Inspur

CN12900 Series

INOS-CN High Availability and Redundancy Guide

Inspur-Cisco Networking Technology Co.,Ltd. provides customers with comprehensive technical support and services. For any assistance, please contact our local office or company headquarters.
Website: http://www.inspur.com/
Technical Support Tel: 400-691-1766
Technical Support Email:icnt_service@inspur.com
Technical Document Support Email: icnt_service@inspur.com
Address: 1036 Langchao Road, Lixia District, Jinan City, Shandong Province
Postal code: 250101

Preface

Objectives

This guide describes main functions of the CN12900 Series. To have a quick grasp of the CN12900 Series, please read this manual carefully.

Versions

The following table lists the product versions related to this document.

Product name	Version
CN12900 Series	

Conventions

Symbol conventions

The symbols that may be found in this document are defined as follows.

Symbol	Description
Warning	Indicates a hazard with a medium or low level of risk which, if not avoided, could result in minor or moderate injury.
Caution	Indicates a potentially hazardous situation that, if not avoided, could cause equipment damage, data loss, and performance degradation, or unexpected results.
Note	Provides additional information to emphasize or supplement important points of the main text.
Tip	Indicates a tip that may help you solve a problem or save time.

General conventions

Convention	Description
Boldface	Names of files, directories, folders, and users are in **boldface**. For example, log in as user **root**.
Italic	Book titles are in *italics*.

Convention	Description
`Lucida Console`	Terminal display is in `Lucida Console`.

Command conventions

Convention	Description
Boldface	The keywords of a command line are in **boldface**.
Italic	Command arguments are in *italics*.
[]	Items (keywords or arguments) in square brackets [ ] are optional.
{ x | y | ... }	Alternative items are grouped in braces and separated by vertical bars. One is selected.
[x | y | ...]	Optional alternative items are grouped in square brackets and separated by vertical bars. One or none is selected.
{ x | y | ... } *	Alternative items are grouped in braces and separated by vertical bars. A minimum of one or a maximum of all can be selected.
[x | y | ...] *	The parameter before the & sign can be repeated 1 to n times.

GUI conventions

Convention	Description
Boldface	Buttons, menus, parameters, tabs, windows, and dialog titles are in **boldface**. For example, click **OK**.
>	Multi-level menus are in boldface and separated by the ">" signs. For example, choose **File** > **Create** > **Folder**.

Keyboard operation

Format	Description
Key	Press the key. For example, press **Enter** and press **Tab**.
Key 1+Key 2	Press the keys concurrently. For example, pressing **Ctrl+C** means the two keys should be pressed concurrently.
Key 1, Key 2	Press the keys in turn. For example, pressing **Alt**, **A** means the two keys should be pressed in turn.

Mouse operation

Action	Description
Click	Select and release the primary mouse button without moving the pointer.
Double-click	Press the primary mouse button twice continuously and quickly without moving the pointer.
Drag	Press and hold the primary mouse button and move the pointer to a certain position.

Change history

Updates between document versions are cumulative. Therefore, the latest document version contains all updates made to previous versions.

Issue 01 (2020-02-24)

Initial commercial release

Contents

Table

Chapter 1 Overview

Inspur INOS-CN is a resilient operating system that is specifically designed for high availability at the network, system, and process level.

This chapter describes high-availability (HA) concepts and features for Inspur INOS-CN devices and includes the following sections:

•About High Availability

•Service-Level High Availability

•System-Level High Availability

•Network-Level High Availability

•Additional Management Tools for Availability

•Software Image

•Virtual Device Contexts

1.1 About High Availability

To prevent or minimize traffic disruption during hardware or software failures, Inspur INOS-CN has these features:

•Redundancy—Inspur INOS-CN HA provides physical and software redundancy at every component level, spanning across the physical, environmental, power, and system software aspects of its architecture.

•Isolation of planes and processes—Inspur INOS-CN HA provides isolation between control and data forwarding planes within the device and between software components so that a failure within one plane or process does not disrupt others.

•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.

•Supervisor stateful switchover—The Inspur CN12904 and CN12908 chassis support an active and standby dual supervisor configuration. State and configuration remain constantly synchronized between the two supervisor modules to provide seamless and stateful switchover in the event of a supervisor module failure.

1.2 Service-Level High Availability

Inspur INOS-CN has a modularized architecture that compartmentalizes components for fault isolation, redundancy, and resource efficiency.

1.2.1 Isolation of Processes

In the Inspur INOS-CN software,independent processes, known as services,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 approach provides a highly fault-tolerant software infrastructure and fault isolation between services.A failure in a service instance (such as BGP) does not affect any other services that are running at that time (such as the Link Aggregation Control Protocol [LACP]). In addition,each instance of a service can run as an independent process, which means that two instances of a routing protocol (for example, two instances of the

Open Shortest Path First [OSPF] protocol) can run as separate processes.

1.2.2 Process Restartability

Inspur INOS-CN processes run in a protected memory space independently from 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, which 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.

1.3 System-Level High Availability

The Inspur CN12900 Series switches are protected from system failure by redundant hardware components and a high-availability software framework.

1.3.1 Physical Redundancy

The Inspur CN12900 Series switches have the following physical redundancies:

•Power Supply Redundancy—To provide redundant power input to the chassis, the Inspur CN12900 Series switches support the following number of power supply modules:

Inspur CN12900 Series Switches	**Maximum Number of Supported Power Supply Modules**
CN12904 switch	4
CN12908 switch	8

•Fan Tray Redundancy—For cooling the system, the Inspur CN12900 Series switches support the following number of fan trays:

•System Controller Redundancy—A pair of redundant system controllers in the Inspur CN12900

Inspur CN12900 Series switches	Maximum Number of Supported Fan Trays
CN12904 and CN12908 switches	3

Series chassis offloads chassis management functions from the supervisor modules. You can have two of the same type or you can mix as follows:

Active	**Standby**	**Ok?**
A	A	Yes
B	B	Yes
A	A+	Yes
B	B+	Yes
A	B	Not unless A is able to failover to B
B	A	Not unless A is able to failover to B
A+	B+	Not unless A+ is able to failover to B+
B+	A+	Not unless A+ is able to failover to B+

Note: Supervisor A and A+ are not supported for CN12904-FM-R fabric modules.

•Supervisor Module Redundancy—The Inspur CN12900 Series chassis support dual supervisor modules to provide redundancy for the control and management plane.

