JUNIPER NETWORKS | Engineering Simplicity

Junos Space Network Management Platform

High Availability and Disaster Recovery Guide

Published
2024-04-24

RELEASE
24.1

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Junos Space Network Management Platform High Availability and Disaster Recovery Guide
24.1

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at https://support.juniper.net/support/eula/. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

About This Guide

Use this guide to get an overview of the design and implementation of high availability in Junos Space. This guide also includes information about steps required to deploy the high availability solution, high availability failover scenario, disaster recovery, switchover between active and standby sites, process to configure disaster recovery, solution for managing disaster recovery, and so on.

1

PART

High Availability

CHAPTER 1

Overview

IN THIS CHAPTER

Junos Space High Availability Overview

Junos Space is designed as a carrier-grade system that provides a complete fault tolerant solution. The set of topics describing Junos Space high availability (HA) provide an overview of the Junos Space high availability design and implementation, as well as all the steps that are required to deploy a high availability solution, from ordering your appliances and preparing a Junos Space high availability cluster, to final deployment.

In order to gain an understanding of the Junos Space high availability solution, we recommend that you read all the Junos Space high availability topics. However, if you are primarily interested in setting up high availability, including the prerequisite steps, see the "Configuring the Junos Space Cluster for High Availability Overview" on page 23 topic. A set of frequently asked questions about Junos Space high availability are also answered in FAQ: Junos Space High Availability.

Junos Space Network Management Platform is available in:

- Virtual appliance for the VMware ESX server or Kernel-based Virtual Machine (KVM) environment

The Junos Space Virtual appliance uses the software build with identical features to provide the complete package including OS, databases, load balancers and JBoss engines. You can cluster multiple appliances together to form a Junos Space cluster, as shown in Figure 1 on page 3.

Figure 1: Deployment of a Junos Space Cluster

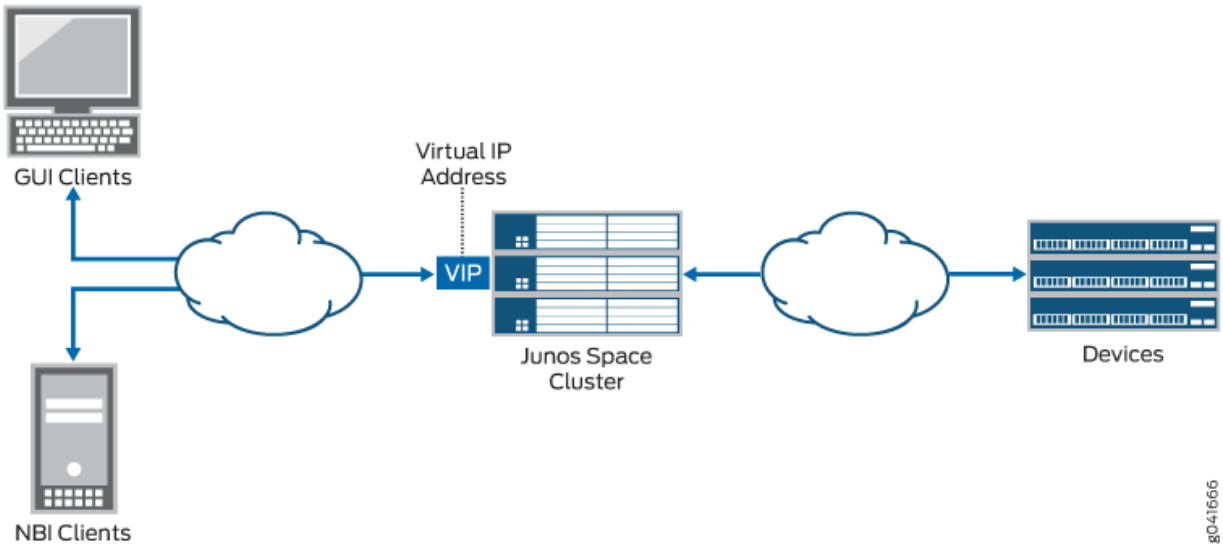

A Junos Space fabric (cluster) contains <only virtual appliances. Each appliance in the cluster is called a *node*. Junos Space cluster architecture also incorporates load balancing across all nodes in the cluster, which becomes the basis for providing scalability for a Junos Space deployment.

A Junos Space high availability solution comprises the following key components:

- Junos Space cluster architecture allows multiple Junos Space Virtual appliances to be connected together to form a single cluster. All services within the cluster are provided through a single virtual IP address that GUI and Northbound Interface (NBI) clients can use. This architecture provides protection against any single point of failure (SPOF) in the cluster. If any node in the cluster fails, all services continue to be available, albeit with reduced capacity.

  Four logical clusters can be formed within the single physical cluster when Junos Space appliances are connected together. For more information, see "Understanding the Logical Clusters Within a Junos Space Cluster" on page 11.

- The Junos Space Appliance contributes significantly to the availability of the overall cluster. For more information, see the No Link Title topic.

- The Watchdog service provides process-level high availability. In the event of any software services failure on a Junos Space appliance, the watchdog service automatically restarts the service.

RELATED DOCUMENTATION

Junos Space High Availability Software Architecture Overview | 4

CHAPTER 2

Understanding the High Availability Software Architecture

IN THIS CHAPTER

Junos Space High Availability Software Architecture Overview

IN THIS SECTION

The Junos Space platform is designed to ensure five-nines availability with a clustered, multi-tiered, distributed architecture comprising the following features:

- Standard browser-based Web 2.0 GUI clients and REST/HTTPS-based NBI clients
- Apache Load Balancer as a top-level load balancer
- JBoss Application Server based on J2EE technology to provide application framework
- MySQL database to manage persistent data
- Cassandra distributed file system to store device image files and files from Junos Space applications

The following sections describe the Junos Space architecture and identify the basic requirements for communication between nodes in a Junos Space cluster:

Junos Space Software Architecture

Figure 2 on page 5 provides a high-level view of the Junos Space software architecture. Junos Space services are accessible to GUI and NBI clients by means of a single virtual IP address for the cluster.

Figure 2: Junos Space Software Architecture

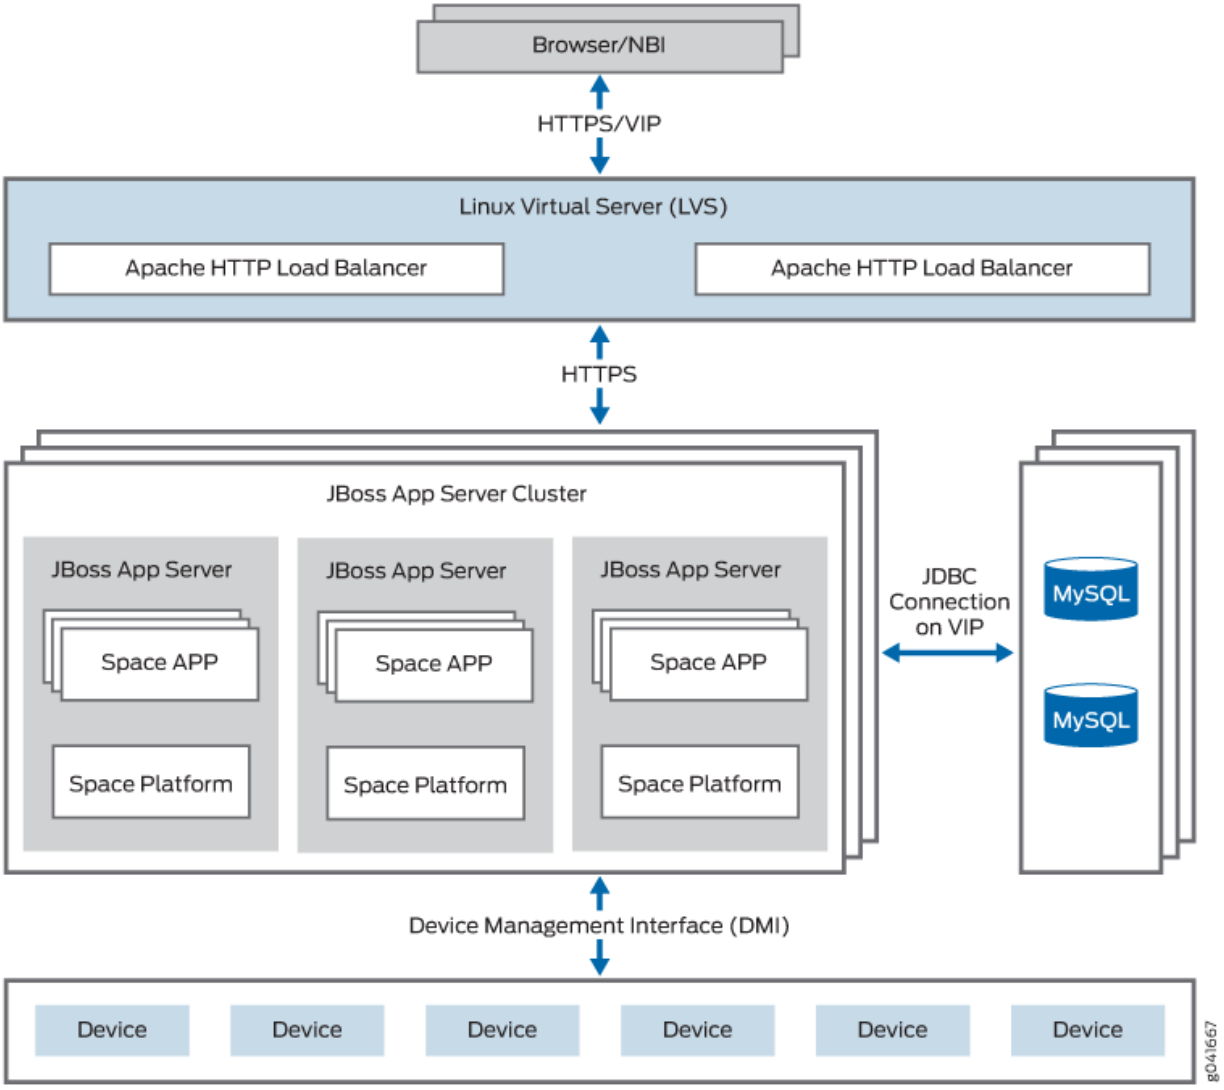

The requests from clients are load-balanced between multiple nodes in the cluster through the Apache HTTP Load Balancer, which is deployed in an active-hot standby configuration on two nodes in the cluster. The load balancer on the node which owns the virtual IP (VIP) address acts as the active instance. If the node which currently owns the VIP address goes down, the other node in the Linux

Virtual Server (LVS) cluster will detect this failure and automatically take over the VIP address. The HTTP requests are load-balanced across all active JBoss servers in the cluster using a round-robin algorithm.

