Cisco Nexus 1000V InterCloud High Availability and Redundancy and Configuration Guide, Release 5.2(1)IC1(1.1)

First Published: June 28, 2013

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Text Part Number: OL-29144-01

CONTENTS

Preface

This preface contains the following sections:

Audience

This publication is for network administrators who configure and maintain Cisco Nexus devices.

This guide is for network and server administrators with the following experience and knowledge:

- An understanding of virtualization
- Using VMM software to create a virtual machine and configure a VMware vSwitch
- Ability to create an account on provider cloud such as Amazon Web Services (AWS).
- Knowledge of VMware vNetwork Distributed Switch is not required.

Document Conventions

Command descriptions use the following conventions:

Convention	Description
bold	Bold text indicates the commands and keywords that you enter literally as shown.
Italic	Italic text indicates arguments for which the user supplies the values.
[x]	Square brackets enclose an optional element (keyword or argument).

Convention	Description
[x | y]	Square brackets enclosing keywords or arguments separated by a vertical bar indicate an optional choice.
{x | y}	Braces enclosing keywords or arguments separated by a vertical bar indicate a required choice.
[x {y | z}]	Nested set of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element.
`variable`	Indicates a variable for which you supply values, in context where italics cannot be used.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.

Examples use the following conventions:

Convention	Description
`screen font`	Terminal sessions and information the switch displays are in screen font.
`boldface screen font`	Information you must enter is in boldface screen font.
italic screen font	Arguments for which you supply values are in italic screen font.
< >	Nonprinting characters, such as passwords, are in angle brackets.
[ ]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

This document uses the following conventions:

Note Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the manual.

Caution Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

Related Documentation for Cisco Nexus 1000V InterCloud

This section lists the documents used with the Cisco Nexus 1000V InterCloud and available on Cisco.com at the following URL:

http://www.cisco.com/en/US/partner/products/ps12904/tsd_products_support_series_home.html

General Information

Cisco Nexus 1000V InterCloud Release Notes

Install and Upgrade

Cisco Nexus 1000V InterCloud Installation Guide

Configuration Guides

Cisco Nexus 1000V InterCloud License Configuration Guide

Cisco Nexus 1000V InterCloud High Availability and Redundancy Configuration Guide

Cisco Nexus 1000V InterCloud Interface Configuration Guide

Cisco Nexus 1000V InterCloud Layer 2 Configuration Guide

Cisco Nexus 1000V InterCloud Port Profile Configuration Guide

Cisco Nexus 1000V InterCloud Security Configuration Guide

Cisco Nexus 1000V InterCloud System Management Configuration Guide

Reference Guides

Cisco Nexus 1000V InterCloud Command Reference

Cisco Nexus 1000V InterCloud Verified Scalability Reference

Cisco Nexus 1000V MIB Quick Reference

Troubleshooting and Alerts

Cisco Nexus 1000V Password Recovery Procedure

Cisco Nexus 1000V Documentation

Cisco Nexus 1000V for VMware vSphere Documentation

http://www.cisco.com/en/US/products/ps9902/tsd_products_support_series_home.html

Cisco Prime Network Services Controller Documentation

http://www.cisco.com/en/US/products/ps13213/tsd_products_support_series_home.html

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to nexus1k-docfeedback@cisco.com. We appreciate your feedback.

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation, at:

http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html

Subscribe to the *What's New in Cisco Product Documentation* as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS version 2.0.