1.4 Network-Level High Availability

Network convergence is optimized by providing tools and functions to make both failover and fallback transparent and fast.

1.4.1 Layer 2 HA Features

Inspur INOS-CN provides these Layer 2 HA features:

•Spanning Tree Protocol (STP) enhancements, such as Bridge Protocol Data Unit (BPDU) Guard, Loop Guard, Root Guard, BPDU Filters, and Bridge Assurance, to guarantee the health of the STP control plane

•Unidirectional Link Detection (UDLD) Protocol

•IEEE 802.3ad link aggregation

Note: Virtual port channels (vPCs) allow you to create redundant physical links between two systems that act as a logical single link.

1.4.2 Layer 3 HA Features

Inspur INOS-CN provides these Layer 3 HA features:

•Nonstop forwarding (NSF) graceful restart extensions for routing protocols

Open Shortest Path First version 2 (OSPFv2), OSPFv3, Intermediate System to Intermediate System (IS-IS), Enhanced Interior Gateway Routing Protocol (EIGRP), and Border Gateway Protocol (BGP) utilize graceful restart extensions to the base protocols to provide nonstop forwarding and least obtrusive routing recovery for those environments.

•Shortest Path First (SPF) optimizations such as link-state advertisement (LSA) pacing and incremental SPF

•Protocol-based periodic refresh

•Millisecond timers for First-Hop Redundancy Protocols (FHRPs) such as the Hot Standby Router Protocol (HSRP) and the Virtual Router Redundancy Protocol (VRRP)

For more information on these Layer 3 routing protocols, see the Inspur CN12900 Series INOS-CN Unicast Routing Configuration Guide.

1.5 Additional Management Tools for Availability

Inspur INOS-CN incorporates several system management tools for monitoring and notification of system availability events.

1.5.1 EEM

Embedded Event Manager (EEM) consists of Event Detectors, the Event Manager, and an Event Manager Policy Engine. Using EEM, you can define policies to take specific actions when the system software recognizes certain events through the Event Detectors. The result is a flexible set of tools to automate many network management tasks and to direct the operation of Inspur INOS-CN to increase availability, collect information, and notify external systems or personnel about critical events.

For information about configuring EEM, see the Inspur CN12900 Series INOS-CN System Management Configuration Guide.

1.5.2 Smart Call Home

Combining GOLD and EEM capabilities, Smart Call Home provides an e-mail-based notification of critical system events. Smart Call Home has message formats that are compatible with pager services, standard e-mail, or XML-based automated parsing applications. You can use this feature to page a network support engineer, e-mail a network operations center.

For information about configuring Smart Call Home, see the Inspur CN12900 Series INOS-CN System Management Configuration Guide

1.6 Virtual Device Contexts

Inspur INOS-CN can segment operating system and hardware resources into virtual device contexts (VDCs) that emulate virtual devices. The Inspur CN12900 Series switches currently do not support multiple VDCs. All switch resources are managed in the default VDC. Service-Level High Availability

This chapter describes the Inspur INOS-CN service restartability for service-level high availability (HA) and includes the following sections:

•About Inspur INOS-CN Service Restarts
•Licensing Requirements
•Restartability Infrastructure
•Process Restartability
•Restarts on Standby Supervisor Services
•Restarts on Switching Module Services
•Troubleshooting Restarts
•Additional References for Service-Level High Availability

1.6.1 About Inspur INOS-CN Service Restarts

The Inspur INOS-CN 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. Inspur INOS-CN 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.

Not all services are designed for a stateful restart. For example, Inspur INOS-CN does not store run-time state information for Layer 3 routing protocols (such as Open Shortest Path First [OSPF] and Routing Information Protocol [RIP]). Their configuration settings are preserved across a restart, but these protocols are designed to rebuild their operational state using information obtained from neighbor routers.

Note: This chapter refers to processes and services interchangeably. A process is considered to be a running instance of a service.

1.6.2 Licensing Requirements

Product	License Requirement
Inspur INOS-CN	The service-level high availability features require no license. Any feature not included in a license package is bundled with the image and is provided for free.

1.6.3 Restartability Infrastructure

Inspur INOS-CN allows stateful restarts of most processes and services. 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 overall system function, service management, and system health monitoring and enforces high-availability policies. The System Manager is responsible for launching, stopping, monitoring, and restarting services as well as initiating and managing the synchronization of service states and supervisor states for a stateful switchover.

Persistent Storage Service

Inspur INOS-CN services use the persistent storage service (PSS) to store and manage the operational run-time information. 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 private (it can be read only by that service) or shared (the information can be read by other services). If the information is shared, the service can specify that it is local (the information can be read only by services on the same supervisor) or global (it can be read by services on either supervisor or on modules). For example, if the PSS information of a service is defined as shared and global, services on other modules can synchronize with the PSS information of the service that runs on the active supervisor.

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

Inspur INOS-CN allows each service to have an associated set of internal HA policies that define how a failed service is 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 is 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—Specifies 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, resulting in a supervisor switchover or restart.

•Minimum lifetime—Specifies the time that a service must run after a restart attempt to consider the restart attempt as successful. The minimum lifetime is no less than 4 minutes.

1.6.4 Process Restartability

Process restartability ensures that a failed service can recover and resume operations without disrupting the data plane or other services. Depending on the service HA policies, previous restart failures, and the health of other services running on the same supervisor, the System Manager determines the action to be taken when a service fails.

The action taken by the System Manager for various failure conditions is described in the following table.

Table 2 **System Manager Action for Various Failure Cases**

Failure	
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)
Supervisor hardware failure	Supervisor reset (single) or switchover (dual)
Kernel failure	Supervisor reset (single) or switchover (dual)
Watchdog timeout	Supervisor reset (single) or switchover (dual)

Types of Process Restarts

A failed service is restarted by one of the methods described in this section, depending on the service's HA implementation and HA policies,

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 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 performs the action specified by the HA policy of the services, 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.Inspur INOS-CN collects the process stack and core for debugging purposes with an option to transfer core files to a remote location.

When a stateful restart occurs, Inspur INOS-CN sends a syslog message of level LOG_ERR. If SNMP traps are enabled, the SNMP agent sends a trap. If the Smart Call Home service is enabled, the service sends an event message.