Active JBoss servers within the cluster provide the application framework for Junos Space applications, including the following services:

- Hosting the applications and associated business logic
- Application-level load balancing within the cluster
- Application monitoring and automatic recovery
- Cluster node monitoring and automatic recovery
- Database services with direct access to MySQL DB through JDBC
- Hosting Device Mediation Logic

Load-Balancing Architecture

A Junos Space cluster is presented with two kinds of loads:

- Incoming requests from GUI and NBI clients
- Communication with managed devices

Junos Space is designed to load-balance incoming requests across all active nodes in the cluster. Requests from GUI and NBI clients arrive as HTTP requests serviced by the active instance of the Apache HTTP load balancer. The load balancer distributes the requests to all active JBoss servers in the cluster using a round-robin algorithm. Sticky sessions are utilized to ensure that all HTTP requests associated with a specific GUI session are served by the same JBoss server during the lifetime of that session. For the purpose of application-level load balancing, JBoss business logic processes complex requests as a set of sub-jobs, which are distributed across multiple nodes in the cluster. For example, a single request to a four-node Space cluster to resynchronize 100 devices is divided into four sub-jobs that are executed on four different nodes, with each node resynchronizing 25 devices. For a detailed overview of load balancing, see the topic "Understanding the Logical Clusters Within a Junos Space Cluster" on page 11.

To perform device-level load balancing, Junos Space employs logic in the Device Mediation Layer (DML) so that device connections are equally distributed across all active nodes in the cluster. Device-level load balancing is performed during device discovery by comparing the number of device connections served by individual nodes and selecting the least loaded node. If any node goes down, all associated device connections are distributed to the remaining active nodes in the cluster, thus preventing a node outage from affecting device connectivity. For a detailed overview of device connectivity management, see the topic "Understanding High Availability Management of DMI Connections" on page 21.

Database Architecture

MySQL Enterprise Edition is used to provide database services for managing persistent data for both platform and applications. MySQL DB servers are running on two nodes in the cluster in active-standby configuration. Database transactions are replicated between the two MySQL servers in near real-time. For information about the MySQL cluster that is formed within each Junos Space cluster, see "Understanding the Logical Clusters Within a Junos Space Cluster" on page 11.

Inter-Node Communication Among Nodes in a Junos Space Cluster

In order to facilitate seamless communication between the nodes in a Space cluster and to achieve optimum performance of the cluster, you need to ensure the following:

- All nodes in a Junos Space cluster are configured with IP addresses inside the same subnet. This is important for the VIP switchover mechanism to work correctly.
- All nodes in a Space cluster are connected by means of a 1-Gbps or 100-Mbps local network with negligible latency.
- JBoss servers within a Junos Space cluster communicate by means of a UDP multicast to form logical clusters.

> **NOTE**: UDP multicast traffic must be allowed within the nodes in the cluster, which also means that you should disable IGMP snooping on the switches that interconnect the cluster or configure them explicitly to allow UDP multicast between the nodes.

RELATED DOCUMENTATION

Junos Space High Availability Overview | 2

Software Components for Junos Space Nodes | 7

Understanding the Logical Clusters Within a Junos Space Cluster | 11

Software Components for Junos Space Nodes

The Junos Space virtual appliance runs the same software stack, as shown in Figure 3 on page 8.

Figure 3: Software Stack on a Junos Space Appliance

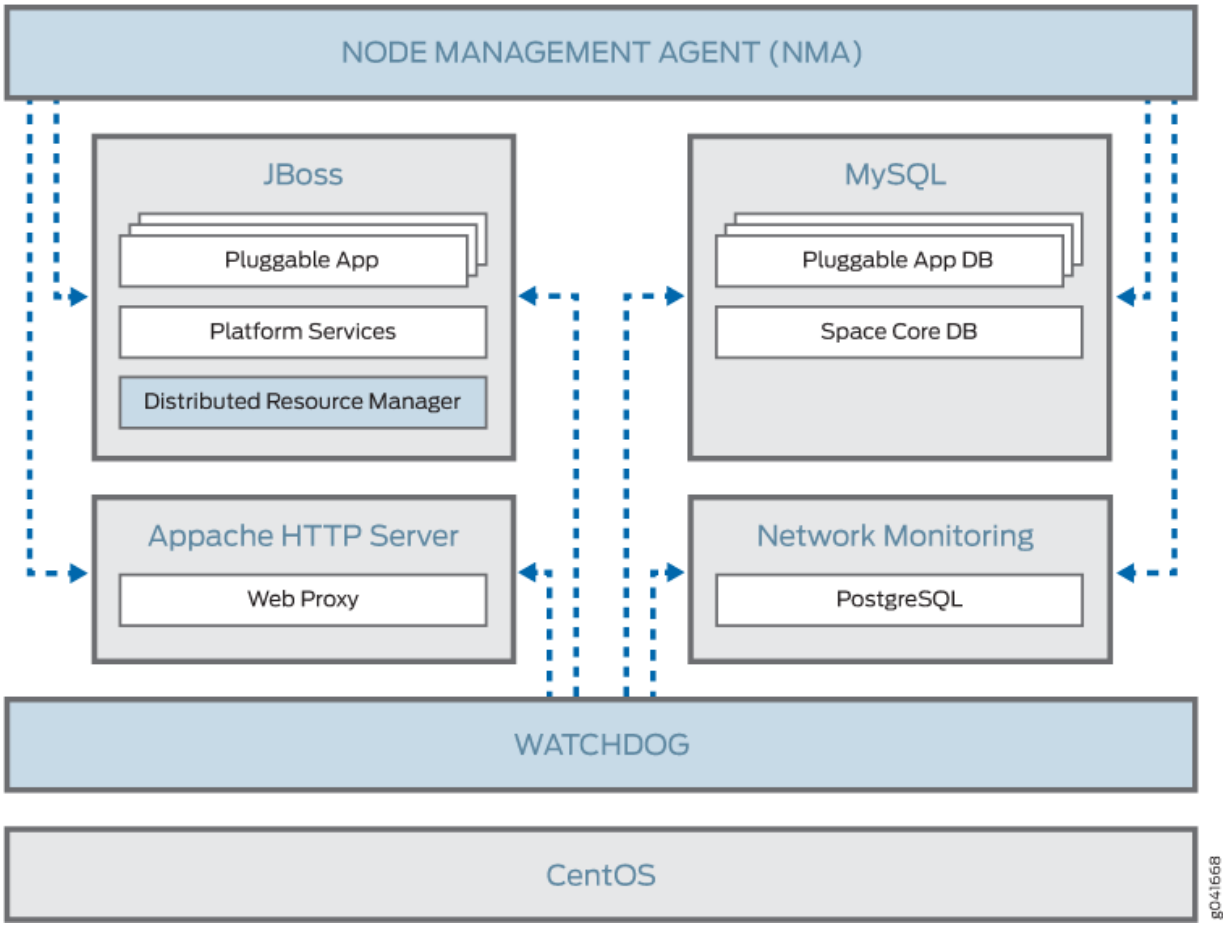

The Junos Space software architecture is based on a combination of the following mature and proven software components:

- CentOS 6.8 distribution is used as the underlying OS of the appliance. CentOS distribution is binary compatible with Red Hat Enterprise Linux (RHEL). Services that are required for Junos Space are leveraged from this distribution, with all other services removed. Junos Space administrators do not need to directly access the Linux components because all operations, administration, and management (OAM) of the platform is performed from the Junos Space user interface or CLI. At the same time, it is important to note that the underlying operating system is an industry-standard distribution with a strong heritage of reliability and security.

- The MySQL Enterprise Edition 5.6 relational database service provides persistent storage for the Junos Space Network Management Platform and all hosted applications. A common database instance stores all persistent data that the Network Management Platform requires. As shown in the preceding illustration, each pluggable application that is installed on the platform has its own unique database instance. All database instances are contained within a single MySQL server, which runs on two nodes in the cluster to form an active-standby cluster. The remaining nodes in the cluster do not run a MySQL server.

- JBoss 7.1 Application Server is the container that hosts the presentation layer, business logic layer, and data access layer of Junos Space platform as well as the hosted applications. One JBoss server runs on each node in the cluster and they all work together as a single load-sharing cluster.
- Apache HTTP Server (version 2.2.34) is the front-end load balancer for all requests coming from GUI and NBI clients. This server runs on two nodes in the cluster which together form an active-standby cluster.

