and rpc.pmfd will delete this tag from its tag list and discontinue retry attempts.

This message is informational; no user action is needed.

442053 :clcomm: Invalid path_manager client_type (%d)

The system attempted to add a client of unknown type to the set of path manager clients.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

442281 :reservation error(%s) - did_get_path() error in other_node_status()

The device fencing program has suffered an internal error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available. Copies of /var/adm/messages from all nodes should be provided for diagnosis. It may be possible to retry the failed operation, depending on the nature of the error. If the message specifies the "node_join" transition, then this node may be unable to access shared devices. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access shared devices. In either case, it may be possible to reacquire access to shared devices by executing "/usr/cluster/lib/sc/run_reserve -c node_join" on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group has failed to start on this node. If another node was available to host the device group, then it should have been started on that node. If desired, it may be possible to switch the device group to this node with the scswitch command. If no other node was available, then the device group will not have been started. The scswitch command may be used to retry the attempt to start the device group. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group has failed. The desired action may be retried.

443271 :clcomm: Pathend: Aborting node because %s for %u ms

The pathend aborted the node for the specified reason.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

443479 :CMM: Quorum device %ld with gdevname %s has %d configured path - Ignoring mis-configured quorum device.

The specified number of configured paths to the specified quorum device is less than two, which is the minimum allowed. This quorum device will be ignored.

Reconfigure the quorum device appropriately.

443746 :resource %s state on node %s change to %s

This is a notification from the rgmd that a resource's state has changed. This may be used by system monitoring tools.

This is an informational message, no user action is needed.

444001 :%s: Call failed, return code=%d

A client was not able to make an rpc connection to a server (rpc.pmfd, rpc.fed or rgmd) to execute the action shown, and was not able to read the rpc error. The rpc error number is shown. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

444144 :clcomm: Cannot change increment

An attempt was made to change the flow control policy parameter that specifies the thread increment level. The flow control system uses this parameter to set the number of threads that are acted upon at one time. This value currently cannot be changed.

No user action required.

445616 :libsecurity: create of rpc handle to program %ld failed, will not retry

A client of the rpc.pmfd, rpc.fed or rgmd server was not able to initiate an rpc connection, after multiple retries. The maximum time allowed for connecting has been exceeded, or the types of rpc errors encountered indicate that there is no point in retrying. An accompanying error message shows the rpc error data. The pmfadm or scha command exits with error. The program number is shown. To find out what program corresponds to this number, use the rpcinfo command. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

446068 :CMM: Node %s (nodeid = %ld) is down.

The specified node has gone down in that communication with it has been lost.

The cause of the failure should be resolved and the node should be rebooted if node failure is unexpected.

446249 :Method <%s> on resource <%s>: authorization error.

An attempted method execution failed, apparently due to a security violation; this error should not occur. This failure is considered a method failure. Depending on which method was being invoked and the Failover_mode setting on the resource, this might cause the resource group to fail over or move to an error state.

Examine other syslog messages occurring at about the same time to see if the problem can be diagnosed. Save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance in diagnosing the problem.

447578 :Duplicated installed nodename when Resource Type <%s> is added.

User has defined duplicated installed node name when creating resource type.

Recheck the installed nodename list and make sure there is no nodename duplication.

447872 :fatal: Unable to reserve %d MBytes of swap space; exiting

The rgmd was unable to allocate a sufficient amount of memory upon starting up. This is a fatal error. The rgmd will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

Make sure that the hardware configuration meets documented minimum requirements. Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

448703 :clcomm: validate_policy: high too small. high %d low %d nodes %d pool %d

The system checks the proposed flow control policy parameters at system startup and when processing a change request. The high server thread level must be large enough to grant the low number of threads to all of the nodes identified in the message for a fixed size resource pool.

No user action required.

448844 :clcomm: inbound_invo::done: state is 0x%x

The internal state describing the server side of a remote invocation is invalid when the invocation completes server side processing.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

448898 :%s.nodes entry in the configuration file must be between 1 and %d.

Illegal value for a node number. Perhaps the system is not booted as part of the cluster.

Make sure the node is booted as part of a cluster.

449159 :clconf: No valid quorum_vote field for node %u

Found the quorum vote field being incorrect while converting the quorum configuration information into quorum table.

Check the quorum configuration information.

449288 :setgid: %s

The rpc.pmfd server was not able to set the group id of a process. The message contains the system error. The server does not perform the action requested by the client, and an error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

449336 :setuid: %s

The rpc.pmfd or rpc.fed server was not able to set the session id, and the system error is shown. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

449344 :setuid: %s

The rpc.pmfd server was not able to set the user id of a process. The message contains the system error. The server does not perform the action requested by the client, and an error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

449907 :scvxvmlg error - mknod(%s) failed

The program responsible for maintaining the VxVM namespace was unable to access the global device namespace. If configuration changes were recently made to VxVM diskgroups or volumes, this node may be unaware of those changes. Recently created volumes may be inaccessible from this node.

Verify that the /global/.devices/node@N (N = this node's node number) is mounted globally and is accessible. If no configuration changes have been recently made to VxVM diskgroups or volumes and all volumes continue to be accessible from this node, then no further action is required. If changes have been made, the device namespace on this node can be updated to reflect those changes by executing "/usr/cluster/lib/dcs/scvxvmlg". If the problem persists, contact your authorized Sun service provider to determine whether a workaround or patch is available.

450173 :Error accessing policy

This message appears when the customer is initializing or changing a scalable services load balancer, by starting or updating a service. The Load_Balancing_Policy is missing.

Add a Load_Balancing_Policy parameter when creating the resource group.

450780 :Error: Unable to create scha_control timestamp file <%s>
for resource <%s>

The rgmd has failed in an attempt to create a file used for the anti-"pingpong" feature. This may prevent the anti-pingpong feature from working, which may permit a resource group to fail over repeatedly between two or more nodes. The failure to create the file might indicate a more serious problem on the node.

Examine other syslog messages occurring around the same time on the same node, to see if the source of the problem can be identified.

451640 :tag %s: stat of command file %s failed

The rpc.fed server checked the command path indicated by the tag, and this check failed, possibly because the path is incorrect. An error message is output to syslog.

Check the path of the command.

452150 :Failed to start the fault monitor.

Process monitor facility has failed to start the fault monitor.

Check whether the system is low in memory or the process table is full and correct these problems. If the error persists, use scswitch to switch the resource group to another node.

452202 :clcomm: sdoor_sendstream::send

This operation should never occur.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

452604 :CMM: Registered key on and acquired quorum device %ld
(gdevname %s).

When this node was booting up, it had found only non-cluster member keys on the specified device. After joining the cluster and having its CCR recovered, this node has been able to register its keys on this device and is its owner.

This is an informational message, no user action is needed.

453919 :Pathprefix is not set for resource group %s.

Resource Group property Pathprefix is not set.

Use scrgadm to set the Pathprefix property on the resource group.

454247 :Error: Unable to create directory <%s> for scha_control
timestamp file

The rgmd is unable to access the directory used for the anti-"pingpong" feature, and cannot create the directory (which should already exist). This may prevent the anti-pingpong feature from working, which may permit a resource group to fail over repeatedly between two or more nodes. The failure to access or create the directory might indicate a more serious problem on the node.

Examine other syslog messages occurring around the same time on the same node, to see if the source of the problem can be identified.

454607 :INTERNAL ERROR: Invalid resource extension property type <%d> on resource <%s>; aborting node

An attempted creation or update of a resource has failed because of invalid resource type data. This may indicate CCR data corruption or an internal logic error in the rgmd. The rgmd will produce a core file and will force the node to halt or reboot.

Use `scrgadm(1M) -pvv` to examine resource properties. If the resource or resource type properties appear to be corrupted, the CCR might have to be rebuilt. If values appear correct, this may indicate an internal error in the rgmd. Re-try the creation or update operation. If the problem recurs, save a copy of the `/var/adm/messages` files on all nodes and contact your authorized Sun service provider for assistance.

454930 :Scheduling class %s not configured

An attempt to change the thread scheduling class failed, because the scheduling class was not configured.

Configure the system to support the desired thread scheduling class.

457114 :fatal: death_ff->arm failed

The daemon specified in the error tag was unable to arm the failfast device. The failfast device kills the node if the daemon process dies either due to hitting a fatal bug or due to being killed inadvertently by an operator. This is a requirement to avoid the possibility of data corruption. The daemon will produce a core file and will cause the node to halt or reboot.

Save a copy of the `/var/adm/messages` files on all nodes, and of the core file generated by the daemon. Contact your authorized Sun service provider for assistance in diagnosing the problem.

457121 :Failed to retrieve the host information for %s: %s.

The data service failed to retrieve the host information.

If the logical hostname and shared address entries are specified in the `/etc/inet/hosts` file, check that the entries are correct. Verify the settings in the `/etc/nsswitch.conf` file include "files" for host lookup. If these are correct, check the health of the name server. For more error information, check the syslog messages.

458091 :CMM: Reconfiguration delaying for %d seconds to allow larger partitions to win race for quorum devices.

In the case of potential split brain scenarios, the CMM allows larger partitions to win the race to acquire quorum devices by forcing the smaller partitions to sleep for a time period proportional to the number of nodes not in that partition.

This is an informational message, no user action is needed.

458530 :Method <%s> on resource <%s>: program file is not executable.

A method pathname points to a file that is not executable. This may have been caused by incorrect installation of the resource type.

Identify registered resource type methods using `scrgadm(1M) -pvv`. Check the permissions on the resource type methods. Reinstall the resource type if necessary, following resource type documentation.

458818 :reservation fatal error(%s) - disk_file not specified

The device fencing program has suffered an internal error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available. Copies of `/var/adm/messages` from all nodes should be provided for diagnosis. It may be possible to retry the failed operation, depending on the nature of the error. If the message specifies the "node_join" transition, then this node may be unable to access shared devices. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access shared devices. In either case, it may be possible to reacquire access to shared devices by executing `"/usr/cluster/lib/sc/run_reserve -c node_join"` on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group has failed to start on this node. If another node was available to host the device group, then it should have been started on that node. If desired, it may be possible to switch the device group to this node with the `scswitch` command. If no other node was available, then the device group will not have been started. The `scswitch` command may be used to retry the attempt to start the device group. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group has failed. The desired action may be retried.

458988 :libcddb: scha_cluster_open failed with %d

Call to initialize a handle to get cluster information failed. The second part of the message gives the error code.

The calling program should handle this error. If it is not recoverable, it will exit.

460027 :Resource <%s> of Resource Group <%s> failed sanity check on node <%s>\n

Message logged for failed scha_control sanity check methods on specific node.

No user action required.

460520 :scvxvmlg fatal error -
dcs_get_service_names_of_class(%s) failed, returned %d

The program responsible for maintaining the VxVM namespace has suffered an internal error. If configuration changes were recently made to VxVM diskgroups or volumes, this node may be unaware of those changes. Recently created volumes may be inaccessible from this node.

If no configuration changes have been recently made to VxVM diskgroups or volumes and all volumes continue to be accessible from this node, then no action is required. If changes have been made, the device namespace on this node can be updated to reflect those changes by executing "/usr/cluster/lib/dcs/scvxvmlg". If the problem persists, contact your authorized Sun service provider to determine whether a workaround or patch is available.

462083 :fatal: Resource <%s> update failed with error <%d>;
aborting node

Rgmd failed to read updated resource from the CCR on this node.

Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

462632 :HA: repl_mgr: exception invalid_repl_prov_state %d

The system did not perform this operation on the primary object.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

464588 :Failed to retrieve the resource group property %s: %s.

An API operation has failed while retrieving the resource group property. Low memory or API call failure might be the reasons.

In case of low memory, the problem will probably be cured by rebooting. If the problem reoccurs, you might need to increase swap space by configuring additional swap devices. Otherwise, if it is API call failure, check the syslog messages from other components. For resource group name and the property name, check the current syslog message.

465065 :Error accessing group

This message appears when the customer is initializing or changing a scalable services load balancer, by starting or updating a service. The specified resource group is invalid.

Check the resource group name specified and make sure that a valid value is used.

468477 :Failed to retrieve the property %s: %s.

API operation has failed in retrieving the cluster property.

For property name, check the syslog message. For more details about API call failure, check the syslog messages from other components.

468732 :Too many modules configured for autopush.

The system attempted to configure a clustering STREAMS module for autopush but too many modules were already configured.

Check in your /etc/iu.ap file if too many modules have been configured to be autopushed on a network adapter. Reduce the number of modules. Use autopush(1m) command to remove some modules from the autopush configuration.

469417 :Failfast: timeout - unit \"%s\"%s.

A failfast client has encountered a timeout and is going to panic the node.

There may be other related messages on this node which may help diagnose the problem. Resolve the problem and reboot the node if node panic is unexpected.

472185 :Failed to retrieve the resource group property %s for %s: %s.

The query for a property failed. The reason for the failure is given in the message.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

473021 :in libsecurity uname sys call failed: %s

A client was not able to make an rpc connection to a server (rpc.pmfd, rpc.fed or rgmd) because the host name could not be obtained. The system error message is shown. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

473460 :Method <%s> on resource <%s>: authorization error: %s.

An attempted method execution failed, apparently due to a security violation; this error should not occur. The last portion of the message describes the error. This failure is considered a method failure. Depending on which method was being invoked and the Failover_mode setting on the resource, this might cause the resource group to fail over or move to an error state.

Correct the problem identified in the error message. If necessary, examine other syslog messages occurring at about the same time to see if the problem can be diagnosed. Save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance in diagnosing the problem.

473653 :Failed to retrieve the resource type handle for %s while querying for property %s: %s.

Access to the object named failed. The reason for the failure is given in the message.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

474690 :clxecd: Error %d from send_fd

clxecd program has encountered a failed fcntl(2) system call. The error message indicates the error number for the failure.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

475398 :Out of memory (memory allocation failed):%s.%s

There is not enough swap space on the system.

Add more swap space. See swap(1M) for more details.

476996 :Failed to unregister callback for NAF0 group %s with tag %s (request failed with %d).

An unexpected error occurred while trying to communicate with the network monitoring daemon (pnmd).

Make sure the network monitoring daemon (pnmd) is running.

477378 :Failed to restart the service.

Restart attempt of the dataservice is failed.

Check the syslog messages that are occurred just before this message to check whether there is any internal error. In case of internal error, contact your Sun service provider. Otherwise, any of the following situations may have happened. 1) Check the Start_timeout and Stop_timeout values and adjust them if they are not appropriate. 2) This might be the result of lack of the system resources. Check whether the system is low in memory or the process table is full and take appropriate action.

478523 :Could not mount "%s" because there was an error (%d) in opening the directory.

While mounting a Cluster file system, the directory on which the mount is to take place could not be opened.

Fix the reported error and retry. The most likely problem is that the directory does not exist - in that case, create it with the appropriate permissions and retry.

479105 :Can not get service status for global service <%s> of path <%s>

Can not get status for the global service. This is a server problem.

Contact your authorized Sun service provider to determine what is the cause of the problem.

479442 :in libsecurity could not allocate memory

A server (rpc.pmfd, rpc.fed or rgmd) was not able to start, or a client was not able to make an rpc connection to the server, probably due to low memory. An error message is output to syslog.

Investigate if the host is low on memory. If not, save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

483160 :Failed to connect to socket: %s.

While determining the health of the resource, process monitor facility has failed to communicate with the resource fault monitor.

Any of the following situations might have occurred. 1) Check whether the fault monitor is running, if not wait for the fault monitor to start. 2) Check whether the fault monitor is disabled, if it is then user can enable the fault monitor, otherwise ignore it. 3) In all other situations, consider it as an internal error. Save /var/adm/messages file and contact your authorized Sun service provider. For more error description check the syslog messages.

484084 :INTERNAL ERROR: non-existent resource <%s> appears in dependency list of resource <%s>

While attempting to execute an operator-requested enable of a resource, the rgmd has found a non-existent resource to be listed in the Resource_dependencies or Resource_dependencies_weak property of the indicated resource. This suggests corruption of the RGM's internal data but is not fatal.

Use scrgadm(1M) -pvv to examine resource group properties. If the values appear corrupted, the CCR might have to be rebuilt. If values appear correct, this may indicate an internal error in the rgmd. Contact your authorized Sun service provider for assistance in diagnosing and correcting the problem.

485464 :clcomm: Failed to allocate simple xdoor server %d

The system could not allocate a simple xdoor server. This can happen when the xdoor number is already in use. This message is only possible on debug systems.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

485759 :transition "%s" failed for cluster "%s"

The mentioned state transition failed for the cluster. udlmctl will exit.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

485942 :(%s) sigprocmask failed: %s (UNIX errno %d)

Call to sigprocmask() failed. The "sigprocmask" man page describes possible error codes. udlmctl will exit.

Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

487022 :The networking components for scalable resource %s have been configured successfully for method %s.

The calls to the underlying scalable networking code succeeded.

This is an informational message, no user action is needed.

487484 :lkcm_reg: lib initialization failed

udlm could not register with cmm because lib initialization failed.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

487574 :Failed to alloc memory

A scha_control call failed because the system has run out of swap space. The system is likely to halt or reboot if swap space continues to be depleted.

Investigate the cause of swap space depletion and correct the problem, if possible.

487827 CCR: Waiting for repository synchronization to finish.

This node is waiting to finish the synchronization of its repository with other nodes in the cluster before it can join the cluster membership.

This is an informational message, generally no user action is needed. If all the nodes in the cluster are hanging at this message for a long time, look for other

messages. The possible cause is the cluster hasn't obtained quorum, or there is CCR metadata missing or invalid. If the cluster is hanging due to missing or invalid metadata, the ccr metadata needs to be recovered from backup.

488276 :in libsecurity write of file %s failed: %s

The rpc.pmfd, rpc.fed or rgmd server was not able to write to a cache file for rpcbind information. The affected component should continue to function by calling rpcbind directly.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

489438 :clcomm: Path %s being drained

A communication link is being removed with another node. The interconnect may have failed or the remote node may be down.

Any interconnect failure should be resolved, and/or the failed node rebooted.

489644 :Could not look up host because host was NULL.

Can't look up the hostname locally in hostfile. The specified host name is invalid.

Check whether the hostname has NULL value. If this is the case, recreate the resource with valid host name. If this is not the reason, treat it as an internal error and contact Sun service provider.

490778 :Local node failed to do affinity switchover to global service <%d> of path <%s>

When prenet_start method of SUNW.HAStorage do affinity switch, it fails. This is a severe problem.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

491081 :resource %s removed.

This is a notification from the rgmd that the operator has deleted a resource. This may be used by system monitoring tools.

This is an informational message, no user action is needed.

491579 :clcomm: validate_policy: fixed size pool low %d must match moderate %d

The system checks the proposed flow control policy parameters at system startup and when processing a change request. The low and moderate server thread levels must be the same for fixed size resource pools.

No user action required.

491694 :Could not %s any ip addresses.

The specified action was not successful for all ip addresses managed by the LogicalHostname resource.

Check the logs for any error messages from pnm. This could be result from the lack of system resources, such as low on memory. Reboot the node if the problem persists.

492603 :launch_fed_prog: fe_method_full_name() failed for program <%s>

The ucmmmd was unable to assemble the full method pathname for the fed program to be launched. This is considered a launch_fed_prog failure.

Examine other syslog messages occurring at about the same time to see if the problem can be identified. Save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance in diagnosing and correcting the problem.

492953 :ORACLE_HOME/bin/lsnrctl not found ORACLE_HOME=%s

Oracle listener binaries not found under ORACLE_HOME. ORACLE_HOME specified for the resource is indicated in the message. HA-Oracle will not be able to manage Oracle listener if ORACLE_HOME is incorrect.

Specify correct ORACLE_HOME when creating resource. If resource is already created, please update resource property "ORACLE_HOME".

494534 :clcomm: per node IP config %s%d:%d (%d): %d.%d.%d.%d failed with %d

The system failed to configure IP communications across the private interconnect of this device and IP address, resulting in the error identified in the message. This happened during initialization. Someone has used the "lo0:1" device before the system could configure it.

If you used "lo0:1", please use another device. Otherwise, Contact your authorized Sun service provider to determine whether a workaround or patch is available.

494913 :pmfd: unknown action (0x%x)

An internal error has occurred in the rpc.pmfd server. This should not happen.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

495284 :dl_attach: DLPI error %u

Could not attach to the physical device. We are trying to open a fast path to the private transport adapters.

Reboot of the node might fix the problem.

495386 :INTERNAL ERROR: %s.

An internal error has occurred.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

495529 :Prog <%s> failed to execute step <%s> - <%s>

ucmmd failed to execute a step.

Examine other syslog messages occurring at about the same time to see if the problem can be identified and if it recurs. Save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance.

495710 :Stopping oracle server using shutdown immediate

Informational message. Oracle server will be stopped using "shutdown immediate" command.

None

496884 :Despite the warnings, the validation of the hostname list succeeded.

While validating the hostname list, non fatal errors have been found.

This is informational message. It is suggested to correct the errors if applicable. For the error information, check the syslog messages that have been encountered before this message.

497795 :gethostbyname() timed out.

The name service could be unavailable.

If the cluster is under load or too much network traffic, increase the timeout value of monitor_check method using scrgadm command. Otherwise, check if name service is configured correctly. Try some commands to query name services, such as ping and nslookup, and correct the problem. If the error still persists, then reboot the node.

497808 :clcomm: Cannot fork() after Galileo initialization

A user level process attempted to fork after Galileo initialization. This is not allowed.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

498582 :Attempt to load %s failed: %s.

A shared address resource was in the process of being created. In order to prepare this node to handle scalable services, the specified kernel module was attempted to be loaded into the system, but failed.

This might be the result from the lack of system resources. Check whether the system is low in memory and take appropriate action (e.g., by killing hung processes). For specific information check the syslog message. After more resources are available on the system, attempt to create shared address resource. If problem persists, reboot.

498711 :Could not initialize the ORB. Exiting.

clxecd program was unable to initialize its interface to the low-level clustering software.

This might occur because the operator has attempted to start clxecd program on a node that is booted in non-cluster mode. If the node is in non-cluster mode, boot it into cluster mode. If the node is already in cluster mode, contact your authorized Sun service provider to determine whether a workaround or patch is available.

499756 :CMM: Node %s: joined cluster.

The specified node has joined the cluster.

This is an informational message, no user action is needed.

499775 :resource group %s added.

This is a notification from the rgmd that a new resource group has been added. This may be used by system monitoring tools.

This is an informational message, no user action is needed.

Message IDs 500000 - 599999

Error Message List

The following list is ordered by the message ID.

501582 :in libsecurity setnetpath failed: %s

The rpc.pmfd, rpc.fed or rgmd server was not able to initiate an rpc connection, because it could not get the network database handle. The server does not start. The rpc error message is shown. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

501733 :scvxvmlg fatal error - _cladm() failed

The program responsible for maintaining the VxVM namespace has suffered an internal error. If configuration changes were recently made to VxVM diskgroups or volumes, this node may be unaware of those changes. Recently created volumes may be inaccessible from this node.

If no configuration changes have been recently made to VxVM diskgroups or volumes and all volumes continue to be accessible from this node, then no action is required. If changes have been made, the device namespace on this node can be updated to reflect those changes by executing "/usr/cluster/lib/dcs/scvxvmlg". If the problem persists, contact your authorized Sun service provider to determine whether a workaround or patch is available.

501763 :recv_request: t_alloc: %s

Call to `t_alloc()` failed. The "t_alloc" man page describes possible error codes. `udlm` will exit and the node will abort.

This is an internal error. Save the contents of `/var/adm/messages`, `/var/cluster/ucmm/ucmm_reconf.log` and `/var/cluster/ucmm/dlm*/logs/*` from all the nodes and contact your Sun service representative.

501917 :process_intention(): IDL exception when communicating to node %d

An inter-node communication failed, probably because a node died.

No action is required; the `rgmd` should recover automatically.

502022 :fatal: joiners_read_ccr: exiting early because of unexpected exception

The low-level cluster machinery has encountered a fatal error. The `rgmd` will produce a core file and will cause the node to halt or reboot.

Save a copy of the `/var/adm/messages` files on all nodes, and of the `rgmd` core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

503048 :NULL value returned for resource group property %s.

NULL value was returned for resource group property.

For the property name check the syslog message. Any of the following situations might have occurred. Different user action is needed for these 1) If a new resource group is created or updated, check whether the value of the property is valid. 2) For all other cases, treat it as an Internal error.

503064 :Method <%s> on resource <%s>: Method timed out.

A `VALIDATE` method execution has exceeded its configured timeout and was killed by the `rgmd`. This in turn will cause the failure of a creation or update operation on a resource or resource group.

Consult resource type documentation to diagnose the cause of the method failure. Other syslog messages occurring just before this one might indicate the reason for the failure. After correcting the problem that caused the method to fail, the operator may retry the resource group update operation.

504363 :ERROR: process_resource: resource <%s> is pending_update but no UPDATE method is registered

A non-fatal internal error has occurred in the `rgmd` state machine.

Since this problem might indicate an internal logic error in the `rgmd`, please save a copy of the `/var/adm/messages` files on all nodes, the output of an `scstat -g`

command, and the output of a `scrgadm -pvv` command. Report the problem to your authorized Sun service provider.

504402 :CMM: Aborting due to stale sequence number. Received a message from node %ld indicating that node %ld has a stale sequence#. node %ld: state "%s", sequence# %lld; node %ld: state "%s", sequence# %lld.

After receiving a message from the specified remote node, the local node has concluded that it has stale state with respect to the remote node, and will therefore abort. The state of a node can get out-of-date if it has been in isolation from the nodes which have majority quorum.

Reboot the node.

505101 :Found another active instance of `clxecd`. Exiting `daemon_process`.

An active instance of `clxecd` program is already running on the node.

This would usually happen if the operator tries to start the `clxecd` program by hand on a node which is booted in cluster mode. If that's not the case, contact your authorized Sun service provider to determine whether a workaround or patch is available.

508671 :mmap: %s

The `rpc.pmfd` server was not able to allocate shared memory for a semaphore, possibly due to low memory, and the system error is shown. The server does not perform the action requested by the client, and `pmfadm` returns error. An error message is also output to `syslog`.