Stateless Restarts

Inspur INOS-CN 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. 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, Inspur INOS-CN sends a syslog message of level LOG_ERR. If SNMP traps are enabled, the SNMP agent sends a trap. If the Smart Call Home service is enabled, the service sends an event message.

Switchovers

If a standby supervisor is available, Inspur INOS-CN performs a supervisor switchover rather than a supervisor restart whenever multiple failures occur at the same time, because these cases are considered unrecoverable on the same supervisor. For example, if more than one HA application fails, that is considered an unrecoverable failure.

1.6.5 Restarts on Standby Supervisor Services

When a service fails on a supervisor that is in the standby state, the System Manager does not apply the HA policies and restarts the service after a delay of 30 seconds. The delay ensures that the active supervisor is not overloaded by repeated standby service failures and synchronizations. If the service being restarted requires synchronization with a service on the active supervisor, the standby supervisor is taken out of hot standby mode until the service is restarted and synchronized. Services that are not restartable cause the standby supervisor to reset.

When a standby service restart occurs, Inspur INOS-CN sends a syslog message of level LOG_ERR. If SNMP traps are enabled, the SNMP agent sends a trap. If the Smart Call Home service is enabled, the service sends an event message.

1.6.6 Restarts on Switching Module Services

When services fail on a switching module or another nonsupervisor module, the recovery action is determined by HA policies for those services. Because service failures on nonsupervisor module services do not require a supervisor switchover, the recovery options are a stateful restart, a stateless restart, or a module reset.

When a nonsupervisor module service restart occurs, Inspur INOS-CN sends a syslog message of level LOG_ERR. If SNMP traps are enabled, the SNMP agent sends a trap. If the Smart Call Home service is enabled, the service sends an event message.

1.6.7 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 the Smart Call Home service is enabled, every service restart generates a Smart Call Home event.

•If SNMP traps are enabled, the SNMP agent sends a trap when a service is restarted.

•When a service failure occurs on a local module, you can view a log of the event by entering the **show processes log** command in that module. The process logs are persistent across supervisor switchovers and resets.

•When a service fails, a system core image file is generated. You can view recent core images by using the **show cores** command on the active supervisor. Core files are not persistent across supervisor switchovers and resets, but you can configure the system to export core files to an external server by using a file transfer utility such as the Trivial File Transfer Protocol (TFTP).

1.6.8 Additional References for Service-Level High Availability

This section describes additional information related to service-level high availability.

Related Documents

Related Topic	Document Title
Inspur INOS-CN fundamentals	Inspur CN12900 Series INOS-CN Fundamentals Configuration Guide

Network-Level High Availability

This chapter describes Inspur INOS-CN network high availability (HA) and includes the following sections:

•About Network-Level High Availability

•Licensing Requirements

•Spanning Tree Protocol

•Virtual Port Channels

•First-Hop Redundancy Protocols

•Nonstop Forwarding in Routing Protocols

•Additional References for Network-Level High Availability

2.1 About Network-Level High Availability

Network-level HA is optimized by tools and functionality that provide failovers and fallbacks transparently and quickly. The features described in this chapter ensure high availability at the network level.

2.2 Licensing Requirements

Product	License Requirement
Inspur INOS-CN	The network-level high availability features require no license. Any feature not included in a license package is bundled with the software image and is provided for free.
BGP	Border Gateway Protocol (BGP) requires the Essentials license package (CN129-ES-M4 and CN129-ES-M8).

2.3 Spanning Tree Protocol

Note: The Spanning Tree Protocol (STP) refers to IEEE 802.1w and IEEE 802.1s. If this publication is referring to the IEEE 802.1D STP, 802.1D is stated specifically.

When you create fault-tolerant internetworks, you must have a loop-free path between all nodes in a network. Multiple active paths between end stations cause loops in the network that result in network devices learning end station MAC addresses on multiple Layer 2 LAN ports. This condition can result in a broadcast storm, which creates an unstable network.

STP provides a loop-free network at the Layer 2 level. Layer 2 LAN ports send and receive STP frames, which are called Bridge Protocol Data Units (BPDUs), at regular intervals. Network devices do not forward these frames but use the frames to determine the network topology and to construct a loop-free path within that topology. Using the spanning tree topology, STP forces redundant data paths into a blocked state. If a network segment in the spanning tree fails and a redundant path exists, the STP algorithm recalculates the spanning tree topology and activates the blocked path.

Inspur INOS-CN also supports the Multiple Spanning Tree Protocol (MSTP). The multiple independent spanning tree topology enabled by MSTP provides multiple forwarding paths for data traffic, enables load balancing, and reduces the number of STP instances required to support a large number of VLANs.

MST incorporates the Rapid Spanning Tree Protocol (RSTP), which allows rapid convergence. MST improves the fault tolerance of the network because a failure in one instance (forwarding path) does not affect other instances (forwarding paths).

Note: You can configure spanning tree parameters only on Layer 2 interfaces; a spanning tree configuration is not allowed on a Layer 3 interface. For information on creating Layer 2 interfaces, see the Inspur CN12900 Series INOS-CN Interfaces Configuration Guide.

For details about STP behavior and configuration, see the Inspur CN12900 Series INOS-CN Layer -2 Configuration Guide.

2.4 Virtual Port Channels

The major limitation in classic port channel communication is that the port channel operates only between two devices. In large networks, the support of multiple devices together is often a design requirement to provide some form of hardware failure alternate path. This alternate path is often connected in a way that would cause a loop, limiting the benefits gained with port channel technology to a single path. To address this limitation, Inspur INOS-CN provides a technology called virtual port channel (vPC). Although a pair of switches acting as a vPC peer endpoint looks like a single logical entity to port channel-attached devices, the two devices that act as the logical port channel endpoint are still two separate devices. This environment combines the benefits of hardware redundancy with the benefits of port channel loop management.

For more information on vPCs, see the see the Inspur CN12900 Series INOS-CN Interfaces Configuration Guide.

2.5 First-Hop Redundancy Protocols

Within a group of two or more routers, first-hop redundancy protocols (FHRPs) allow a transparent failover of the first-hop IP router. Inspur INOS-CN supports the following FHRPs:

•Hot Standby Router Protocol (HSRP)—HSRP provides first-hop routing redundancy for IP hosts on Ethernet networks configured with a default gateway IP address. An HSRP router group of two or more routers chooses an active gateway and a standby gateway. The active gateway routes packets while the standby gateway remains idle until the active gateway fails or when preset conditions are met.