The following software components or services also play a significant role in the overall management of a Junos Space cluster:

- Distributed Resource Manager (DRM)—DRM is deployed as a service inside the JBoss application server, just like all other services provided by Network Management Platform and the hosted applications. You can think of DRM as the server-side component that you interact with when you navigate to the **Network Management Platform > Administration > Fabric** workspace in the Junos Space user interface. DRM works together with the Node Management Agent to fulfill the following responsibilities:
  - Managing the Junos Space cluster—DRM implements the business logic for adding and removing nodes in the cluster and monitors the overall health of the cluster.
  - Managing the logical clusters in the cluster—The logical clusters within the physical cluster formed by the Junos Space nodes include the Apache Load Balancer cluster, JBoss cluster, and Database cluster. DRM implements the business logic to add and remove nodes in these logical clusters and monitors their status. The logical clusters are described in detail in "Understanding the Logical Clusters Within a Junos Space Cluster" on page 11.
- Node Management Agent (NMA)—NMA runs on each node in the cluster and is deployed as a set of CGI scripts run by an Apache HTTP daemon. NMA has the following responsibilities:
  - Monitor system resource usage on the node and the health of various services running on the node.
  - Start and stop services on the node based on requests from DRM.
  - Manage the configuration files for various services running on the node.
  - Manage installation, uninstallation, and upgrades of pluggable applications as well as upgrade of the Network Management Platform software on the node.
- Watchdog—The watchdog service (jmp-watchdog) runs on each node in the cluster to ensure that required services on the node are running. Every second, the watchdog checks that the required services are running and if the watchdog detects that a service is down, it restarts the service.

RELATED DOCUMENTATION

Junos Space High Availability Overview | 2

Junos Space High Availability Software Architecture Overview | 4

Understanding the Logical Clusters Within a Junos Space Cluster | 11

CHAPTER 3

Understanding the Junos Space Cluster (Fabric) Architecture

IN THIS CHAPTER

Understanding the Logical Clusters Within a Junos Space Cluster

IN THIS SECTION

You can connect multiple Junos Space Virtual appliances together to form a Junos Space cluster. Figure 4 on page 12 shows the logical clusters (Apache Load Balancer cluster, the JBoss cluster, and MySQL cluster) that are formed within each Junos Space cluster.

Figure 4: Junos Space Logical Clusters

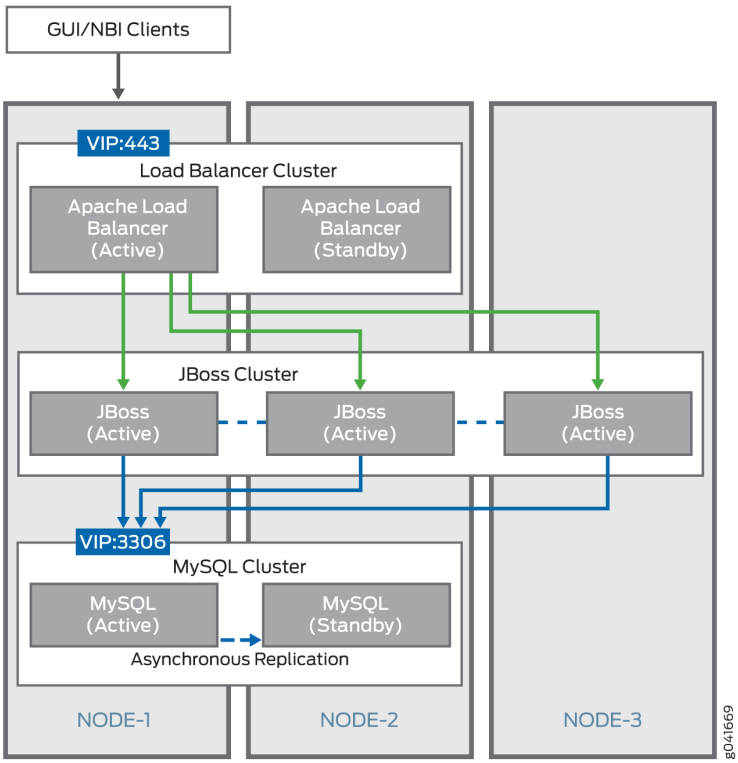

Apache Load-Balancer Cluster

The Apache HTTP server, with the mod_proxy load-balancer module enabled, runs on two nodes in the cluster at any given time. These servers form an active-standby logical cluster. They both listen on the TCP port 443 for HTTP requests from GUI and NBI clients. All clients use the virtual IP (VIP) address of the cluster to access its services. At any time, the VIP address is owned by only one node in the cluster. Hence, the Apache HTTP server on the node that owns the VIP address receives all HTTP requests from GUI and NBI clients and acts as the active load-balancer server, whereas the other server acts as the standby. A round-robin load-balancing algorithm is used to distribute requests to JBoss servers running on all nodes in the cluster. The load-balancer also employs session-stickiness to ensure that all HTTP requests from a user session are sent to the same node in the cluster. To achieve this, the server sets a cookie named JSESSIONID. The value of this cookie identifies the specific node in the cluster that serves requests that belong to this user session. All additional requests contain this cookie and the load-balancer forwards the request to the JBoss server that runs on the node that this cookie identifies.

If the Apache HTTP server on a node goes down, the server is automatically restarted by the watchdog service on that node. If this node owns the VIP address, then the GUI and NBI clients might experience a brief service outage until the Apache HTTP server is restarted. However, this outage lasts only a few seconds (typically, two seconds) and is hardly noticed by the clients. On the other hand, if the Apache HTTP server goes down on the node that does not currently own the VIP address, no side -effects are

noticed by any clients or any other components. The watchdog service restarts the server and the server comes back up in about two seconds.

JBoss Cluster

The JBoss application server runs on all nodes except dedicated database nodes in the Junos Space cluster. The nodes form a single all-active logical cluster and the load-balancer server (described previously) distributes the load across all the nodes. Even if one or more of the JBoss servers in the cluster fails, the application logic still continues to be accessible from the surviving nodes. JBoss servers on all nodes are started with the same configuration and use UDP multicast to detect each other and form a single cluster. JBoss also uses UDP multicast for session replication and caching services across all the nodes.

NOTE: The JBoss server does not run on Fault Monitoring and Performance Monitoring (FMPM) nodes and hosted virtual machines.

When the JBoss server on a node goes down, other nodes in the JBoss cluster detect this change and automatically reconfigure themselves to remove the failed node from the cluster. The time taken by other cluster members to detect a failed JBoss server depends on whether the JBoss server process crashed abnormally or is unresponsive. In the former case, cluster members detect the failure immediately (around two seconds) because their TCP connections to the crashed JBoss server are closed by the operating system. In the latter case, cluster members detect the failure in about 52 seconds. If a JBoss server crashes, the JBoss server is restarted automatically by the watchdog service (jmp-watchdog) running on the node. When the JBoss server comes back up, the JBoss server is automatically discovered by other cluster members and added to the cluster. The JBoss server then synchronizes its cache from the other nodes in the cluster. The typical restart time for the JBoss server is two to five minutes, but it can take more time depending on the number of applications installed, the number of devices being managed, the number of DMI schema versions installed, and so forth.

One JBoss server in the cluster always acts as the primary of the cluster. The main purpose of the primary designation is to host services that are deployed as cluster-wide singletons (HA singletons)—for example, services that must be deployed on only one server in the cluster at any time. Junos Space uses a several services of this type, including the Job Poller service, which provides a single timer for scheduling jobs across the cluster, and the Distributed Resource Manager (DRM) service, which monitors and manages the nodes in the cluster. These services are deployed only on the JBoss server that is designated as the primary.

NOTE: This does not mean that the primary does not host other services. Non-cluster singleton services are also hosted on the primary server. Junos Space is configured such that the first JBoss

server that comes up in the cluster becomes the primary. If the primary server goes down, other members in the JBoss cluster detect this and elect a new primary.

MySQL Cluster

The MySQL server runs on two nodes in the Junos Space cluster at any given time. These nodes form a logical active-standby cluster and both nodes listen on TCP port 3306 for database requests from JBoss servers. By default, JBoss servers are configured to use the Virtual IP (VIP) address of the cluster to access database services. At any time, the VIP address is owned by only one node in the cluster. Thus, the MySQL server on the node that owns the VIP address receives all database requests from the JBoss server, which acts as the active database server while the other server acts as the standby.

If you want to improve the performance of Junos Space Network Management Platform and Junos Space applications, you can add two Junos Space nodes to run as dedicated database nodes. When you add any two Junos Space nodes as the primary and secondary database nodes, the MySQL server is moved to the two dedicated database nodes and is disabled on the first two nodes of the Junos Space cluster. This frees system resources on the Junos Space active VIP node, improving the performance of the node.

JBoss servers use a separate database virtual IP (VIP) address to access database services on dedicated database nodes. You specify the VIP address for the database when you add nodes as dedicated database nodes to the Junos Space cluster. This VIP address is owned by the node designated the primary database node. The MySQL server on the primary database node acts as the active database server, and the server on the secondary database node acts as the standby. Figure 5 on page 15 shows the logical clusters (Apache Load Balancer cluster, the JBoss cluster, and MySQL cluster) that are formed within a Junos Space cluster when you have dedicated database nodes as part of the Junos Space cluster.