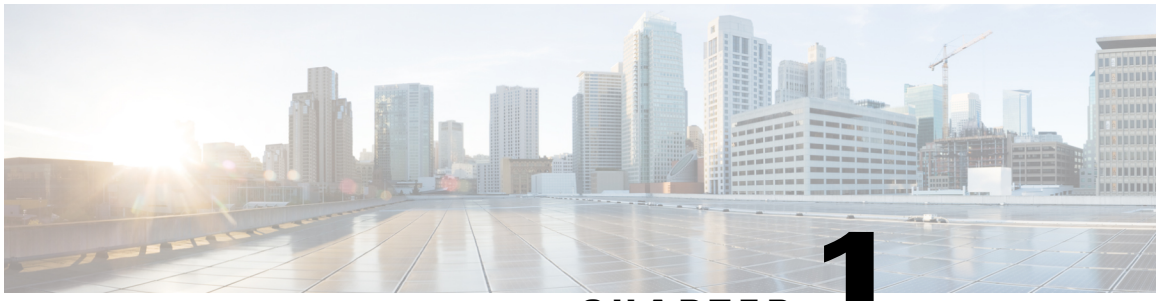

CHAPTER 1

Overview

This chapter contains the following sections:

Information About High Availability

The purpose of high availability (HA) is to limit the impact of failures—both hardware and software— within a system. The Cisco NX-OS operating system is designed for high availability at the network, system, and service levels.

The following Cisco NX-OS features minimize or prevent traffic disruption in the event of a failure:

- Redundancy—Redundancy at every aspect of the software architecture.
- Isolation of processes—Isolation between software components to prevent a failure within one process disrupting other processes.
- Restartability—Most system functions and services are isolated so that they can be restarted independently after a failure while other services continue to run. In addition, most system services can perform stateful restarts, which allow the service to resume operations transparently to other services.

System Components

The Cisco Nexus 1000V InterCloud consists of the following components:

- Cisco Prime Network Controller
- Cisco Nexus 1000V InterCloud Virtual Supervisor Module (VSMs)
- InterCloud Switch
- InterCloud Extender
- Virtual Ethernet modules (VEMs) that are represented as modules within the VSM
- A remote management component such as VMware vCenter Server.

See the *Cisco Nexus 1000V InterCloud Installation Guide* for more information about the various system components.

InterCloud Link

The Cisco Nexus 1000V InterCloud provides a system level high availability solution via redundant InterCloud Links. An InterCloud Link is a secure connection between an enterprise and public cloud. A secure Layer 2 tunnel connects the InterCloud Extender and InterCloud Switch, thereby extending the enterprise network into the cloud. The InterCloud Link is considered a single unit consisting of the InterCloud Extender in the enterprise and InterCloud Switch in the public cloud.

Service-Level High Availability

Isolation of Processes

The Cisco NX-OS software has independent processes, known as services, that perform a function or set of functions for a subsystem or feature set. Each service and service instance runs as an independent, protected process. This way of operating provides a highly fault-tolerant software infrastructure and fault isolation between services. A failure in a service instance does not affect any other services that are running at that time. Additionally, each instance of a service can run as an independent process, which means that two instances of a routing protocol can run as separate processes.

Process Restartability

Cisco NX-OS processes run in a protected memory space independently of each other and the kernel. This process isolation provides fault containment and enables rapid restarts. Process restartability ensures that process-level failures do not cause system-level failures. In addition, most services can perform stateful restarts. These stateful restarts allow a service that experiences a failure to be restarted and to resume operations transparently to other services within the platform and to neighboring devices within the network.

System Level High Availability

The Cisco Nexus InterCloud 1000V supports redundant InterCloud Links, a primary and a secondary, running as an HA pair. Dual InterCloud Links operate in an active/standby capacity in which only one of the InterCloud

Links is active at any given time, while the other acts as a standby backup. The InterCloud Links are configured as primary or secondary during the Cisco Nexus 1000V InterCloud installation.

Redundancy Manager

Redundancy Manager is the service on the Cisco Nexus 1000V InterCloud that manages the high availability feature within and among gateways to provide a system-level high availability solution. Redundancy Manager within each gateway communicates with its peer gateways to ensure the system is in a healthy state.