Investigate if the machine is running out of memory. If this is not the case, save the `/var/adm/messages` file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

508687 :monitor_check: method <%s> failed on resource <%s> in resource group <%s> on node <%s>, exit code <%d>

In `scha_control`, `monitor_check` method of the resource failed on specific node.

No action is required, this is normal phenomenon of `scha_control`, which launches the corresponding `monitor_check` method of the resource on all candidate nodes and looks for a healthy node which passes the test. If a healthy node is found, `scha_control` will let the node take over the resource group. Otherwise, `scha_control` will just exit early.

509069 :CMM: Halting because this node has no configuration info about node %ld which is currently configured in the cluster and running.

The local node has no configuration information about the specified node. This indicates a misconfiguration problem in the cluster. The `/etc/cluster/ccr/infrastructure` table on this node may be out of date with respect to the other nodes in the cluster.

Correct the misconfiguration problem or update the `infrastructure` table if out of date, and reboot the nodes. To update the table, boot the node in non-cluster (-x) mode, restore the table from the other nodes in the cluster or backup, and boot the node back in cluster mode.

509136 :Probe failed.

Fault monitor was unable to perform complete health check of the service.

1) Fault monitor would take appropriate action (by restarting or failing over the service.). 2) Data service could be under load, try increasing the values for `Probe_timeout` and `Thororugh_probe_interval` properties. 3) If this problem continues to occur, look at other messages in `syslog` to determine the root cause of the problem. If all else fails reboot node.

510659 :Failover %s data services must be in a failover resource group.

The Scalable resource property for the data service was set to `FALSE`, which indicates a failover resource, but the corresponding data service resource group is not a failover resource group. Failover resources of this resource type must reside in a failover resource group.

Decide whether this resource is to be scalable or failover. If scalable, set the Scalable property value to `TRUE`. If failover, leave Scalable set to `FALSE` and create this resource in a failover resource group. A failover resource group has its resource group property `RG_mode` set to `Failover`.

511177 :clcomm: solaris xdoor door_info failed

A `door_info` operation failed. Refer to the "door_info" man page for more information.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

511749 :rgm_launch_method: failed to get method <%s> timeout value from resource <%s>

Due to an internal error, the `rgmd` was unable to obtain the method timeout for the indicated resource. This is considered a method failure. Depending on which method was being invoked and the `Failover_mode` setting on the resource, this might cause the resource group to fail over or move to an error state.

Examine other `syslog` messages occurring at about the same time to see if the problem can be identified. Save a copy of the `/var/adm/messages` files on all

nodes and contact your authorized Sun service provider for assistance in diagnosing the problem.

511810 :Property <%s> does not exist in SUNW.HAStorage

Property set in SUNW.HAStorage type resource is not defined in SUNW.HAStorage.

Check /var/adm/message and see what property name is used. Correct it according to the definition in SUNW.HAStorage.

511917 :clcomm: orbdata: unable to add to hash table

The system records object invocation counts in a hash table. The system failed to enter a new hash table entry for a new object type.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

512422 :WARNING: unknown msg (type %d) was picked up by a lkcm_act, returning LKCM_NOOP

Warning for unknown message picked up during udlm state update.

None.

513538 :scvxvmlg error - mkdirp(%s) failed

The program responsible for maintaining the VxVM namespace was unable to access the global device namespace. If configuration changes were recently made to VxVM diskgroups or volumes, this node may be unaware of those changes. Recently created volumes may be inaccessible from this node.

Verify that the /global/.devices/node@N (N = this node's node number) is mounted globally and is accessible. If no configuration changes have been recently made to VxVM diskgroups or volumes and all volumes continue to be accessible from this node, then no further action is required. If changes have been made, the device namespace on this node can be updated to reflect those changes by executing "/usr/cluster/lib/dcs/scvxvmlg". If the problem persists, contact your authorized Sun service provider to determine whether a workaround or patch is available.

514156 :PNM: shutting down

The PNM daemon (pnmd) is shutting down, disabling all monitoring and failover of adapters.

This message is informational; no user action is needed.

514731 :Failed to kill listener process for %s

Failed to kill listener processes.

None

515583 :%s is not a valid IP address.

Validation method has failed to validate the ip addresses. The mapping for the given ip address in the local host files can't be done: the specified ip address is invalid.

Invalid hostnames/ip addresses have been specified while creating the resource. Recreate the resource with valid hostnames.

516407 :reservation warning(%s) - MHIOCGRP_INKEYS error will
retry in %d seconds

The device fencing program has encountered errors while trying to access a device. The failed operation will be retried

This is an informational message, no user action is needed.

517009 :lkcm_act: invalid handle was passed %s %d

Handle for communication with udlmctl is invalid.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

517343 :clexecd: Error %d from pipe

clexecd program has encountered a failed pipe(2) system call. The error message indicates the error number for the failure.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

517363 :clconf: Unrecognized property type

Found the unrecognized property type in the configuration file.

Check the configuration file.

518018 :CMM: Node being aborted from the cluster.

This node is being excluded from the cluster.

Node should be rebooted if required. Resolve the problem according to other messages preceding this message.

518291 :Warning: Failed to check if scalable service group %s
exists: %s.

A call to the underlying scalable networking code failed.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

520231 :Unable to set the number of threads for the FED RPC threadpool.

The rpc.fed server was unable to set the number of threads for the RPC threadpool. This happens while the server is starting up.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

520982 :CMM: Preempting node %ld from quorum device %s failed with error %d.

This node was unable to preempt the specified node from the quorum device, indicating that the partition to which the local node belongs has been preempted and will abort. If a cluster gets divided into two or more disjoint subclusters, exactly one of these must survive as the operational cluster. The surviving cluster forces the other subclusters to abort by grabbing enough votes to grant it majority quorum. This is referred to as preemption of the losing subclusters.

There may be other related messages that may indicate why the partition to which the local node belongs has been preempted. Resolve the problem and reboot the node.

521538 :monitor_check: set_env_vars() failed for resource <%s>, resource group <%s>

During execution of a scha_control(1HA,3HA) function, the rgmd was unable to set up environment variables for method execution, causing a MONITOR_CHECK method invocation to fail. This in turn will prevent the attempted failover of the resource group from its current master to a new master.

Examine other syslog messages occurring at about the same time to see if the problem can be identified. Save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance in diagnosing and correcting the problem.

522315 :resource group %s on node %s state change to RG_PENDING_ONLINE

This is a notification from the rgmd that the resource group is being brought online on the indicated node. This may be used by system monitoring tools.

This is an informational message, no user action is needed.

522480 :RGM state machine returned error %d

An error has occurred on this node while attempting to execute the rgmd state machine.

Examine other syslog messages occurring at about the same time to see if the problem can be identified. Save a copy of the `/var/adm/messages` files on all nodes and contact your authorized Sun service provider for assistance in diagnosing the problem.

522779 :IP address (hostname) string %s in property %s, entry %d could not be resolved to an IP address.

The IP address (hostname) string within the named property in the message did not resolve to a real IP address.

Change the IP address (hostname) string within the entry in the property to one that does resolve to a real IP address. Make sure the syntax of the entry is correct.

523302 :fatal: thr_keycreate: %s (UNIX errno %d)

The rgmd failed in a call to thr_keycreate(3T). The error message indicates the reason for the failure. The rgmd will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

Save a copy of the `/var/adm/messages` files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

523643 :INTERNAL ERROR: %s

An internal error has occurred.

Save a copy of the `/var/adm/messages` files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

525628 :CMM: Cluster has reached quorum.

Enough nodes are operational to obtain a majority quorum; the cluster is now moving into operational state.

This is an informational message, no user action is needed.

526403 :ff_open: %s

A server (rpc.pmfd or rpc.fed) was not able to establish a link to the failfast device, which ensures that the host aborts if the server dies. The error message is shown. An error message is output to syslog.

Save the `/var/adm/messages` file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

526413 :Failed to verify that all NAF0 groups are in a stable state. Assuming this node cannot respond to client requests.

The state of the NAF0 groups on the node could not be determined.

Make sure all adapters and cables are working. Look in the `/var/adm/messages` file for message from the network monitoring daemon (pnmd).

526492 :Service object [%s, %s, %d] removed from group "%s"

A specific service known by its unique name SAP (service access point), the three-tuple, has been deleted in the designated group.

This is an informational message, no user action is needed.

526846 :Daemon <%s> is not running.

The HA-NFS fault monitor detected that the specified daemon is no longer running.

No action. The fault monitor would restart the daemon. If it doesn't happen, reboot the node.

527795 :clexecd: setrlimit returned %d

clexecd program has encountered a failed setrlimit() system call. The error message indicates the error number for the failure.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

528020 :CCR: Remove table %s failed.

The CCR failed to remove the indicated table.

The failure can happen due to many reasons, for some of which no user action is required because the CCR client in that case will handle the failure. The cases for which user action is required depends on other messages from CCR on the node, and include: If it failed because the cluster lost quorum, reboot the cluster. If the root file system is full on the node, then free up some space by removing unnecessary files. If the root disk on the afflicted node has failed, then it needs to be replaced. If the cluster repository is corrupted as indicated by other CCR messages, then boot the offending node(s) in -x mode to restore the cluster repository from backup. The cluster repository is located at `/etc/cluster/ccr/`.

528566 :Method <%s> on resource <%s>, resource group <%s>, is_frozen=<%d>: Method timed out.

A method execution has exceeded its configured timeout and was killed by the rgmd. Depending on which method was being invoked and the Failover_mode setting on the resource, this might cause the resource group to fail over or move to an error state.

Consult resource type documentation to diagnose the cause of the method failure. Other syslog messages occurring just before this one might indicate the reason for the failure. After correcting the problem that caused the method to fail,

the operator may choose to issue an `scswitch(1M)` command to bring resource groups onto desired primaries. Note, if the indicated value of `is_frozen` is 1, this might indicate an internal error in the `rgmd`. Please save a copy of the `/var/adm/messages` files on all nodes, and report the problem to your authorized Sun service provider.

529131 :Method <%s> on resource <%s>: RPC connection error.

An attempted method execution failed, due to an RPC connection problem. This failure is considered a method failure. Depending on which method was being invoked and the `Failover_mode` setting on the resource, this might cause the resource group to fail over or move to an error state; or it might cause an attempted edit of a resource group or its resources to fail.

Examine other syslog messages occurring around the same time on the same node, to see if the cause of the problem can be identified. If the same error recurs, you might have to reboot the affected node. After the problem is corrected, the operator may choose to issue an `scswitch(1M)` command to bring resource groups onto desired primaries, or re-try the resource group update operation.

529191 :cllexecd: Sending fd to workerd returned %d. Exiting.

There was some error in setting up interprocess communication in the `cllexecd` program.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

529407 :resource group %s state on node %s change to %s

This is a notification from the `rgmd` that a resource group's state has changed. This may be used by system monitoring tools.

This is an informational message, no user action is needed.

530064 :reservation error(%s) - do_enfailfast() error for disk %s

The device fencing program has encountered errors while trying to access a device. All retry attempts have failed.

For the user action required by this message, see the user action for message 192619.

530492 :fatal: ucmm_initialize() failed

The daemon indicated in the message tag (`rgmd` or `ucmmd`) was unable to initialize its interface to the low-level cluster membership monitor. This is a fatal error, and causes the node to be halted or rebooted to avoid data corruption. The daemon produces a core file before exiting.

Save a copy of the /var/adm/messages files on all nodes, and of the core file generated by the daemon. Contact your authorized Sun service provider for assistance in diagnosing the problem.

530603 :Warning: Scalable service group for resource %s has already been created.

It was not expected that the scalable services group for the named resource existed.

Rebooting all nodes of the cluster will cause the scalable services group to be deleted.

531148 :fatal: thr_create stack allocation failure: %s (UNIX error %d)

The rgmd was unable to create a thread stack, most likely because the system has run out of swap space. The rgmd will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

Rebooting the node has probably cured the problem. If the problem recurs, you might need to increase swap space by configuring additional swap devices. See swap(1M) for more information.

531989 :Prog <%s> step <%s>: authorization error: %s.

An attempted program execution failed, apparently due to a security violation; this error should not occur. The last portion of the message describes the error. This failure is considered a program failure.

Correct the problem identified in the error message. If necessary, examine other syslog messages occurring at about the same time to see if the problem can be diagnosed. Save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance in diagnosing the problem.

532454 :file specified in USER_ENV parameter %s does not exist

"User_env" property was set when configuring the resource. File specified in "User_env" property does not exist or is not readable. File should be specified with fully qualified path.

Specify existing file with fully qualified file name when creating resource. If resource is already created, please update resource property "User_env".

532654 :The -c or -u flag must be specified for the %s method.

The arguments passed to the function unexpectedly omitted the given flags.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

532980 :clcomm: Pathend %p: deferred task not allowed in state %d

The system maintains state information about a path. A deferred task is not allowed in this state.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

534512 :in libsecurity svc_tp_create failed for transport %s

A server (rpc.pmfd, rpc.fed or rgmd) was not able to start because it could not create a rpc handle for the network specified. The rpc error message is shown. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

534826 :clxecd: Error %d from start_failfast_server

clxecd program could not enable one of the mechanisms which causes the node to be shutdown to prevent data corruption, when clxecd program dies.

To avoid data corruption, system will halt or reboot the node. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

535044 :Creation of resource <%s> failed because none of the nodes on which VALIDATE would have run are currently up

In order to create a resource whose type has a registered VALIDATE method, the rgmd must be able to run VALIDATE on at least one node. However, all of the candidate nodes are down. "Candidate nodes" are either members of the resource group's Nodelist or members of the resource type's Installed_nodes list, depending on the setting of the resource's Init_nodes property.

Boot one of the resource group's potential masters and retry the resource creation operation.

535181 :Host %s is not valid.

Validation method has failed to validate the ip addresses.

Invalid hostnames/ip addresses have been specified while creating resource. Recreate the resource with valid hostnames. Check the syslog message for the specific information.

535182 :in libsecurity NETPATH=%s

A server (rpc.pmfd, rpc.fed or rgmd) was not able to start because it could not establish a rpc connection for the network specified, because it couldn't find any

transport. This happened because either there are no available transports at all, or there are but none is a loopback. The NETPATH environment variable is shown. This error message is informational, and appears together with other messages appropriate for this situation. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

535886 :Could not find a mapping for %s in %s. It is recommended that a mapping for %s be added to %s.

No mapping was found in the local hosts file for the specified ip address.

Applications may use hostnames instead of ip addresses. It is recommended to have a mapping in the hosts file. Add an entry in the hosts file for the specified ip address.

536091 :Failed to retrieve the cluster handle while querying for property %s: %s.

Access to the object named failed. The reason for the failure is given in the message.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

536838 :clconf: Node %d is not in the cluster

Executed a program on the node that is not a current member of the cluster.

The specified node needs to be rebooted.

537175 :CMM: node %s (nodeid: %ld, incarnation #: %ld) has become reachable.

The cluster can communicate with the specified node. A node becomes reachable before it is declared up and having joined the cluster.

This is an informational message, no user action is needed.

537352 :reservation error(%s) - do_scsi3_preemptandabort() error for disk %s

The device fencing program has encountered errors while trying to access a device. All retry attempts have failed.

For the user action required by this message, see the user action for message 192619.

537380 :Invalid option -%c for the validate method.

Invalid option is passed to validate call back method.

This is an internal error. Contact your authorized Sun service provider for assistance in diagnosing and correcting the problem.

537498 :Invalid value was returned for resource property %s for %s.

The value returned for the named property was not valid.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

537607 :Not found clexecd on node %d for %d seconds. Retrying
...

Could not find clexecd to execute the program on a node. Indicated retry times.

This is an informational message, no user action is needed.

540274 :got unexpected exception %s

An inter-node communication failed with an unknown exception.

Examine syslog output for related error messages. Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file (if any). Contact your authorized Sun service provider for assistance in diagnosing the problem.

540376 :Unable to change the directory to %s: %s. Current directory is /.

Callback method is failed to change the current directory . Now the callback methods will be executed in "/", so the core dumps from this callbacks will be located in "/".

No user action needed. For detailed error message, check the syslog message.

541180 :Sun udlmlib library called with unknown option: "%c"

Unknown option used while starting up Oracle unix dlm.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

541206 :Couldn't read deleted directory: error (%d)

The file system is unable to create temporary copies of deleted files.

Mount the affected file system as a local file system, and ensure that there is no file system entry with name "._" at the root level of that file system. Alternatively, run fsck on the device to ensure that the file system is not corrupt.

541818 :Service group "%s" created

The service group by that name is now known by the scalable services framework.

This is an informational message, no user action is needed.

543720 :Desired Primaries is %d. It should be 1.

Invalid value for Desired Primaries property.

Invalid value is set for Desired Primaries property. The value should be 1. Reset the property value using scrgadm(1M).

544252 :Method <%s> on resource <%s>: Execution failed: no such method tag.

An internal error has occurred in the rpc.fed daemon which prevents method execution. This is considered a method failure. Depending on which method was being invoked and the Failover_mode setting on the resource, this might cause the resource group to fail over or move to an error state, or it might cause an attempted edit of a resource group or its resources to fail.

Examine other syslog messages occurring at about the same time to see if the problem can be identified. Save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance in diagnosing the problem. Re-try the edit operation.

544380 :Failed to retrieve the resource type handle: %s.

An API operation on the resource type has failed.

For the resource type name, check the syslog tag. For more details, check the syslog messages from other components. If the error persists, reboot the node.

544592 :PCSENTRY: %s

The rpc.pmfd server was not able to monitor a process, and the system error is shown. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

546856 :CCR: Could not find the CCR transaction manager.

The CCR data server could not find the CCR transaction manager in the cluster.

Reboot the cluster. Also contact your authorized Sun service provider to determine whether a workaround or patch is available.

547057 :thr_sigsetmask: %s

A server (rpc.pmfd or rpc.fed) was not able to establish a link with the failfast device because of a system error. The error message is shown. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

547385 :dl_bind: bad ACK header %u

An unexpected error occurred. The acknowledgment header for the bind request (to bind to the physical device) is bad. We are trying to open a fast path to the private transport adapters.

Reboot of the node might fix the problem.

549709 :Local node isn't in replica nodelist of service <%s> with path <%s>. No affinity switchover can be done

Local node does not support the replica of the service.

No user action required.

550143 :resource %s state on node %s change to R_OFFLINE

This is a notification from the rgmd that the resource has been brought offline on the indicated node. This may be used by system monitoring tools.

This is an informational message, no user action is needed.

550471 :Failed to initialize the cluster handle: %s.

An API operation has failed while retrieving the cluster information.

This may be solved by rebooting the node. For more details about API failure, check the messages from other components.

551094 :reservation warning(%s) - Unable to open device %s, will retry in %d seconds

The device fencing program has encountered errors while trying to access a device. The failed operation will be retried

This is an informational message, no user action is needed.

551139 :Failed to initialize scalable services group: Error %d.

The data service in scalable mode was unable to register itself with the cluster networking.

There may be prior messages in syslog indicating specific problems. Reboot the node if unable to correct the situation.

551436 :libsecurity: clnt_authenticate failed

A client of the rpc.pmfd, rpc.fed or rgmd server was not able to initiate an rpc connection, because it failed the authentication process. The pmfadm or scha command exits with error. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

553652 :fatal: cannot create thread to wake up President

The rgmd was unable to create a thread upon starting up. This is a fatal error. The rgmd will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

Make sure that the hardware configuration meets documented minimum requirements. Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

556466 :clexecd: dup2 of stdout returned with errno %d while exec"ing (%s). Exiting.

clexecd program has encountered a failed dup2(2) system call. The error message indicates the error number for the failure.

The clexecd program will exit and the node will be halted or rebooted to prevent data corruption. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

558742 :Resource group %s is online on more than one node.

The named resource group should be online on only one node, but it is actually online on more than one node.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

559550 :Error in opening /etc/vfstab: %s

Failed to open /etc/vfstab. Error message is followed.

Check with system administrator and make sure /etc/vfstab is properly defined.

559614 :Resource <%s> of Resource Group <%s> failed monitor check on node <%s>\n

Message logged for failed scha_control monitor check methods on specific node.

No user action required.

560047 :UNIX DLM version (%d) and SUN Unix DLM library version (%d): compatible.

The Unix DLM is compatible with the installed version of lubudlm.

None.

560781 :Tag %s: could not allocate history.

The rpc.pmfd server was not able to allocate memory for the history of the tag shown, probably due to low memory. The process associated with the tag is stopped and pmfadm returns error.

Investigate if the machine is running out of memory. If this is not the case, save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

562397 :Failfast: %s.

A failfast client has encountered a deferred panic timeout and is going to panic the node. This may happen if a critical userland process, as identified by the message, dies unexpectedly.

Check for core files of the process after rebooting the node and report these to your authorized Sun service provider.

563343 :resource type %s updated.

This is a notification from the rgmd that the operator has edited a property of a resource type. This may be used by system monitoring tools.

This is an informational message, no user action is needed.

563847 :INTERNAL ERROR: POSTNET_STOP method is not registered for resource <%s>

A non-fatal internal error has occurred in the rgmd state machine.

Since this problem might indicate an internal logic error in the rgmd, please save a copy of the /var/adm/messages files on all nodes, the output of an scstat -g command, and the output of a scradm -pvv command. Report the problem to your authorized Sun service provider.

564771 :Error in reading /etc/vfstab: getvfsent() returns <%d>

Error in reading /etc/vfstab. The return code of getvfsent() is followed.

Check with system administrator and make sure /etc/vfstab is properly defined.

565159 : "pmfadm -s": Error signaling <%s>: %s

An error occurred while rpc.pmfd attempted to send a signal to one of the processes of the given tag. The reason for the failure is also given. The signal was sent as a result of a "pmfadm -s" command.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

565198 :did subpath %s created for instance %d.

Informational message from scdidadm.

No user action required.

565438 :svc_run returned

The rpc.pmfd server was not able to run, due to an rpc error. This happens while the server is starting up, at boot time. The server does not come up, and an error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

565884 :tag %s: command file %s is not executable

The rpc.fed server checked the command indicated by the tag, and this check failed because the command is not executable. An error message is output to syslog.

Check the permission mode of the command, make sure that it is executable.

566781 :ORACLE_HOME %s does not exist

Directory specified as ORACLE_HOME does not exist. ORACLE_HOME property is specified when creating Oracle_server and Oracle_listener resources.

Specify correct ORACLE_HOME when creating resource. If resource is already created, please update resource property "ORACLE_HOME".

567374 :Failed to stop %s.

Sun Cluster failed to stop the application.

Use process monitor facility (pmfadm (1M)) with -L option to retrieve all the tags that are running on the server. Identify the tag name for the application in this resource. This can be easily identified as the tag ends in the string ".svc" and contains the resource group name and the resource name. Then use pmfadm (1M) with -s option to stop the application. This problem may occur when the cluster is under load and Sun Cluster cannot stop the application within the timeout period specified. You may consider increasing the Stop_timeout property. If the error still persists, then reboot the node.

567610 :PARAMTER_FILE %s does not exist

Oracle parameter file (typically init.ora) specified in property "Parameter_file" does not exist or is not readable.

Please make sure that "Parameter_file" property is set to the existing Oracle parameter file. Reissue command to create/update the resource using correct "Parameter_file".

567819 :clcomm: Fixed size resource_pool short server threads:
pool %d for client %d total %d

The system can create a fixed number of server threads dedicated for a specific purpose. The system expects to be able to create this fixed number of threads. The system could fail under certain scenarios without the specified number of threads. The server node creates these server threads when another node joins the cluster. The system cannot create a thread when there is inadequate memory.

There are two possible solutions. Install more memory. Alternatively, reduce memory usage. Application memory usage could be a factor, if the error occurs when a node joins an operational cluster and not during cluster startup.

568162 :Unable to create failfast thread

A server (rpc.pmfd or rpc.fed) was not able to start because it was not able to create the failfast thread, which ensures that the host aborts if the server dies. An error message is output to syslog.

Investigate if the host is low on memory. If not, save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

568314 :Failed to remove node %d from scalable service group %s:
%s.

A call to the underlying scalable networking code failed.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

570802 :fatal: Got error <%d> trying to read CCR when disabling
monitor of resource <%s>; aborting node

Rgmd failed to read updated resource from the CCR on this node.

Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

571642 :ucm_callback for cmmreturn generated exception %d

ucmm callback for step cmmreturn failed.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

571734 :Validation failed. ORACLE_SID is not set

ORACLE_SID property for the resource is not set. HA-Oracle will not be able to manage Oracle server if ORACLE_SID is incorrect.

Specify correct ORACLE_SID when creating resource. If resource is already created, please update resource property "ORACLE_SID".

571825 :Stopping listener %s.

Informational message. HA-Oracle will be stopping Oracle listener.

None

571950 :Fault monitor detected error %s: %ld Action=%s : %s

Fault monitor has detected an error. Error detected by fault monitor and action taken by fault monitor is indicated in message.

None

572955 :host %s: client is null

The rgm is not able to obtain an rpc client handle to connect to the rpc.fed server on the named host. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

573690 :reservation error(%s) - did_get_num_paths() error in is_scsi3_disk(), returned %d

The device fencing program has suffered an internal error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available. Copies of /var/adm/messages from all nodes should be provided for diagnosis. It may be possible to retry the failed operation, depending on the nature of the error. If the message specifies the "node_join" transition, then this node may be unable to access shared devices. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access shared devices. In either case, it may be possible to reacquire access to shared devices by executing "/usr/cluster/lib/sc/run_reserve -c node_join" on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group has failed to start on this node. If another node was available to host the device group, then it should have been started on that node. If desired, it may be possible to switch the device group to this node with the scswitch command. If no other node was available, then the device group will not have been started. The scswitch command may be used to retry the attempt to start the device group. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group has failed. The desired action may be retried.

574542 :clexecd: fork1 returned %d. Exiting.

clexecd program has encountered a failed fork1(2) system call. The error message indicates the error number for the failure.

If the error number is 12 (ENOMEM), install more memory, increase swap space, or reduce peak memory consumption. If error number is something else, contact your authorized Sun service provider to determine whether a workaround or patch is available.

574675 :nodeid of ctxp is bad: %d

nodeid in the context pointer is bad.

None. udln takes appropriate action.