Many host implementations do not support any dynamic router discovery mechanisms but can be configured with a default router. Running a dynamic router discovery mechanism on every host is not feasible for a number of reasons, including administrative overhead, processing overhead, and security issues. HSRP provides failover services to these hosts.

•Virtual Router Redundancy Protocol (VRRP)—VRRP dynamically assigns responsibility for one or more virtual routers to the VRRP routers on a LAN, which allows several routers on a multi-access link to use the same virtual IP address. A VRRP router is configured to run VRRP with one or more other routers attached to a LAN. One router is elected as the virtual router master, while the other routers act as backups if the virtual router master fails.

For configuration details about FHRPs, see the see the Inspur CN12900 Series INOS-CN Unicast Configuration Guide.

2.6 Nonstop Forwarding in Routing Protocols

Inspur INOS-CN provides a multilevel high-availability architecture. Open Shortest Path First version 2

(OSPFv2) supports stateful restart, which is also referred to as nonstop routing (NSR). If OSPFv2 experiences problems, it attempts to restart from its previous runtime state. The neighbors would not register any neighbor event in this case.

If the first restart is not successful and another problem occurs, OSPFv2 attempts a graceful restart. A graceful restart, or nonstop forwarding (NSF), allows OSPFv2 to remain in the data forwarding path through a process restart. When OSPFv2 needs to do a graceful restart, it first sends a link-local opaque (type 9) link-state advertisement (LSA), called a grace LSA. (For more information about opaque LSAs, see the Inspur CN12900 Series INOS-CN Unicast Routing Configuration Guide.) The restarting of the OSPFv2 platform is called NSF capable. The grace LSA includes a grace period, which is a specified time that the neighbor OSPFv2 interfaces hold onto the LSAs from the restarting OSPFv2 interface. (Typically, OSPFv2 tears down the adjacency and discards all LSAs from a down or restarting OSPFv2 interface.) The participating neighbors, which are called NSF helpers, keep all LSAs that originate from the restarting OSPFv2 interface as if the interface were still adjacent. When the restarting OSPFv2 interface is operational again, it rediscovers its neighbors, establishes adjacency, and starts sending its LSA updates again. At this point, the NSF helpers recognize that graceful restart has finished.

Scenarios where a stateful restart is used:

•First recovery attempt after a process experiences problems.

•User-initiated switchover using the **system switchover** command.

•Active supervisor removal.

•Active supervisor reload using the **reload module** active-sup command.

Scenarios where graceful restart is used:

•Second recovery attempt after a process experiences problems within a 4-minute interval.

•Manual restart of the process using the **restart {ospfv3 | ospf}** command.

Note: For more information on nonstop routing in routing protocols, see the Inspur CN12900 Series INOS-CN Unicast Routing Configuration Guide.

2.7 Additional References for Network-Level High Availability

This section describes additional information related to network-level high availability.

Related Documents

Related Topic	Document Title
Graceful restart	Inspur CN12900 Series INOS-CN Unicast Routing Configuration Guide.

System-Level High Availability

This chapter describes the Inspur INOS-CN high availability (HA) system and application restart operations and includes the following sections:

•About Inspur INOS-CN System-Level High Availability

•Physical Redundancy

•Fabric Module and Line Card Compatibility

•Supervisor Restarts and Switchovers

•Displaying HA Status Information

•Additional References for System-Level High Availability

3.1 About Inspur INOS-CN System-Level High Availability

Inspur INOS-CN system-level HA mitigates the impact of hardware or software failures and is supported by the following features:

•Redundant hardware components:

•Power supply

•Fan tray

•Switch fabric

•System controller

•Supervisor

For details about physical requirements and redundant hardware components, see the Hardware Installation Guide for your specific Inspur CN12900 Series chassis.

•HA software features:

•Nonstop forwarding (NSF)—For details about nonstop forwarding, also known as graceful restart, see the Inspur CN12900 Series INOS-CN Unicast Routing Configuration Guide.

•Embedded Event Manager (EEM)—For details about configuring EEM, see the Inspur CN12900 Series INOS-CN System Management Configuration Guide. Smart Call Home—For details about configuring Smart Call Home, see the Inspur CN12900 Series INOS-CN System Management Configuration Guide.

3.2 Physical Redundancy

The Inspur CN12904 and CN12908 chassis include the following physical redundancies:

•Power supply

•Fan tray

•Switch fabric

•System controller

•Supervisor module

For additional details about physical redundancies, see the Hardware Installation Guide for your specific

Inspur CN12900 Series INOS-CN chassis.

3.2.1 Power Supply Redundancy

The Inspur CN12904 chassis supports up to four power supply modules, and the Inspur CN12908 chassis supports up to eight power supply modules. Each Inspur CN12900 Series switch AC or DC power supply module can deliver up to 3 kW.

Note: You can use the Inspur CN12900 Series switches with all AC power supplies, all DC power supplies, or a mix of AC and DC power supplies to provide the required power for the device.

The power subsystem allows the power supplies to be configured in one of the available redundancy modes. By installing more modules, you can ensure that the failure of one module does not disrupt system operations. You can replace the failed module while the system is operating. For information on power supply module installation and replacement, see the Hardware Installation Guide for your specific Inspur CN12900 Series chassis.

3.2.2 Power Modes

Each of the power redundancy modes imposes different power budgeting and allocation models, which deliver varying usable power yields and capacities. For more information about power budgeting, usable capacity, planning requirements, and redundancy configuration, see the Hardware Installation Guide for your specific Inspur CN12900 Series chassis.

The available power supply redundancy modes are described in the following table.

Table 3 Power Redundancy Modes

Redundancy Mode	**Descriptio**
Combined (nonredundant)	This mode does not provide power redundancy. The available power is the total power capacity of all power supplies.
insrc-redundant (grid redundancy)	This mode provides grid redundancy when you connect half of the power supplies to one grid and the other half of the power supplies to the second grid. The available power is the amount of power available through a grid. To enable grid redundancy, you must connect the power supplies to the correct power grid slots. For example, on the Inspur CN12908 switch, slots 1, 2, 3, and 4 are in grid A, and slots 5, 6, 7, and 8 are in grid B. To configure and operate in grid redundancy mode, you must connect half of your power supplies to the slots in grid A and the rest of your power supplies to the slots in grid B. For more information on power grid slot assignments for your power supplies, see the Hardware Installation Guide for your specific Inspur INOS-CN Series platform.
ps-redundant (N+1 redundancy)	This mode provides an extra power supply if an active power supply goes down. One power supply of all the available power supplies is considered an extra power supply, and the total available power is the amount provided by the active power supply units.