InterCloud Extender High Availability Handshake

In the Cisco Nexus 1000V InterCloud HA model, the InterCloud Link consisting of InterCloud Extender in the enterprise and the InterCloud Switch in a provider cloud are together considered a single unit. In an HA deployment, a second, standby InterCloud Link exists to minimize the impact of a failure on the active InterCloud Link.

In HA mode, two InterCloud Links are deployed in the system. During instantiation, the Cisco Prime Network Services Controller assigns an HA role (either primary or secondary) to each InterCloud Extender. Once configured with their local and peer information and the site-to-site tunnel has been established, the InterCloud Extenders perform a handshake over UDP port 9984. Upon initial handshake, the InterCloud Extenders use their HA roles to determine which should move to the active or standby states. Normally, the InterCloud Extender with the Primary HA role will move active and the InterCloud Extender with the Secondary HA role will move standby.

It is possible, however, that a Secondary InterCloud Extender moves active if it cannot communicate with the Primary InterCloud Extender during initial handshake. If this occurs, the Primary will move to the standby state once it has come up and performed the handshake with the Secondary InterCloud Extender.

If the active InterCloud Link fails, the standby InterCloud Link moves to active state and the failed InterCloud Link is rebooted and moves to the standby state.

InterCloud Extender High Availability Heartbeats

After the handshake has occurred between the primary and secondary InterCloud Extenders, they exchange heartbeats to share data and ensure the system is healthy. Similar to the handshake, the heartbeats are sent and received on UDP port 9984.

The heartbeats include useful information such as local and peer states, control flags, and tunnel status that allow the Redundancy Manager on each InterCloud Extender to make intelligent decisions regarding the health of the system as a whole.

The following intervals apply when sending heartbeat messages.

Interval	Description
5 seconds	Interval at which heartbeat requests are sent.
35 seconds	Interval after which missed heartbeats indicate degraded communication on the management interface so that heartbeats are also sent through the site-to-site secure tunnel and InterCloud Switches.

Interval	Description
300 seconds	Interval after which the standby InterCloud Link will reload should a WAN connectivity issue occur. This protects against the possibility of a failure occurring on both InterCloud Extenders that results in the false detection of a WAN connectivity issue.
240 seconds	Interval after which Tunnel Manager declares the site-to-site secure tunnel destroyed due to missed heartbeats. In standalone mode, both the InterCloud Extender and InterCloud Switch will be rebooted. In HA mode, a switchover will occur and the failed InterCloud Link will reboot and come back as standby.

Management and Tunnel Interface Redundancy

The active InterCloud Extender sends a heartbeat request to the standby InterCloud Extender who then sends a reply. If the standby InterCloud Extender does not receive a heartbeat request for 30 seconds, it will explicitly send a request to the active InterCloud Extender. If no response is received, the standby InterCloud Extender sends a heartbeat request through the InterCloud Switch in its InterCloud Link who then forwards to its HA peer InterCloud Switch in the active InterCloud Link and finally to the intended active InterCloud Extender.

If a response is received, the InterCloud Extenders will print logs describing the detection of a connectivity issue between InterCloud Extenders. If no response is received, the standby InterCloud Extender will initiate a switchover. As a result, the standby InterCloud Link will move to the active state and the failed InterCloud Link will be rebooted and come up in the standby state.

Note If the control connection between any of the gateways is broken, the redundant heartbeat mechanism will fail. If two redundant heartbeat requests fail across 5 seconds, the source InterCloud Extender will consider its HA peer InterCloud Extender failed.

Split Brain Resolution

When connectivity issues exist among InterCloud Extenders, Intercloud Switches, and the Cisco Prime Network Services Controller, it is possible that both InterCloud Extenders take the active role. This condition is called active-active or split-brain condition. When the communication is restored within the system, the InterCloud Extenders exchange information to decide which would have a lesser impact on the system, if rebooted.