575545 :fatal: rgm_chg_freeze: INTERNAL ERROR: invalid value of rgl_is_frozen <%d> for resource group <%s>

The in-memory state of the rgmd has been corrupted due to an internal error. The rgmd will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

575853 :libsecurity: create of rpc handle to program %ld failed, will keep trying

A client of the rpc.pmfd, rpc.fed or rgmd server was not able to initiate an rpc connection. The maximum time allowed for connecting (1 hr) has not been reached yet, and the pmfadm or scha command will retry to connect. An accompanying error message shows the rpc error data. The program number is shown. To find out what program corresponds to this number, use the rpcinfo command. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

576196 :clcomm: error loading kernel module: %d

The loading of the cl_comm module failed.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

576744 :INTERNAL ERROR: Invalid resource property type <%d> on resource <%s>

An attempted creation or update of a resource has failed because of invalid resource type data. This may indicate CCR data corruption or an internal logic error in the rgmd.

Use scrgadm(1M) -pvv to examine resource properties. If the resource or resource type properties appear to be corrupted, the CCR might have to be

rebuilt. If values appear correct, this may indicate an internal error in the rgmd. Re-try the creation or update operation. If the problem recurs, save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance.

577140 :clcomm: Exception during unmarshal_receive

The server encountered an exception while unmarshalling the arguments for a remote invocation. The system prints the exception causing this error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

579190 :INTERNAL ERROR: resource group <%s> state <%s> node <%s> contains resource <%s> in state <%s>

The rgmd has discovered that the indicated resource group's state information appears to be incorrect. This may prevent any administrative actions from being performed on the resource group.

Since this problem might indicate an internal logic error in the rgmd, please save a copy of the /var/adm/messages files on all nodes, the output of an scstat -g command, and the output of a scrgadm -pvv command. Report the problem to your authorized Sun service provider.

579235 :Method <%s> on resource <%s> terminated abnormally

A resource method terminated without using an exit(2) call. The rgmd treats this as a method failure.

Consult resource type documentation, or contact the resource type developer for further information.

579987 :Error binding "%s" in the name server. Exiting.

clxecd program was unable to start because of some problems in the low-level clustering software.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

580163 :reservation warning(%s) - MHIOCTKOWN error will retry in %d seconds

The device fencing program has encountered errors while trying to access a device. The failed operation will be retried

This is an informational message, no user action is needed.

580416 :Cannot restart monitor: Monitor is not enabled.

An update operation on the resource would have been restarted the fault monitor. But, the monitor is currently disabled for the resource.

This is informational message. Check whether the monitor is disabled for the resource. If not, consider it as an internal error and contact your authorised Sun service provider.

581180 :launch_validate: call to rpc.fed failed for resource
<%s>, method <%s>

The rgmd failed in an attempt to execute a VALIDATE method, due to a failure to communicate with the rpc.fed daemon. If the rpc.fed process died, this might lead to a subsequent reboot of the node. Otherwise, this will cause the failure of a creation or update operation on a resource or resource group.

Examine other syslog messages occurring at about the same time to see if the problem can be identified. Re-try the creation or update operation. If the problem recurs, save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance.

581376 :clcomm: solaris xdoor: too much reply data

The reply from a user level server will not fit in the available space.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

581413 :Daemon %s is not running.

HA-NFS fault monitor checks the health of statd, lockd, mountd and nfsd daemons on the node. It detected that one these are not currently running.

No action. The monitor would restart these.

581902 :(%s) invalid timeout "%d"

Invalid timeout value for a method.

Make sure udln.conf file has correct timeouts for methods.

582651 :tag %s: does not belong to caller

The user sent a suspend/resume command to the rpc.fed server for a tag that was started by a different user. An error message is output to syslog.

Check the tag name.

582757 :No PDT Fastpath thread.

The system has run out of resources that is required to create a thread. The system could not create the Fastpath thread that is required for cluster networking.

If cluster networking is required, add more resources (most probably, memory) and reboot.

583138 :dfstab not readable

HA-NFS fault monitor failed to read dfstab when it detected that dfstab has been modified.

Make sure the dfstab file exists and has read permission set appropriately. Look at the prior syslog messages for any specific problems and correct them.

583224 :Rebooting this node because daemon %s is not running.

The rpcbind daemon on this node is not running.

No action. Fault monitor would reboot the node. Also see message id 804791.

583542 :clcomm: Pathend: would abort node because %s for %u ms

The system would have aborted the node for the specified reason if the check for send thread running was enabled.

No user action is required.

583563 :fatal: rgm_run_state: internal error: bad state <%d> for resource group <%s>

An internal error has occurred. The rgmd will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

584386 :PENDING_OFFLINE: bad resource state <%s> (%d) for resource <%s>

The rgmd state machine has discovered a resource in an unexpected state on the local node. This should not occur and may indicate an internal logic error in the rgmd.

Look for other syslog error messages on the same node. Save a copy of the /var/adm/messages files on all nodes, and report the problem to your authorized Sun service provider.

586298 :clcomm: unknown type of signals message %d

The system has received a signals message of unknown type.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

586344 :clcomm: unable to unbind %s from name server

The name server would not unbind the specified entity.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

589719 :Issuing failover request.

This is informational message. We are about to call API function to request for failover. In case of failure, follow the syslog messages after this message.

No user action is needed.

589817 :clcomm: nil_sendstream::send

The system attempted to use a "send" operation for a local invocation. Local invocations do not use a "send" operation.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

590263 :Online check Error %s : %ld

Error detected when checking ONLINE status of RDBMS. Error number is indicated in message. This can be because of RDBMS server problems or configuration problems.

Check RDBMS server using vendor provided tools. If server is running properly, this can be fault monitor set-up error.

590700 :ALERT_LOG_FILE %s doesn't exist

File specified in resource property "Alert_log_file" does not exist. HA-Oracle requires correct Alert Log file for fault monitoring.

Check "Alert_log_file" property of the resource. Specify correct Oracle Alert Log file when creating resource. If resource is already created, please update resource property Alert_log_file".

592285 :clexecd: getrlimit returned %d

clexecd program has encountered a failed getrlimit(2) system call. The error message indicates the error number for the failure.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

593330 :Resource type name is null.

This is an internal error. While attempting to retrieve the resource information, null value was retrieved for the resource type name.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

594629 :Failed to stop the fault monitor.

Process monitor facility has failed to stop the fault monitor.

Use pmfadm(1M) with -L option to retrieve all the tags that are running on the server. Identify the tag name for the fault monitor of this resource. This can be easily identified, as the tag ends in string ".mon" and contains the resource group name and the resource name. Then use pmfadm(1M) with -s option to stop the fault monitor. If the error still persists, then reboot the node.

594675 :reservation warning(%s) - MHI0CGRP_REGISTER error will retry in %d seconds

The device fencing program has encountered errors while trying to access a device. The failed operation will be retried

This is an informational message, no user action is needed.

595101 :t_sndudata in send_reply: %s

Call to t_sndudata() failed. The "t_sndudata" man page describes possible error codes. udlm will try to resend the message. abort.

None.

595686 :%s is %d for %s. It should be 1.

The named property has an unexpected value.

Change the value of the property to be 1.

596447 :UNIX DLM is asking for a reconfiguration to recover from a communication error.

A reconfiguration has been requested by udlm.

None.

596604 :clcomm: solookup on routing socket failed with error = %d

The system prepares IP communications across the private interconnect. A lookup operation on the routing socket failed.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

597171 :Unexpected early exit while performing: "%s"

clexecd program got an error while executing the program indicated in the error message.

Please check the error message. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

597239 :The weight portion of %s at position %d in property %s is not a valid weight. The weight should be an integer between %d and %d.

The weight noted does not have a valid value. The position index, which starts at 0 for the first element in the list, indicates which element in the property list was invalid.

Give the weight a valid value.

597381 :setrlimit before exec: %s

rpc.pmfd was unable to set the number of file descriptors before executing a process.

Look for other syslog error messages on the same node. Save a copy of the /var/adm/messages files on all nodes, and report the problem to your authorized Sun service provider.

598087 :PCWSTOP: %s

The rpc.pmfd server was not able to monitor a process, and the system error is shown. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

598259 :scvxvmlg fatal error - ckmode received unknown mode %d

The program responsible for maintaining the VxVM namespace has suffered an internal error. If configuration changes were recently made to VxVM diskgroups or volumes, this node may be unaware of those changes. Recently created volumes may be inaccessible from this node.

If no configuration changes have been recently made to VxVM diskgroups or volumes and all volumes continue to be accessible from this node, then no action is required. If changes have been made, the device namespace on this node can be updated to reflect those changes by executing "/usr/cluster/lib/dcs/scvxvmlg". If the problem persists, contact your authorized Sun service provider to determine whether a workaround or patch is available.

598540 :clcomm: solaris xdoor: completed invo: door_return returned, errno = %d

An unusual but harmless event occurred. System operations continue unaffected.

No user action is required.

598554 :launch_validate_method: getlocalhostname() failed for resource <%s>, resource group <%s>, method <%s>

The rgmd was unable to obtain the name of the local host, causing a VALIDATE method invocation to fail. This in turn will cause the failure of a creation or update operation on a resource or resource group.

Examine other syslog messages occurring at about the same time to see if the problem can be identified. Re-try the creation or update operation. If the problem recurs, save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance.

598979 :tag %s: already suspended

The user sent a suspend command to the rpc.fed server for a tag that is already suspended. An error message is output to syslog.

Check the tag name.

599430 :Failed to retrieve the resource property %s: %s.

An API operation has failed while retrieving the resource property. Low memory or API call failure might be the reasons.

In case of low memory, the problem will probably be cured by rebooting. If the problem reoccurs, you might need to increase swap space by configuring additional swap devices. Otherwise, if it is API call failure, check the syslog messages from other components. For the resource name and property name, check the current syslog message.

599558 :SIOCLIFADDIF of %s failed: %s.

Specified system operation failed

This is an internal error. Contact your authorized Sun service provider with the following information. 1) Saved copy of /var/adm/messages file. 2) Output of "ifconfig -a" command.

Message IDs 600000 - 699999

Error Message List

The following list is ordered by the message ID.

600967 :Could not allocate buffer for DBMS log messages: %m

Fault monitor could not allocate memory for reading RDBMS log file. As a result of this error, fault monitor will not scan errors from log file. However it will continue fault monitoring.

Check if syetem is low on memory. If problem persists, please stop and start the fault monitor.

601901 :Failed to retrieve the resource property %s for %s: %s.

The query for a property failed. The reason for the failure is given in the message.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

603096 :resource %s disabled.

This is a notification from the rgmd that the operator has disabled a resource. This may be used by system monitoring tools.

This is an informational message, no user action is needed.

603192 :For global service <%s> of path <%s>, local node is less preferred than node <%d>. But affinity switch over may still be done.

A service is switched to a less preferred node due to affinity switchover of SUNW.HAStorage prenet_start method.

Check which configuration can gain more performance benefit, either to leave the service on its most preferred node or let the affinity switchover take effect. Using scswitch(1m) to switch it back if necessary.

603479 :This node can be a primary for scalable resource %s, but there is no NAFO group defined on this node. A NAFO group must be created on this node.

The node does not have a NAFO group defined.

Any adapters on the node which are connected to the public network should be put under NAFO control by placing them in a NAFO group. See the pnmset(1M) man page for details.

604153 :clcomm: Path %s errors during initiation

Communication could not be established over the path. The interconnect may have failed or the remote node may be down.

Any interconnect failure should be resolved, and/or the failed node rebooted.

605301 :lkcm_sync: invalid handle was passed %s%d

Invalid handle passed during lockstep execution.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

606203 :Couldn't get the root vnode: error (%d)

The file system is corrupt or was not mounted correctly.

Run fsck, and mount the affected file system again.

606467 :CMM: Initialization for quorum device %s failed with error EACCES. Will retry later.

This node is not able to access the specified quorum device because the node is still fenced off. An attempt will be made to access the quorum device again after the node's CCR has been recovered.

This is an informational message, no user action is needed.

607054 :%s not found.

Could not find the binary to startup udln.

Make sure the unix dlm package is installed properly.

607613 :transition "%s" timed out for cluster, as did attempts to reconfigure.

Step transition failed. udlmctl will exit.

Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

607678 :clconf: No valid quorum_resv_key field for node %u

Found the quorum_resv_key field being incorrect while converting the quorum configuration information into quorum table.

Check the quorum configuration information.

608202 :scha_control: resource group <%s> was frozen on Global_resources_used within the past %d seconds; exiting

A scha_control call has failed with a SCHA_ERR_CHECKS error because the resource group has a non-null Global_resources_used property, and a global device group was failing over within the indicated recent time interval. The resource fault probe is presumed to have failed because of the temporary unavailability of the device group. A properly-written resource monitor, upon getting the SCHA_ERR_CHECKS error code from a scha_control call, should sleep for awhile and restart its probes.

No user action is required. Either the resource should become healthy again after the device group is back online, or a subsequent scha_control call should succeed in failing over the resource group to a new master.

608453 :failfast disarm error: %d

Error during a failfast device disarm operation.

None.

608876 :PCRUN: %s

The rpc.pmfd server was not able to monitor a process, and the system error is shown. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

609118 :Error creating deleted directory: error (%d)

While mounting this file system, PXFS was unable to create some directories that it reserves for internal use.

If the error is 28(ENOSPC), then mount this FS non-globally, make some space, and then mount it globally. If there is some other error, and you are unable to

correct it, contact your authorized Sun service provider to determine whether a workaround or patch is available.

612049 :resource <%s> in resource group <%s> depends on disabled network address resource <%s>

An enabled application resource was found to implicitly depend on a network address resource that is disabled. This error is non-fatal but may indicate an internal logic error in the rgmd.

Look for other syslog error messages on the same node. Save a copy of the /var/adm/messages files on all nodes, and report the problem to your authorized Sun service provider.

612124 :Volume configuration daemon not running.

Volume manager is not running.

Bring up the volume manager.

612931 :Unable to get device major number for %s driver: %s.

System was unable to translate the given driver name into device major number.

Check whether the /etc/name_to_major file is corrupted. Reboot the node if problem persists.

613522 :clexecd: Error %d from poll. Exiting.

clexecd program has encountered a failed poll(2) system call. The error message indicates the error number for the failure.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

613896 :INTERNAL ERROR: process_resource: Resource <%s> is R_BOOTING in PENDING_OFFLINE or PENDING_DISABLED resource group

The rgmd is attempting to bring a resource group offline on a node where BOOT methods are still being run on its resources. This should not occur and may indicate an internal logic error in the rgmd.

Look for other syslog error messages on the same node. Save a copy of the /var/adm/messages files on all nodes, and report the problem to your authorized Sun service provider.

615120 :fatal: unknown scheduling class "%s"

An internal error has occurred. The daemon indicated in the message tag (rgmd or ucmmmd) will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

Save a copy of the /var/adm/messages files on all nodes, and of the core file generated by the daemon. Contact your authorized Sun service provider for assistance in diagnosing the problem.

617643 :Unable to fork(): %s.

Upon a NAFO failure, the system was unable to take any action, because it failed to fork another process.

This might be the result from the lack of the system resources. Check whether the system is low in memory or the process table is full, and take appropriate action. For specific error information check the syslog message.

617917 :Initialization failed. Invalid command line %s %s

Unable to process parameters passed to the call back method. This is an internal error.

Please report this problem.

618107 :Path %s initiation encountered errors, errno = %d.
Remote node may be down or unreachable through this path.

Communication with another node could not be established over the path.

Any interconnect failure should be resolved, and/or the failed node rebooted.

618466 :Unix DLM no longer running

UNIX DLM is expected to be running, but is not. This will result in a udlmstep1 failure.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

618585 :clexecd: getmsg returned %d. Exiting.

clexecd program has encountered a failed getmsg(2) system call. The error message indicates the error number for the failure.

The clexecd program will exit and the node will be halted or rebooted to prevent data corruption. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

618637 :The port number %d from entry %s in property %s was not found in config file <%s>.

All entries in the list property must have port numbers that correspond to ports configured in the configuration file. The port number from the list entry does not correspond to a port in the configuration file.

Remove the entry or change its port number to correspond to a port in the configuration file.

618764 :fe_set_env_vars() failed for Resource <%s>, resource group <%s>, method <%s>

The rgmd was unable to set up environment variables for a method execution, causing the method invocation to fail. Depending on which method was being invoked and the Failover_mode setting on the resource, this might cause the resource group to fail over or move to an error state.

Examine other syslog messages occurring at about the same time to see if the problem can be identified. Save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance in diagnosing and correcting the problem.

619213 :t_alloc (recv_request) failed with error %d

Call to t_alloc() failed. The "t_alloc" man page describes possible error codes. ucmmmd will exit and the node will abort.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

619312 :"%s" restarting too often ... sleeping %d seconds.

The tag shown, run by rpc.pmfd server, is restarting and exiting too often. This means more than once a minute. This can happen if the application is restarting, then immediately exiting for some reason, then the action is executed and returns OK (0), which causes the server to restart the application. When this happens, the rpc.pmfd server waits for up to 1 min before it restarts the application. An error message is output to syslog.

Examine the state of the application, try to figure out why the application doesn't stay up, and yet the action returns OK.

620204 :Failed to start scalable service.

Unable to configure service for scalability.

The start method on this node will fail. Sun Cluster resource management will attempt to start the service on some other node.

621264 :clconf: Not found clexecd on node %d for %d seconds. Retrying ...

Could not find clexecd to execute the program on a node. Indicated retry times.

No action required. This is informational message.

621686 :CCR: Invalid checksum length %d in table %s, expected %d.

The checksum of the indicated table has a wrong size. This causes the consistency check of the indicated table to fail.

Boot the offending node in -x mode to restore the indicated table from backup or other nodes in the cluster. The CCR tables are located at /etc/cluster/ccr/.

622387 :constchar*fmt

Function definition. Please ignore

None

623528 :clcomm: Unregister of adapter state proxy failed

The system failed to unregister an adapter state proxy.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

623759 :svc_setschedprio: Could not lookup RT (real time) scheduling class info: %s

The server was not able to determine the scheduling mode info, and the system error is shown. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

624447 :fatal: sigaction: %s (UNIX errno %d)

The rgmd has failed to initialize signal handlers by a call to sigaction(2). The error message indicates the reason for the failure. The rgmd will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

627610 :clconf: Invalid clconf_obj type

An invalid clconf_obj type has been encountered while converting an clconf_obj type to group name. Valid objtypes are "CL_CLUSTER", "CL_NODE", "CL_ADAPTER", "CL_PORT", "CL_BLACKBOX", "CL_CABLE", "CL_QUORUM_DEVICE".

This is an unrecoverable error, and the cluster needs to be rebooted. Also contact your authorized Sun service provider to determine whether a workaround or patch is available.

628203 :in libsecurity could not find any tcp transport

A client was not able to make an rpc connection to a server (rpc.pmfd, rpc.fed or rgmd) because it could not find a tcp transport. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

628771 :CCR: Can't read CCR metadata.

Reading the CCR metadata failed on this node during the CCR data server initialization.

There may be other related messages on this node, which may help diagnose the problem. For example: If the root disk on the afflicted node has failed, then it needs to be replaced. If the cluster repository is corrupted, then boot this node in -x mode to restore the cluster repository from backup or other nodes in the cluster. The cluster repository is located at /etc/cluster/ccr/.

630653 :Failed to initialize DCS

There was a fatal error while this node was booting.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

631408 :PCSET: %s

The rpc.pmfd server was not able to monitor a process, and the system error is shown. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

631429 :huge address size %d

Size of MAC address in acknowledgment of the bind request exceeds the maximum size allotted. We are trying to open a fast path to the private transport adapters.

Reboot of the node might fix the problem.

631648 :Retrying to retrieve the resource group information.

An update to cluster configuration occurred while resource group properties were being retrieved

Ignore the message.

634957 :thr_keycreate failed in init_signal_handlers

The ucmmmd failed in a call to thr_keycreate(3T). ucmmmd will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

Save a copy of the /var/adm/messages files on all nodes and of the ucmm core. Contact your authorized Sun service provider for assistance in diagnosing the problem.

636848 :Failed to get all NAFO groups (request failed with %d).

An unexpected error occurred while trying to communicate with the network monitoring daemon (pnmd).

Make sure the network monitoring daemon (pnmd) is running.

637677 :(%s) t_alloc: tli error: %s

Call to t_alloc() failed. The "t_alloc" man page describes possible error codes. udlmctl will exit.

Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

640029 :PENDING_ONLINE: bad resource state <%s> (%d) for resource <%s>

The rgmd state machine has discovered a resource in an unexpected state on the local node. This should not occur and may indicate an internal logic error in the rgmd.

Look for other syslog error messages on the same node. Save a copy of the /var/adm/messages files on all nodes, and report the problem to your authorized Sun service provider.

640087 :udlmctl: incorrect comand line

udlmctl will not startup because of incorrect command line options.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

640090 :CMM: Initialization for quorum device %s failed with error %d.

The initialization of the specified quorum device failed with the specified error, and this node will ignore this quorum device.

There may be other related messages on this node which may indicate the cause of this problem. Refer to the quorum disk repair section of the administration guide for resolving this problem.

640484 :clconf: No valid voteount field for quorum device %d

Found the votecount field for the quorum device being incorrect while converting the quorum configuration information into quorum table.

Check the quorum configuration information.

640799 :pmf_alloc_thread: ENOMEM

The rpc.pmfd server was not able to allocate a new monitor thread, probably due to low memory. As a consequence, the rpc.pmfd server was not able to monitor a process. An error message is output to syslog.

Investigate if the machine is running out of memory. If this is not the case, save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

643472 :fatal: Got error <%d> trying to read CCR when enabling resource <%s>; aborting node

Rgmd failed to read updated resource from the CCR on this node.

Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

643802 :Resource group is online on more than one node.

An internal error has occurred. Resource group should be online on only one node.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

644140 :Fault monitor is not running.

Sun cluster tried to stop the fault monitor for this resource, but the fault monitor was not running. This is most likely because the fault monitor was unable to start.

Look for prior syslog messages relating to starting of fault monitor and take corrective action. No other action needed

644850 :File %s is not readable: %s.

Unable to open the file in read only mode.

Make sure the specified file exists and have correct permissions. For the file name and details, check the syslog messages.

645665 :resource group %s state on node %s change to RG_OFFLINE

This is a notification from the rgmd that a resource group's state has changed. This may be used by system monitoring tools.

This is an informational message, no user action is needed.

646815 :PCUNSET: %s

The rpc.pmfd server was not able to monitor a process, and the system error is shown. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

646950 :clcomm: Path %s being cleaned up

A communication link is being removed with another node. The interconnect may have failed or the remote node may be down.

Any interconnect failure should be resolved, and/or the failed node rebooted.

647339 :(%s) scan of dlmmmap failed on "%s", idx =%d

Failed to scan dlmmmap.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

647673 :scvxvmlg error - dcs_get_service_parameters() failed, returned %d

The program responsible for maintaining the VxVM namespace has suffered an internal error. If configuration changes were recently made to VxVM diskgroups or volumes, this node may be unaware of those changes. Recently created volumes may be inaccessible from this node.

If no configuration changes have been recently made to VxVM diskgroups or volumes and all volumes continue to be accessible from this node, then no action is required. If changes have been made, the device namespace on this node can be updated to reflect those changes by executing "/usr/cluster/lib/dcs/scvxvmlg". If the problem persists, contact your authorized Sun service provider to determine whether a workaround or patch is available.

648339 :Failed to retrieve ip addresses configured on adapter %s.

System was attempting to list all the ip addresses configured on the specified adapter, but it was unable to do that.

Check the messages that are logged just before this message for possible causes. For more help, contact your authorized Sun service provider with the following information. Output of /var/adm/messages file and the output of "ifconfig -a" command .

649584 :Modification of resource group <%s> failed because none of the nodes on which VALIDATE would have run for resource <%s> are currently up

Before it will permit the properties of a resource group to be edited, the rgmd runs the VALIDATE method on each resource in the group for which a VALIDATE method is registered. For each such resource, the rgmd must be able to run VALIDATE on at least one node. However, all of the candidate nodes are down. "Candidate nodes" are either members of the resource group's Nodelist or members of the resource type's Installed_nodes list, depending on the setting of the resource's Init_nodes property.

Boot one of the resource group's potential masters and retry the resource creation operation.

649860 :RGM isn't failing resource group <%s> off of node <%d>, because no current or potential master is healthy enough

A scha_control(1HA,3HA) GIVEOVER attempt failed on all potential masters, because no candidate node was healthy enough to host the resource group.

Examine other syslog messages on all cluster members that occurred about the same time as this message, to see if the problem that caused the MONITOR_CHECK failure can be identified. Repair the condition that is preventing any potential master from hosting the resource.

650276 :Failed to get port numbers from config file <%s>.

An error occurred while parsing the configuration file to extract port numbers.

Check that the configuration file path exists and is accessible. Check that port keywords and values exist in the file.

650390 :Validation failed. init.ora file does not exist: %s

Oracle Parameter file has not been specified. Default parameter file indicated in the message does not exist. Cannot start Oracle server.

Please make sure that parameter file exists at the location indicated in message or specify "Parameter_file" property for the resource. Clear START_FAILED flag on the resource and bring the resource online.

650825 :Method <%s> on resource <%s> terminated due to receipt of signal <%d>

A resource method was terminated by a signal, most likely resulting from an operator-issued kill(1). The method is considered to have failed.

No action is required. The operator may choose to issue an scswitch(1M) command to bring resource groups onto desired primaries, or re-try the administrative action that was interrupted by the method failure.

650932 :malloc failed for ipaddr string

Call to malloc failed. The "malloc" man page describes possible reasons.

Install more memory, increase swap space or reduce peak memory consumption.

651093 :reservation message(%s) - Fencing node %d from disk %s

The device fencing program is taking access to the specified device away from a non-cluster node.

This is an informational message, no user action is needed.

651327 :Failed to delete scalable service group %s: %s.

A call to the underlying scalable networking code failed.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

652399 :Ignoring the SCHA_ERR_SEQID while retrieving %s

An update to the cluster configuration tables occurred while trying to retrieve certain cluster related information. However, the update does not affect the property that is being retrieved.

Ignore the message

653062 :Syntax error on line %s in dfstab file.

The specified share command is incorrect.

Correct the share command using the dfstab(4) man pages.

653183 :Unable to create the directory %s: %s. Current directory is /.