Use the **power redundancy-mode {combined | insrc_redundant | ps-redundant}** command to specify one of these power modes.

3.2.3 Fan Tray Redundancy

The Inspur CN12900 Series switches contain redundant system fan trays for cooling the system. For the number of supported fan trays per chassis, see Physical Redundancy on page 5.

The fan speeds are variable and depend on the temperature of the ASICs in the system. If fans are removed or go bad, the other fan modules can start running at a higher speed to compensate for the missing or failed fans. If the system temperature increases above the thresholds, the system shuts down.

•If a single fan fails within a fan tray, the fan speed of the other fans in the tray does not increase.

•If multiple fans fail within a fan tray, the fan speed increases to 100% on all the fan trays.

•If an entire fan tray is removed, the fan speed for the other two fan trays increases to 100% as soon as the tray is removed.

•If multiple fan trays are removed and not replaced within 2 minutes, the device reloads. When the device comes back, if it still detects the multiple fan tray failure, it reloads again after 2 minutes. If desired, you can use EEM to overwrite this policy.

•If a fan tray fails, leave the failed unit in place to ensure proper airflow until you can replace it. The fan trays are hot swappable, but you must replace one fan tray at a time. Otherwise, the device reboots after 2 minutes if multiple fan trays are missing.

Note: There is no time limit for replacing a single fan tray, but to ensure proper airflow, replace the fan tray as soon as possible.

3.2.4 Switch Fabric Redundancy

Inspur INOS-CN provides switching fabric availability through redundant switch fabric module implementation. You can configure a single Inspur CN12904 or CN12908 chassis with one to six switch fabric modules for capacity and redundancy. Each line card installed in the system automatically connects to and uses all functionally of the installed switch fabric modules. A failure of a switch fabric module triggers an automatic reallocation and balancing of traffic across the remaining active switch fabric modules. Replacing the failed fabric module reverses this process. After you insert the replacement fabric module and bring it online, traffic is again redistributed across all installed fabric modules and redundancy is restored.

Fabric modules are hot swappable. Hot swapping can temporarily disrupt traffic. To prevent the disruption of traffic when you hot-swap fabric modules, use the **poweroff module** slot-number command before you remove a fabric module and the **no poweroff module** slot-number command after you reinsert the fabric module.

The maximum bandwidth allowed per CN129-X636C-R line card requires five fabric modules (CN12904-FM-R or CN12908-FM-R), and the maximum bandwidth allowed per CN129-X636Q-R line card requires four fabric modules (CN12904-FM-R or CN12908-FM-R). Additional fabric modules will provide additional redundancy for these line cards. The maximum bandwidth allowed per CN129-X6136YC-R line card requires six CN12904-FM-R fabric modules for redundancy.

3.2.5 Line Card Failures

To keep line cards powered down whenever they fail or crash, use the **system module failure-action shutdown** command to prevent the cards from rebooting. This command is useful if your topology is configured for network-level redundancy and you want to prevent a second disruption from occurring in the network because a line card is trying to come up.

You can use the **show module** module command to verify that the line card has been powered down. If desired, use the **no poweroff module** module command to manually bring the line card back up.

```
switch (config) # system module failure-action shutdown
2018 August 12 23:31:51 switch %$ VDC-1 %$ %SYSMGR-SLOT1-2-SERVICE_CRASHED:
Service "ipfib" (PID 2558) hasn't caught signal 11 (core will be saved).
2018 August 12 23:32:25 switch %$ VDC-1 %$ %PLATFORM-2-MOD_PWRDN:
Module 1 powered down (Serial number SAL1815Q1DP)
switch(config)# show module 1
Mod Ports Module-Type Model Status
--- ----- ------------------------------------- --------------- ----------
1 36 36x40G Ethernet Module CN129-X636Q-R powered-dn
switch (config) # no poweroff module 1
2018 August 12 23:34:31 switch %$ VDC-1 %$ %PLATFORM-2-PFM_MODULE_POWER_ON:
Manual power-on of Module 1 from Command Line Interface
2018 August 12 23:34:31 switch %$ VDC-1 %$ %PLATFORM-2-MOD_DETECT:
Module 1 detected (Serial number SAL1815Q1DP) Module-Type
36x40G Ethernet Module Model CN129-X636Q-R
2018 August 12 23:34:31 switch %$ VDC-1 %$ %PLATFORM-2-MOD_
```

3.2.6 1System Controller Redundancy

Two redundant system controllers in the Inspur CN12904 and CN12908 chassis offload chassis management functions from the supervisor modules. The controllers are responsible for managing power supplies and fan trays and act as a central point for the Gigabit Ethernet Out-of-Band Channel (EOBC) between the supervisors, fabric modules, and line cards.

3.3 Supervisor Module Redundancy

The Inspur CN12904 and CN12908 chassis support dual supervisor modules to provide 1+1 redundancy for the control and management plane. A dual supervisor configuration operates in an active or standby capacity in which only one of the supervisor modules is active at any given time, while the other acts as a standby backup. The state and configuration remain constantly synchronized between the two supervisor modules to provide stateful switchover in the event of a supervisor module failure.

An Inspur INOS-CN generic online diagnostics (GOLD) subsystem and additional monitoring processes on the supervisor trigger a stateful failover to the redundant supervisor when the processes detect unrecoverable critical failures, service restartability errors, kernel errors, or hardware failures.

If a supervisor-level unrecoverable failure occurs, the currently active failed supervisor triggers a switchover. The standby supervisor becomes the new active supervisor and uses the synchronized state and configuration while the failed supervisor is reloaded. If the failed supervisor is able to reload and pass self-diagnostics, it initializes, becomes the new standby supervisor, and then synchronizes its operating state with the newly active unit.

3.3.1 Supervisor Modules

Two supervisor modules are available for the Inspur CN12900 Series switches: Supervisor A (SUP A) and Supervisor B (SUP B). The following table lists the differences between the two modules.