A split-brain is not possible only due to InterCloud Extender connectivity loss because when the standby InterCloud Extender moves to active due to heartbeat failure, it sends a request to the Cisco Prime Network Services Controller to reboot the failed InterCloud Link. Once the failed InterCloud Link comes up it will move to the standby state.

A split-brain scenario is possible if there is a connectivity loss between the Cisco Prime Network Services Controller and the standby InterCloud Extender who is moving active. In this situation, the gateways in the failed InterCloud Link will not be rebooted, creating a active-active scenario.

If an active-active scenario occurs, the following parameters are considering during handshake resolution:

- Heartbeats Sent - The InterCloud Extender with a greater number of heartbeats sent within some threshold will remain active. If the difference in heartbeats sent is insignificant, the resolution will occur based on HA role.
- HA role - The InterCloud Extender with HA role Primary will remain active.

The InterCloud Link which moves to the standby state will be rebooted because many of the platform components do not support an active to standby state transition. The rebooted InterCloud Link will move to standby after it performs the handshake with the active InterCloud Extender.

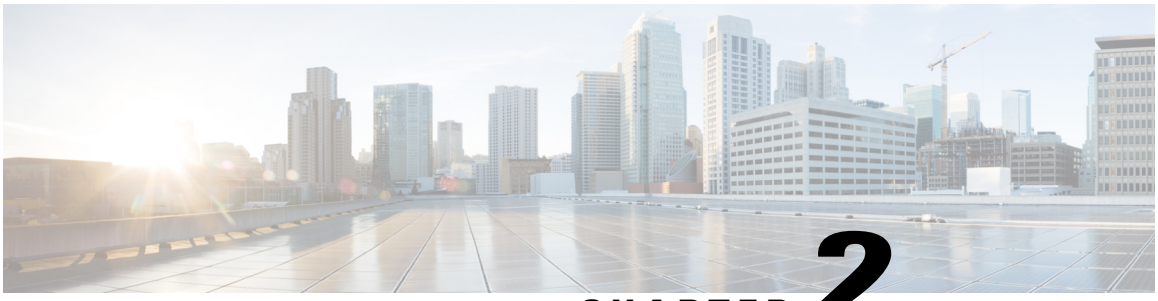

CHAPTER 2

Understanding Service-Level High Availability

This chapter contains the following sections:

Information About Cisco NX-OS Service Restarts

The Cisco NX-OS service restart features allow you to restart a faulty service without restarting the supervisor to prevent process-level failures from causing system-level failures. You can restart a service depending on current errors, failure circumstances, and the high-availability policy for the service. A service can undergo either a stateful or stateless restart. Cisco NX-OS allows services to store run-time state information and messages for a stateful restart. In a stateful restart, the service can retrieve this stored state information and resume operations from the last checkpoint service state. In a stateless restart, the service can initialize and run as if it had just been started with no prior state.

Restartability Infrastructure

Cisco NX-OS allows stateful restarts of most processes and services. The back-end management and orchestration of processes, services, and applications within a platform are handled by a set of high-level system-control services.

System Manager

The System Manager directs the overall system function, service management, and system health monitoring. The System Manager is responsible for launching, stopping, monitoring, and restarting services.

Persistent Storage Service

Cisco NX-OS services use the persistent storage service (PSS) to store and manage the operational run-time information and configuration of platform services. The PSS component works with system services to recover states in the event of a service restart. PSS functions as a database of state and run-time information, which allows services to make a checkpoint of their state information whenever needed. A restarting service can recover the last known operating state that preceded a failure, which allows for a stateful restart.

Each service that uses PSS can define its stored information as one of the following:

- Private—It can be read only by that service.
- Shared—The information can be read by other services.

The service can specify that it is one of the following:

- Local—The information can be read only by services on the same supervisor.
- Global—It can be read by services on either supervisor or on modules.