Callback method is failed to create the directory specified. Now the callback methods will be executed in "/", so the core dumps from this callbacks will be located in "/".

No user action needed. For detailed error message, check the syslog message.

654520 :INTERNAL ERROR: rgm_run_state: bad state <%d> for resource group <%s>

The rgmd state machine on this node has discovered that the indicated resource group's state information is corrupted. The state machine will not launch any methods on resources in this resource group. This may indicate an internal logic error in the rgmd.

Other syslog messages occurring before or after this one might provide further evidence of the source of the problem. If not, save a copy of the /var/adm/

messages files on all nodes, and (if the rgmd crashes) a copy of the rgmd core file, and contact your authorized Sun service provider for assistance.

656721 :clexecd: %s: sigdelset returned %d. Exiting.

clexecd program has encountered a failed sigdelset(3C) system call. The error message indicates the error number for the failure.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

656795 :CMM: Unable to bind <%s> to nameserver.

An instance of the userland CMM encountered an internal initialization error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

657560 :CMM: Reading reservations from quorum device %s failed with error %d.

The specified error was encountered while trying to read reservations on the specified quorum device.

There may be other related messages on this and other nodes connected to this quorum device that may indicate the cause of this problem. Refer to the quorum disk repair section of the administration guide for resolving this problem.

658555 :Retrying to retrieve the resource information.

An update to cluster configuration occurred while resource properties were being retrieved

Ignore the message.

659665 :kill -KILL: %s

The rpc.fed server is not able to stop a tag that timed out, and the error message is shown. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

659827 :CCR: Can't access CCR metadata on node %s errno = %d.

The indicated error occurred when CCR is trying to access the CCR metadata on the indicated node. The errno value indicates the nature of the problem. errno values are defined in the file /usr/include/sys/errno.h. An errno value of 28(ENOSPC) indicates that the root file system on the node is full. Other values of errno can be returned when the root disk has failed(EIO).

There may be other related messages on the node where the failure occurred. These may help diagnose the problem. If the root file system is full on the node,

then free up some space by removing unnecessary files. If the root disk on the afflicted node has failed, then it needs to be replaced. If the cluster repository is corrupted, boot the indicated node in -x mode to restore it from backup. The cluster repository is located at /etc/cluster/ccr/.

660332 :launch_validate: fe_set_env_vars() failed for resource <%s>, resource group <%s>, method <%s>

The rgmd was unable to set up environment variables for method execution, causing a VALIDATE method invocation to fail. This in turn will cause the failure of a creation or update operation on a resource or resource group.

Examine other syslog messages occurring at about the same time to see if the problem can be identified. Re-try the creation or update operation. If the problem recurs, save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance.

660974 :file specified in USER_ENV %s does not exist

"User_env" property was set when configuring the resource. File specified in "User_env" property does not exist or is not readable. File should be specified with fully qualified path.

Specify existing file with fully qualified file name when creating resource. If resource is already created, please update resource property "User_env".

661614 :Method <%s> failed on resource <%s> in resource group <%s>, exit code <%d>

A resource method exited with a non-zero exit code; this is considered a method failure. Depending on which method is being invoked and the Failover_mode setting on the resource, this might cause the resource group to fail over or move to an error state.

Consult resource type documentation to diagnose the cause of the method failure. Other syslog messages occurring just before this one might indicate the reason for the failure. After correcting the problem that caused the method to fail, the operator may choose to issue an scswitch(1M) command to bring resource groups onto desired primaries.

663089 :clexecd: %s: sigwait returned %d. Exiting.

clexecd program has encountered a failed sigwait(3C) system call. The error message indicates the error number for the failure.

The clexecd program will exit and the node will be halted or rebooted to prevent data corruption. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

663293 :reservation error(%s) - do_status() error for disk %s

The device fencing program has encountered errors while trying to access a device. All retry attempts have failed.

The action which failed is a scsi-2 ioctl. These can fail if there are scsi-3 keys on the disk. To remove invalid scsi-3 keys from a device, use "scdidadm -R" to repair the disk (see scdidadm man page for details). If there were no scsi-3 keys present on the device, then this error is indicative of a hardware problem, which should be resolved as soon as possible. Once the problem has been resolved, the following actions may be necessary: If the message specifies the "node_join" transition, then this node may be unable to access the specified device. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access the device. In either case, access can be reacquired by executing "/usr/cluster/lib/sc/run_reserve -c node_join" on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group may have failed to start on this node. If the device group was started on another node, it may be moved to this node with the scswitch command. If the device group was not started, it may be started with the scswitch command. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group may have failed. If so, the desired action may be retried.

663835 :in libsecurity creat of file %s failed: %s

The rpc.pmfd, rpc.fed or rgmd server was not able to create a cache file for rpcbind information. The affected component should continue to function by calling rpcbind directly.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

663851 :Failover %s data services must have exactly one value for extension property %s.

Failover data services must have one and only one value for Confdir_list.

Create a failover resource group for each configuration file.

663897 :clcomm: Endpoint %p: %d is not an endpoint state

The system maintains information about the state of an Endpoint. The Endpoint state is invalid.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

663943 :Quorum: Unable to reset node information on quorum disk.

This node was unable to reset some information on the quorum device. This will lead the node to believe that its partition has been preempted. This is an internal error. If a cluster gets divided into two or more disjoint subclusters, exactly one

of these must survive as the operational cluster. The surviving cluster forces the other subclusters to abort by grabbing enough votes to grant it majority quorum. This is referred to as preemption of the losing subclusters.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

665015 :Scalable service instance [%s,%s,%d] registered on node %s.

The specified scalable service has been registered on the specified node. Now, the gif node can redirect packets for the specified service to this node.

This is an informational message, no user action is needed.

665195 :INTERNAL ERROR: rebalance: invalid node name in Nodelist of resource group <%s>

An internal error has occurred in the rgmd. This error may prevent the rgmd from bringing the affected resource group online.

Look for other syslog error messages on the same node. Save a copy of the /var/adm/messages files on all nodes, and report the problem to your authorized Sun service provider.

665931 :Initialization error. CONNECT_STRING is NULL

Error occurred in monitor initialization. Monitor is unable to get resource property "Connect_string".

Check syslog messages for errors logged from other system modules. Check the resource configuration and value of "Connect_string" property. Check syslog messages for errors logged from other system modules. Stop and start fault monitor. If error persists then disable fault monitor and report the problem.

666391 :clcomm: invalid invocation result status %d

An invocation completed with an invalid result status.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

666443 :unix DLM already running

UNIX DLM is already running. Another dlm will not be started.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

666603 :clxecd: Error %d in fcntl(F_GETFD). Exiting.

clexecd program has encountered a failed fcntl(2) system call. The error message indicates the error number for the failure.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

667020 :Invalid shared path

HA-NFS fault monitor detected that one or more shared paths in dfstab are invalid paths.

Make sure all paths in dfstab are correct. Look at the prior syslog messages for any specific problems and correct them.

669026 :fcntl(F_SETFD) failed in close_on_exec

A fcntl operation failed. The "fcntl" man page describes possible error codes.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

670799 :CMM: Registering reservation key on quorum device %s failed with error %d.

The specified error was encountered while trying to place the local node's reservation key on the specified quorum device. This node will ignore this quorum device.

There may be other related messages on this and other nodes connected to this quorum device that may indicate the cause of this problem. Refer to the quorum disk repair section of the administration guide for resolving this problem.

671954 :waitpid: %s

The rpc.pmfd server was not able to wait for a process. The message contains the system error. The server does not perform the action requested by the client, and an error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

672019 :Stop method failed. Error: %d.

Stop method failed, while attempting to restart the data service.

Check the Stop_timeout and adjust it if it is not appropriate. For the detailed explanation of failure, check the syslog messages that occurred just before this message.

672372 :dl_attach: bad ACK header %u

Could not attach to the physical device. We are trying to open a fast path to the private transport adapters.

Reboot of the node might fix the problem.

674359 :load balancer deleted

This message indicates that the service group has been deleted.

This is an informational message, no user action is needed.

674415 :svc_restore_priority: %s

The rpc.pmfd or rpc.fed server was not able to run the application in the correct scheduling mode, and the system error is shown. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

674848 :fatal: Failed to read CCR

The rgmd is unable to read the cluster configuration repository. This is a fatal error. The rgmd will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

675776 :Stopped the fault monitor.

The fault monitor for this data service was stopped successfully.

No action needed.

676141 :in libsecurity could not copy host name

A client was not able to make an rpc connection to a server (rpc.pmfd, rpc.fed or rgmd) because the host name could not be saved, probably due to low memory. An error message is output to syslog.

Investigate if the host is low on memory. If not, save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

676379 :NAFO group %s has status %s. Assuming this node cannot respond to client requests.

The state of the NAFO group named is degraded.

Make sure all adapters and cables are working. Look in the /var/adm/messages file for message from the network monitoring daemon (pnmd).

676478 :kill -TERM: %s

The rpc.fed server is not able to kill a tag that timed out, and the error message is shown. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

676558 :WARNING: Global_resources_used property of resource group <%s> is set to non-null string, assuming wildcard

The Global_resources_used property of the resource group was set to a specific non-null string. The only supported settings of this property in the current release are null ("") or wildcard ("*").

No user action is required; the rgmd will interpret this value as wildcard. This means that method timeouts for this resource group will be suspended while any device group temporarily goes offline during a switchover or failover. This is usually the desired setting, except when a resource group has no dependency on any global device service or pxfs file system.

677278 :No network address resource in resource group.

A resource has no associated network address.

For a failover data service, add a network address resource to the resource group. For a scalable data service, add a network resource to the resource group referenced by the RG_dependencies property.

677785 :PNM: nafo%d: could not failover static routes in %s

A failover has happened for the named NAFO group, but the static routes commands contained in the named file cannot be successfully executed. Some static routes therefore could not be restored after an adapter failover.

Check that the named file has the right permission to be executed by the PNM daemon (pnmd), and that the file consists of valid route(1M) commands. Execute the file manually to restore the static routes. If no static routes need to be failed over explicitly, remove the named file.

678041 :lkcm_sync: cm_reconfigure failed: %s

ucmm reconfiguration failed.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

678319 :(%s) getenv of "%s" failed.

Failed to get the value of an environmental variable. udlm will fail to go through a transition.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

678755 :dl_bind: DL_BIND_ACK protocol error

Could not bind to the physical device. We are trying to open a fast path to the private transport adapters.

Reboot of the node might fix the problem.

679912 :uaddr2taddr: %s

Call to uaddr2taddr() failed. The "uaddr2taddr" man page describes possible error codes. udlm will exit and the node will abort and panic.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

680437 :Start method failed. Error: %d.

Restart of the data service failed.

Check the syslog messages that are occurred just before this message to check whether there is any internal error. In case of internal error, contact your Sun service provider. Otherwise, any of the following situations may have happened. 1) Check the Start_timeout value and adjust it if it is not appropriate. 2) Check whether the application's configuration is correct. 3) This might be the result of lack of the system resources. Check whether the system is low in memory or the process table is full and take appropriate action.

681547 :fatal: Method <%s> on resource <%s>: Received unexpected result <%d> from rpc.fed, aborting node

A serious error has occurred in the communication between rgmd and rpc.fed. The rgmd will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

683997 :Failed to retrieve the resource group property %s: %s

Unable to retrieve the resource group property.

For the property name and the reason for failure, check the syslog message. For more details about the api failure, check the syslog messages from the RGM.

684753 :store_binding: <%s> bad bind type <%d>

During a name server binding store an unknown binding type was encountered.

No action required. This is informational message.

684895 :Failed to validate scalable service configuration: Error %d.

An error was detected in the Load_balancing_weights property for the data service.

Use the scrgadm command to change the Load_balancing_weights property to a valid value.

685886 :Failed to communicate: %s.

While determining the health of the data service, fault monitor is failed to communicate with the process monitor facility.

This is internal error. Save /var/adm/messages file and contact your authorized Sun service provider. For more details about error, check the syslog messages.

686220 :Node attempted to join with invalid version message

Initial handshake message from a cluster node did not have a valid format.

Check if all cluster nodes are running the same version of the clustering software.

687543 :shutdown abort did not succeed.

HA-Oracle failed to shutdown Oracle server using "shutdown abort".

Examine log files and syslog messages to determine the cause of failure.

687929 :daemon %s did not respond to null rpc call: %s.

HA-NFS fault monitor failed to ping an nfs daemon.

No action required. The fault monitor will restart the daemon if necessary.

688163 :clexecd: pipe returned %d. Exiting.

clexecd program has encountered a failed pipe(2) system call. The error message indicates the error number for the failure.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

689538 :Listener %s did not stop.(%s)

Failed to start Oracle listener using "lsnrctl" command. HA-Oracle will attempt to kill listener process.

None

689887 :Failed to stop the process with: %s. Retry with SIGKILL.

Process monitor facility is failed to stop the data service. It is reattempting to stop the data service.

This is informational message. Check the Stop_timeout and adjust it, if it is not appropriate value.

689989 :Invalid device group name <%s> supplied

The diskgroup name defined in SUNW.HAStorage type resource is invalid

Check and set the correct diskgroup name in extension property "ServicePaths" of SUNW.HAStorage type resource.

690463 :Cannot bring server online on this node.

Oracle server is running but it cannot be brought online on this node. START method for the resource has failed.

Check if Oracle server can be started manually. Examine the log files and setup. Clear START_FAILED flag on the resource and bring the resource online.

691736 :CMM: Quorum device %ld (%s) with vote count = %d removed.

The specified quorum device with the specified vote count has been removed from the cluster. A quorum device being placed in maintenance state is equivalent to it being removed from the quorum subsystem's perspective, so this message will be logged when a quorum device is put in maintenance state as well as when it is actually removed.

This is an informational message, no user action is needed.

696186 :This list element in System property %s has an invalid port number: %s.

The system property that was named does not have a valid port number.

Change the value of the property to use a valid port number.

697026 :did instance %d created.

Informational message from scdidadm.

No user action required.

697108 :t_sndudata in send_reply failed.

Call to t_sndudata() failed. The "t_sndudata" man page describes possible error codes. ucmmmd will exit and the node will abort.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

697588 :Nodeid must be less than %d. Nodeid passed: "%s"

Incorrect nodeid passed to Oracle unix dlm. Oracle unix dlm will not start.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

698512 :Directory %s is not readable: %s.

The specified path doesn't exist or is not readable

Consult the HA-NFS configuration guide on how to configure the dfstab_name> file for HA-NFS resources.

698526 :scvxvmlg error - service %s has service_class %s, not %s, ignoring it

The program responsible for maintaining the VxVM namespace has suffered an internal error. If configuration changes were recently made to VxVM diskgroups or volumes, this node may be unaware of those changes. Recently created volumes may be inaccessible from this node.

If no configuration changes have been recently made to VxVM diskgroups or volumes and all volumes continue to be accessible from this node, then no action is required. If changes have been made, the device namespace on this node can be updated to reflect those changes by executing "/usr/cluster/lib/dcs/scvxvmlg". If the problem persists, contact your authorized Sun service provider to determine whether a workaround or patch is available.

698744 :scvxvmlg error - lstat(%s) failed with errno %d

The program responsible for maintaining the VxVM namespace was unable to access the global device namespace. If configuration changes were recently made to VxVM diskgroups or volumes, this node may be unaware of those changes. Recently created volumes may be inaccessible from this node.

Verify that the /global/.devices/node@N (N = this node's node number) is mounted globally and is accessible. If no configuration changes have been recently made to VxVM diskgroups or volumes and all volumes continue to be accessible from this node, then no further action is required. If changes have been made, the device namespace on this node can be updated to reflect those changes by executing "/usr/cluster/lib/dcs/scvxvmlg". If the problem persists, contact your authorized Sun service provider to determine whether a workaround or patch is available.

699689 :(%s) poll failed: %s (UNIX errno %d)

Call to poll() failed. The "poll" man page describes possible error codes. udlmctl will exit.

Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

Message IDs 700000 - 799999

Error Message List

The following list is ordered by the message ID.

700161 :Fault monitor is already running.

The resource's fault monitor is already running.

This is an internal error. Save the /var/adm/messages file from all the nodes.
Contact your authorized Sun service provider.

700321 :exec() of %s failed: %m.

The exec() system call failed for the given reason.

Verify that the pathname given is valid.

702368 :Failed to register callback for NAF0 group %s with tag
%s and callback command %s (request failed with %d).

An unexpected error occurred while trying to communicate with the network
monitoring daemon (pnmd).

Make sure the network monitoring daemon (pnmd) is running.

703476 :clcomm: unable to create desired unref threads

The system was unable to create thread that deal with no longer needed objects.
The system fails to create threads when memory is not available. This message
can be generated by the inability of either the kernel or a user level process. The

kernel creates unref threads when the cluster starts. A user level process creates threads when it initializes.

Take steps to increase memory availability. The installation of more memory will avoid the problem with a kernel inability to create threads. For a user level process problem: install more memory, increase swap space, or reduce the peak work load.

703553 :Resource group name or resource name is too long.

Process monitor facility is failed to execute the command. Resource group name or resource name is too long for the process monitor facility command.

Check the resource group name and resource name. Give short name for resource group or resource .

704795 :in libsecurity could not negotiate uid on any transport in NETPATH

A server (rpc.pmfd, rpc.fed or rgmd) was not able to start because it could not establish a rpc connection for the network specified, because it couldn't find any transport. This happened because either there are no available transports at all, or there are but none is a loopback. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

705163 :load balancer thread failed to start for %s

The system has run out of resources that is required to create a thread. The system could not create the load balancer thread.

The service group is created with the default load balancing policy. If rebalancing is required, free up resources by shutting down some processes. Then delete the service group and re-create it.

705379 :resource %s state on node %s change to R_ONLINE_UNMON

This is a notification from the rgmd that a resource has changed state to online-not-monitored on the given node. This may be used by system monitoring tools.

This is an informational message, no user action is needed.

705629 :clutil: Can't allocate hash table

The system attempted unsuccessfully to allocate a hash table. There was insufficient memory.

Install more memory, increase swap space, or reduce peak memory consumption.

706314 :clexecd: Error %d from open(/dev/zero). Exiting.

clxecd program has encountered a failed open(2) system call. The error message indicates the error number for the failure.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

707301 :lkcm_reg: Unix DLM version (???) and the OSD library version (%d) are not compatible. Unix DLM versions acceptable to this library are: %d

UNIX DLM and Oracle DLM are not compatible. Compatible versions will be printed as part of this message.

Check installation procedure to make sure you have the correct versions of Oracle DLM and Unix DLM. Contact Sun service representative if versions cannot be resolved.

707421 :%s: Cannot create a thread.

Solaris has run out of its limit on threads. Either too many clients are requesting a service, causing many threads to be created at once or system is overloaded with processes.

Reduce system load by reducing number of requestors of this service or halting other processes on the system.

707881 :clcomm: thread_create failed for autom_thread

The system could not create the needed thread, because there is inadequate memory.

There are two possible solutions. Install more memory. Alternatively, reduce memory usage. Since this happens during system startup, application memory usage is normally not a factor.

707948 :launching method <%s> for resource <%s>, resource group <%s>, timeout <%d> seconds

RGM has invoked a callback method for the named resource, as a result of a cluster reconfiguration, scha_control GIVEOVER, or scswitch.

This is an informational message, no user action is needed.

710143 :Failed to add node %d to scalable service group %s: %s.

A call to the underlying scalable networking code failed.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

711956 :open /dev/ip failed: %s.

System was attempting to open the specified device, but was unable to do so.

This might be the result of lack of the system resources. Check whether the system is low in memory and take appropriate action. For specific error information check the syslog message.

712367 :clcomm: Endpoint %p: deferred task not allowed in state %d

The system maintains information about the state of an Endpoint. A deferred task is not allowed in this state.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

714838 :reservation fatal error(%s) - Unable to open name file "%s", errno %d

The device fencing program has suffered an internal error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available. Copies of /var/adm/messages from all nodes should be provided for diagnosis. It may be possible to retry the failed operation, depending on the nature of the error. If the message specifies the "node_join" transition, then this node may be unable to access shared devices. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access shared devices. In either case, it may be possible to reacquire access to shared devices by executing "/usr/cluster/lib/sc/run_reserve -c node_join" on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group has failed to start on this node. If another node was available to host the device group, then it should have been started on that node. If desired, it may be possible to switch the device group to this node with the scswitch command. If no other node was available, then the device group will not have been started. The scswitch command may be used to retry the attempt to start the device group. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group has failed. The desired action may be retried.

715958 :Method <%s> on resource <%s> stopped due to receipt of signal <%d>

A resource method was stopped by a signal, most likely resulting from an operator-issued kill(1). The method is considered to have failed.

The operator must kill the stopped method. The operator may then choose to issue an scswitch(1M) command to bring resource groups onto desired primaries, or re-try the administrative action that was interrupted by the method failure.

716253 :launch_fed_prog: fe_set_env_vars() failed for program <%s>, step <%s>

The ucmm server was not able to get the locale environment. An error message is output to syslog.

Investigate if the host is running out of memory. If not save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

717089 :NAFO group %s has unknown status %d. Skipping this NAFO group.

The status of the NAFO group is not among the set of statuses that is known.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

719114 :Failed to parse key/value pair from command line for %s.

The validate method for the scalable resource network configuration code was unable to convert the property information given to a usable format.

Verify the property information was properly set when configuring the resource.

719497 :clcomm: path_manager using RT lwp rather than clock interrupt

The system has been built to use a real time thread to support path_manager heart beats instead of the clock interrupt.

No user action is required.

719735 :resource %s state on node %s change to R_PENDING_BOOT

This is a notification from the rgmd that the BOOT method is running on the given resource, on a node that has recently joined the cluster. This may be used by system monitoring tools.

This is an informational message, no user action is needed.

721252 :cm2udlm: cm_getclustmbname: %s

Could not create a structure for communication with the cluster monitor process.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

721881 :dl_attach: kstr_msg failed %d error

Could not attach to the private interconnect.

Reboot of the node might fix the problem.

722270 :fatal: cannot create state machine thread

The rgmd was unable to create a thread upon starting up. This is a fatal error. The rgmd will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

Make sure that the hardware configuration meets documented minimum requirements. Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

722904 :Failed to open the resource group handle: %s.

An API operation has failed while retrieving the resource group property. Low memory or API call failure might be the reasons.

In case of low memory, the problem will probably be cured by rebooting. If the problem reoccurs, you might need to increase swap space by configuring additional swap devices. Otherwise, if it is API call failure, check the syslog messages from other components. For resource group name and the property name, check the current syslog message.

722984 :call to rpc.fed failed for resource <%s>, resource group <%s>, method <%s>

The rgmd failed in an attempt to execute a method, due to a failure to communicate with the rpc.fed daemon. Depending on which method was being invoked and the Failover_mode setting on the resource, this might cause the resource group to fail over or move to an error state. If the rpc.fed process died, this might lead to a subsequent reboot of the node.

Examine other syslog messages occurring at about the same time to see if the problem can be identified. Save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance in diagnosing the problem.

724035 :Failed to connect to %s secure port %d.

An error occurred while the fault monitor was trying to connect to a secure port specified in the Port_list property for this resource.

Check to make sure that the Port_list property is correctly set to the same port number that the Netscape Directory Server is running on.

726417 :read %d for %sport

Could not get the port information from config file udln.conf.

Check to make sure udln.conf file exists and has entry for udln.port. If everything looks normal and the problem persists, contact your Sun service representative.

727065 :CMM: Enabling failfast on quorum device %s failed with error %d.

An attempt to enable failfast on the specified quorum device failed with the specified error.

Check if the specified quorum disk has failed. This message may also be logged when a node is booting up and has been preempted from the cluster, in which case no user action is necessary.

727160 :msg of wrong version %d, expected %d

udlmctl received an illegal message.

None. udlm will handle this error.

728216 :reservation error(%s) - did_get_path() error

The device fencing program has suffered an internal error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available. Copies of /var/adm/messages from all nodes should be provided for diagnosis. It may be possible to retry the failed operation, depending on the nature of the error. If the message specifies the "node_join" transition, then this node may be unable to access shared devices. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access shared devices. In either case, it may be possible to reacquire access to shared devices by executing "/usr/cluster/lib/sc/run_reserve -c node_join" on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group has failed to start on this node. If another node was available to host the device group, then it should have been started on that node. If desired, it may be possible to switch the device group to this node with the scswitch command. If no other node was available, then the device group will not have been started. The scswitch command may be used to retry the attempt to start the device group. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group has failed. The desired action may be retried.

728425 :INTERNAL ERROR: bad state <%s> (%d) for resource group <%s> in rebalance()

An internal error has occurred in the rgmd. This may prevent the rgmd from bringing the affected resource group online.

Look for other syslog error messages on the same node. Save a copy of the /var/adm/messages files on all nodes, and report the problem to your authorized Sun service provider.

728928 :CCR: Can't access table %s on node %s errno = %d.

The indicated error occurred when CCR was tried to access the indicated table on the nodes in the cluster. The errno value indicates the nature of the problem. errno values are defined in the file /usr/include/sys/errno.h. An errno value of

28(ENOSPC) indicates that the root files system on the node is full. Other values of errno can be returned when the root disk has failed(EIO).

There may be other related messages on the node where the failure occurred. They may help diagnose the problem. If the root file system is full on the node, then free up some space by removing unnecessary files. If the root disk on the afflicted node has failed, then it needs to be replaced. If the indicated table was accidentally removed, boot the indicated node in -x mode to restore the indicated table from backup. The CCR tables are located at /etc/cluster/ccr/.

729152 :clxecd: Error %d from F_SETFD. Exiting.

clxecd program has encountered a failed fcntl(2) system call. The error message indicates the error number for the failure.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

730190 :scvxvmlg error - found non device-node or non link %, directory not removed

The program responsible for maintaining the VxVM namespace had detected suspicious entries in the global device namespace.

The global device namespace should only contain diskgroup directories and volume device nodes for registered diskgroups. The specified path was not recognized as either of these and should be removed from the global device namespace.

730685 :PCSTATUS: %s

The rpc.pmfd server was not able to monitor a process, and the system error is shown. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

730782 :Failed to update scalable service: Error %d.

Update to a property related to scalability was not successfully applied to the system.

Use scswitch to try to bring resource offline and online again on this node. If the error persists, reboot the node and contact your Sun service representative.