Figure 5: Junos Space Logical Clusters with Dedicated Database Nodes

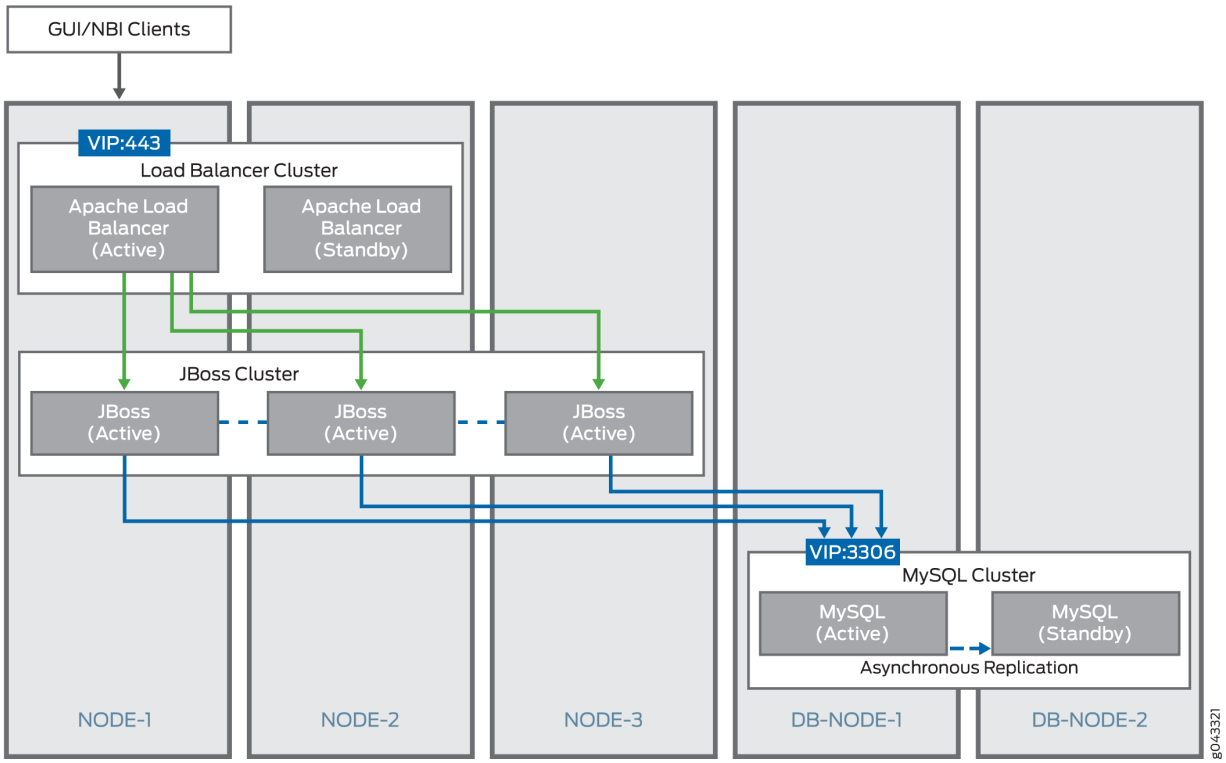

MySQL servers on each of the nodes are configured with unique server IDs. The primary-/backup relationship is also configured symmetrically on the nodes so that the server on the first node is configured with the second node as the primary; and the server on the second node is configured with the first node as the primary. Thus, both nodes are capable of acting as a backup to the other, and the server running on the node that owns the VIP address acts as the primary at any time, which ensures that the primary-backup relationship switches dynamically as the VIP ownership switches from one node to the other. All transactions committed on the active (primary) server are replicated to the standby (backup) server in near real time, by means of the asynchronous replication solution [2] provided by MySQL, which is based on the binary logging mechanism. The MySQL server operating as the primary (the source of the database changes) writes updates and changes as “events” to the binary log. The information in the binary log is stored in different logging formats according to the database changes that are recorded. The backup server is configured to read the binary log from the primary and to execute all the events in the binary log on the backup's local database.

If the MySQL server on a node goes down, the server is restarted automatically by the watchdog service on that node. When restarted, the MySQL server should come up within 20 to 60 seconds. If this node owns the VIP address, JBoss might experience a brief database outage for this 20 to 60 second duration. Any requests that require database access fail during this period. On the other hand, if the MySQL server goes down on the node that does not currently own the VIP address, there are no side- effects noticed by JBoss. The watchdog service restarts the server and the server comes back up in less than

one minute. After the server is back up, it resynchronizes with the primary in the background and the resynchronization time depends on the number of changes that occurred during the outage.

Cassandra Cluster

Starting in Release 15.2R2, Cassandra cluster is an optional logical cluster that you can include within the Junos Space cluster. The Cassandra cluster is formed when there are two or more dedicated Cassandra nodes or two or more JBoss nodes with the Cassandra service running, or a combination of both, within the Junos Space fabric. You can choose to run the Cassandra service on none or all of the nodes in the fabric except dedicated database nodes and FMPM nodes. The Cassandra service running on Junos Space nodes provides a distributed file system to store device images and files from Junos Space applications (such as Juniper Message Bundle [JMB] generated by Service Now and RRD files generated by Network Director). If there are no Cassandra nodes in the fabric, device image files and Junos Space application files are stored in the MySQL database. Figure 6 on page 17 shows the logical clusters (Apache Load Balancer cluster, JBoss cluster, MySQL cluster, and Cassandra cluster) that are formed within a Junos Space cluster when you have Cassandra nodes as part of the Junos Space cluster.

Figure 6: Junos Space Logical Clusters Including the Cassandra Cluster

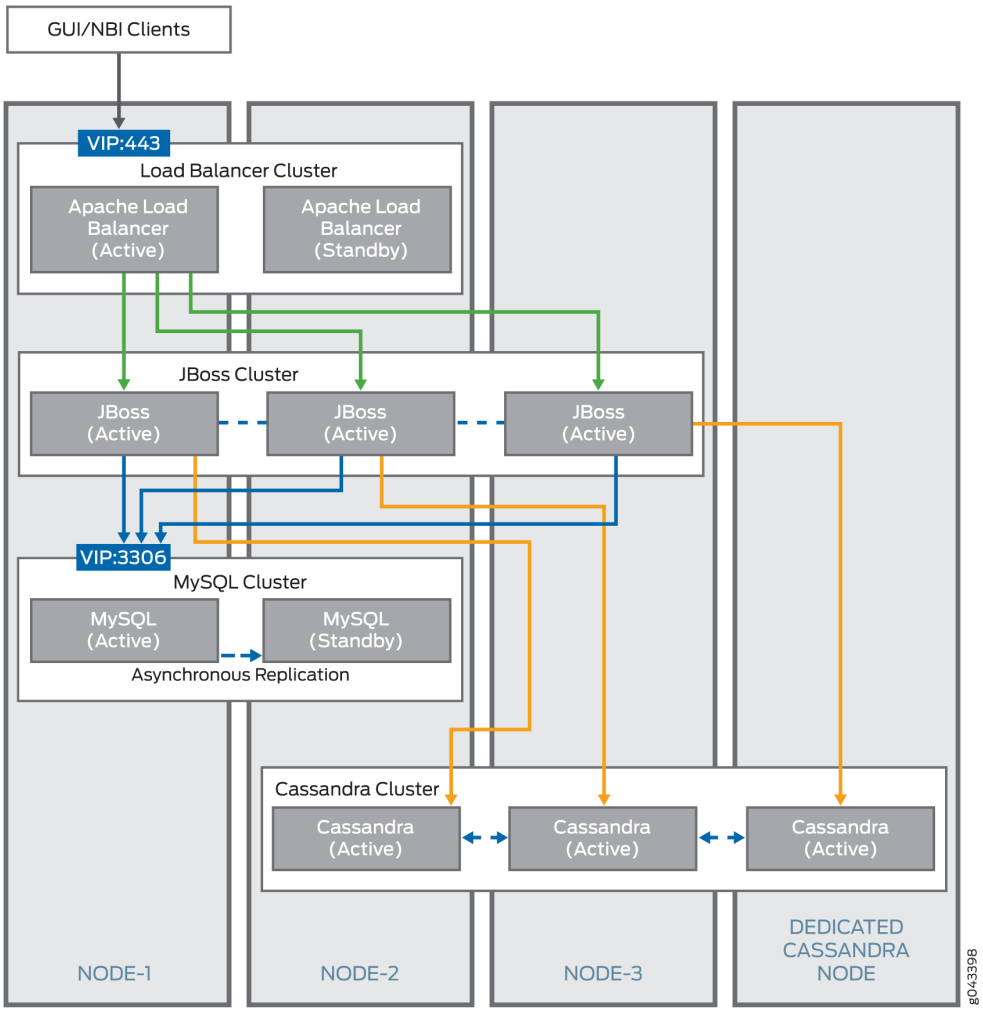

The Cassandra service runs on all the Cassandra nodes in an active-active configuration with real-time replication of the Cassandra database. All the files uploaded to the Cassandra database are copied to all the nodes in the Cassandra cluster. JBoss servers send requests to the Cassandra nodes in the Cassandra cluster in a round-robin manner and access the nodes by using the IP address (of the eth0 interface) of the respective Cassandra node.

If any Cassandra node goes down, Junos Space Platform cannot upload files to or delete files from the Cassandra database until the node that is down is deleted from the fabric. If all existing Cassandra nodes are deleted, the files stored in the Cassandra database are lost.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
15.2R2	Starting in Release 15.2R2, Cassandra cluster is an optional logical cluster that you can include within the Junos Space cluster.

RELATED DOCUMENTATION

Understanding Virtual IP Availability Within a Junos Space Cluster | 18

Understanding High Availability Nodes in a Cluster | 20

Configuring the Junos Space Cluster for High Availability Overview | 23

Understanding Virtual IP Availability Within a Junos Space Cluster

Junos Space must ensure that the virtual IP (VIP) address is always available on one of the nodes in the cluster. This is essential for the HA solution because if the VIP address becomes unavailable, the entire cluster becomes unavailable to all user interface clients and NBI clients. To protect against this scenario, Junos Space uses the heartbeat service (version 2.1.3 to version 3) provided by the Linux-HA project to ensure that the VIP address is always available on one of the nodes in the cluster. For information about the Linux-HA project, see the Linux HA User Guide.

Figure 7 on page 19 shows the heartbeat service that runs on two nodes in the cluster, which together form a Linux HA cluster.