Message and Transaction Service

The message and transaction service (MTS) is a high-performance interprocess communications (IPC) message broker that specializes in high-availability semantics. MTS handles message routing and queuing between services on and across modules and between supervisors. MTS facilitates the exchange of messages such as event notification, synchronization, and message persistency between system services and system components. MTS can maintain persistent messages and logged messages in queues for access even after a service restart.

HA Policies

Cisco NX-OS allows each service to have an associated set of internal HA policies that define how a failed service will be restarted. Each service can have four defined policies—a primary and secondary policy when two supervisors are present, and a primary and secondary policy when only one supervisor is present. If no HA policy is defined for a service, the default HA policy to be performed upon a service failure will be a switchover if two supervisors are present or a supervisor reset if only one supervisor is present.

Each HA policy specifies three parameters:

- Action to be performed by the System Manager:
  - Stateful restart
  - Stateless restart
  - Supervisor switchover (or restart)
- Maximum retries—The number of restart attempts to be performed by the System Manager. If the service has not restarted successfully after this number of attempts, the HA policy is considered to have failed, and the next HA policy is used. If no other HA policy exists, the default policy is applied, which results in a supervisor switchover or restart.

- Minimum lifetime—The time that a service must run after a restart attempt in order to consider the restart attempt as successful. The minimum lifetime is no less than four minutes.

Process Restartability

Cisco NX-OS processes run in a protected memory space independently of each other and the kernel. This process isolation provides fault containment and enables rapid restarts. Process restartability ensures that process-level failures do not cause system-level failures. In addition, most services can perform stateful restarts. These stateful restarts allow a service that experiences a failure to be restarted and to resume operations transparently to other services within the platform and to neighboring devices within the network.

A failed service is restarted by different methods depending on the service's HA implementation and HA policies.

The following table describes the action taken by the System Manager for various failure conditions.

Failure	Action
Service/process exception	Service restart
Service/process crash	Service restart
Unresponsive service/process	Service restart
Repeated service failure	Supervisor reset (single) or switchover (dual)
Unresponsive System Manager	Supervisor reset (single) or switchover (dual)
Kernel failure	Supervisor reset (single) or switchover (dual)
Watchdog timeout	Supervisor reset (single) or switchover (dual)

Stateful Restarts

When a restartable service fails, it is restarted on the same supervisor. If the new instance of the service determines that the previous instance was abnormally terminated by the operating system, the service then determines whether a persistent context exists. The initialization of the new instance attempts to read the persistent context to build a run-time context that makes the new instance appear like the previous one. After the initialization is complete, the service resumes the tasks that it was performing when it stopped. During the restart and initialization of the new instance, other services are unaware of the service failure. Any messages that are sent by other services to the failed service are available from the MTS when the service resumes.

Whether or not the new instance survives the stateful initialization depends on the cause of the failure of the previous instance. If the service is unable to survive a few subsequent restart attempts, the restart is considered as failed. In this case, the System Manager executes the action specified by the service's HA policy, forcing either a stateless restart, no restart, or a supervisor switchover or reset.

During a successful stateful restart, there is no delay while the system reaches a consistent state. Stateful restarts reduce the system recovery time after a failure.

The events before, during, and after a stateful restart are as follows:

1. The running services make a checkpoint of their run-time state information to the PSS.
2. The System Manager monitors the health of the running services that use heartbeats.
3. The System Manager restarts a service instantly when it crashes or hangs
4. After restarting, the service recovers its state information from the PSS and resumes all pending transactions.
5. If the service does not resume a stable operation after multiple restarts, the System Manager initiates a reset or switchover of the supervisor.
6. Cisco NX-OS will collect the process stack and core for debugging purposes with an option to transfer core files to a remote location.

When a stateful restart occurs, Cisco NX-OS sends a syslog message of level LOG_ERR. If SNMP traps are enabled, the SNMP agent sends a trap.