730956 :%d entries found in property %s. For a nonsecure Netscape Directory Server instance %s should have exactly one entry.

Since a nonsecure Netscape Directory Server instance only listens on a single port, the list property should only have a single entry. A different number of entries was found.

Change the number of entries to be exactly one.

732069 :dl_attach: DL_ERROR_ACK protocol error

Could not attach to the physical device.

Check the documentation for the driver associated with the private interconnect. It might be that the message returned is too small to be valid.

732822 :clconf: Invalid group name

An invalid group name has been encountered while converting a group name to clconf_obj type. Valid group names are "cluster", "nodes", "adapters", "ports", "blackboxes", "cables", and "quorum_devices".

This is an unrecoverable error, and the cluster needs to be rebooted. Also contact your authorized Sun service provider to determine whether a workaround or patch is available.

733367 :lkcm_act: %s: %s cm_reconfigure failed

ucmm reconfiguration failed.

None if the next reconfiguration succeeds. If not, save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

734057 :clcomm: Duplicate TypeId"s: %s : %s

The system records type identifiers for multiple kinds of type data. The system checks for type identifiers when loading type information. This message identifies two items having the same type identifiers. This checking only occurs on debug systems.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

734832 :clutil: Created insufficient threads in threadpool

There was insufficient memory to create the desired number of threads.

Install more memory, increase swap space, or reduce peak memory consumption.

734890 :pthread_detach: %s

The rpc.pmfd server was not able to detach a thread, possibly due to low memory. The message contains the system error. The server does not perform the action requested by the client, and an error message is output to syslog.

Investigate if the machine is running out of memory. If all looks correct, save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

737104 :Received unexpected result <%d> from rpc.fed, aborting node

This node encountered an unexpected error while communicating with other cluster nodes during a cluster reconfiguration. The ucmmmd will produce a core file and will cause the node to halt or reboot.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

737125 :INTERNAL ERROR: PENDING_OFF_STOP_FAILED or ERROR_STOP_FAILED in rebalance()

An internal error has occurred in the rgmd. This may prevent the rgmd from bringing affected resource groups online.

Look for other syslog error messages on the same node. Save a copy of the /var/adm/messages files on all nodes, and report the problem to your authorized Sun service provider.

737445 :rebalance: not attempting to start resource group <%s> on node <%s> because this resource group has already failed to start on this node %d or more times in the past %d seconds

The rgmd is preventing "ping-pong" failover of the resource group, i.e., repeated failover of the resource group between two or more nodes.

The time interval in seconds that is mentioned in the message can be adjusted by using scrgadm(1M) to set the Pingpong_interval property of the resource group.

738465 :Malformed adapter specification %s.

Failed to retrieve the nafo information. The given adapter specification is invalid.

Check whether the adapter specification is in the form of nafogroup@nodename. If not, recreate the resource with the properly formatted adapter information.

738847 :clxeccd: unable to create failfast object.

clxeccd problem could not enable one of the mechanisms which causes the node to be shutdown to prevent data corruption, when clxeccd program dies.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

739356 :warning: cannot store start_failed timestamp for resource group <%s>: time() failed, errno <%d> (%s)

The specified resource group failed to come online on some node, but this node is unable to record that fact due to the failure of the time(2) system call. The consequence of this is that the resource group may continue to pingpong between nodes for longer than the Pingpong_interval property setting.

Examine other syslog messages occurring around the same time on the same node, to see if the cause of the problem can be identified. If the same error recurs, you might have to reboot the affected node.

739653 :Port number %d is listed twice in property %s, at entries %d and %d.

The port number in the message was listed twice in the named property, at the list entry locations given in the message. A port number should only appear once in the property.

Specify the property with only one occurrence of the port number.

740373 :Failed to get the scalable service related properties for resource %s.

An unexpected error occurred while trying to collect the properties related to scalable networking for the named resource.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

740972 :in fe_set_env_vars setlocale failed

The rgmd server was not able to get the locale environment, while trying to connect to the rpc.fed server. An error message is output to syslog.

Investigate if the host is running out of memory. If not save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

742337 :Node %d is in the %s for resource %s, but property %s identifies resource %s which cannot host an address on node %d.

All IP addresses used by this resource must be configured to be available on all nodes that the scalable resource can run on.

Either change the resource group nodelist to exclude the nodes that cannot host the SharedAddress IP address, or select a different network resource whose IP address will be available on all nodes where this scalable resource can run.

743362 :could not read failfast mode, using panic

/opt/SUNWudlm/etc/udlm.conf did not have an entry for failfast mode. Default mode of "panic" will be used.

None.

745100 :CMM: Quorum device %ld has been changed from %s to %s.

The name of the specified quorum device has been changed as indicated. This can happen if while this node was down, the previous quorum device was

removed from the cluster, and a new one was added (and assigned the same id as the old one) to the cluster.

This is an informational message, no user action is needed.

747567 :Unable to complete any share commands.

None of the paths specified in the dfstab. file were shared successfully.

The prenet_start method would fail. Sun Cluster resource management would attempt to bring the resource on-line on some other node. Manually check that the paths specified in the dfstab. file are correct.

749409 :clcomm: validate_policy: high not enough. high %d low %d in c %d nodes %d pool %d

The system checks the proposed flow control policy parameters at system startup and when processing a change request. For a variable size resource pool, the high server thread level must be large enough to allow all of the nodes identified in the message join the cluster and receive a minimal number of server threads.

No user action required.

749958 :CMM: Unable to create %s thread.

The CMM was unable to create its specified thread and the system can not continue. This is caused by inadequate memory on the system.

Add more memory to the system. If that does not resolve the problem, contact your authorized Sun service provider to determine whether a workaround or patch is available.

751934 :scswitch: rgm_change_mastery() failed with NOREF, UNKNOWN, or invalid error on node %d

An inter-node communication failed with an unknown exception while the rgmd was attempting to execute an operator-requested switch of the primaries of a resource group, or was attempting to "fail back" a resource group onto a node that just rejoined the cluster. This will cause the attempted switching action to fail.

Examine other syslog messages occurring around the same time on the same node, to see if the cause of the problem can be identified. If the switch was operator-requested, retry it. If the same error recurs, you might have to reboot the affected node. Since this problem might indicate an internal logic error in the clustering software, please save a copy of the /var/adm/messages files on all nodes, the output of an scstat -g command, and the output of a scrgadm -pvv command. Report the problem to your authorized Sun service provider.

753898 :resource group %s state on node %s change to RG_OFF_PENDING_BOOT

This is a notification from the rgmd that BOOT methods are running on resources in the given resource group, on a node that has recently joined the cluster. This may be used by system monitoring tools.

This is an informational message, no user action is needed.

754521 :Property %s does not have a value. This property must have exactly one value.

The property named does not have a value specified for it.

Set the property to have exactly one value.

754848 :The property %s must contain at least one SharedAddress network resource.

The named property must contain at least one SharedAddress.

Specify a SharedAddress resource for this property.

756369 :NAFO group %s has failed, so scalable resource %s in resource group %s may not be able to respond to client requests. A request will be issued to relocate resource %s off of this node.

The named NAFO group has failed, so the node may not be able to respond to client requests. It would be desirable to move the resource to another node that has functioning NAFO groups. A request will be issued on behalf of this resource to relocate the resource to another node.

Check the status of the NAFO group on the node. Try to fix the adapters in the NAFO group.

756650 :Failed to set the global interface node to %d for IP %s: %s.

A call to the underlying scalable networking code failed.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

757581 :Failed to stop daemon %s.

The HA-NFS implementation was unable to stop the specified daemon.

The resource could be in a STOP_FAILED state. If the failover mode is set to HARD, the node would get automatically rebooted by the SunCluster resource management. If the Failover_mode is set to SOFT or NONE, please check that the specified daemon is indeed stopped (by killing it by hand, if necessary). Then clear the STOP_FAILED status on the resource and bring it on-line again using the scswitch command.

760001 :(%s) netconf error: cannot get transport info for "ticlts" %s

Call to getnetconfig failed and udlmctl could not get network information. udlmctl will exit.

Make sure the internconnect does not have any problems. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm/*logs/* from all the nodes and contact your Sun service representative.

761076 :dl_bind: DL_ERROR_ACK protocol error

Could not bind to the physical device. We are trying to open a fast path to the private transport adapters.

Reboot of the node might fix the problem.

762902 :Failed to restart fault monitor.

The resource property that was updated needed the fault monitor to be restarted in order for the change to take effect, but the attempt to restart the fault monitor failed.

Look at the prior syslog messages for specific problems. Correct the errors if possible. Look for the process _probe operating on the desired resource (indicated by the argument to "-R" option). This can be found from the command: ps -ef | egrep _probe | grep "\-R " Send a kill signal to this process. If the process does not get killed and restarted by the process monitor facility, reboot the node.

763929 :HA: rm_service_thread_create failed

The system could not create the needed thread, because there is inadequate memory.

There are two possible solutions. Install more memory. Alternatively, reduce memory usage.

765395 :clcomm: RT class not configured in this system

Sun Cluster requires that the real time thread scheduling class be configured in the kernel.

Configure Solaris with the RT thread scheduling class in the kernel.

766093 :IP address (hostname) and Port pairs %s%c%d%c%s and %s%c%d%c%s in property %s, at entries %d and %d, effectively duplicate each other. The port numbers are the same and the resolved IP addresses are the same.

The two list entries at the named locations in the named property have port numbers that are identical, and also have IP address (hostname) strings that

resolve to the same underlying IP address. An IP address (hostname) string and port entry should only appear once in the property.

Specify the property with only one occurrence of the IP address (hostname) string and port entry.

767363 :CMM: Disconnected from node %ld; aborting using %s rule.

Due to a connection failure between the local and the specified node, the local node must be halted to avoid a "split brain" configuration. The CMM used the specified rule to decide which node to fail. Rules are: rebootee: If one node is rebooting and the other was a member of the cluster, the node that is rebooting must abort. quorum: The node with greater control of quorum device votes survives and the other node aborts. node number: The node with higher node number aborts.

The cause of the failure should be resolved and the node should be rebooted if node failure is unexpected.

767488 :reservation fatal error(UNKNOWN) - Command not specified

The device fencing program has suffered an internal error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available. Copies of /var/adm/messages from all nodes should be provided for diagnosis. It may be possible to retry the failed operation, depending on the nature of the error. If the message specifies the "node_join" transition, then this node may be unable to access shared devices. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access shared devices. In either case, it may be possible to reacquire access to shared devices by executing "/usr/cluster/lib/sc/run_reserve -c node_join" on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group has failed to start on this node. If another node was available to host the device group, then it should have been started on that node. If desired, it may be possible to switch the device group to this node with the scswitch command. If no other node was available, then the device group will not have been started. The scswitch command may be used to retry the attempt to start the device group. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group has failed. The desired action may be retried.

767629 :lkcm_reg: Unix DLM version (%d) and the OSD library version (%d) are not compatible. Unix DLM versions acceptable to this library are: %d

Unix DLM and Oracle DLM are not compatible. Compatible versions will be printed as part of this message.

Check installation procedure to make sure you have the correct versions of Oracle DLM and Unix DLM. Contact Sun service representative if versions cannot be resolved.

767858 :in libsecurity unknown security type %d

This is an internal error which shouldn't occur. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

770355 :fatal: received signal %d

The daemon indicated in the message tag (rgmd or ucmmmd) has received a SIGTERM signal, possibly caused by an operator-initiated kill(1) command. The daemon will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

The operator must use scswitch(1M) and shutdown(1M) to take down a node, rather than directly killing the daemon.

770675 :monitor_check: fe_method_full_name() failed for resource <%s>, resource group <%s>

During execution of a scha_control(1HA,3HA) function, the rgmd was unable to assemble the full method pathname for the MONITOR_CHECK method. This is considered a MONITOR_CHECK method failure. This in turn will prevent the attempted failover of the resource group from its current master to a new master.

Examine other syslog messages occurring at about the same time to see if the problem can be identified. Save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance in diagnosing and correcting the problem.

770776 :INTERNAL ERROR: process_resource: Resource <%s> is R_BOOTING in PENDING_ONLINE resource group

The rgmd is attempting to bring a resource group online on a node where BOOT methods are still being run on its resources. This should not occur and may indicate an internal logic error in the rgmd.

Look for other syslog error messages on the same node. Save a copy of the /var/adm/messages files on all nodes, and report the problem to your authorized Sun service provider.

771340 :fatal: Resource group <%s> update failed with error <%d>; aborting node

Rgmd failed to read updated resource group from the CCR on this node.

Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

772294 :%s requests reconfiguration in step %s

Return status at the end of a step execution indicates that a reconfiguration is required.

None.

772395 :shutdown immediate did not succeed. (%s)

Failed to shutdown Oracle server using "shutdown immediate" command.

Examine "Stop_timeout" property of the resource and increase "Stop_timeout" if Oracle server takes long time to shutdown. and if you don't wish to use "shutdown abort" for stopping Oracle server.

773078 :Error in configuration file lookup (%s, ...): %s

Could not read configuration file udlm.conf.

Make sure udlm.conf exists under /opt/SUNWudlm/etc and has the correct permissions.

773366 :thread create for hb_threadpool failed

The system was unable to create thread used for heartbeat processing.

Take steps to increase memory availability. The installation of more memory will avoid the problem with a kernel inability to create threads. For a user level process problem: install more memory, increase swap space, or reduce the peak work load.

774752 :reservation error(%s) - do_scsi3_inresv() error for disk %s

The device fencing program has encountered errors while trying to access a device. All retry attempts have failed.

For the user action required by this message, see the user action for message 192619.

776199 :(%s) reconfigure: cm error %s

ucmm reconfiguration failed.

None if the next reconfiguration succeeds. If not, save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

776339 :INTERNAL ERROR: postpone_stop_r: meth type <%d>

A non-fatal internal error has occurred in the rgmd state machine.

Since this problem might indicate an internal logic error in the rgmd, please save a copy of the /var/adm/messages files on all nodes, the output of an scstat -g command, and the output of a scrgadm -pvv command. Report the problem to your authorized Sun service provider.

778629 :ERROR: MONITOR_STOP method is not registered for ONLINE resource <%s>

A non-fatal internal error has occurred in the rgmd state machine.

Since this problem might indicate an internal logic error in the rgmd, please save a copy of the /var/adm/messages files on all nodes, the output of an scstat -g command, and the output of a scrgadm -pvv command. Report the problem to your authorized Sun service provider.

779073 :in fe_set_env_vars malloc of env_name[%d] failed

The rgmd server was not able to allocate memory for an environment variable, while trying to connect to the rpc.fed server, possibly due to low memory. An error message is output to syslog.

Investigate if the host is running out of memory. If not save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

779089 :Could not start up DCS client because we could not contact the name server.

There was a fatal error while this node was booting.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

780283 :clcomm: Exception in coalescing region - Lost data

While supporting an invocation, the system wanted to combine buffers and failed. The system identifies the exception prior to this message.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

780792 :Failed to retrieve the resource type information.

A Sun cluster dataservice has failed to retrieve the resource type's property information. Low memory or API call failure might be the reasons.

In case of low memory, the problem will probably be cured by rebooting. If the problem reoccurs, you might need to increase swap space by configuring

additional swap devices. Otherwise, if it is API call failure, check the syslog messages from other components.

781731 :Failed to retrieve the cluster handle: %s.

An API operation has failed while retrieving the cluster information.

This may be solved by rebooting the node. For more details about API failure, check the messages from other components.

782111 :This list element in System property %s is missing a protocol: %s.

The system property that was named does not have a valid format. The value of the property must include a protocol.

Add a protocol to the property value.

782694 :The value returned for property %s for resource %s was invalid.

An unexpected value was returned for the named property.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

783130 :Failed to retrieve the node id for node %s: %s.

API operation has failed while retrieving the node id for the given node.

Check whether the node name is valid. For more information about API call failure, check the messages from other components.

783581 :scvxvmlg fatal error - clconf_lib_init failed, returned %d

The program responsible for maintaining the VxVM namespace has suffered an internal error. If configuration changes were recently made to VxVM diskgroups or volumes, this node may be unaware of those changes. Recently created volumes may be inaccessible from this node.

If no configuration changes have been recently made to VxVM diskgroups or volumes and all volumes continue to be accessible from this node, then no action is required. If changes have been made, the device namespace on this node can be updated to reflect those changes by executing "/usr/cluster/lib/dcs/scvxvmlg". If the problem persists, contact your authorized Sun service provider to determine whether a workaround or patch is available.

784560 :resource %s status on node %s change to %s

This is a notification from the rgmd that a resource's fault monitor status has changed.

This is an informational message, no user action is needed.

785101 :transition "%s" failed for cluster "%s": unknown code %d

The mentioned state transition failed for the cluster because of an unexpected command line option. udlmctl will exit.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

785154 :Could not look up IP because IP was NULL.

The mapping for the given ip address in the local host files can't be done: the specified ip address is NULL.

Check whether the ip address has NULL value. If this is the case, recreate the resource with valid host name. If this is not the reason, treat it as an internal error and contact Sun service provider.

786412 :reservation fatal error(UNKNOWN) - clconf_lib_init()
error, returned %d

The device fencing program has suffered an internal error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available. Copies of /var/adm/messages from all nodes should be provided for diagnosis. It may be possible to retry the failed operation, depending on the nature of the error. If the message specifies the "node_join" transition, then this node may be unable to access shared devices. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access shared devices. In either case, it may be possible to reacquire access to shared devices by executing "/usr/cluster/lib/sc/run_reserve -c node_join" on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group has failed to start on this node. If another node was available to host the device group, then it should have been started on that node. If desired, it may be possible to switch the device group to this node with the scswitch command. If no other node was available, then the device group will not have been started. The scswitch command may be used to retry the attempt to start the device group. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group has failed. The desired action may be retried.

786765 :Failed to get host names from resource %s.

The networking information for the resource could not be retrieved.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

787063 :Error in getting parameters for global service <%s> of path <%s>: %s

Can not get information of global service.

Save a copy of /var/adm/messages and contact your authorized Sun service provider to determine what is the cause of the problem.

787529 :NAFO group %s has status %s. The status of the NAFO group will be checked again in %d seconds.

The NAFO group named is in a transition state. The status will be checked again.

This is an informational message, no user action is needed.

788145 :gethostbyname() failed: %s.

gethostbyname() failed with unexpected error.

Check if name servcie is configured correctly. Try some commands to query name serves, such as ping and nslookup, and correct the problem. If the error still persists, then reboot the node.

789223 :lkcm_sync: caller is not registered

udlm is not registered with ucmm.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

789460 :monitor_check: call to rpc.fed failed for resource <%s>, resource group <%s>, method <%s>

A scha_control(1HA,3HA) GIVEOVER attempt failed, due to a failure of the rgmd to communicate with the rpc.fed daemon. If the rpc.fed process died, this might lead to a subsequent reboot of the node. Otherwise, this will prevent a resource group on the local node from failing over to an alternate primary node

Examine other syslog messages occurring at about the same time to see if the problem can be identified and if it recurs. Save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance.

791495 :Unregistered syscall (%d)

An internal error has ocured. This should not happen. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

792109 :Unable to set number of file descriptors.

rpc.pmfd was unable to set the number of file descriptors used in the RPC server.

Look for other syslog error messages on the same node. Save a copy of the /var/adm/messages files on all nodes, and report the problem to your authorized Sun service provider.

792338 :The property %s must contain at least one value.

The named property does not have a legal value.

Assign the property a value.

792967 :Unable to parse configuration file.

While parsing the Netscape configuration file an error occurred in while either reading the file, or one of the fields within the file.

Make sure that the appropriate configuration file is located in its default location with respect to the Confdir_list property.

794220 :switchback: bad nodename <%s>

The rgmd encountered a bad node name in the Nodelist of a resource group it was trying to failback. This might indicate corruption of CCR data or rgmd in-memory state.

Use scrgadm(1M) -pvv to examine resource group properties. If the values appear corrupted, the CCR might have to be rebuilt. If values appear correct, this may indicate an internal error in the rgmd. Please contact your authorized Sun service provider for assistance in diagnosing and correcting the problem.

794535 :clcomm: Marshal Type mismatch. Expecting type %d got type %d

When MARSHAL_DEBUG is enabled, the system tags every data item marshalled to support an invocation. This reports that the current data item in the received message does not have the expected type. The received message format is wrong.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

795047 :Stop fault monitor using pmfadm failed. tag %s error=%d

Failed to stop fault monitor will be stopped using Process Monitoring Facility (PMF), with the tag indicated in message. Error returned by PMF is indicated in message.

Stop fault monitor processes. Please report this problem.

795062 :Stop fault monitor using pmfadm failed. tag %s error=%s

Failed to stop fault monitor will be stopped using Process Monitoring Facility (PMF), with the tag indicated in message. Error returned by PMF is indicated in message.

Stop fault monitor processes. Please report this problem.

795381 :t_open: %s

Call to t_open() failed. The "t_open" man page describes possible error codes. udlm exits and the node will abort and panic.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

796491 :Error in checking file system mount point for path <%s>

The file system mount point can not be mapped to global service correctly.

Check the definition of mountpoint path in extension property "ServicePaths" of SUNW.HAStorage type resource and make sure they are for global file system.

796536 :Password file %s is not readable: %s

For the secure server to run, a password file named keypass is required. This file could not be read, which resulted in an error when trying to start the Data Service.

Create the keypass file and place it under the Confdir_list path for this resource. Make sure that the file is readable.

796771 :check_for_ccrdata failed malloc of size %d

Call to malloc failed. The "malloc" man page describes possible reasons.

Install more memory, increase swap space or reduce peak memory consumption.

797604 :CMM: Connectivity of quorum device %ld (%s) has been changed from 0x%llx to 0x%llx.

The number of configured paths to the specified quorum device has been changed as indicated. The connectivity information is depicted as bitmasks.

This is an informational message, no user action is needed.

798060 :Error opening procfs status file <%s> for tag <%s>: %s

The rpc.pmf server was not able to open a procfs status file, and the system error is shown. procfs status files are required in order to monitor user processes.

Investigate if the machine is running out of memory. If this is not the case, save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

798175 :sema_wait: %s

The rpc.pmfd server was not able to act on a semaphore. The message contains the system error. The server does not perform the action requested by the client, and an error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

798514 :Starting fault monitor. pmf tag %s

Informational message. Fault monitor is being started under control of Process Monitoring Facility (PMF), with the tag indicated in message.

None

798658 :Failed to get the resource type name: %s.

While retrieving the resource information, API operation has failed to retrieve the resource type name.

This is internal error. Contact your authorized Sun service provider. For more error description, check the syslog messages.

799348 :INTERNAL ERROR: MONITOR_START method is not registered for resource <%s>

A non-fatal internal error has occurred in the rgmd state machine.

Since this problem might indicate an internal logic error in the rgmd, please save a copy of the /var/adm/messages files on all nodes, the output of an scstat -g command, and the output of a scrgadm -pvv command. Report the problem to your authorized Sun service provider.

799426 :clcomm: can't ifkconfig private interface: %s:%d cmd %d error %d

The system failed to configure private network device for IP communications across the private interconnect of this device and IP address, resulting in the error identified in the message.

Ensure that the network interconnect device is supported. Otherwise, Contact your authorized Sun service provider to determine whether a workaround or patch is available.

Message IDs 800000 - 899999

Error Message List

The following list is ordered by the message ID.

802295 :monitor_check: resource group <%s> changed while running MONITOR_CHECK methods

An internal error has occurred in the locking logic of the rgmd, such that a resource group was erroneously allowed to be edited while a failover was pending on it, causing the scha_control call to return early with an error. This in turn will prevent the attempted failover of the resource group from its current master to a new master. This should not occur and may indicate an internal logic error in the rgmd.

Examine other syslog messages occurring at about the same time to see if the problem can be identified. Save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance in diagnosing and correcting the problem.

803391 :Could not validate the settings in %s. It is recommended that the settings for host lookup consult`files before a name server.

Validation callback method has failed to validate the hostname list. There may be syntax error in the nsswitch.conf file.

Check for the following syntax rules in the nsswitch.conf file. 1) Check if the lookup order for "hosts" has "files". 2) "cluster" is the only entry that can come before "files". 3) Everything in between "[" and "]" is ignored. 4) It is illegal to

have any leading whitespace character at the beginning of the line; these lines are skipped. Correct the syntax in the `nsswitch.conf` file and try again.

803649 :Failed to check whether the resource is a logical host resource.

While retrieving the IP addresses from the network resources in the resource group, the attempt to check whether the resource is a logical host resource or not has failed.

Internal error or API call failure might be the reasons. Check the error messages that occurred just before this message. If there is internal error, contact your authorized Sun service provider. For API call failure, check the syslog messages from other components. For the resource name and resource group name, check the syslog tag.

803719 :host %s failed, and `clnt_spcREATEerror` returned NULL

The rgm is not able to establish an rpc connection to the `rpc.fed` server on the host shown, and the rpc error could not be read. An error message is output to syslog.

Save the `/var/adm/messages` file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

804658 :clxecd: close returned %d while exec"ing (%s). Exiting.

clxecd program has encountered a failed `close(2)` system call. The error message indicates the error number for the failure.

The clxecd program will exit and the node will be halted or rebooted to prevent data corruption. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

804791 :A warm restart of rpcbind may be in progress.

The HA-NFS probe detected that the rpcbind daemon is not running, however it also detected that a warm restart of rpcbind is in progress.

If a warm restart is indeed in progress, ignore this message. Otherwise, check to see if the rpcbind daemon is running. If not, reboot the node. If the rpcbind process is not running, the HA-NFS probe would reboot the node itself if the `Failover_mode` on the resource is set to `HARD`.

804820 :clcomm: path_manager failed to create RT lwp (%d)

The system failed to create a real time thread to support path manager heart beats.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

805788 :reservation fatal error(%s) - service_name not specified

The device fencing program has suffered an internal error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available. Copies of /var/adm/messages from all nodes should be provided for diagnosis. It may be possible to retry the failed operation, depending on the nature of the error. If the message specifies the "node_join" transition, then this node may be unable to access shared devices. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access shared devices. In either case, it may be possible to reacquire access to shared devices by executing "/usr/cluster/lib/sc/run_reserve -c node_join" on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group has failed to start on this node. If another node was available to host the device group, then it should have been started on that node. If desired, it may be possible to switch the device group to this node with the scswitch command. If no other node was available, then the device group will not have been started. The scswitch command may be used to retry the attempt to start the device group. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group has failed. The desired action may be retried.