	Supervisor A	**Supervisor B**	**Supervisor A+**	**Supervisor B+**
CPU	4 core, 1.8 GHz	6 core, 2.1 GHz	4 core, 1.8 GHz	6 core, 1.9 GHz
Memory	16 GB	24 GB	16 GB	32 GB
SSD storage	64 GB	256 GB	256 GB	256 GB

SUP A and SUP B are not compatible and should not be installed in the same chassis, except for migration purposes. For dual supervisor systems, you should install either two SUP A modules or two SUP B modules (and not a combination of the two) to ensure supervisor module redundancy.

In a dual supervisor system, Inspur INOS-CN checks the memory size of both the active and standby supervisors. If the memory size is different for each supervisor (because both SUP A and SUP B are installed), a message appears instructing you to replace SUP A with a second SUP B.

To migrate from SUP A to SUP B, insert SUP B into the device and enter the system switchover command. SUP B becomes the active supervisor, and SUP A becomes the standby supervisor, which is not a supported configuration. A warning message appears every hour until you remove SUP A or replace it with a second SUP B.

3.3.2 Fabric Module and Line Card Compatibility

Inspur CN12900 Series switches support only one type of fabric module per chassis.

Only the following fabric module and line card combinations are supported:

Fabric Module	**Supported Line Cards**
CN12904-FM-R CN12908-FM-R	CN129-X636C-R, CN129-X636Q-R, and CN129-X6136YC-R line cards only

3.4 Supervisor Restarts and Switchovers

3.4.1 Restarts on Single Supervisors

In a system with only one supervisor, when all HA policies have been unsuccessful in restarting a service, the supervisor restarts. The supervisor and all services reset and start with no prior state information.

3.4.2 Restarts on Dual Supervisors

When a supervisor-level failure occurs in a system with dual supervisors, the System Manager performs a switchover rather than a restart to maintain stateful operation. In some cases, however, a switchover might not be possible at the time of the failure. For example, if the standby supervisor module is not in a stable standby state, a restart rather than a switchover is performed.

3.4.3 Switchovers on Dual Supervisors

A dual supervisor configuration allows nonstop forwarding (NSF) with a stateful switchover (SSO) when a supervisor-level failure occurs. The two supervisors operate in an active/standby capacity in which only one of the supervisor modules is active at any given time, while the other acts as a standby backup. The two supervisors constantly synchronize the state and configuration in order to provide a seamless and stateful switchover of most services if the active supervisor module fails.

3.4.4 Switchover Characteristics

An HA switchover has the following characteristics:

•It is stateful (nondisruptive) because control traffic is not affected.

•It does not disrupt data traffic because the switching modules are not affected.

•Switching modules are not reset.

3.4.5 Switchover Mechanisms

Switchovers occur by one of the following two mechanisms:

•The active supervisor module fails and the standby supervisor module automatically takes over.

•You manually initiate a switchover from an active supervisor module to a standby supervisor module.

When a switchover process begins, another switchover process cannot be started on the same switch until a stable standby supervisor module is available.

3.4.6 Switchover Failures

Supervisor switchovers are generally hitless and occur without traffic loss. If for some reason a switchover does not complete successfully, the supervisors reset. A reset prevents loops in the Layer 2 network if the network topology was changed during the switchover. For optimal performance of this recovery function, we recommend that you do not change the Spanning Tree Protocol (STP) default timers.

If three system-initiated switchovers occur within 20 minutes, all nonsupervisor modules shut down to prevent switchover cycling. The supervisors remain operational to allow you to collect system logs before resetting the switch.

3.4.7 Manually Initiating a Switchover

To manually initiate a switchover from an active supervisor module to a standby supervisor module, use the system switchover command. After you run this command, you cannot start another switchover process on the same system until a stable standby supervisor module is available.

Note: If the standby supervisor module is not in a stable state (ha-standby), a manually initiated switchover is not performed.

To ensure that an HA switchover is possible, use the show system redundancy status command or the show module command. If the command output displays the ha-standby state for the standby supervisor module, you can manually initiate a switchover.

3.4.8 Switchover Guidelines

Follow these guidelines when performing a switchover:

•When you manually initiate a switchover, it takes place immediately.

•A switchover can be performed only when two supervisor modules are functioning in the switch.

•The modules in the chassis must be functioning.

3.4.9 Verifying Switchover Possibilities

This section describes how to verify the status of the switch and the modules before a switchover.

•Use the **show system redundancy status** command to ensure that the system is ready to accept a switchover.

•Use the **show module** command to verify the status (and presence) of a module at any time. A sample output of the **show module** command follows:

```
switch# show module
Mod Ports           Module-Type                      Model         Status
--- ----- ------------------------------------- -------------------- ---------
1    36   36p 100G Ethernet Module              CN129-X636C-R       ok
2    36   36p 40G Ethernet Module               CN129-X636Q-R       ok
3    52   16x10G + 32x10/25G + 4x100G Module    CN129-X6136YC-R     ok
22   0    Fabric Module                         CN12908-FM-R        ok
24   0    Fabric Module                         CN12908-FM-R        ok
26   0    Fabric Module                         CN12908-FM-R        ok
27   0    Supervisor Module                                         powered-up
28   0    Supervisor Module                     CN129-SUP-B+       active *
29   0    System Controller                     CN129-SC-A          standby
30   0    System Controller                     CN129-SC-A          active

Mod  Sw                    Hw    Slot
---  ----------------------- ------ ----
1    9.2(1)C               1.0    LC1
2    9.2(1)C               1.0    LC2
3    9.2(1)C               1.0    LC3
22   7.0(3)IHA8(0.96)      1.0    FM2
24   7.0(3)IHA8(0.96)      1.0    FM4
26   7.0(3)IHA8(0.96)      1.0    FM6
28   7.0(3)IHA8(0.96)      1.0    SUP2
29   7.0(3)IHA8(0.96)      1.0    SC1
30   7.0(3)IHA8(0.96)      1.0    SC2


Mod  MAC-Address(es)                        Serial-Num
---  --------------------------------------  ----------
1    00-04-31-10-33-2a to 00-04-31-10-33-cb  JAE21310AAX
2    00-04-31-10-0b-d4 to 00-04-31-10-0c-75  JAE213005GB
3    00-04-31-10-4a-f8 to 00-04-31-10-4b-3f  JAE22210C6E
22   NA                                      JAE220600A8
24   NA                                      JAE2206009J
```

```
26   NA                                     JAE2206009Z
28   00-04-31-10-0c-62 to 00-04-31-10-0c-73  FOC22191BRY
29   NA                                     FOC22040SFK
30   NA                                     FOC22046U37

Mod  Online Diag Status
---  ------------------
1    Pass
2    Pass
3    Pass
22   Pass
24   Pass
26   Pass
28   Pass
29   Pass
30   Pass

* this terminal session
switch#
```

The Status column in the output should display an OK status for switching modules and an active or ha-standby status for supervisor modules.