Figure 7: Heartbeat Service on a Linux High Availability Cluster

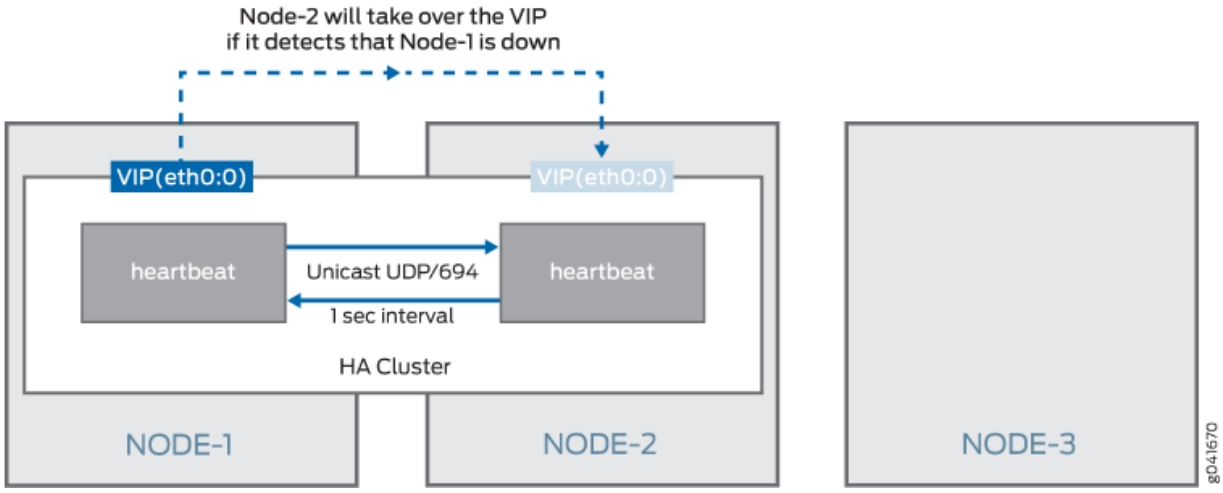

The heartbeat service is configured symmetrically on both nodes to send a heartbeat message to the other node at a 1-second interval. Unicast messages to UDP port 694 are used to send the heartbeat messages. If a node misses 10 consecutive heartbeat messages from the other node, it will consider the other node as dead and initiate a failover to take ownership of the protected resource. The protected resource in this case is the VIP address of the cluster. When failover occurs, the virtual IP address is obtained using a method known as IP address takeover (for more information, see IP Address Take Over) whereby the newly activated node configures the VIP address on one of its interfaces (eth0:0 is used in Junos Space for this) and sends gratuitous ARP packets for the VIP address. All hosts on the network should receive these ARP packets and, from this point forward, send subsequent packets for the VIP address to this node. When the node that currently owns the VIP address crashes, an automatic failover of the VIP address to the other node in the cluster occurs in a little more than 10 seconds. When the crashed node comes back up (for example, in the case of a reboot), it joins the HA cluster and acts as the standby node. In other words, an automatic failback of the VIP address does not happen.

NOTE: The 10 seconds that it takes Junos Space to detect a failed node is applicable when the node crashes or becomes nonresponsive. However, in cases where the node is shut down or rebooted, or if the heartbeat service on the node is stopped by the Junos Space administrator, a message is sent to the heartbeat service on the other node and VIP failover occurs almost instantaneously.

In the case of dedicated database nodes, the database VIP address failover happens in a similar manner to ensure database high availability.

RELATED DOCUMENTATION

Understanding the Logical Clusters Within a Junos Space Cluster | 11

Understanding High Availability Nodes in a Cluster | 20

Configuring the Junos Space Cluster for High Availability Overview | 23

Understanding High Availability Nodes in a Cluster

A Junos Space cluster must include at least two nodes to achieve high availability (HA). If the cluster includes more than two nodes, the availability of the cluster does not increase, but the amount of load that the cluster can handle increases with each node added to the cluster. So at any given time, only two nodes in the cluster provide HA to the whole cluster. By default, these two nodes alone (referred to as the HA nodes in the cluster) form the Linux HA cluster, the Apache Load Balancer cluster, and the MySQL cluster. If you have added dedicated database nodes to the cluster, the MySQL cluster is formed by the primary and secondary database nodes.

By default, the first two nodes added to the cluster function as the HA nodes. In the topic "Understanding the Logical Clusters Within a Junos Space Cluster" on page 11, the example shows that the first two nodes (Node-1 and Node-2) are HA nodes. If you were to delete Node-1 or Node-2 from the **Network Management Platform > Administration > Fabric** workspace, the system checks to see if other nodes in the cluster are available to replace the deleted HA node. The system then displays the list of capable nodes (only Node-3 in the example), which you can select. After you confirm the selected node, the Distributed Resource Manager (DRM) service adds the node to the HA cluster by sending requests to the Node Management Agent (NMA) running on the newly selected node. The following actions are initiated on the node added to the HA cluster:

- Apache HTTP server with the mod_proxy load balancer is started on the node and the node is configured with all JBoss nodes as members.
- If there are no dedicated database nodes in the cluster, the database from the MySQL server on the other HA node in the cluster is copied and the MySQL server is started on the node. This server is configured as a backup of the other MySQL server in the cluster and it resynchronizes with the primary in the background. The existing MySQL server is also reconfigured to act as a backup of this new server to ensure a symmetric primary/backup configuration on both.

When you add dedicated database nodes to the Junos Space cluster, you add two nodes together as the primary and secondary database nodes to form the MySQL cluster. The database is copied from the active HA node to the two database nodes and is disabled on the HA nodes. If you were to delete one of the database nodes from the cluster, the other database node is designated the primary database node. The system checks whether non-HA nodes in the cluster are available to replace the deleted database node and displays the list of nodes you can select to replace the deleted node.

After you select a node, the Distributed Resource Manager (DRM) service adds the node to the MySQL cluster by sending requests to the Node Management Agent (NMA) running on the newly selected node.

The following actions are initiated on the node added to the MySQL cluster:

- The database from the MySQL server on the primary database node in the cluster is copied and the MySQL server is started on the newly-added secondary database node. This server is configured as a backup of the MySQL server on the primary database node and it resynchronizes with the primary in the background. The existing MySQL server on the primary database node is also reconfigured to act as a backup of this new server on the secondary database node to ensure a symmetric primary/backup configuration on both.
- The JBoss server is stopped on the newly added database node.

In addition to the three default logical clusters, if you have a Cassandra cluster as part of the Junos Space fabric, the files uploaded to Cassandra are copied to all the Cassandra nodes that are part of the Cassandra cluster. Hence, if one Cassandra node fails, the files from the failed node are not lost. However, Junos Space Platform cannot upload files to or delete files in the Cassandra database until the node that failed is deleted.

If the Cassandra service is enabled on an HA node and that node goes down, and if you want to run the Cassandra service on the newly added HA node, you must manually enable and start the Cassandra service on the node. When the last node with the Cassandra service running is deleted, the files stored in the Cassandra database are lost.

RELATED DOCUMENTATION

Understanding the Logical Clusters Within a Junos Space Cluster | 11

Configuring the Junos Space Cluster for High Availability Overview | 23

Understanding High Availability Management of DMI Connections

Junos Space maintains a persistent device management interface (DMI) connection with each managed device and supports the following types of DMI connections:

- Space-initiated (default)—A TCP connection from a JBoss server process on a node to the SSH port (22 by default) on the device.
- Device-initiated—A TCP connection from the device to port 7804 on a JBoss server process on a node.

To load balance DMI connections, all connections are distributed across all the nodes in a Junos Space cluster. A device keepalive monitor sends a heartbeat message to devices every 40 seconds. If there is no reply for 15 minutes, the device keepalive monitor marks the connection status of the device as Down.

A device connection monitor scans the connection status of all devices with space-initiated connections. If the monitor detects that the connection status of a device is Down, it attempts to reconnect to the device. If this first attempt fails, a second attempt is made after 30 minutes. Because each reconnect attempt is performed from a node in the cluster that is the least loaded in terms of the number of devices managed, the device might get reconnected from a different node in the cluster after a connection failure.

When devices are discovered using device-initiated connection mode, the device management IP address of all nodes in the Junos Space cluster gets configured in the outbound SSH stanza on the device. The device will keep trying to connect to one of these IP addresses until one succeeds. The device is responsible for detecting any failures on the connection and for reconnecting to another node in the cluster. For more information, see the *Junos XML Management Protocol Guide*.

If a JBoss server process crashes or is stopped, or if the node running the process is shut down, all the DMI connections that it maintains are migrated to another node in the cluster. When this JBoss server comes up, these DMI connections are not automatically migrated back to the JBoss server because it is available for any new devices that are being discovered. At present, there is no way to migrate DMI connections back to this original JBoss server, which can result in poor load balancing of DMI connections if there are not many new devices to be discovered.

RELATED DOCUMENTATION

Understanding High Availability Nodes in a Cluster | **20**

CHAPTER 4

Configuring High Availability Overview

IN THIS CHAPTER

Configuring the Junos Space Cluster for High Availability Overview

IN THIS SECTION

This topic provides an overview of the key steps required to configure a Junos Space cluster as a carrier-grade system with all high-availability capabilities enabled.

Requirements

You can choose Virtual Appliances for setting up a Junos Space cluster.

For a cluster of Virtual Appliances, the following recommendations apply for the underlying virtualization infrastructure on which the appliances are deployed:

- Use VMware ESX server 4.0 or later or VMware ESXi server 4.0, 5.0, 5.1, 5.5, or 6.0 or a kernel-based virtual machine (KVM) server on qemu-kvm (KVM) Release 0.12.1.2-2/448.el6 or later (which is on CentOS Release 6.5) that can support a virtual machine.