806365 :monitor_check: getlocalhostname() failed for resource
<%s>, resource group <%s>

While attempting to process a scha_control(1HA,3HA) call, the rgmd failed in an attempt to obtain the hostname of the local node. This is considered a MONITOR_CHECK method failure. This in turn will prevent the attempted failover of the resource group from its current master to a new master.

Examine other syslog messages occurring at about the same time to see if the problem can be identified. Save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance in diagnosing and correcting the problem.

806618 :Resource group name is null.

This is an internal error. While attempting to retrieve the resource information, null value was retrieved for the resource group name.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

806902 :clutil: Could not create lwp during respawn

There was insufficient memory to support this operation.

Install more memory, increase swap space, or reduce peak memory consumption.

807249 :CMM: Node %s (nodeid = %d) with votecount = %d removed.

The specified node with the specified votecount has been removed from the cluster.

This is an informational message, no user action is needed.

808746 :Node id %d is higher than the maximum node id of %d in the cluster.

In one of the scalable networking properties, a node id was encountered that was higher than expected.

Verify that the nodes listed in the scalable networking properties are still valid cluster members.

809322 :Couldn't create deleted directory: error (%d)

The file system is unable to create temporary copies of deleted files.

Mount the affected file system as a local file system, and ensure that there is no file system entry with name "._" at the root level of that file system. Alternatively, run fsck on the device to ensure that the file system is not corrupt.

809329 :No adapter for node %s.

No NAFO group has been specified for this node.

If this error message has occurred during resource creation, supply valid adapter information and retry it. If this message has occurred after resource creation, remove the LogicalHostname resource and recreate it with the correct NAFO group for each node which is a potential master of the resource group.

809858 :ERROR: method <%s> timeout for resource <%s> is not an integer

The indicated resource method timeout, as stored in the CCR, is not an integer value. This might indicate corruption of CCR data or rgmd in-memory state. The method invocation will fail; depending on which method was being invoked and the Failover_mode setting on the resource, this might cause the resource group to fail over or move to an error state.

Use scstat(1M) -g and scrgadm(1M) -pvv to examine resource properties. If the values appear corrupted, the CCR might have to be rebuilt. If values appear correct, this may indicate an internal error in the rgmd. Contact your authorized Sun service provider for assistance in diagnosing and correcting the problem.

809956 :PCSEXIT: %s

The rpc.pmfd server was not able to monitor a process, and the system error is shown. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

809985 :Elements in Confdir_list and Port_list must be 1-1 mapping

The Confdir_list and Port_list properties must contain the same number of entries, thus maintaining a 1-1 mapping between the two.

Using the appropriate scrgadm command, configure this resource to contain the same number of entries in the Confdir_list and the Port_list properties.

810551 :fatal: Unable to bind president to nameserver

The low-level cluster machinery has encountered a fatal error. The rgmd will produce a core file and will cause the node to halt or reboot to avoid the possibility of data corruption.

Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

811254 :VALIDATE failed on resource <%s>, resource group <%s>

The resource's VALIDATE method exited with a non-zero exit code. This indicates that an attempted update of a resource or resource group is invalid.

Examine syslog messages occurring just before this one to determine the cause of validation failure. Re-try the update.

811463 :match_online_key failed strdup for (%s)

Call to strdup failed. The "strdup" man page describes possible reasons.

Install more memory, increase swap space or reduce peak memory consumption.

812706 :dl_attach: DL_OK_ACK protocol error

Could not attach to the private interconnect interface.

Reboot of the node might fix the problem.

813831 :reservation warning(%s) - MHIOCSTATUS error will retry in %d seconds

The device fencing program has encountered errors while trying to access a device. The failed operation will be retried

This is an informational message, no user action is needed.

813866 :Property %s has no hostnames for resource %s.

The named property does not have any hostnames set for it.

Re-create the named resource with one or more hostnames.

813977 :Node %d is listed twice in property %s.

The node in the message was listed twice in the named property.

Specify the property with only one occurrence of the node.

813990 :Started the HA-NFS system fault monitor.

The HA-NFS system fault monitor was started successfully.

No action required.

814232 :fork() failed: %m.

The fork() system call failed for the given reason.

If system resources are not available, consider rebooting the node.

814905 :Could not start up DCS client because major numbers on this node do not match the ones on other nodes. See /var/adm/messages for previous errors.

Some drivers identified in previous messages do not have the same major number across cluster nodes, and devices owned by the driver are being used in global device services.

Look in the /etc/name_to_major file on each cluster node to see if the major number for the driver matches across the cluster. If a driver is missing from the /etc/name_to_major file on some of the nodes, then most likely, the package the driver ships in was not installed successfully on all nodes. If this is the case, install that package on the nodes that don't have it. If the driver exists on all nodes but has different major numbers, see the documentation that shipped with this product for ways to correct this problem.

815551 :System property %s with value %s has an empty list element.

The system property that was named does not have a value for one of its list elements.

Assign the property to have a value where all list elements have values.

816002 :The port number %d from entry %s in property %s was not for a nonsecure port.

The Netscape Directory Server instance has been configured as nonsecure, but the port number given in the list property is for a secure port.

Remove the the entry from the list or change its port number to correspond to a nonsecure port.

816578 :Node %u attempting to join cluster has incompatible cluster software. %s not compatible with %s

A node is attempting to join the cluster but it is either using an incompatible software version or is booted in a different mode (32-bit vs. 64-bit).

Ensure that all nodes have the same clustering software installed and are booted in the same mode.

817592 :HA: rma::admin_impl failed to bind

An HA framework component failed to register with the name server.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

818821 :Value %d is listed twice in property %s.

The value listed occurs twice in the named property.

Specify the property with only one occurrence of the value.

818824 :HA: rma::reconf can't talk to RM

An HA framework component failed to register with the Replica Manager.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

818836 :Value %s is listed twice in property %s.

The value listed occurs twice in the named property.

Specify the property with only one occurrence of the value.

819642 :fatal: unable to register RPC service; aborting node

The rgmd was unable to start up successfully because it failed to register an RPC service. It will produce a core file and will force the node to halt or reboot.

If rebooting the node doesn't fix the problem, examine other syslog messages occurring at about the same time to see if the problem can be identified and if it recurs. Save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance.

819721 :Failed to start %s.

Sun Cluster could not start the application. It would attempt to start the service on another node if possible.

1) Check prior syslog messages for specific problems and correct them. 2) This problem may occur when the cluster is under load and Sun Cluster cannot start the application within the timeout period specified. You may consider increasing the Start_timeout property. 3) If the resource was unable to start on any node, resource would be in START_FAILED state. In this case, use scswitch to bring the resource ONLINE on this node. 4) If the service was successfully started on another node, attempt to restart the service on this node using scswitch. 5) If the above steps do not help, disable the resource using scswitch. Check to see that

the application can run outside of the Sun Cluster framework. If it cannot, fix any problems specific to the application, until the application can run outside of the Sun Cluster framework. Enable the resource using `scswitch`. If the application runs outside of the Sun Cluster framework but not in response to starting the data service, contact your authorized Sun service provider for assistance in diagnosing the problem.

820394 :Cannot check online status. Server processes are not running.

HA-Oracle could not check online status of Oracle server. Oracle server processes are not running.

Examine "Connect_string" property of the resource. Make sure that user id and password specified in connect string are correct and permissions are granted to user for connecting to the server. Check whether Oracle server can be started manually. Examine the log files and setup.

821304 :Failed to retrieve the resource group information.

A Sun cluster dataservice has failed to retrieve the resource group property information. Low memory or API call failure might be the reasons.

In case of low memory, the problem will probably be cured by rebooting. If the problem reoccurs, you might need to increase swap space by configuring additional swap devices. Otherwise, if it is API call failure, check the syslog messages from other components.

821781 :Fencing shared disk groups: %s

A reservation failfast will be set so nodes which share these disk groups will be brought down if they are fenced off by other nodes.

None.

822385 :Failed to retrieve process monitor facility tag.

Failed to create the tag that is used to register with the process monitor facility.

Check the syslog messages that occurred just before this message. In case of internal error, save the `/var/adm/messages` file and contact authorized Sun service provider.

824468 :Invalid probe values.

The values for system defined properties `Retry_count` and `Retry_interval` are not consistent with the property `Thorough_Probe_Interval`.

Change the values of the properties to satisfy the following relationship:
`Thorough_Probe_Interval * Retry_count <= Retry_interval`.

824550 :clcomm: Invalid flow control parameters

The flow control policy is controlled by a set of parameters. These parameters do not satisfy guidelines. Another message from `validay_policy` will have already identified the specific problem.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

824861 :Resource %s named in property %s is not a SharedAddress resource.

The resource given for the named property is not a SharedAddress resource. All resources for that property must be SharedAddress resources.

Specify only SharedAddresses for the named property.

825274 :idl_scha_control_checkall(): IDL Exception on node <%d>

During a failover attempt, the `scha_control` function was unable to check the health of the indicated node, because of an error in inter-node communication. This was probably caused by the death of the indicated node during `scha_control` execution. The RGM will still attempt to master the resource group on another node, if available.

No action is required; the `rgmd` should recover automatically. Identify what caused the node to die by examining syslog output. The syslog output might indicate further remedial actions.

826050 :Failed to retrieve the cluster property %s for %s: %s.

The query for a property failed. The reason for the failure is given in the message.

Save a copy of the `/var/adm/messages` files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

826747 :reservation error(%s) - do_scsi3_inkeys() error for disk %s

The device fencing program has encountered errors while trying to access a device. All retry attempts have failed.

For the user action required by this message, see the user action for message 192619.

827525 :reservation message(%s) - Fencing other node from disk %s

The device fencing program is taking access to the specified device away from a non-cluster node.

This is an informational message, no user action is needed.

828170 :CCR: Unrecoverable failure during updating table %s.

CCR encountered an unrecoverable error while updating the indicated table on this node.

The node needs to be rebooted. Also contact your authorized Sun service provider to determine whether a workaround or patch is available.

828283 :clconf: No memory to read quorum configuration table

Could not allocate memory while converting the quorum configuration information into quorum table.

This is an unrecoverable error, and the cluster needs to be rebooted. Also contact your authorized Sun service provider to determine whether a workaround or patch is available.

828407 :WARNING: lkcm_sync failed: unknown message type %d

An message of unknown type was sent to udln. This will be ignored.

None.

828474 :resource group %s property changed.

This is a notification from the rgmd that the operator has edited a property of a resource group. This may be used by system monitoring tools.

This is an informational message, no user action is needed.

828739 :transition "%s" timed out for cluster, forcing reconfiguration.

Step transition failed. A reconfiguration will be initiated.

Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

829117 :scha_control GIVEOVER failed. error %d

Fault monitor had detected problems in RDBMS server. Attempt to switchover resource to another node failed. Error returned by API call scha_control is indicated in the message.

None.

829262 :Switchover (%s) error: cannot find clexecd

The file system specified in the message could not be hosted on the node the message came from. Check to see if the user program "clexecd" is running on that node.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

829384 :INTERNAL ERROR: launch_method: state machine attempted to launch invalid method <%s> (method <%d>) for resource <%s>; aborting node

An internal error occurred when the rgmd attempted to launch an invalid method for the named resource. The rgmd will produce a core file and will force the node to halt or reboot.

Look for other syslog error messages on the same node. Save a copy of the /var/adm/messages files on all nodes, and report the problem to your authorized Sun service provider.

830211 :Failed to accept connection on socket: %s.

While determining the health of the data service, fault monitor is failed to communicate with the process monitor facility.

This is internal error. Save /var/adm/messages file and contact your authorized Sun service provider. For more details about error, check the syslog messages.

831036 :Service object [%s, %s, %d] created in group "%s"

A specific service known by its unique name SAP (service access point), the three-tuple, has been created in the designated group.

This is an informational message, no user action is needed.

833229 :Couldn't remove deleted directory file, "%s" error: (%d)

The file system is unable to create temporary copies of deleted files.

Mount the affected file system as a local file system, and ensure that there is no file system entry with name "." at the root level of that file system. Alternatively, run fsck on the device to ensure that the file system is not corrupt.

833970 :clcomm: getrlimit(RLIMIT_NOFILE): %s

During Galileo initialization within this user process, the getrlimit call failed with the specified error.

Read the man page for getrlimit for a more detailed description of the error.

837169 :Starting listener %s.

Informational message. HA-Oracle will be starting Oracle listener.

None

837211 :Resource is already online.

While attempting to restart the resource, error has occurred. The resource is already online.

This is an internal error. Save the `/var/adm/messages` file from all the nodes. Contact your authorized Sun service provider.

837752 :Failed to retrieve the resource group handle for %s while querying for property %s: %s.

Access to the object named failed. The reason for the failure is given in the message.

Save a copy of the `/var/adm/messages` files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

837760 :monitored processes forked failed (errno=%d)

The `rpc.pmfd` server was not able to start (fork) the application, probably due to low memory, and the system error number is shown. An error message is output to `syslog`.

Investigate if the machine is running out of memory. If this is not the case, save the `/var/adm/messages` file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

839641 :t_alloc (reqp): %s

Call to `t_alloc()` failed. The "t_alloc" man page describes possible error codes. `udlm` will exit and the node will abort.

This is an internal error. Save the contents of `/var/adm/messages`, `/var/cluster/ucmm/ucmm_reconf.log` and `/var/cluster/ucmm/dlm*/logs/*` from all the nodes and contact your Sun service representative.

839649 :t_alloc (resp): %s

Call to `t_alloc()` failed. The "t_alloc" man page describes possible error codes. `udlm` will exit and the node will abort.

This is an internal error. Save the contents of `/var/adm/messages`, `/var/cluster/ucmm/ucmm_reconf.log` and `/var/cluster/ucmm/dlm*/logs/*` from all the nodes and contact your Sun service representative.

839936 :Some ip addresses may not be plumbed.

Some of the ip addresses managed by the `LogicalHostname` resource were not successfully brought on-line on this node.

Use `ifconfig` command to make sure that the ip addresses are indeed absent. Check for any error message before this error message for a more precise reason for this error. Use `scswitch` to move the resource group to some other node.

840542 :OFF_PENDING_BOOT: bad resource state <%s> (%d) for resource <%s>

The rgmd state machine has discovered a resource in an unexpected state on the local node. This should not occur and may indicate an internal logic error in the rgmd.

Look for other syslog error messages on the same node. Save a copy of the /var/adm/messages files on all nodes, and report the problem to your authorized Sun service provider.

840619 :Invalid value was returned for resource group property
%s for %s.

The value returned for the named property was not valid.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

840696 :DNS database directory %s is not readable: %s

The DNS database directory is not readable. This may be due to the directory not existing or the permissions not being set properly.

Make sure the directory exists and has read permission set appropriately. Look at the prior syslog messages for any specific problems and correct them.

841616 :CMM: This node has been preempted from quorum device %s.

This node's reservation key was on the specified quorum device, but is no longer present, implying that this node has been preempted by another cluster partition. If a cluster gets divided into two or more disjoint subclusters, exactly one of these must survive as the operational cluster. The surviving cluster forces the other subclusters to abort by grabbing enough votes to grant it majority quorum. This is referred to as preemption of the losing subclusters.

There may be other related messages that may indicate why quorum was lost. Determine why quorum was lost on this node, resolve the problem and reboot this node.

841719 :listener %s is not running. restart limit reached.
Stopping fault monitor.

Listener is not running. Listener monitor has reached the restart limit specified in "Retry_count" and "Retry_interval" properties. Listener monitor will be stopped.

Check Oracle listener setup. Please make sure that Listener_name specified in the resource property is configured in listener.ora file. Check "Host" property of listener in listener.ora file. Examine log file and syslog messages for additional information. Stop and start listener monitor.

841875 :remote node died

An inter-node communication failed because another cluster node died.

No action is required. The cluster will reconfigure automatically. Examine syslog output on the rebooted node to determine the cause of node death.

842313 :clxecd: Sending fd on common channel returned %d.
Exiting.

clxecd program has encountered a failed fcntl(2) system call. The error message indicates the error number for the failure.

The node will halt or reboot itself to prevent data corruption. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

842382 :fcntl: %s

A server (rpc.pmfd or rpc.fed) was not able to execute the action shown, and the process associated with the tag is not started. The error message is shown.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

842712 :clcomm: solaris xdoor door_create failed

A door_create operation failed. Refer to the "door_create" man page for more information.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

843093 :fatal: Got error <%d> trying to read CCR when enabling
monitor of resource <%s>; aborting node

Rgmd failed to read updated resource from the CCR on this node.

Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

843978 :Socket creation failed: %s.

System is unable to create a socket.

This might be the result from the lack of system resources. Check whether the system is low in memory and take appropriate action. For specific error information check the syslog message.

843983 :CMM: Node %s: attempting to join cluster.

The specified node is attempting to become a member of the cluster.

This is an informational message, no user action is needed.

845866 :Failover attempt failed: %s.

The failover attempt of the resource is rejected or encountered an error.

For more detailed error message, check the syslog messages. Check whether the Pingpong_interval has appropriate value. If not, adjust it using scrgadm(1M). Otherwise, use scswitch to switch the resource group to a healthy node.

846053 :Fast path enable failed on %s%d, could cause path timeouts

DLPI fast path could not be enabled on the device.

Check if the right version of the driver is in use.

846376 :fatal: Got error <%d> trying to read CCR when making resource group <%s> unmanaged; aborting node

Rgmd failed to read updated resource from the CCR on this node.

Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

846420 :CMM: Nodes %ld and %ld are disconnected from each other; node %ld will abort using %s rule.

Due to a connection failure between the two specified non-local nodes, one of the nodes must be halted to avoid a "split brain" configuration. The CMM used the specified rule to decide which node to fail. Rules are: rebootee: If one node is rebooting and the other was a member of the cluster, the node that is rebooting must abort. quorum: The node with greater control of quorum device votes survives and the other node aborts. node number: The node with higher node number aborts.

The cause of the failure should be resolved and the node should be rebooted if node failure is unexpected.

846813 :Switchover (%s) error (%d) converting to primary

The file system specified in the message could not be hosted on the node the message came from.

Check /var/adm/messages to make sure there were no device errors. If not, contact your authorized Sun service provider to determine whether a workaround or patch is available.

847065 :Failed to start listener %s.

Failed to start Oracle listener.

Check Oracle listener setup. Please make sure that Listener_name specified in the resource property is configured in listener.ora file. Check "Host" property of listener in listener.ora file. Examine log file and syslog messages for additional information.

847124 :getnetconfignt: %s

call to getnetconfignt in udlm port setup failed.udlm fails to start and the node will eventually panic.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

847496 :CMM: Reading reservation keys from quorum device %s failed with error %d.

The specified error was encountered while trying to read reservation keys on the specified quorum device.

There may be other related messages on this and other nodes connected to this quorum device that may indicate the cause of this problem. Refer to the quorum disk repair section of the administration guide for resolving this problem.

847809 :Must be in cluster to start %s

Machine on which this command or daemon is running is not part of a cluster.

Run the command on another machine or make the machine is part of a cluster by following appropriate steps.

847916 :(%s) netdir error: uaddr2taddr: %s

Call to uaddr2taddr() failed. The "uaddr2taddr" man page describes possible error codes. udlmctl will exit.

Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

847978 :reservation fatal error(UNKNOWN) - cluster_get_quorum_status() error, returned %d

The device fencing program has suffered an internal error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available. Copies of /var/adm/messages from all nodes should be provided for diagnosis. It may be possible to retry the failed operation, depending on the nature of the error. If the message specifies the "node_join" transition, then this node may be unable to access shared devices. If the failure occurred during the "release_shared_scsi2" transition, then a node which was

joining the cluster may be unable to access shared devices. In either case, it may be possible to reacquire access to shared devices by executing `"/usr/cluster/lib/sc/run_reserve -c node_join"` on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group has failed to start on this node. If another node was available to host the device group, then it should have been started on that node. If desired, it may be possible to switch the device group to this node with the `scswitch` command. If no other node was available, then the device group will not have been started. The `scswitch` command may be used to retry the attempt to start the device group. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group has failed. The desired action may be retried.

848033 :SharedAddress online.

The status of the sharedaddress resource is online.

This is informational message. No user action required.

848652 :CMM aborting.

The node is going down due to a decision by the cluster membership monitor.

This message is preceded by other messages indicating the specific cause of the abort, and the documentation for these preceding messages will explain what action should be taken. The node should be rebooted if node failure is unexpected.

848881 :Received notice that NAFO group %s has failed.

The status of the named NAFO group has become degraded. If possible, the scalable resources currently running on this node with monitoring enabled will be relocated off of this node, if the NAFO group stays in a degraded state.

Check the status of the NAFO group on the node. Try to fix the adapters in the NAFO group.

848943 :clconf: No valid gdevname field for quorum device %d

Found the gdevname field for the quorum device being incorrect while converting the quorum configuration information into quorum table.

Check the quorum configuration information.

852212 :reservation message(%s) - Taking ownership of disk %s away from non-cluster node

The device fencing program is taking access to the specified device away from a non-cluster node.

This is an informational message, no user action is needed.

852497 :scvxvmlg error - readlink(%s) failed

The program responsible for maintaining the VxVM namespace was unable to access the global device namespace. If configuration changes were recently made to VxVM diskgroups or volumes, this node may be unaware of those changes. Recently created volumes may be inaccessible from this node.

Verify that the `/global/.devices/node@N` (N = this node's node number) is mounted globally and is accessible. If no configuration changes have been recently made to VxVM diskgroups or volumes and all volumes continue to be accessible from this node, then no further action is required. If changes have been made, the device namespace on this node can be updated to reflect those changes by executing `"/usr/cluster/lib/dcs/scvxvmlg"`. If the problem persists, contact your authorized Sun service provider to determine whether a workaround or patch is available.

852615 :reservation error(%s) - Unable to gain access to device "%s"

The device fencing program has encountered errors while trying to access a device.

Another cluster node has fenced this node from the specified device, preventing this node from accessing that device. Access should have been reacquired when this node joined the cluster, but this must have experienced problems. If the message specifies the "node_join" transition, this node will be unable to access the specified device. If the failure occurred during the "make_primary" transition, then this will be unable to access the specified device and a device group containing the specified device may have failed to start on this node. An attempt can be made to acquire access to the device by executing `"/usr/cluster/lib/sc/run_reserve -c node_join"` on this node. If a device group failed to start on this node, the `scswitch` command can be used to start the device group on this node if access can be reacquired. If the problem persists, please contact your authorized Sun service provider to determine whether a workaround or patch is available.

853478 :Received non interrupt heartbeat on %s - path timeouts are likely.

Solaris Clustering requires network drivers to deliver heartbeat messages in the interrupt context. A heartbeat message has unexpectedly arrived in non interrupt context.

Check if the right version of the driver is in use.

854468 :failfast arm error: %d

Error during failfast device arm operation.

This is an internal error. Save the contents of `/var/adm/messages`, `/var/cluster/ucmm/ucmm_reconf.log` and `/var/cluster/ucmm/dlm*/logs/*` from all the nodes and contact your Sun service representative.

854792 :clcomm: error in copyin for cl_change_threads_min

The system failed a copy operation supporting a flow control state change.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

854894 :No LogicalHostname resource in resource group.

The probe method for this data service could not find a LogicalHostname resource in the same resource group as the data service.

Use scrgadm to configure the resource group to hold both the data service and the LogicalHostname.

856492 :waitpid() failed: %m.

The waitpid() system call failed for the given reason.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

856919 :INTERNAL ERROR: process_resource: resource group <%s> is pending_methods but contains resource <%s> in STOP_FAILED state

During a resource creation, deletion, or update, the rgmd has discovered a resource in STOP_FAILED state. This may indicate an internal logic error in the rgmd, since updates are not permitted on the resource group until the STOP_FAILED error condition is cleared.

Look for other syslog error messages on the same node. Save a copy of the /var/adm/messages files on all nodes, and report the problem to your authorized Sun service provider.

857573 :scvxvmlg error - rmdir(%s) failed

The program responsible for maintaining the VxVM namespace was unable to access the global device namespace. If configuration changes were recently made to VxVM diskgroups or volumes, this node may be unaware of those changes. Recently created volumes may be inaccessible from this node.

Verify that the /global/.devices/node@N (N = this node's node number) is mounted globally and is accessible. If no configuration changes have been recently made to VxVM diskgroups or volumes and all volumes continue to be accessible from this node, then no further action is required. If changes have been made, the device namespace on this node can be updated to reflect those changes by executing "/usr/cluster/lib/dcs/scvxvmlg". If the problem persists, contact your authorized Sun service provider to determine whether a workaround or patch is available.

857792 :UNIX DLM initiating cluster abort.

Due to an error encountered, unix dlm is initiating an abort.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

859126 :System property %s is empty.

The system property that was named does not have a value.

Assign the property a value.

862493 :in libsecurity could not register on any transport in NETPATH

A server (rpc.pmfd, rpc.fed or rgmd) was not able to start because it could not establish a rpc connection for the network specified, because it couldn't find any transport. An error message is output to syslog. This happened because either there are no available transports at all, or there are but none is a loopback.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

862716 :sema_init: %s

The rpc.pmfd server was not able to initialize a semaphore, possibly due to low memory, and the system error is shown. The server does not perform the action requested by the client, and pmfadm returns error. An error message is also output to syslog.

Investigate if the machine is running out of memory. If this is not the case, save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

865635 :lkcm_act: caller is not registered

udlm is not currently registered with ucmm.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

866624 :clcomm: validate_policy: threads_low not big enough low %d pool %d

The system checks the proposed flow control policy parameters at system startup and when processing a change request. The low server thread level must not be less than twice the thread increment level for resource pools whose number threads varies dynamically.

No user action required.

867059 :Could not shutdown replica for device service (%s). Some file system replicas that depend on this device service may

already be shutdown. Future swichovers to this device service will not succeed unless this node is rebooted.

See message.

If mounts or node reboots are on at the time this message was displayed, wait for that activity to complete, and then retry the command to shutdown the device service replica. If not, then contact your authorized Sun service provider to determine whether a workaround or patch is available.

869406 :Failed to communicate with server %s port %d: %s.