•Use the **show boot auto-copy** command to verify the configuration of the auto-copy feature and if an auto-copy to the standby supervisor module is in progress. Sample outputs of the **show boot auto-copy** command are as follows:

```
switch#           show           boot           auto-copy
Auto-copy         feature        is             enabled

switch#           show           boot           auto-copy list
No file currently being auto-copied
```

3.4.10 Replacing the Active Supervisor Module in a Dual Supervisor System

You can nondisruptively replace the active supervisor module in a dual supervisor system.

Procedure

1.switch # **system switchover**

Initiates a manual switchover to the standby supervisor.

Note: Wait until the switchover completes and the standby supervisor becomes active.

2.switch# **reload module** slot-number **force**

Boots the supervisor module replacement immediately.

Note: If you do not force the boot, the replacement supervisor module should be booted by the active supervisor module 6 minutes after insertion. For information on replacing a supervisor module, see the Hardware Installation Guide for your specific Inspur CN12900 Series chassis.

3.switch# **copy bootflash:**inos-cn-image **bootflash:**inos-cn-image

Copies the inos-cn image from the active supervisor module to the standby supervisor module.

4.switch# **configure terminal**

Enters global configuration mode.

5.switch (config)# **boot inos-cn bootflash:**inos-cn-image [sup-number]

Configures the standby supervisor boot variables.

6.switch(config)# **copy running-config startup-config**

Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to replace the active supervisor module in a dual supervisor system:

```
switch#
switch# config
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# system switchover
switch(config)#
Inspur-Cisco Networking Technology Co. Ltd.
BIOS Ver: 5.30
RC Revision:  02.03.00

Memory Information:
MRC Revision:00.50.00
Total  DRAM: 32768 MB
Memory TOLM: 80000000
PCIE   BASE: 80000000      Size : 10000000
PCI32  BASE: 90000000      Limit: FBFFFFFF
PCI64  BASE: 380000000000      Limit: 383FFFFFFFFF
UC    START: 380000000000      End  : 384000000000
ME Operational Firmware Version: 06:3.0.3.27

DIMM Information:
Clock Speed: 1067MHz
Socket: 0x0 Channel: 0x0 Number: 0x0 Presence: Yes Size: 16GB
Socket: 0x0 Channel: 0x1 Number: 0x0 Presence: Yes Size: 16GB

Detected Inspur IOFPGA
Booting from Primary Bios
Code Signing Results: 0x0
Booting from Upgrade FPGA
```

```
FPGA Revision                      :0x2
FPGA ID                            :0x11543991
FPGA Date                          :0x20180615
Power Debug Register1              :0xbaadbeef
Power Debug Register2              :0xbaadbeef
Reset Cause Register               :0x2
Boot Ctrl Register                 :0xe0ff
FPGA Update Status                 :0x0
EventLog  Register1                :0xc2048000
EventLog  Register2                :0xff7effff

Board Type:Sup-B+

Bootable Disk is detected. Device Name: Micron_M600_MTFDDAT256MBF

Board type  1
INSPUR SUP-B+
IOFPGA @ 0xe8000000
SLOT_ID @ 0x1c
Standalone chassis
check_bootmode: grub:Continue grub
Trying to read config file /boot/grub/menu.lst.local from (hd0,5)
Filesystem type is ext2fs,partition type 0x83
Trying to read config file/boot/grub/menu.lst.local from (hd0,4)
Filesystem type is ext2fs,partition type 0x83
Auto boot configuration file is absent.
No autoboot or failed autoboot. falling to loader


             Loader Version 5.30
loader >

*******************************************************************
current standby sup
*******************************************************************

switch(standby)#
switch(standby)# 2018 Aug 27 18:45:36 switch %$ VDC-1 %$ Aug 27 18:45:36 %KERN-2-SYSTEM_MSG: [2269.991378] Switchover started by redundancy driver-kernel
2018 Aug 27 18:45:36 switch %$ VDC-1 %$ %SYSMGR-2-HASWITCHOVER_PRE_START: This supervisor is becoming active (pre-start phase).
2018 Aug 27 18:45:36 switch %$ VDC-1 %$ %SYSMGR-2-HASWITCHOVER_START: Supervisor 27 is becoming active.
2018 Aug 27 18:45:36 switch %$ VDC-1 %$ %USER-2-SYSTEM_MSG: sysmgr-dispatch(0t1371) status 1-kim
```

2018 Aug 27 18:45:36 switch %$ VDC-1 %$ %PLATFORM-2-PS_ABSENT: Power supply 4 is absent/shutdown, ps-redundancy might be affected

2018 Aug 27 18:45:36 switch %$ VDC-1 %$ %PLATFORM-2-PS_ABSENT: Power supply 5 is absent/shutdown, ps-redundancy might be affected

2018 Aug 27 18:45:37 switch %$ VDC-1 %$ netstack: Registration with cli server complete

2018 Aug 27 18:45:37 switch %$ VDC-1 %$ %SYSMGR-2-SWITCHOVER_OVER: Switchover completed.

2018 Aug 27 18:45:37 switch %$ VDC-1 %$ %USER-2-SYSTEM_MSG: KIM Switchover active cb - kim

2018 Aug 27 18:45:37 switch %$ VDC-1 %$ vlan_mgr: sending stp with mode change - 1

2018 Aug 27 18:45:38 switch %$ VDC-1 %$ %PLATFORM-2-MODULE_EJECTOR_POLICY_ENABLED: All Ejectors closed for module 1. Ejector based shutdown enabled

2018 Aug 27 18:45:38 switch %$ VDC-1 %$ %PLATFORM-2-MODULE_EJECTOR_POLICY_ENABLED: All Ejectors closed for module 2. Ejector based shutdown enabled

2018 Aug 27 18:45:38 switch %$ VDC-1 %$ %PLATFORM-2-MODULE_EJECTOR_POLICY_ENABLED: All Ejectors closed for module 3. Ejector based shutdown enabled