- Deploy the two Junos Space Virtual Appliances (JSVA) on two separate servers.
- Each server must be able to dedicate 4 vCPUs or 2.66 GHz or more, 32 GB RAM, and sufficient hard disk for the Junos Space Virtual Appliance that it hosts.
- The servers should have similar fault tolerance features as the Junos Space appliance: dual redundant power supplies connected to two separate power circuits, RAID array of hard disks for storage, and hot-swappable fans.

> **NOTE**: For more information on the requirements for the virtual appliance, refer to the *Deploying a Junos Space Virtual Appliance on a VMware ESXi Server* and *Deploying a Junos Space Virtual Appliance on a KVM Server* topics in the *Junos Space Virtual Appliance* documentation.

If you choose Junos Space appliances, you need to choose two instances of the corresponding SKUs for the appliance that you are using. In addition, order a second power supply module for each appliance in order to provide the redundant power supply module for each appliance.

Preparation

We recommend you use the following guidelines as you prepare a Junos Space cluster for high availability:

- The Junos Space cluster architecture allows you to dedicate two Junos Space nodes solely for MySQL database functions. Dedicated database nodes can free up system resources such as CPU time and memory utilization on the Junos Space VIP node, thereby improving the performance of the Junos Space VIP node. If you decide to add dedicated database nodes to the Junos Space cluster, in the first instance you must add two nodes together as the primary and secondary database nodes, enabling database high availability by default.
- Junos Space Platform enables you to run the Cassandra service on dedicated nodes with only the Cassandra service running or on nodes with the JBoss server running. When the Cassandra service is started on any of the nodes, device images and files from Junos Space applications are moved from the MySQL database to the Cassandra database, thereby improving the performance of the MySQL database. If you want to ensure redundancy for files stored in the Cassandra database, you must ensure that the Cassandra service is running on two or more nodes that together form the Cassandra cluster.
- A Junos Space Virtual appliance utilizes two Ethernet interfaces: eth0 and eth3. The eth0 interface is used for all inter-node communication within the cluster and also for communication between GUI and NBI clients and the cluster. The eth3 interface can be configured as the device management interface, in which case, all communication between the cluster and the managed devices occur over this interface. If the eth3 interface is not configured, all device communication also takes place over

the eth0 interface. So, you must first decide whether or not to use eth3 as the device management interface. If you choose to use eth3, you should use eth3 for all appliances in the same cluster.

- You also must decide on the following networking parameters to be configured on the Junos Space appliances:
  - IP address and subnet mask for the interface “eth0”, the default gateway address, and the address of one or more name servers in the network.
  - IP address and subnet mask for the interface “eth3” if you choose to use a separate device management interface.
  - The virtual IP address to use for the cluster, which should be an address in the same subnet as the IP address assigned to the “eth0” interface.

    If you decide to add dedicated database nodes, you must choose a separate virtual IP (VIP) address in the same subnet as the VIP address of the Junos Space cluster. This database VIP address must be in the same subnet as the IP address assigned to the eth0 Ethernet interface and must be different from the VIP address of the Junos Space cluster nodes.
  - NTP server settings from which to synchronize the appliance’s time.
- The IP address that you assign to each Junos Space node in the cluster and the virtual IP address for the cluster must be in the same subnet. This is required for the IP address takeover mechanism to function correctly.

  NOTE: Strictly speaking, you can choose to deploy the non-HA nodes in a different subnet. However, doing so will cause a problem if one of the HA nodes goes down and you want to promote one of the other nodes as an HA node. So, we recommend that you configure eth0 on all nodes in the same subnet.
- Because JBoss servers on all the nodes communicate using UDP multicast to form and manage the JBoss cluster, you must ensure that UDP multicast is enabled in the network where you deploy the cluster nodes. You must also disable IGMP snooping on the switches interconnecting the cluster, or configure them explicitly to allow UDP multicast between the nodes.

Configuring the First Node in the Cluster

After you power on the appliance and connect to its console, Junos Space displays a menu-driven command-line interface (CLI) that you use to specify the initial configuration of the appliance. To complete this initial configuration, you specify the following parameters:

- IP address and subnet mask for the interface “eth0”

- IP address of the default gateway
- IP address of the name server
- IP address and subnet mask for the interface “eth3”, if you choose to configure a cluster as described in the topic "Understanding the Logical Clusters Within a Junos Space Cluster" on page 11.
- Whether this appliance being added to an existing cluster. Choose “n” to indicate that this is the first node in the cluster.
- The virtual IP address that the cluster will use.
- NTP server settings from which to synchronize the appliance’s time.
- Maintenance mode user ID and password.

> **NOTE**: Make note of the user ID and password that you specify for maintenance mode, as you will need this ID and password to perform Network Management Platform software upgrades and database restoration.

For detailed step-by-step instructions on configuring the appliance for initial deployment, refer to the Junos Space appliance documentation. After you have completed the initial configuration, all Junos Space services are started on the appliance and you can log in to the Network Management Platform User Interface from the virtual IP address assigned to it. At this stage, you have a single node cluster with no HA, which you can see by navigating to the **Network Management Platform > Administration> Fabric** workspace.

Adding a Second Node to the Cluster

In order to add a second node to the cluster, you must first configure the second appliance using its console. The process is identical to that of the first appliance except that you need to choose “y” when it you are prompted to specify whether this appliance will be added to an existing cluster. Make sure that the IP address you assign to this node is in the same subnet as the first node. You must also ensure its uniformity in using a separate device management interface (eth3). If you chose to use eth3 for the first node, choose the same for all additional nodes in the cluster.

After you configure the second appliance, you can log in to the Network Management Platform user interface of the first node at its virtual IP address to add the node to the cluster from the **Network Management Platform > Administration > Fabric > Add Fabric Node** workspace. To add the node to the cluster, specify the IP address assigned to the eth0 interface of the new node, assign a name for the new node, and (optionally) schedule the date and time to add the node. The Distributed Resource Manager (DRM) service running on the first node contacts Node Management Agent (NMA) on the new node to make necessary configuration changes and add it to the cluster. The DRM service also ensures that

required services are started on this node. After the new node joins the cluster, you can monitor its status from the **Network Management Platform > Administration > Fabric** workspace.

For more information about adding nodes to an existing cluster from the Junos Space Platform UI, see *Fabric Management Overview* (in the *Junos Space Network Management Platform* Workspaces User Guide).

Adding Additional Nodes to a Cluster

The process for adding additional nodes is identical to the process for adding the second node. However, these additional nodes do not participate in any of the HA clusters in the fabric, unless explicitly promoted to that role if another HA node is removed, or if they are added as dedicated database nodes to form the MySQL cluster.

For more information about adding nodes to an existing cluster from the Junos Space Platform UI, see *Fabric Management Overview* (in the *Junos Space Network Management Platform* Workspaces User Guide).

Removing Nodes from a Cluster

If a node has failed and needs to be replaced, you can easily remove the node from the cluster. Navigate to the **Network Management Platform > Administration > Fabric** workspace, select the node you want to remove, and choose the **Delete Node** action. If the node being deleted is an HA node, the system will check if other nodes in the cluster can be elected as the replacement for the HA node being deleted. The system then shows the list of capable nodes (only Node-3 in this example) and allows you to choose from the available nodes. The process is described in "Understanding High Availability Nodes in a Cluster" on page 20.

If the node being deleted is a database node, the system checks whether other nodes in the cluster can replace the database node being deleted. If there are nodes present that are capable of replacing the deleted node, the system displays the list of capable nodes and allows you to choose from the available nodes.

For more information about deleting nodes from the cluster, see *Deleting a Node from the Junos Space Fabric* (in the *Junos Space Network Management Platform* Workspaces User Guide).

RELATED DOCUMENTATION

Understanding High-Availability Failover Scenarios | 28

CHAPTER 5

High Availability Failover Scenarios

IN THIS CHAPTER

Understanding High-Availability Failover Scenarios

IN THIS SECTION

The following sections describe possible high-availability failure scenarios: how a failure is detected, what recovery action to take, and if applicable, the impact on the system caused by the failure.

Active VIP Node Crashes

Detection

The heartbeat service running on a standby VIP node detects a crash within 10 seconds of not receiving any heartbeat messages from its peer. The JBoss clustering mechanism enables JBoss servers on other nodes to detect that the JBoss server on the failed node is unresponsive, in about 52 seconds.

Recovery

The standby node immediately takes over the VIP address.

Device connections served by the failed node are migrated to the remaining nodes in the cluster. This process starts in about one minute after the JBoss cluster members detect that the JBoss server on the failed node is down. The time it takes for the process to complete depends on the number of device connections to be migrated, the load on the remaining nodes, and so on. Typically, the process is completed within a few minutes.

Impact

The VIP address becomes unavailable for about 10 seconds until it is taken over by the standby node. The GUI or API client access during this period encounters transient errors. In addition, any SNMP traps sent by the devices to the VIP address during this interval are lost.

Device connectivity is down for a few minutes for devices whose connections were being served by the JBoss server on the failed node.

Any jobs that were in progress on the failed node are marked as failed and the reason is indicated.

NOTE: Starting from Junos Space Network Management Platform 21.1R1, to perform a manual fail over instead of reboot, run the below commands in VIP node CLI:

- `systemctl restart corosync`
- `systemctl restart pacemaker`

Standby VIP Node Crashes

Detection

The JBoss clustering mechanism enables JBoss servers on the other nodes to detect that the JBoss server on the failed node is unresponsive in about 52 seconds.

Recovery

Device connections served by the failed node are migrated to the remaining nodes in the cluster. This process starts in about one minute after the JBoss cluster members detect that the JBoss server on the failed node is down. The process completion time depends on the number of device connections to be migrated, the load on the remaining nodes, and so on. Typically, this process is completed within a few minutes.