Stateless Restarts

Cisco NX-OS infrastructure components manage stateless restarts. During a stateless restart, the System Manager identifies the failed process and replaces it with a new process. The service that failed does not maintain its run-time state upon the restart, so the service can either build the run-time state from the running configuration, or if necessary, exchange information with other services to build a run-time state.

When a stateless restart occurs, Cisco NX-OS sends a syslog message of level LOG_ERR. If SNMP traps are enabled, the SNMP agent sends a trap.

Troubleshooting Restarts

When a service fails, the system generates information that can be used to determine the cause of the failure. The following sources of information are available:

- Every service restart generates a syslog message of level LOG_ERR.
- If SNMP traps are enabled, the SNMP agent sends a trap when a service is restarted
- When a service failure occurs , the event is logged. To view the log, use the **show processes log** command in that module. The process logs are persistent across switchovers and resets.
- When a service fails, a system core image file is generated. You can view recent core images by entering the **show cores** command on the gateway. Core files are not persistent across switchovers and resets, but you can configure the system to export core files to an external server using a file transfer utility such as the Trivial File Transfer Protocol (TFTP).

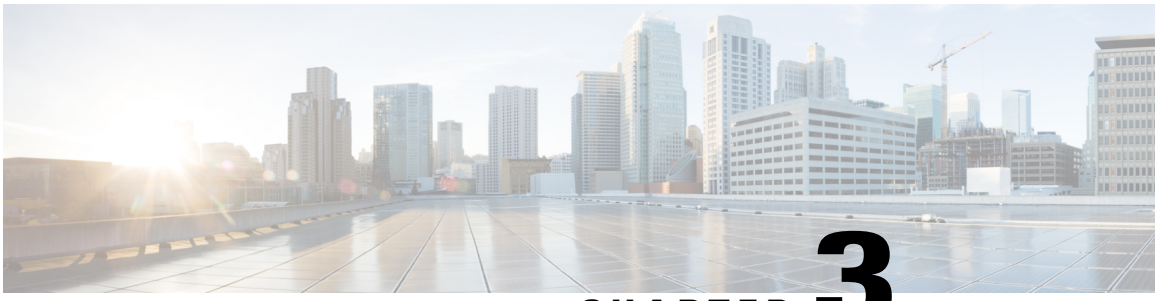

CHAPTER 3

Configuring System-Level High Availability

This chapter contains the following sections:

Information About System-Level High Availability

Configuring System-Level High Availability

Changing the InterCloud Link High Availability Deployment Mode

The High Availability feature can be enabled on a previously deployed standalone InterCloud Link. In doing so, the previously standalone InterCloud Link is given the HA role primary and remains active after the feature is enabled. A standby InterCloud Link is created and paired with the existing InterCloud Link.

Caution HA to standalone deployment mode change is not supported in this release.

Procedure

Step 1 Login to the Cisco Prime Network Services Controller.

Step 2 Navigate to the **InterCloud Management** > **InterCloud Link** .

Step 3 Click on VPCs and select the InterCloud Link to be modified.

Step 4 Check the **Enable HA** checkbox. In the Extend Network to Cloud wizard, enter the information to create the additional, standby InterCloud Link. See the *Cisco Prime Network Services Controller 3.0 User Guide* for more information.
Once completed, the standby InterCloud Link will deploy and HA will be enabled.

High Availability Commands

Use one of the following commands to display High Availability related information or perform a switchover:

Command	Source	Purpose
show intercloud redundancy configuration	InterCloud Switch InterCloud Extender	Display the HA configuration and status of the local system and remote peers.
show intercloud redundancy statistics heartbeat	InterCloud Extender	Display handshake and heartbeat related statistics between ICX's in an HA deployment.
show intercloud redundancy statistics platform	InterCloud Switch InterCloud Extender	Display platform related statistic for a gateway.
intercloud redundancy switchover	Active InterCloud Extender (HA)	Initiate a switchover in HA deployment.
show processes	InterCloud Switch InterCloud Extender Cisco Nexus 1000V InterCloud VSM	Displays the state of all processes and the start count of all the processes.
show module service-module	Cisco Nexus 1000V InterCloud VSM	Displays information about available InterCloud Extender's and InterCloud Switche's in the system.