The data service fault monitor probe was trying to read from or write to the service specified and failed. Sun Cluster will attempt to correct the situation by either doing a restart or a failover of the data service. The problem may be due to an overloaded system or other problems, causing a timeout to occur before communications could be completed.

If this problem is due to an overloaded system, you may consider increasing the Probe_timeout property.

870181 :Failed to retrieve the resource handle for %s while querying for property %s: %s.

Access to the object named failed. The reason for the failure is given in the message.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

870317 :INTERNAL ERROR: START method is not registered for resource <%s>

A non-fatal internal error has occurred in the rgmd state machine.

Since this problem might indicate an internal logic error in the rgmd, please save a copy of the /var/adm/messages files on all nodes, the output of an scstat -g command, and the output of a scrgadm -pvv command. Report the problem to your authorized Sun service provider.

870566 :clutil: Scheduling class %s not configured

An attempt to change the thread scheduling class failed, because the scheduling class was not configured.

Configure the system to support the desired thread scheduling class.

871642 :Validation failed. Invalid command line %s %s

Unable to process parameters passed to the call back method. This is an internal error.

Please report this problem.

872086 :Service is degraded.

Probe is detected a failure in the data service. Probe is setting resource's status as degraded.

Wait for the fault monitor to restart the data service. Check the syslog messages and configuration of the data service.

872599 :Error in getting service name for device path <%s>

Can not map the device path to a valid global service name.

Check the path passed into extension property "ServicePaths" of SUNW.HAStorage type resource.

872839 :Resource is already stopped.

Sun Cluster attempted to stop the resource, but found it already stopped.

No user action required.

874879 :clcomm: Path %s being deleted

A communication link is being removed with another node. The interconnect may have failed or the remote node may be down.

Any interconnect failure should be resolved, and/or the failed node rebooted.

875171 :clcomm: Pathend %p: %d is not a pathend state

The system maintains state information about a path. The state information is invalid.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

875345 :None of the shared paths in file %s are valid.

All the paths specified in the dfstab._name> file are invalid.

Check that those paths are valid. This might be a result of the underlying disk failure in an unavailable file system. The monitor_check method would thus fail and the HA-NFS resource would not be brought online on this node. However, it is advisable that the file system be brought online soon.

875595 :CMM: Shutdown timer expired. Halting.

The node could not complete its shutdown sequence within the halt timeout, and is aborting to enable another node to safely take over its services.

This is an informational message, no user action is needed.

875796 :CMM: Reconfiguration callback timed out; node aborting.

One or more CMM client callbacks timed out and the node will be aborted.

There may be other related messages on this node which may help diagnose the problem. Resolve the problem and reboot the node if node failure is unexpected. If unable to resolve the problem, contact your authorized Sun service provider to determine whether a workaround or patch is available.

875939 :ERROR: Failed to initialize callbacks for
Global_resources_used, error code <%d>

The rgmd encountered an error while trying to initialize the Global_resources_used mechanism on this node. This is not considered a fatal error, but probably means that method timeouts will not be suspended while a device service is failing over. This could cause unneeded failovers of resource groups when device groups are switched over.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem. This error might be cleared by rebooting the node.

876090 :fatal: must be superuser to start %s

The rgmd can only be executed by the super-user.

This probably occurred because a non-root user attempted to start the rgmd manually. Normally, the rgmd is started automatically when the node is booted.

876324 :CCR: CCR transaction manager failed to register with the
cluster HA framework.

The CCR transaction manager failed to register with the cluster HA framework.

This is an unrecoverable error, and the cluster needs to be rebooted. Also contact your authorized Sun service provider to determine whether a workaround or patch is available.

876834 :Could not start server

HA-Oracle failed to start Oracle server. Syslog messages and log file will provide additional information on possible reasons of failure.

Check whether Oracle server can be started manually. Examine the log files and setup.

877905 :ff_ioctl: %s

A server (rpc.pmfd or rpc.fed) was not able to arm or disarm the failfast device, which ensures that the host aborts if the server dies. The error message is shown. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

878089 :fatal: realloc: %s (UNIX error %d)

The rgmd failed to allocate memory, most likely because the system has run out of swap space. The rgmd will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

The problem was probably cured by rebooting. If the problem recurs, you might need to increase swap space by configuring additional swap devices. See swap(1M) for more information.

878135 :WARNING: udln_update_from_saved_msg

There is no saved message to update udln.

None. This is a warning only.

879301 :reservation error(%s) - clconf_do_execution() error

The device fencing program has suffered an internal error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available. Copies of /var/adm/messages from all nodes should be provided for diagnosis. It may be possible to retry the failed operation, depending on the nature of the error. If the message specifies the "node_join" transition, then this node may be unable to access shared devices. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access shared devices. In either case, it may be possible to reacquire access to shared devices by executing "/usr/cluster/lib/sc/run_reserve -c node_join" on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group has failed to start on this node. If another node was available to host the device group, then it should have been started on that node. If desired, it may be possible to switch the device group to this node with the scswitch command. If no other node was available, then the device group will not have been started. The scswitch command may be used to retry the attempt to start the device group. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group has failed. The desired action may be retried.

879511 :reservation fatal error(%s) - service_class not specified

The device fencing program has suffered an internal error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available. Copies of /var/adm/messages from all nodes should be provided for diagnosis. It may be possible to retry the failed operation, depending on the nature of the error. If the message specifies the "node_join"

transition, then this node may be unable to access shared devices. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access shared devices. In either case, it may be possible to reacquire access to shared devices by executing "/usr/cluster/lib/sc/run_reserve -c node_join" on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group has failed to start on this node. If another node was available to host the device group, then it should have been started on that node. If desired, it may be possible to switch the device group to this node with the scswitch command. If no other node was available, then the device group will not have been started. The scswitch command may be used to retry the attempt to start the device group. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group has failed. The desired action may be retried.

880317 :scvxvmlg fatal error - %s does not exist, VxVM not installed?

The program responsible for maintaining the VxVM namespace was unable to access the local VxVM device namespace. If configuration changes were recently made to VxVM diskgroups or volumes, this node may be unaware of those changes. Recently created volumes may be inaccessible from this node.

If VxVM is used to manage shared device, it must be installed on all cluster nodes. If VxVM is installed on this node, but the local VxVM namespace does not exist, VxVM may have to be re-installed on this node. If VxVM is installed on this node and the local VxVM device namespace does exist, the namespace management can be manually run on this node by executing "/usr/cluster/lib/dcs/scvxvmlg" on this node. If the problem persists, please contact your authorized Sun service provider to determine whether a workaround or patch is available. If VxVM is not being used on this cluster, then no user action is required.

880835 :pmf_search_children: Error stopping <%s>: %s

An error occurred while rpc.pmf attempted to send a KILL signal to one of the processes of the given tag. The reason for the failure is also given. rpc.pmf attempted to kill the process because a previous error occurred while creating a monitor process for the process to be monitored.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

884114 :clcomm: Adapter %s constructed

A network adapter has been initialized.

No action required.

884482 :clconf: Quorum device ID %ld is invalid. The largest supported ID is %ld

Found the quorum device ID being invalid while converting the quorum configuration information into quorum table.

Check the quorum configuration information.

884823 :Prog <%s> step <%s>: stat of program file failed.

A step points to a file that is not executable. This may have been caused by incorrect installation of the package.

Identify the program for the step. Check the permissions on the program. Reinstall the package if necessary.

884979 :(%s) aborting, but got a message of type %d

Going through udlm abort and received an unexpected message of the mentioned type.

None.

887666 :clcomm: sxdoor: op %d fcntl failed: %s

A user level process is unmarshalling a door descriptor and creating a new door. The specified operation on the fcntl operation fails. The "fcntl" man page describes possible error codes.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

887669 :clcomm: coalesce_region request(%d) > MTUsize(%d)

While supporting an invocation, the system wanted to create one buffer that could hold the data from two buffers. The system cannot create a big enough buffer. After generating another system error message, the system will panic. This message only appears on debug systems.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

888259 :clcomm: Path %s being deleted and cleaned

A communication link is being removed with another node. The interconnect may have failed or the remote node may be down.

Any interconnect failure should be resolved, and/or the failed node rebooted.

889884 :scha_control RESTART failed. error %d

Fault monitor had detected problems in RDBMS server. Attempt to restart RDBMS server on the same node failed. Error returned by API call sch_a_control is indicated in the message.

None.

890129 :dl_attach: DL_ERROR_ACK access error

Could not attach to the physical device. We are trying to open a fast path to the private transport adapters.

Reboot of the node might fix the problem

890927 :HA: repl_mgr_impl: thr_create failed

The system could not create the needed thread, because there is inadequate memory.

There are two possible solutions. Install more memory. Alternatively, reduce memory usage.

891030 :resource group %s state on node %s change to
RG_OFF_BOOTED

This is a notification from the rgmd that a resource group has completed running its resources" BOOT methods on the given node. This may be used by system monitoring tools.

This is an informational message, no user action is needed.

891362 :scha_resource_open error (%d)

Error occurred in API call scha_resource_open.

Check syslog messages for errors logged from other system modules. Stop and start fault monitor. If error persists then disable fault monitor and report the problem.

891462 :in libsecurity caller is %d, not the desired uid %d

A server (rpc.pmfd, rpc.fed or rgmd) refused an rpc connection from a client because it has the wrong uid. The actual and desired uids are shown. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

891920 :rgm_clear_util called on resource <%s> with unknown flag
<%d>

An internal rgmd error has occurred while attempting to carry out an operator request to clear an error flag on a resource. The attempted clear action will fail.

Since this problem might indicate an internal logic error in the rgmd, please save a copy of the /var/adm/messages files on all nodes, the output of an scstat -g

command, and the output of a `scrgadm -pvv` command. Report the problem to your authorized Sun service provider.

892183 :libsecurity: NULL RPC to program %ld failed will not
retry %s

A client of the `rpc.pmfd`, `rpc.fed` or `rgmd` server was not able to initiate an rpc connection, because it could not execute a test rpc call. The program will not retry because the time limit of 1 hr was exceeded. The message shows the specific rpc error. The program number is shown. To find out what program corresponds to this number, use the `rpcinfo` command. An error message is output to `syslog`.

Save the `/var/adm/messages` file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

893095 :Service <%s> with path <%s> is not available.
Retrying...

The service is not available yet. `prenet_start` method of `SUNW.HASStorage` is still testing and waiting.

Not user action is required.

894418 :reservation warning(%s) - Found invalid key, preempting

The device fencing program has discovered an invalid scsi-3 key on the specified device and is removing it.

This is an informational message, no user action is needed.

894711 :Could not resolve "%s" in the name server. Exiting.

`clexecd` program was unable to start due to an error in registering itself with the low-level clustering software.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

895149 :(%s) t_open: tli error: %s

Call to `t_open()` failed. The "t_open" man page describes possible error codes. `udlmctl` will exit.

Save the contents of `/var/adm/messages`, `/var/cluster/ucmm/ucmm_reconf.log` and `/var/cluster/ucmm/dlm*/logs/*` from all the nodes and contact your Sun service representative.

895159 :clcomm: solaris xdoor dup failed: %s

A dup operation failed. The "dup" man page describes possible error codes.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

895821 :INTERNAL ERROR: cannot get nodeid for node <%s>

The scha_control function is unable to obtain the node id number for one of the resource group's potential masters. This node will not be considered a candidate destination for the scha_control giveover.

Try issuing an scstat(1M) -n command and see if it successfully reports status for all nodes. If not, then the cluster configuration data may be corrupted. If so, then there may be an internal logic error in the rgmd. In either case, please save copies of the /var/adm/messages files on all nodes, and report the problem to your authorized Sun service provider.

896275 :CCR: Ignoring override field for table %s on joining node %s.

The override flag for a table indicates that the CCR should use this copy as the final version when the cluster is coming up. If the cluster already has a valid copy while the indicated node is joining the cluster, then the override flag on the joining node is ignored.

This is an informational message, no user action is needed.

896441 :Unknown scalable service method code: %d.

The method code given is not a method code that was expected.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

896799 :INTERNAL ERROR: resource group <%s> is PENDING_BOOT or ERROR_STOP_FAILED, but contains no resources

The operator is attempting to delete the indicated resource group. Although the group is empty of resources, it was found to be in an unexpected state. This will cause the resource group deletion to fail.

Use scswitch(1M) -z to switch the resource group offline on all nodes, then retry the deletion operation. Since this problem might indicate an internal logic error in the rgmd, please save a copy of the /var/adm/messages files on all nodes, and report the problem to your authorized Sun service provider.

897348 :%s: must be run in secure mode using -S flag

rpc.sccheckd should always be invoked in secure mode. If this message shows up, someone has modified configuration files that affects server startup.

Reinstall cluster packages or contact your service provider.

898001 :launch_fed_prog: getlocalhostname() failed for program <%s>

The ucmmmd was unable to obtain the name of the local host. Launching of a method failed.

Examine other syslog messages occurring at about the same time to see if the problem can be identified. Save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance in diagnosing the problem.

898738 :Aborting node because pm_tick delay of %lld ms exceeds %lld ms

The system is unable to send heartbeats for a long time. (This is half of the minimum of timeout values of all the paths. If the timeout values for all the paths is 10 secs then this value is 5 secs.) There is probably heavy interrupt activity causing the clock thread to get delayed, which in turn causes irregular heartbeats. The node is aborted because it is considered to be in "sick" condition and it is better to abort this node instead of causing other nodes (or the cluster) to go down.

Check to see what is causing high interrupt activity and configure the system accordingly.

899278 :Retry_count exceeded in Retry_interval

Fault monitor has detected problems in RDBMS server. Number of restarts through fault monitor exceed the count specified in "Retry_count" parameter in "Retry_interval". Database server is unable to survive on this node. Switching over the resource group to other node.

Please check the RDBMS setup and server configuration.

899305 :clexecd: Daemon exiting because child died.

Child process in the clexecd program is dead.

If this message is seen when the node is shutting down, ignore the message. If that's not the case, the node will halt or reboot itself to prevent data corruption. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

899648 :Failed to process the resource information.

A Sun cluster data service is unable to retrieve the resource property information. Low memory or API call failure might be the reasons.

In case of low memory, the problem will probably be cured by rebooting. If the problem reoccurs, you might need to increase swap space by configuring additional swap devices. Otherwise, if it is API call failure, check the syslog messages from other components.

899776 :ERROR: scha_control() was called on resource group <%s>, resource <%s> before the RGM started

This message most likely indicates that a program called `scha_control(1ha,3ha)` before the RGM had started up. Normally, `scha_control` is called by a resource monitor to request failover or restart of a resource group. If the RGM had not yet started up on the cluster, no resources or resource monitors should have been running on any node. The `scha_control` call will fail with a `SCHA_ERR_CLRECONF` error.

On the node where this message appeared, confirm that `rgmd` was not yet running (i.e., the cluster was just booting up) when this message was produced. Find out what program called `scha_control`. If it was a customer-supplied program, this most likely represents an incorrect program behavior which should be corrected. If there is no such customer-supplied program, or if the cluster was not just starting up when the message appeared, contact your authorized Sun service provider for assistance in diagnosing the problem.

Message IDs 900000 - 999999

Error Message List

The following list is ordered by the message ID.

900102 :Failed to retrieve the resource type property %s: %s.

An API operation has failed while retrieving the resource type property. Low memory or API call failure might be the reasons.

In case of low memory, the problem will probably be cured by rebooting. If the problem reoccurs, you might need to increase swap space by configuring additional swap devices. Otherwise, if it is API call failure, check the syslog messages from other components.

900206 :parameter "%s%s" must be an integer "%s". Using default value of %d.

Using a default value for a parameter.

None.

900675 :cluster volume manager shared access mode enabled

Message indicating shared access availability of the volume manager.

None.

900843 :Retrying to retrieve the cluster information.

An update to cluster configuration occurred while cluster properties were being retrieved

Ignore the message.

900954 :fatal: Unable to open CCR

The rgmd was unable to open the cluster configuration repository (CCR). The rgmd will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

901030 :clconf: Data length is more than max supported length in clconf_file_io

In reading configuration data through CCR FILE interface, found the data length is more than max supported length.

Check the CCR configuraton information.

902721 :Switching over resource group using scha_control
GIVEOVER

Fault monitor has detected problems in RDBMS server. Fault monitor has determined that RDBMS server cannot be restarted on this node. Attempt will be made to switchover the resource to any other node, if a healthy node is available.

Check the cause of RDBMS failure.

903007 :txmit_common: udp is null!

Can not transmit a message and communicate with udlmctl because the address to send to is null.

None.

903200 :PNM: nafo%d: could not start monitoring

An internal resource problem has stopped PNM from performing monitoring of the named NAFO group. Fault detection and failover for adapters in the group are therefore disabled.

Send a KILL (9) signal to the pnmd daemon. Because pnmd is under PMF control, it will be restarted automatically. If the problem persists, restart the node with scswitch -S and shutdown(1M).

903317 :Entry at position %d in property %s was invalid.

An invalid entry was found in the named property. The position index, which starts at 0 for the first element in the list, indicates which element in the property list was invalid.

Make sure the property has a valid value.

903734 :Failed to create lock directory %s: %s.

This network resource failed to create a directory in which to store lock files. These lock files are needed to serialize the running of the same callback method on the same adapter for multiple resources.

This might be the result of a lack of system resources. Check whether the system is low in memory and take appropriate action. For specific error information, check the syslog message.

905023 :clexecd: dup2 of stderr returned with errno %d while exec"ing (%s). Exiting.

clexecd program has encountered a failed dup2(2) system call. The error message indicates the error number for the failure.

The clexecd program will exit and the node will be halted or rebooted to prevent data corruption. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

905591 :Could not determine volume configuration daemon mode

Could not get information about the volume manager daemon mode.

Check if the volume manager has been started up right.

905720 :Failed to get NAFO status for group %s (request failed with %d).

A query to get the state of a NAFO group failed. This may cause a method failure to occur.

Make sure the network monitoring daemon (pnmd) is running. Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

908240 :in libsecurity realloc failed

A server (rpc.pmfd, rpc.fed or rgmd) was not able to start, or a client was not able to make an rpc connection to the server, probably due to low memory. An error message is output to syslog.

Investigate if the host is low on memory. If not, save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

908387 :ucm_callback for step %d generated exception %d

ucmm callback for a step failed.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

908591 :Failed to stop fault monitor.

An attempt was made to stop the fault monitor and it failed. There may be prior messages in syslog indicating specific problems.

If there are prior messages in syslog indicating specific problems, these should be corrected. If that doesn't resolve the issue, the user can try the following. Use process monitor facility (pmfadm (1M)) with -L option to retrieve all the tags that are running on the server. Identify the tag name for the fault monitor of this resource. This can be easily identified as the tag ends in string ".mon" and contains the resource group name and the resource name. Then use pmfadm (1M) with -s option to stop the fault monitor. This problem may occur when the cluster is under load and Sun Cluster cannot stop the fault monitor within the timeout period specified. You may consider increasing the Monitor_Stop_timeout property. If the error still persists, then reboot the node.

908716 :fatal: Aborting this node because method <%s> failed on resource <%s> and Failover_mode is set to HARD

A STOP method has failed or timed out on a resource, and the Failover_mode property of that resource is set to HARD.

No action is required. This is normal behavior of the RGM. With Failover_mode set to HARD, the rgmd reboots the node to force the resource group offline so that it can be switched onto another node. Other syslog messages that occurred just before this one might indicate the cause of the STOP method failure.

912352 :Could not unplumb some ip addresses.

Some of the ip addresses managed by the LogicalHostname resource were not successfully brought offline on this node.

Use the ifconfig command to make sure that the ip addresses are indeed absent. Check for any error message before this error message for a more precise reason for this error. Use scswitch command to move the resource group to a different node. If problem persists, reboot.

913654 :%s not specified on command line.

Required property not included in a scrgadm command.

Reissue the scrgadm command with the required property and value.

914519 :Error when sending message to child %m

Error occurred when communicating with fault monitor child process. Child process will be stopped and restarted.

If error persists, then disable the fault monitor and resport the problem.

914866 :Unable to complete some unshare commands.

HA-NFS postnet_stop method was unable to complete the unshare(1M) command for some of the paths specified in the dfstab file.

The exact pathnames which failed to be unshared would have been logged in earlier messages. Run those unshare commands by hand. If problem persists, reboot the node.

917591 :fatal: Resource type <%s> update failed with error <%d>; aborting node

Rgmd failed to read updated resource type from the CCR on this node.

Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

918018 :load balancer for group "%s" released

This message indicates that the service group has been deleted.

This is an informational message, no user action is needed.

919417 :Failed to pre-allocate <%d> Mbyte swap space

The pmfd, fed, or other program was not able to allocate swap space. This means that the machine is low in swap space. The server does not come up, and an error message is output to syslog.

Investigate if the machine is running out of swap. If this is not the case, save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

919740 :WARNING: error in translating address (%s) for nodeid %d

Could not get address for a node.

Make sure the node is booted as part of a cluster.

919860 :scvxvmlg warning - %s does not link to %s, changing it

The program responsible for maintaining the VxVM device namespace has discovered inconsistencies between the VxVM device namespace on this node and the VxVM configuration information stored in the cluster device configuration system. If configuration changes were made recently, then this message should reflect one of the configuration changes. If no changes were made recently or if this message does not correctly reflect a change that has been made, the VxVM device namespace on this node may be in an inconsistent state. VxVM volumes may be inaccessible from this node.

If this message correctly reflects a configuration change to VxVM diskgroups then no action is required. If the change this message reflects is not correct, then the information stored in the device configuration system for each VxVM diskgroup should be examined for correctness. If the information in the device configuration system is accurate, then executing `"/usr/cluster/lib/dcs/scvxxvmlg"` on this node should restore the device namespace. If the information stored in the device configuration system is not accurate, it must be updated by executing `"/usr/cluster/bin/scconf -c -D name=diskgroup_name"` for each VxVM diskgroup with inconsistent information.

920103 :created %d threads to handle resource group switchback;
desired number = %d

The rgmd was unable to create the desired number of threads upon starting up. This is not a fatal error, but it may cause RGM reconfigurations to take longer because it will limit the number of tasks that the rgmd can perform concurrently.

Make sure that the hardware configuration meets documented minimum requirements. Examine other syslog messages on the same node to see if the cause of the problem can be determined.

920272 :CMM: Open failed for quorum device %ld with gdevname %s.

The open operation on the specified quorum device failed, and this node will ignore the quorum device.

The quorum device has failed or the path to this device may be broken. Refer to the quorum disk repair section of the administration guide for resolving this problem.

922085 :INTERNAL ERROR CMM: Memory allocation error.

The CMM failed to allocate memory during initialization.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

922363 :resource %s status msg on node %s change to <%s>

This is a notification from the rgmd that a resource's fault monitor status message has changed.

This is an informational message, no user action is needed.

923184 :CMM: Scrub failed for quorum device %s.

The scrub operation for the specified quorum device failed, due to which this quorum device will not be added to the cluster.

There may be other related messages on this node that may indicate the cause of this problem. Refer to the disk repair section of the administration guide for

resolving this problem. After the problem has been resolved, retry adding the quorum device.

923618 :Prog <%s>: unknown command.

An internal error in ucmmmd has prevented it from successfully executing a program.

Save a copy of /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

923712 :CCR: Table %s on joining node %s has the same version but different checksum as the copy in the current membership. The table on the joining node will be replaced by the one in the current membership.

The indicated table on the joining node has the same version but different contents as the one in the current membership. It will be replaced by the one in the current membership.

This is an informational message, no user action is needed.

923959 :Warning: out of swap space; cannot store start-failed timestamp for resource group <%s>

The specified resource group failed to come online on some node, but this node is unable to record that fact due to insufficient swap space. The system is likely to halt or reboot if swap space continues to be depleted.

Investigate the cause of swap space depletion and correct the problem, if possible.

925953 :reservation error(%s) - do_scsi3_register() error for disk %s

The device fencing program has encountered errors while trying to access a device. All retry attempts have failed.

For the user action required by this message, see the user action for message 192619.

926201 :RGM aborting

A fatal error has occurred in the rgmd. The rgmd will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

926550 :scha_cluster_get failed for (%s) with %d

Call to get cluster information failed. The second part of the message gives the error code.

The calling program should handle this error. If it is not recoverable, it will exit.

926749 :Resource group nodelist is empty.

Empty value was specified for the nodelist property of the resource group.

Any of the following situations might have occurred. Different user action is required for these different scenarios. 1) If a new resource is created or updated, check the value of the nodelist property. If it is empty or not valid, then provide valid value using `scrgadm(1M)` command. 2) For all other cases, treat it as an Internal error. Contact your authorized Sun service provider.

927753 :Fault monitor does not exist or is not executable

Fault monitor program specified in support file is not executable or does not exist. Recheck your installation.

Please report this problem.

927846 :fatal: Received unexpected result <%d> from rpc.fed, aborting node

A serious error has occurred in the communication between `rgmd` and `rpc.fed` while attempting to execute a `VALIDATE` method. The `rgmd` will produce a core file and will force the node to halt or reboot to avoid the possibility of data corruption.

Save a copy of the `/var/adm/messages` files on all nodes, and of the `rgmd` core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

928382 :CCR: Failed to read table %s on node %s.

The CCR failed to read the indicated table on the indicated node. The CCR will attempt to recover this table from other nodes in the cluster.

This is an informational message, no user action is needed.

928455 :clcomm: Couldn't write to routing socket: %d

The system prepares IP communications across the private interconnect. A write operation to the routing socket failed.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

929252 :Failed to start HA-NFS system fault monitor.

Process monitor facility has failed to start the HA-NFS system fault monitor.

Check whether the system is low in memory or the process table is full and correct these problems. If the error persists, use `scswitch` to switch the resource group to another node.

929712 :Share path %s: file system %s is not global.

The specified share path exists on a file system which is neither mounted via `/etc/vfstab` nor is a global file system.

Share paths in HA-NFS must satisfy this requirement.

930851 :ERROR: process_resource: resource <%s> is pending_fini but no FINI method is registered

A non-fatal internal error has occurred in the `rgmd` state machine.

Since this problem might indicate an internal logic error in the `rgmd`, please save a copy of the `/var/adm/messages` files on all nodes, the output of an `scstat -g` command, and the output of a `scrgadm -pvv` command. Report the problem to your authorized Sun service provider.