Impact

Device connectivity is down for a few minutes for devices whose connections were being served by the JBoss server on the failed node.

Any jobs that were in progress on the failed node are marked as failed and the reason is indicated.

eth0 on the Active VIP Node Goes Down

Detection

The heartbeat service running on the standby VIP node detects the crash within 10 seconds of not receiving any heartbeat messages from its peer. The JBoss clustering mechanism enables JBoss servers on the other nodes to detect that the JBoss server on the failed node is unresponsive, in about 52 seconds.

Recovery

The standby node immediately takes over the VIP address.

Device connections served by the failed node are migrated to the remaining nodes in the cluster. This process starts in about one minute after the JBoss cluster members detect that the JBoss server on the failed node is down. The time it takes for the process to complete depend on the number of device connections to be migrated, the load on the remaining nodes, and so on. Typically, the process is completed within a few minutes.

Impact

The VIP address becomes unavailable for about 10 seconds until it is taken over by the standby node. The GUI or API client access during this period encounters transient errors. In addition, any SNMP traps sent by the devices to the VIP address during this interval are lost.

Device connectivity is down for a few minutes for the devices whose connections were being served by the JBoss server on the failed node.

Any jobs that were in progress on the failed node are marked as failed and the reason is indicated.

eth0 on the Standby VIP Node Goes Down

Detection

The JBoss clustering mechanism enables JBoss servers on the other nodes to detect that the JBoss server on the failed node is unresponsive in about 52 seconds.

Recovery

Device connections served by the failed node are migrated to the remaining nodes in the cluster. This process starts in about one minute after the JBoss cluster members detect that the JBoss server on the failed node is down. The process completion time depends on the number of device connections to be migrated, the load on the remaining nodes, and so on. Typically, this process is completed within a few minutes.

Impact

Device connectivity is down for a few minutes for the devices whose connections were being served by the JBoss server on the failed node.

Any jobs that were in progress on the failed node are marked as failed and the reason is indicated.

A Non-VIP Node Crashes

Detection

The JBoss clustering mechanism enables JBoss servers on the other nodes to detect that the JBoss server on the failed node is unresponsive in about 52 seconds.

Recovery

Device connections served by the failed node are migrated to the remaining nodes in the cluster. This process starts in about one minute after the JBoss cluster members detect that the JBoss server on the failed node is down. The time it takes for the process to complete depends on the number of device connections to be migrated, the load on the remaining nodes, and so on. Typically, this process is completed in a few minutes.

Impact

Device connectivity is down for a few minutes for devices whose connections were served by the JBoss server on the failed node. Any jobs that were in progress on the failed node are marked as failed and the reason is indicated.

eth0 on a Non-VIP Node Goes Down

Detection

The JBoss clustering mechanism enables JBoss servers on the other nodes to detect that the JBoss server on the failed node is unresponsive in about 52 seconds.

Recovery

Device connections served by the failed node are migrated to the remaining nodes in the cluster. This process starts in about 1one minute after the JBoss cluster members detect that the JBoss server on the failed node is down. The process completion time depends on the number of device connections to be migrated, the load on the remaining nodes, and so on. Typically, this process is completed in, a few minutes.

Impact

Device connectivity is down for a few minutes for the devices whose connections were being served by the JBoss server on the failed node.

Any jobs that were in progress on the failed node are marked as failed and the reason is indicated.

eth3 on a Non-VIP Node Goes Down

Detection

The device keepalive monitor detects that all device connections served by this node are down in 15 minutes and marks the connection status of these devices as Down.

Recovery

For connections initiated by Junos Space , Junos Space attempts to reconnect with these devices. Each attempt is made from the cluster node that is determined to be the least loaded in terms of the number of devices it manages. If other nodes in the cluster are significantly less loaded than this node, according to this load-balancing check, reconnection attempts are made from those nodes and they succeed. In this case, connectivity for these devices comes back up in a few minutes. If this node happens to be the least loaded, then all reconnection attempts are made from this node and these attempts continue to fail as long as eth3 remains down.

In the case of device-initiated connections, the device detects a connection failure in about 15 minutes, and then reconnects with another node in the cluster in the next few seconds.

Impact

Device connectivity is down for devices whose connections were being served by this node. Connectivity might be down for 15 minutes (best case) or until eth3 is brought back up (worst case). In addition, the outage time might vary from device to device depending on which node is chosen to

attempt a reconnection for that device. In the case of device-initiated connections, the outage lasts for a little more than 15 minutes.

eth3 on the Active VIP Node Goes Down

Detection

The device keepalive monitor detects that all device connections served by this node are down in 15 minutes and marks the connection status of these devices as Down.

Recovery

For Jconnections initiated by Junos Space, Junos Space attempts to reconnect with these devices. Each attempt is made from the cluster node that is determined to be the least loaded in terms of the number of devices it manages. If other nodes in the cluster are significantly less loaded than this node, according to this load-balancing check, reconnection attempts are made from those nodes and they succeed. In this case, connectivity for these devices comes back up in a few minutes. If this node happens to be the least loaded, then all reconnection attempts are made from this node and these attempts continue to fail as long as eth3 remains down.

In the case of device-initiated connections, the device detects a connection failure in about 15 minutes and then reconnects with another node in the cluster in the next few seconds.

Impact

Device connectivity is down for the devices whose connections were being served by this node. Connectivity might be down for 15 minutes (best case) or until eth3 is brought back up (worst case). In addition, the outage time might vary from device to device depending on which node is chosen to attempt a reconnection for that device. In the case of device-initiated connections, the outage lasts for a little more than 15 minutes.

JBoss Server on a Node Goes Down

Detection

When the JBoss server on a node goes down, other nodes in the JBoss cluster detect the failure in about two seconds) because their TCP connections to the failed JBoss server are closed by the operating system.

Recovery

Device connections served by the failed JBoss server are migrated to the other nodes in the cluster. This process starts in about one minute after the JBoss cluster members detect that the JBoss server on the failed node is down. The time it takes for the process to complete depends on the number of device connections to be migrated, the load on the remaining nodes, and so on. Typically, the process is completed within a few minutes.

The watchdog service (jmp-watchdog) running on the node detects that the JBoss server is down and restarts it automatically. When the JBoss server comes back up, it is automatically discovered by other cluster members and added to the cluster. It then synchronize its cache from the other nodes in the cluster. The typical restart time for JBoss is two to five minutes. However, it can take more time depending on the number of applications installed, the number of devices being managed, the number of DMI schema versions installed, and so on.

Impact

Device connectivity is down for a few minutes for devices whose connections were being served by the JBoss server that went down.

Any jobs that were in progress on the crashed JBoss server are marked as failed and the reason is indicated.

MySQL Server on the Active VIP Node Goes Down

Detection

If the MySQL server on a node goes down, the watchdog service detects the down MySQL server on that active node in about one to two seconds.

Recovery

The watchdog service immediately restarts the MySQL server on the node. When restarted, the MySQL server comes up in around 20 to 60 seconds.

Impact

The MySQL server on the VIP node is the active database servicing all requests from all JBoss servers in the cluster. This effectively means that a brief database outage could be experienced by JBoss on all nodes for this duration (20 to 60 seconds). Any requests that require database access fail during this period. This results in failures encountered by GUI or API clients on their requests, which internally require database access during this period. This also results in failures of jobs that require database access during this period.

MySQL Server on the Standby VIP Node Goes Down

Detection

If the MySQL server on a node goes down, the watchdog service detects the down MySQL server on that standby node in about one to two seconds.

Recovery

The watchdog service immediately restarts the MySQL server on the node. When restarted, it takes around 20 to 60 seconds for the MySQL server to come up. After it is back up, this server

resynchronizes with the primary server in the background and the resynchronization time depends on the number of changes that happened during the outage.

Impact

Since the MySQL server on the standby VIP node is not accessed by JBoss, its downtime does not cause any adverse impact that is noticed by the rest of the system or users of the system.

Primary Database Node Crashes

Detection

The heartbeat service running on the secondary database node detects the crash within 10 seconds of not receiving any heartbeat messages from the primary database node.

Recovery

The database VIP address is transferred to the secondary database node within 10 to 20 seconds. The JBoss servers on other nodes can access the database after the database VIP address is taken over by the secondary database node.

Impact

The database VIP address becomes unavailable for about 10 to 20 seconds until it is taken over by the secondary database node.The MySQL server on the primary database node is the active database servicing all requests from all JBoss servers in the cluster. This effectively means that a brief database outage could be experienced by JBoss on all nodes for this duration (20 to 60 seconds). Any requests that require database access fail during this period. This results in failures encountered by GUI and API clients on their requests that internally require database access during this period. This also results in failures of jobs that require database access during this period.

Secondary Database Node Crashes

Detection

The heartbeat service running on the primary database node detects the crash within 10 seconds of not receiving any heartbeat messages from the secondary database node.

Recovery

The node can be deleted and a new node can be added to the Junos Space cluster as a secondary database node to maintain database high availability.

Impact

Because the MySQL server on the secondary database node is not accessed by JBoss, its downtime does not cause any adverse impact that is noticed by the rest of the system or users of the system.

MySQL Server on the Primary Database Node Goes Down

Detection

If the MySQL server on a node goes down, the watchdog service detects the down MySQL server on that active node in about one to two seconds.

Recovery

The watchdog service immediately restarts the MySQL server on the node. When restarted, the MySQL server comes up in around 20 to 60 seconds.

Impact

The MySQL server on the primary database node is the active database servicing all requests from all JBoss servers in the cluster. This effectively means that a brief database outage could be experienced by JBoss on all nodes for this duration (20 to 60 seconds). Any requests that require database access fail during this period. This results in failures encountered by GUI and API clients on their requests that internally require database access during this period. This also results in failures of jobs that require database access during this period.