show intercloud redundancy configuration

```
switch# show intercloud redundancy config

Redundancy Manager Information:
```

```
Cluster Node Count: 3

Local Node:
state : Active
HA mode : High Availability
uuid : A65AD6DC-D80F-0D5D-3341-58AEA0D0938C
version : 521111
cluster_id : 4
priority : Secondary
type : InterCloud Extender
ipaddr [mgmt] : 10.193.73.171

Tunnel Peer:
state : Active
uuid : 92566D7A-E0BD-9977-84E9-78037EE4BC94
type : InterCloud Switch
ipaddr [public]: 107.21.132.239

HA Peer:
state : Standby
uuid : AC5A0B56-51CF-397A-7529-CC0920BC87A3
type : InterCloud Extender
ipaddr [mgmt] : 10.193.73.174
```

show intercloud redundancy statistics heartbeat

```
switch# show intercloud redundancy statistics heartbeat

HA Manager Heartbeat Stats:

Heartbeat Frequency (s) : 5
Heartbeat Timeout (s) : 30

rx_handshake_pkts : 2
tx_handshake_pkts : 4

rx_heartbeat_pkts : 143
tx_heartbeat_pkts : 143

rx_drops_invalid_src_addr : 0
rx_drops_wrong_cluster : 0
rx_drops_queue_full : 0
rx_drops_not_ready : 0
rx_drops_wrong_version : 0
rx_unknown_pkts : 0

WAN Timeout (s) : 300
WAN HB Count : 0
```

show intercloud redundancy statistics platform

```
switch# show intercloud redundancy statistics platform

HA Manager Platform Stats:

rx_cncl : 2
rx_cncl_inval : 0
tx_cncl : 2
tx_cncl_err : 0

rx_cnc_state_push_req : 0
rx_cnc_state_push_req_inval : 0
rx_cnc_state_push_rsp : 2
rx_cnc_state_push_rsp_inval : 0
tx_cnc_state_push_req : 2
tx_cnc_state_push_req_err : 0
tx_cnc_state_push_rsp : 0
tx_cnc_state_push_rsp_err : 0
```

```
tx_cnc_state_push_req_timeouts: 0
```

High Availability Troubleshooting

Configuring Redundancy Manager Event-Logs

The Redundancy Manager event-logs can be configured using the following command:

Command	Source	Purpose
[**no**] **event-log redundancy-mgr** { **trace** | **info** | **error** } [**terminal**]	InterCloud Switch InterCloud Extender	Configures the Redundancy Manager event logs.

There are three levels of event-logs: Trace, Info, and Error (most critical). Info and Error event logs are enabled by default.

The optional parameter **terminal** will display the event-logs in real time on the terminal.