2018 Aug 27 18:45:38 switch %$ VDC-1 %$ %PLATFORM-2-MODULE_EJECTOR_POLICY_ENABLED: All Ejectors closed for module 22. Ejector based shutdown enabled

2018 Aug 27 18:45:38 switch %$ VDC-1 %$ %PLATFORM-2-MODULE_EJECTOR_POLICY_ENABLED: All Ejectors closed for module 24. Ejector based shutdown enabled

2018 Aug 27 18:45:38 switch %$ VDC-1 %$ %PLATFORM-2-MODULE_EJECTOR_POLICY_ENABLED: All Ejectors closed for module 26. Ejector based shutdown enabled

2018 Aug 27 18:45:38 switch %$ VDC-1 %$ %PLATFORM-2-FANMOD_FAN_OK: Fan module 1 (Fan1(sys_fan1) fan) ok

2018 Aug 27 18:45:38 switch %$ VDC-1 %$ %PLATFORM-2-FANMOD_FAN_OK: Fan module 2 (Fan2(sys_fan2) fan) ok

2018 Aug 27 18:45:38 switch %$ VDC-1 %$ %PLATFORM-2-FANMOD_FAN_OK: Fan module 3 (Fan3(sys_fan3) fan) ok

2018 Aug 27 18:45:39 switch %$ VDC-1 %$ %ASCII-CFG-2-CONF_CONTROL: System ready

2018 Aug 27 18:48:11 switch %$ VDC-1 %$ %VMAN-2-GENERAL_FAILURE: Virtual Service [guestshell+]::Activate::Unable to activate container::'3' restarts were attempted, but failed to activate virtual-service

2018 Aug 27 18:48:17 switch %$ VDC-1 %$ %VMAN-2-ACTIVATION_STATE: Failed to deactivate virtual service 'guestshell+'

2018 Aug 27 19:06:13 switch %$ VDC-1 %$ %SYSMGR-2-STANDBY_BOOT_FAILED: Standby supervisor failed to boot up.

2018 Aug 27 19:12:09 switch %$ VDC-1 %$ %USBHSD-STANDBY-2-MOUNT: logflash: online

Inactive timeout reached, logging out.

User Access Verification

switch 登录: admin

密码:

Login incorrect

```
User Access Verification
登录: admin
密码:

Network Operating System-Cloud Network (NOS-CN) Software
for technical assistance contact helpline: 400-691-1766
Copyright (C) 2018, User subject to license from copyright owner.
All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under their own
licenses, such as open source.  This software is provided "as is," and unless
otherwise stated, there is no warranty, express or implied, including but not
limited to warranties of merchantability and fitness for a particular purpose.
Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or
GNU General Public License (GPL) version 3.0  or the GNU
Lesser General Public License (LGPL) Version 2.1 or
Lesser General Public License (LGPL) Version 2.0.
A copy of each such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://opensource.org/licenses/gpl-3.0.html and
http://www.opensource.org/licenses/lgpl-2.1.php and
http://www.gnu.org/licenses/old-licenses/library.txt.
You may use this software subject to the ICNT end user license agreement
which can be found at http://www.icntnetworks.com

switch#
```

Replacing the Standby Supervisor Module in a Dual Supervisor System

You can nondisruptively replace the standby supervisor module in a dual supervisor system.

Procedure

1. switch# **reload module** slot-number **force**

Boots the supervisor module replacement immediately.

Note: If you do not force the boot, the replacement supervisor module should be booted by the active supervisor module 6 minutes after insertion. For information on replacing a supervisor module, see the Hardware Installation Guide for your specific Inspur CN12900 Series chassis.

2. switch# **copy bootflash:**inos-cn-image **bootflash:**inos-cn-image

Copies the software image from the active supervisor module to the standby supervisor module.

3. switch# **configure terminal**

Enters global configuration mode.

4. switch (config)# **boot inos-cn bootflash:**inos-cn-image [sup-number]

Configures the standby supervisor boot variables.

5. switch(config)# **copy running-config startup-config**

Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to replace the standby supervisor module in a dual supervisor system:

```
switch# reload module 27 force
switch# copy bootflash:CN12900-dk9.6.1.2.I3.1.bin bootflash:CN12900-dk9.6.1.2.I3.1.bin
switch# config terminal
switch# boot inos-cn bootflash:CN12900-dk9.6.1.2.I3.1.bin sup-1
switch# copy running-config startup-config
```

3.5 Displaying HA Status Information

Use the **show system redundancy status** command to view the HA status of the system.

```
switch# show system redundancy status
Redundancy mode
---------------
      administrative: HA
        operational: HA
This supervisor (sup-1)
-----------------------
    Redundancy state:   Active
    Supervisor state:   Active
      Internal state:   Active with HA standby
Other supervisor (sup-2)
------------------------
    Redundancy state:   Standby
    Supervisor state:   HA standby
      Internal state:   HA standby
```

The following conditions identify when automatic synchronization is possible:

•If the internal state of one supervisor module is Active with HA standby and the other supervisor module is ha-standby, the system is operationally HA and can perform automatic synchronization.

•If the internal state of one of the supervisor modules is none, the system cannot perform automatic synchronization.

The following table lists the possible values for the redundancy states.

Table 4 Redundancy States

State	Description
Not present	The supervisor module is not present or is not plugged into the chassis.
Initializing	The diagnostics have passed, and the configuration is being downloaded.
Active	The active supervisor module and the switch are ready to be configured.
Standby	A switchover is possible.

Failed	The system detects a supervisor module failure on initialization and automatically attempts to power-cycle the module three times. After the third attempt, it continues to display a failed state.
Offline	The supervisor module is intentionally shut down for debugging purposes.
At BIOS	The system has established a connection with the supervisor, and the supervisor module is performing diagnostics.
Unknown	The system is in an invalid state. If it persists, call TAC.
HA standby	The HA switchover mechanism in the standby supervisor module is enabled.
Active with no standby	A switchover is impossible.

The following table lists the possible values for the supervisor module states.

Table 5 Supervisor States

The following table lists the possible values for the internal redundancy states.

Table 6 Internal States

State	**Description**
Active	The active supervisor module in the switch is ready to be configured.
HA standby	A switchover is possible.
Offline	The system is intentionally shut down for debugging purposes.
Unknown	The system is in an invalid state and requires a support call to TAC.

3.6 Additional References for System-Level High Availability

This section describes additional information related to system-level high availability.