932126 :CCR: Quorum regained.

The cluster lost quorum at sometime after the cluster came up, and the quorum is regained now.

This is an informational message, no user action is needed.

936306 :svc_setschedprio: Could not setup RT (real time) scheduling parameters: %s

The server was not able to set the scheduling mode parameters, and the system error is shown. An error message is output to `syslog`.

Save the `/var/adm/messages` file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

937669 :CCR: Failed to update table %s.

The CCR data server failed to update the indicated table.

There may be other related messages on this node, which may help diagnose the problem. If the root file system is full on the node, then free up some space by removing unnecessary files. If the root disk on the afflicted node has failed, then it needs to be replaced.

938189 :rpc.fed: program not registered or server not started

The `rpc.fed` daemon was not correctly initialized, or has died. This has caused a step invocation failure, and may cause the node to reboot.

Check that the Solaris Clustering software has been installed correctly. Use `ps(1)` to check if `rpc.fed` is running. If the installation is correct, the reboot should restart all the required daemons including `rpc.fed`.

938618 :Couldn't create deleted subdir %s: error (%d)

While mounting this file system, PXFS was unable to create some directories that it reserves for internal use.

If the error is 28(ENOSPC), then mount this FS non-globally, make some space, and then mount it globally. If there is some other error, and you are unable to correct it, contact your authorized Sun service provider to determine whether a workaround or patch is available.

938836 :invalid value for parameter "%sfailfast": "%s". Using default value of "panic"

`/opt/SUNWudlm/etc/udlm.conf` did not have a valid entry for failfast mode. Default mode of "panic" will be used.

None.

939374 :CCR: Failed to access cluster repository during synchronization. ABORT node.

This node failed to access its cluster repository when it first came up in cluster mode and tried to synchronize its repository with other nodes in the cluster.

This is usually caused by an unrecoverable failure such as disk failure. There may be other related messages on this node, which may help diagnose the problem. If the root disk on the afflicted node has failed, then it needs to be replaced. If the root disk is full on this node, boot the node into non-cluster mode and free up some space by removing unnecessary files.

941367 :open failed: %s

Failed to open `/dev/console`. The "open" man page describes possible error codes.

None. `ucmmd` will exit.

941693 :"%s" Failed to stay up.

The tag shown, being run by the `rpc.pmfd` server, has exited. Either the user has decided to stop monitoring this process, or the process exceeded the number of retries. An error message is output to `syslog`.

This message is informational; no user action is needed.

942307 :NAFO group %s has status %s.

The specified NAFO group is not in functional state. Logical host resource can't be started without a functional NAFO.

LogicalHostname resource will not be brought online on this node. Check the messages(pnmd errors) that encountered just before this message for any NAFO or adapter problem. Correct the problem and rerun the scrgadm command.

944068 :clcomm: validate_policy: invalid relationship moderate
%d high %d pool %d

The system checks the proposed flow control policy parameters at system startup and when processing a change request. The moderate server thread level cannot be higher than the high server thread level.

No user action required.

944121 :Incorrect permissions set for %s

This file does not have the expected default execute permissions.

Reset the permissions to allow execute permissions using the chmod command.

948847 :ucm_callback for start_trans generated exception %d

ucmm callback for start transition failed. Step may have timedout.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

949148 :"%s" requeued

The tag shown has exited and was restarted by the rpc.pmfd server. An error message is output to syslog.

This message is informational; no user action is needed.

949565 :reservation error(%s) - do_scsi2_tkown() error for disk
%s

The device fencing program has encountered errors while trying to access a device. All retry attempts have failed.

The action which failed is a scsi-2 ioctl. These can fail if there are scsi-3 keys on the disk. To remove invalid scsi-3 keys from a device, use "scdidadm -R" to repair the disk (see scdidadm man page for details). If there were no scsi-3 keys present on the device, then this error is indicative of a hardware problem, which should be resolved as soon as possible. Once the problem has been resolved, the following actions may be necessary: If the message specifies the "node_join" transition, then this node may be unable to access the specified device. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access the device. In either case, access

can be reacquired by executing `"/usr/cluster/lib/sc/run_reserve -c node_join"` on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group may have failed to start on this node. If the device group was started on another node, it may be moved to this node with the `scswitch` command. If the device group was not started, it may be started with the `scswitch` command. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group may have failed. If so, the desired action may be retried.

949937 :Out of memory.

The dataservice has failed to allocate memory, most likely because the system has run out of swap space.

The problem will probably be cured by rebooting. If the problem reoccurs, you might need to increase swap space by configuring additional swap devices. See `swap(1M)` for more information.

950747 :resource %s monitor disabled.

This is a notification from the `rgmd` that the operator has disabled monitoring on a resource. This may be used by system monitoring tools.

This is an informational message, no user action is needed.

950760 :reservation fatal error(%s) - get_resv_lock() error

The device fencing program has suffered an internal error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available. Copies of `/var/adm/messages` from all nodes should be provided for diagnosis. It may be possible to retry the failed operation, depending on the nature of the error. If the message specifies the "node_join" transition, then this node may be unable to access shared devices. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access shared devices. In either case, it may be possible to reacquire access to shared devices by executing `"/usr/cluster/lib/sc/run_reserve -c node_join"` on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group has failed to start on this node. If another node was available to host the device group, then it should have been started on that node. If desired, it may be possible to switch the device group to this node with the `scswitch` command. If no other node was available, then the device group will not have been started. The `scswitch` command may be used to retry the attempt to start the device group. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group has failed. The desired action may be retried.

951501 :CCR: Could not initialize CCR data server.

The CCR data server could not initialize on this node. This usually happens when the CCR is unable to read its metadata entries on this node. There is a CCR data server per cluster node.

There may be other related messages on this node, which may help diagnose this problem. If the root disk failed, it needs to be replaced. If there was cluster repository corruption, then the cluster repository needs to be restored from backup or other nodes in the cluster. Boot the offending node in -x mode to restore the repository. The cluster repository is located at /etc/cluster/ccr/.

951634 :INTERNAL ERROR CMM: clconf_get_quorum_table() returned error %d.

The node encountered an internal error during initialization of the quorum subsystem object.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

951733 :Incorrect usage: %s.

The usage of the program was incorrect for the reason given.

Use the correct syntax for the program.

952237 :Method <%s>: unknown command.

An internal error has occurred in the interface between the rgmd and fed daemons. This in turn will cause a method invocation to fail. This should not occur and may indicate an internal logic error in the rgmd.

Look for other syslog error messages on the same node. Save a copy of the /var/adm/messages files on all nodes, and report the problem to your authorized Sun service provider.

952465 :HA: exception adding secondary

A failure occurred while attempting to add a secondary provider for an HA service.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

953642 :Server is not running. Calling shutdown abort to clear shared memory (if any)

Informational message. Oracle server is not running. However if Oracle processes are aborted without clearing shared memory, it can cause problems when starting Oracle server. Clearing leftover shared memory if any.

None

954497 :clcomm: Unable to find %s in name server

The specified entity is unknown to the name server.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

957086 :Prog <%s> failed to execute step <%s> - error=<%d>

ucmmd failed to execute a step.

Examine other syslog messages occurring at about the same time to see if the problem can be identified and if it recurs. Save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance.

958425 :clcomm: Cannot fork1() after Galileo initialization

A user level process attempted to fork1 after Galileo initialization. This is not allowed.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

958832 :INTERNAL ERROR: monitoring is enabled, but MONITOR_STOP method is not registered for resource <%s>

A non-fatal internal error has occurred in the rgmd state machine.

Since this problem might indicate an internal logic error in the rgmd, please save a copy of the /var/adm/messages files on all nodes, the output of an scstat -g command, and the output of a scrgadm -pvv command. Report the problem to your authorized Sun service provider.

958888 :clcomm: Failed to allocate simple xdoor client %d

The system could not allocate a simple xdoor client. This can happen when the xdoor number is already in use. This message is only possible on debug systems.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

959384 :Possible syntax error in hosts entry in %s.

Validation callback method has failed to validate the hostname list. There may be syntax error in the nsswitch.conf file.

Check for the following syntax rules in the nsswitch.conf file. 1) Check if the lookup order for "hosts" has "files". 2) "cluster" is the only entry that can come before "files". 3) Everything in between "[" and "]" is ignored. 4) It is illegal to have any leading whitespace character at the beginning of the line; these lines are skipped. Correct the syntax in the nsswitch.conf file and try again.

959610 :Property %s should have only one value.

A multi-valued (comma-separated) list was provided to the `scrgadm` command for the property, while the implementation supports only one value for this property.

Specify a single value for the property on the `scrgadm` command.

960308 :clcomm: Pathend %p: remove_path called twice

The system maintains state information about a path. The `remove_path` operation is not allowed in this state.

No user action is required.

960344 :ERROR: process_resource: resource <%s> is pending_init but no INIT method is registered

A non-fatal internal error has occurred in the `rgmd` state machine.

Since this problem might indicate an internal logic error in the `rgmd`, please save a copy of the `/var/adm/messages` files on all nodes, the output of an `scstat -g` command, and the output of a `scrgadm -pvv` command. Report the problem to your authorized Sun service provider.

960862 :(%s) sigaction failed: %s (UNIX errno %d)

The `udlm` has failed to initialize signal handlers by a call to `sigaction(2)`. The error message indicates the reason for the failure.

Save a copy of the `/var/adm/messages` files on all nodes, and of the `rgmd` core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

960932 :Switchover (%s) error: failed to fsck disk

The file system specified in the message could not be hosted on the node the message came from because an `fsck` on the file system revealed errors.

Unmount the `PXFS` file system (if mounted), `fsck` the device, and then mount the `PXFS` file system again.

961551 :Signal %d terminated the scalable service configuration process.

An unexpected signal caused the termination of the program that configures the networking components for a scalable resource. This premature termination will cause the scalable service configuration to be aborted for this resource.

Save a copy of the `/var/adm/messages` files on all nodes. If a core file was generated, submit the core to your service provider. Contact your authorized Sun service provider for assistance in diagnosing the problem.

962746 :Usage: %s [-c|-u] -R -T -G [-r sys_def_prop=values ...]
[-x ext_prop=values ...].

Incorrect arguments are passed to the callback methods.

This is an internal error. Contact your authorized Sun service provider for assistance in diagnosing and correcting the problem.

963465 :fatal: rpc_control() failed to set automatic MT mode;
aborting node

The rgmd failed in a call to rpc_control(3N). This error should never occur. If it did, it would cause the failure of subsequent invocations of scha_cmds(1HA) and scha_calls(3HA). This would most likely lead to resource method failures and prevent RGM reconfigurations from occurring. The rgmd will produce a core file and will force the node to halt or reboot.

Examine other syslog messages occurring at about the same time to see if the source of the problem can be identified. Save a copy of the /var/adm/messages files on all nodes and contact your authorized Sun service provider for assistance in diagnosing the problem. Reboot the node to restart the clustering daemons.

963755 :lkcm_cfg: caller is not registered

udlm is not registered with ucmm.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

964072 :Unable to resolve %s.

The data service has failed to resolve the host information.

If the logical host and shared address entries are specified in the /etc/inet/hosts file, check that these entries are correct. If this is not the reason, then check the health of the name server. For more error information, check the syslog messages.

964083 :t_open (open_cmd_port) failed

Call to t_open() failed. The "t_open" man page describes possible error codes. ucmmmd will exit and the node will abort.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

964399 :udlm seq no (%d) does not match library"s (%d).

Mismatch in sequence numbers between udlm and the library code is causing an abort.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

964521 :Failed to retrieve the resource handle: %s.

An API operation on the resource has failed.

For the resource name, check the syslog tag. For more details, check the syslog messages from other components. If the error persists, reboot the node.

965873 :CMM: Node %s (nodeid = %d) with vote count = %d added.

The specified node with the specified vote count has been added to the cluster.

This is an informational message, no user action is needed.

966245 :%d entries found in property %s. For a secure Netscape Directory Server instance %s should have one or two entries.

Since a secure Netscape Directory Server instance can listen on only one or two ports, the list property should have either one or two entries. A different number of entries was found.

Change the number of entries to be either one or two.

966416 :This list element in System property %s has an invalid protocol: %s.

The system property that was named does not have a valid protocol.

Change the value of the property to use a valid protocol.

966842 :in libsecurity unknown security flag %d

This is an internal error which shouldn't happen. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

967050 :Validation failed. Listener binaries not found
ORACLE_HOME=%s

Oracle listener binaries not found under ORACLE_HOME. ORACLE_HOME specified for the resource is indicated in the message. HA-Oracle will not be able to manage Oracle listener if ORACLE_HOME is incorrect.

Specify correct ORACLE_HOME when creating resource. If resource is already created, please update resource property "ORACLE_HOME".

967970 :Modification of resource <%s> failed because none of the nodes on which VALIDATE would have run are currently up

In order to change the properties of a resource whose type has a registered VALIDATE method, the rgmd must be able to run VALIDATE on at least one node. However, all of the candidate nodes are down. "Candidate nodes" are either members of the resource group's Nodelist or members of the resource type's Installed_nodes list, depending on the setting of the resource's Init_nodes property.

Boot one of the resource group's potential masters and retry the resource change operation.

968557 :Could not unplug any ip addresses.

Failed to unplug any ip addresses. The resource cannot be brought offline. Node will be rebooted by Sun cluster.

Check the syslog messages from other components for possible root cause. Save a copy of /var/adm/messages and contact Sun service provider for assistance in diagnosing and correcting the problem.

968853 :scha_resource_get error (%d) when reading system property %s

Error occurred in API call scha_resource_get.

Check syslog messages for errors logged from other system modules. Stop and start fault monitor. If error persists then disable fault monitor and report the problem.

969008 :t_alloc (open_cmd_port-T_ADDR) %d

Call to t_alloc() failed. The "t_alloc" man page describes possible error codes. ucmmmd will exit and the node will abort.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

969827 :Failover attempt has failed.

The failover attempt of the resource is rejected or encountered an error.

For more detailed error message, check the syslog messages. Check whether the Pingpong_interval has appropriate value. If not, adjust it using scrgadm(1M). Otherwise, use scswitch to switch the resource group to a healthy node.

970912 :execve: %s

The rpc.pmfd server was not able to exec a new process, possibly due to bad arguments. The message contains the system error. The server does not perform the action requested by the client, and an error message is output to syslog.

Investigate that the file path to be executed exists. If all looks correct, save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

971233 :Property %s is not set.

The property has not been set by the user and must be.

Reissue the scrgadm command with the required property and value.

971412 :Error in getting global service name for path <%s>

The path can not be mapped to a valid service name.

Check the path passed into extension property "ServicePaths" of SUNW.HAStorage type resource.

972580 :CCR: Highest epoch is < 0, highest_epoch = %d.

The epoch indicates the number of times a cluster has come up. It should not be less than 0. It could happen due to corruption in the cluster repository.

Boot the cluster in -x mode to restore the cluster repository on all the members of the cluster from backup. The cluster repository is located at /etc/cluster/ccr/.

972610 :fork: %s

The rgmd, rpc.pmfd or rpc.fed daemon was not able to fork a process, possibly due to low swap space. The message contains the system error. This can happen while the daemon is starting up (during the node boot process), or when executing a client call. If it happens when starting up, the daemon does not come up. If it happens during a client call, the server does not perform the action requested by the client.

Investigate if the machine is running out of swap space. If this is not the case, save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

973615 :Node %s: weight %d

The load balancer set the specified weight for the specified node.

This is an informational message, no user action is needed.

973933 :resource %s added.

This is a notification from the rgmd that the operator has created a new resource. This may be used by system monitoring tools.

This is an informational message, no user action is needed.

974106 :lkcm_parm: caller is not registered

udlm is not registered with ucmm.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

974664 :HA: no valid secondary provider in rmm - aborting

This node joined an existing cluster. Then all of the other nodes in the cluster died before the HA framework components on this node could be properly initialized.

This node must be rebooted.

976495 :fork failed: %s

Failed to run the "fork" command. The "fork" man page describes possible error codes.

Some system resource has been exceeded. Install more memory, increase swap space or reduce peak memory consumption.

978125 :in libsecurity setnetconfig failed when initializing the server: %s - %s

A server (rpc.pmfd, rpc.fed or rgmd) was not able to start because it could not establish a rpc connection for the network specified. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

978534 :%s: lookup failed.

Could not get the hostname for a node. This could also be because the node is not booted as part of a cluster.

Make sure the node is booted as part of a cluster.

978829 :t_bind, did not bind to desired addr

Call to t_bind() failed. The "t_bind" man page describes possible error codes. udlm will exit and the node will abort.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

979343 :Error: duplicate prog <%s> launched step <%s>

Due to an internal error, uccmd has attempted to launch the same step by duplicate programs. uccmd will reject the second program and treat it as a step failure.

Save a copy of the /var/adm/messages files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

979803 :CMM: Node being shut down.

This node is being shut down.

This is an informational message, no user action is needed.

980307 :reservation fatal error(%s) - Illegal command

The device fencing program has suffered an internal error.

Contact your authorized Sun service provider to determine whether a workaround or patch is available. Copies of /var/adm/messages from all nodes should be provided for diagnosis. It may be possible to retry the failed operation, depending on the nature of the error. If the message specifies the "node_join" transition, then this node may be unable to access shared devices. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access shared devices. In either case, it may be possible to reacquire access to shared devices by executing "/usr/cluster/lib/sc/run_reserve -c node_join" on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group has failed to start on this node. If another node was available to host the device group, then it should have been started on that node. If desired, it may be possible to switch the device group to this node with the scswitch command. If no other node was available, then the device group will not have been started. The scswitch command may be used to retry the attempt to start the device group. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group has failed. The desired action may be retried.

980477 :LogicalHostname online.

The status of the logicalhost resource is online.

This is informational message. No user action required.

980681 :clconf: CSR removal failed

While executing task in clconf and modifying the state of proxy, received failure from trying to remove CSR.

This is informational message. No user action required.

980942 :CMM: Cluster doesn't have operational quorum yet;
waiting for quorum.

Not enough nodes are operational to obtain a majority quorum; the cluster is waiting for more nodes before starting.

If nodes are booting, wait for them to finish booting and join the cluster. Boot nodes that are down.

981739 CCR: Updating invalid table %s.

This joining node carries a valid copy of the indicated table with override flag set while the current cluster membership doesn't have a valid copy of this table. This node will update its copy of the indicated table to other nodes in the cluster.

This is an informational message, no user action is needed.

981931 :INTERNAL ERROR: postpone_start_r: meth type <%d>

A non-fatal internal error has occurred in the rgmd state machine.

Since this problem might indicate an internal logic error in the rgmd, please save a copy of the /var/adm/messages files on all nodes, the output of an scstat -g command, and the output of a scrgadm -pvv command. Report the problem to your authorized Sun service provider.

983305 :clconf: Failed to open table infrastructure in unregister_callback

Failed to open table infrastructure in unregistered clconf callback with CCR. Table infrastructure not found.

Check the table infrastructure.

984704 :reset_rg_state: unable to change state of resource group <%s> on node <%d>; assuming that node died

The rgmd was unable to reset the state of the specified resource group to offline on the specified node, presumably because the node died.

Examine syslog output on the specified node to determine the cause of node death. The syslog output might indicate further remedial actions.

985111 :lkcm_reg: illegal %s value

Cluster information that is being used during udln registration with ucmm is incorrect.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

986190 :Entry at position %d in property %s with value %s is not a valid node identifier or node name.

The value given for the named property has an invalid node specified for it. The position index, which starts at 0 for the first element in the list, indicates which element in the property list was invalid.

Specify a valid node for the property.

986197 :reservation fatal error(%s) - malloc() error, errno %d

The device fencing program has been unable to allocate required memory.

Memory usage should be monitored on this node and steps taken to provide more available memory if problems persist. Once memory has been made available, the following steps may need to be taken: If the message specifies the "node_join" transition, then this node may be unable to access shared devices. If the failure occurred during the "release_shared_scsi2" transition, then a node which was joining the cluster may be unable to access shared devices. In either case, access to shared devices can be reacquired by executing `"/usr/cluster/lib/sc/run_reserve -c node_join"` on all cluster nodes. If the failure occurred during the "make_primary" transition, then a device group has failed to start on this node. If another node was available to host the device group, then it should have been started on that node. The device group can be switched back to this node if desired by using the `scswitch` command. If no other node was available, then the device group will not have been started. The `scswitch` command may be used to start the device group. If the failure occurred during the "primary_to_secondary" transition, then the shutdown or switchover of a device group has failed. The desired action may be retried.

986466 :clxecd: stat of "%s" failed

clxecd problem failed to stat the directory indicated in the error message.

Make sure the directory exists. If the problem persists, contact your authorized Sun service provider to determine whether a workaround or patch is available.

987455 :in libsecurity weak Unix authorization failed

A server (rgmd) refused an rpc connection from a client because it failed the Unix authentication. This happens if a caller program using scha public api, either in its C form or its CLI form, is not running as root. An error message is output to syslog.

Check that the calling program using the scha public api is running as root. If the program is running as root, save the `/var/adm/messages` file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

987601 :scvxvmlg error - opendir(%s) failed

The program responsible for maintaining the VxVM namespace was unable to access the global device namespace. If configuration changes were recently made to VxVM diskgroups or volumes, this node may be unaware of those changes. Recently created volumes may be inaccessible from this node.

Verify that the `/global/.devices/node@N` (N = this node's node number) is mounted globally and is accessible. If no configuration changes have been recently made to VxVM diskgroups or volumes and all volumes continue to be

accessible from this node, then no further action is required. If changes have been made, the device namespace on this node can be updated to reflect those changes by executing `"/usr/cluster/lib/dcs/scvxvmlg"`. If the problem persists, contact your authorized Sun service provider to determine whether a workaround or patch is available.

988416 :t_sndudata (2) in send_reply: %s

Call to `t_sndudata()` failed. The "t_sndudata" man page describes possible error codes.

None.

988719 :Warning: Unexpected result returned while checking for the existence of scalable service group %s: %d.

A call to the underlying scalable networking code failed.

Save a copy of the `/var/adm/messages` files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

989693 :thr_create failed

Could not create a new thread. The "thr_create" man page describes possible error codes.

Some system resource has been exceeded. Install more memory, increase swap space or reduce peak memory consumption.

989846 :ERROR: unpack_rg_seq(): rgname_to_rg failed <%s>

Due to an internal error, the `rgmd` was unable to find the specified resource group data in memory.

Save a copy of the `/var/adm/messages` files on all nodes. Contact your authorized Sun service provider for assistance in diagnosing the problem.

990215 :HA: repl_mgr: exception while invoking RMA reconf object

An unrecoverable failure occurred in the HA framework.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

990418 :received signal %d

The daemon indicated in the message tag (`rgmd` or `ucmmd`) has received a signal, possibly caused by an operator-initiated `kill(1)` command. The signal is ignored.

The operator must use `scswitch(1M)` and `shutdown(1M)` to take down a node, rather than directly killing the daemon.

991108 :uaddr2taddr (open_cmd_port) failed

Call to uaddr2taddr() failed. The "uaddr2taddr" man page describes possible error codes. ucmmmd will exit and the node will abort.

Save the files /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

991130 :pthread_create: %s

The rpc.pmfd server was not able to allocate a new thread, probably due to low memory, and the system error is shown. This can happen when a new tag is started, or when monitoring for a process is set up. If the error occurs when a new tag is started, the tag is not started and pmfadm returns error. If the error occurs when monitoring for a process is set up, the process is not monitored. An error message is output to syslog.

Investigate if the machine is running out of memory. If this is not the case, save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

991800 :in libsecurity transport %s is not a loopback transport

A server (rpc.pmfd, rpc.fed or rgmd) refused an rpc connection from a client because the named transport is not a loopback. An error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

991864 :putenv: %s

The rpc.pmfd server was not able to change environment variables. The message contains the system error. The server does not perform the action requested by the client, and an error message is output to syslog.

Save the /var/adm/messages file. Contact your authorized Sun service provider to determine whether a workaround or patch is available.

992912 :clexecd: thr_sigsetmask returned %d. Exiting.

clexecd program has encountered a failed thr_sigsetmask(3THR) system call. The error message indicates the error number for the failure.

Contact your authorized Sun service provider to determine whether a workaround or patch is available.

992998 :clconf: CSR registration failed

While executing task in clconf and modifying the state of proxy, received failure from registering CSR.

This is informational message. No user action required.

995026 :lkcm_cfg: invalid handle was passed %s %d

Handle for communication with udlmctl during a call to return the current DLM configuration is invalid.

This is an internal error. Save the contents of /var/adm/messages, /var/cluster/ucmm/ucmm_reconf.log and /var/cluster/ucmm/dlm*/logs/* from all the nodes and contact your Sun service representative.

995339 :Restarting using scha_control RESTART

Fault monitor has detected problems in RDBMS server. Attempt will be made to restart RDBMS server on the same node.

Check the cause of RDBMS failure.

996075 :fatal: Unable to resolve %s from nameserver

The low-level cluster machinery has encountered a fatal error. The rgmd will produce a core file and will cause the node to halt or reboot to avoid the possibility of data corruption.

Save a copy of the /var/adm/messages files on all nodes, and of the rgmd core file. Contact your authorized Sun service provider for assistance in diagnosing the problem.

996369 :PNM: successfully started

The PNM daemon has successfully started. Adapter monitoring and failover are enabled.

This message is informational; no user action is needed.

996897 :Method <%s> on resource <%s>: stat of program file failed.

The rgmd was unable to access the indicated resource method file. This may be caused by incorrect installation of the resource type.

Consult resource type documentation; [re-]install the resource type, if necessary.

996902 :Stopped the HA-NFS system fault monitor.

The HA-NFS system fault monitor was stopped successfully.

No action required.

997689 :IP address %s is an IP address in resource %s and in resource %s.

The same IP address is being used in two resources. This is not a correct configuration.

Delete one of the resources that is using the duplicated IP address.

998022 :Failed to restart the service: %s.

Restart attempt of the dataservice has failed.

Check the syslog messages that are occurred just before this message to check whether there is any internal error. In case of internal error, contact your Sun service provider. Otherwise, any of the following situations may have happened. 1) Check the Start_timeout and Stop_timeout values and adjust them if they are not appropriate. 2) This might be the result of lack of the system resources. Check whether the system is low in memory or the process table is full and take appropriate action.