event-log redundancy-mgr info terminal

```
switch# event-log redundancy-mgr info terminal
switch# event-log redundancy-mgr error terminal

switch# Thu Jun 27 14:37:25 2013 90000 usec hamgr_config_cg(394):Received configuration
from PA
Thu Jun 27 14:37:25 2013 90000 usec hamgr_config_cg(410):Received cgu_info configuration
Thu Jun 27 14:37:25 2013 90000 usec hamgr_cguinfo_to_nodeinfo(772):cgu_info: uuid
[AC5A0B56-51CF-397A-7529-CC0
920BC87A3], cluster_id [4], ha_role [1], opcode [1] , ip [10.193.73.174]
Thu Jun 27 14:37:25 2013 90000 usec hamgr_cginfo_process_action(538):Received action ADD
Thu Jun 27 14:37:25 2013 90000 usec hamgr_node_add(404):Using MGMT IP address from config
Thu Jun 27 14:37:25 2013 90000 usec hamgr_node_add(407):Local node configured successfully
Thu Jun 27 14:37:25 2013 90000 usec hamgr_pss_config_write(556):Writing pss config for info
 type HAMGR_NODE_INFO_LOCAL
Thu Jun 27 14:37:25 2013 90000 usec hamgr_pss_config_write(585):node_config pss does not
exist, creating uri volatile:/dev/shm/hamgr_node_configs
Thu Jun 27 14:37:25 2013 90000 usec hamgr_sockets_update(230):Updating sockets if needed...
Thu Jun 27 14:37:25 2013 90000 usec hamgr_sockets_update(271):Sockets not enabled, tunnel
must be up
Thu Jun 27 14:37:25 2013 100000 usec hamgr_config_cg(394):Received configuration from PA
Thu Jun 27 14:37:25 2013 100000 usec hamgr_config_cg(436):Received peer_info configuration
Thu Jun 27 14:37:25 2013 100000 usec hamgr_peerinfo_to_nodeinfo(909):Printing peer_info
Thu Jun 27 14:37:25 2013 100000 usec hamgr_peerinfo_to_nodeinfo(911):peer_info : uuid :
A65AD6DC-D80F-0D5D-3341-58AEA0D0938C
Thu Jun 27 14:37:25 2013 100000 usec hamgr_peerinfo_to_nodeinfo(912):peer_info : ip_addr
(int) : 2873737482
Thu Jun 27 14:37:25 2013 100000 usec hamgr_peerinfo_to_nodeinfo(914):peer_info : ip_addr
(str) : 10.193.73.171
Thu Jun 27 14:37:25 2013 110000 usec hamgr_peerinfo_to_nodeinfo(915):peer_info : opcode :
1
Thu Jun 27 14:37:25 2013 110000 usec hamgr_cginfo_process_action(538):Received action ADD
Thu Jun 27 14:37:25 2013 110000 usec hamgr_node_list_add(744):Peer node added successfully
Thu Jun 27 14:37:25 2013 110000 usec hamgr_pss_config_write(556):Writing pss config for
info type HAMGR_NODE_INFO_PEER
```

Displaying Redundancy Manager Event-Logs

The Redundancy Manager event-logs can be displayed using the following command:

Command	Source	Purpose
show system internal event-log redundancy-mgr	InterCloud Switch InterCloud Extender	Displays the Redundancy Manager event logs.

Note Since event logs are stored in a ring buffer, older logs may be overwritten by newer logs.

show system internal event-log redundancy-mgr

```
switch# show system internal event-log redundancy-mgr

1) Event:E_MTS_RX, length:60, at 100000 usecs after Thu Jun 27 07:42:27 2013
[REQ] Opc:MTS_OPC_SDWRAP_DEBUG_DUMP(1530), Id:0X00006803, Ret:SUCCESS
Src:0x00000101/2682, Dst:0x00000101/1240, Flags:None
HA_SEQNO:0X00000000, RRtoken:0x00006803, Sync:UNKNOWN, Payloadsize:216
Payload:
0x0000: 01 00 2f 74 6d 70 2f 64 62 67 64 75 6d 70 35 33


2) Event:E_MTS_RX, length:60, at 600000 usecs after Thu Jun 27 07:42:11 2013
[REQ] Opc:MTS_OPC_SDWRAP_DEBUG_DUMP(1530), Id:0X000060E6, Ret:SUCCESS
Src:0x00000101/2678, Dst:0x00000101/1240, Flags:None
HA_SEQNO:0X00000000, RRtoken:0x000060E6, Sync:UNKNOWN, Payloadsize:216
Payload:
0x0000: 01 00 2f 74 6d 70 2f 64 62 67 64 75 6d 70 35 33
```

Feature History for System-Level High Availability

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

Feature Name	Releases	Feature Information
System -Level High Availability	5.2(1)IC1(1.1)	This feature was introduced.