MySQL Server on the Secondary Database Node Goes Down

Detection

If the MySQL server on a node goes down, the watchdog service detects the down MySQL server on that standby node in about one to two seconds.

Recovery

The watchdog service immediately restarts the MySQL server on the node. When restarted, it takes around 20 to 60 seconds for the MySQL server to come up. After it is back up, this server resynchronizes with the primary database node in the background. The resynchronization time depends on the number of changes that happened during the outage.

Impact

Because the MySQL server on the secondary database node is not accessed by JBoss, its downtime does not cause any adverse impact that is noticed by the rest of the system or users of the system.

Apache HTTP Server on the Active VIP Node Goes Down

Detection

If the Apache HTTP server on a node goes down, the watchdog service detects the down HTTP server on that node in about one to two seconds.

Recovery

The watchdog service immediately restarts the Apache HTTP server on the node and it becomes ready for service in one second.

Impact

A brief service outage could be experienced by GUI and NBI clients until the Apache HTTP server is restarted. However, this outage is only for a few seconds (typically, two seconds) and is hardly noticed by the clients.

Apache HTTP Server on the Standby VIP Node Goes Down

Detection

If the Apache HTTP server on a node goes down, the watchdog service detects the down HTTP server on that node in about one to two seconds.

Recovery

The watchdog service immediately restarts the Apache HTTP Server on the node and it becomes ready for service in one second.

Impact

No impact.

Dedicated Cassandra Node Crashes

Detection

If the Cassandra node goes down, the watchdog service detects that the Cassandra service is down on that node in about one to two seconds.

Recovery

The Cassandra node that is down must be deleted from the fabric.

Impact

Files cannot be uploaded to or deleted from the Cassandra database until the node that is down is deleted from the fabric.

Cassandra Service on a JBoss Node Goes Down

Detection

If the Cassandra service on a JBoss node goes down, the watchdog service detects that the Cassandra service is down on that node in about one to two seconds.

Recovery

The Cassandra service on the node must be disabled.

Impact

Files cannot be uploaded to or deleted from the Cassandra database until the Cassandra service is disabled on the node.

RELATED DOCUMENTATION

Configuring the Junos Space Cluster for High Availability Overview | 23

Disaster Recovery Overview | 40

2

PART

Disaster Recovery

CHAPTER 6

Disaster Recovery Solution

IN THIS CHAPTER

Disaster Recovery Overview

IN THIS SECTION

A Junos Space cluster allows you to maintain high availability and scalability in your network management solution. However, because all nodes in a cluster need to be within the same subnet, they are typically deployed in the same data center or within the same campus. But you can easily recover a cluster from a disaster at a location by mirroring the original Junos Space installation on a cluster to another cluster at a geographically different location. So if the main Junos Space site fails due to a disaster such as an earthquake, the other site can take over. Hence, the physical installation of the

disaster recovery setup is typically a set of two geographically separate clusters: the active or main site (that is, the local site) and the standby or backup site (that is, the remote site).

When the basic connectivity requirements and prerequisites are met (refer to "Prerequisites to Configure Disaster Recovery" on page 43 and "Connectivity Requirements to Configure Disaster Recovery" on page 44), data from the cluster at the active site is replicated to the cluster at the standby site in near realtime.

The data in the MySQL databases is replicated asynchronously from the active site to the standby site over an SSL connection. MySQL data between the disaster recovery sites is encrypted using self-signed SSL certificates that are generated when disaster recovery is initialized. CA root certificate, CRLs, user certificates, scripts, device images, archived audit logs, and information about scheduled jobs are replicated to the standby site during the real-time data replication to the standby site. The configuration and round-robin database (RRD) files are synchronized periodically by using Secure Copy Protocol (SCP) from the active site to the standby site.

The disaster recovery watchdog, an in-built Junos Space mechanism, monitors the integrity of database replication across sites. All other services (such as JBoss, Apache, and so on) do not run on the standby site until the active site fails over to the standby site. A failover to the standby site is automatically initiated when the active site is down. A device arbitration algorithm is used to determine which site should be the active site to prevent a split-brain scenario where both sites try to be active. For information about the device arbitration algorithm, see "Failure Detection by Using the Device Arbitration Algorithm" on page 46.

The following sections describe the connectivity requirements for the disaster recovery process, failure-detection mechanisms, and the disaster recovery commands:

Disaster Recovery Solution

After you configure and initiate the disaster recovery process between an active site and a standby site, asynchronous replication of MySQL database between the sites is initiated. Configuration and RRD files are backed up to the standby site through SCP at defined time intervals.

The disaster recovery process does not perform real-time replication of the Cassandra database to the standby site or monitor the Cassandra service running on the Junos Space nodes.

During the normal operation of the disaster recovery solution, the GUI and API users and the managed devices are connected to the active site for all network management services. The connectivity between the standby site and managed devices is disabled as long as the active site is functional. When the active site becomes unavailable due to a disaster, the standby site becomes operational. At this time, all services on the standby site are started and the connectivity between the standby site and managed devices is established.

Figure 8 on page 42 displays the disaster recovery solution.

Figure 8: Disaster Recovery Solution

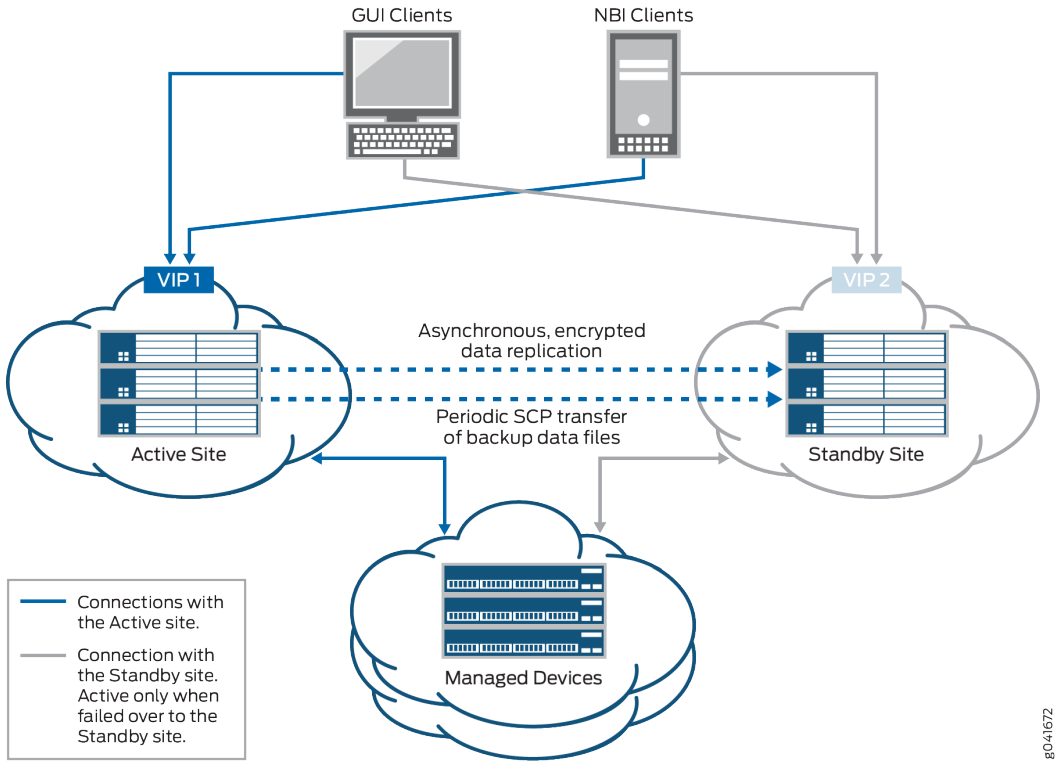

The disaster recovery watchdog process is initiated at the VIP node of both the active and standby sites to monitor the health of the replication process and detect when the remote site goes down. The disaster recovery watchdog at the local site checks whether there are connectivity issues between both sites (by pinging the nodes at the remote site) and whether the sites are connected to arbiter devices (if you use the device arbitration algorithm).

The disaster recovery watchdog at a site performs the following tasks to confirm connectivity with the remote site and arbiter devices:

- Ping the VIP address of the remote site at a regular configurable interval. The default value for the interval is 30 seconds.

  For each ping, expect a reply within a configurable timeout interval. The default value for the timeout interval is 5 seconds.

- If the local site fails to receive a reply within the timeout interval, the disaster recovery watchdog retries the ping for a configurable number of times. By default, the number of retries is 4.

- If all the retries fail, the disaster recovery watchdog at the local site concludes that the VIP address of the remote site is not reachable.

However, the disaster recovery watchdog does not conclude that the remote site is down because the remote site may be switching over the VIP address to a standby node due to a local switchover.

- To consider the possibility of a VIP address switchover, the disaster recovery watchdog pings the IP addresses of the other load-balancer nodes at the remote site. If the ping to any of the nodes receives a reply, the disaster recovery watchdog concludes that the remote site is still up.

  If the ping to the nodes fails, the disaster recovery watchdog does not conclude that the remote site is down. Instead, the disaster recovery watchdog considers the possibility of connectivity issues between the sites. Both sites will try to become active.

- To prevent both sites from trying to become active, the disaster recovery watchdog initiates the device arbitration algorithm and determines whether a failover is required.

  A failover is initiated only if the percentage of arbiter devices managed by the active site falls below the failover threshold. Then the active site becomes the standby site and the standby site becomes the active site.

  If the percentage of arbiter devices is above the failover threshold, the standby site remains standby and the active site remains active. The percentage of arbiter devices managed by the active site is configurable and its default value is 50%.

The failover is initiated if the following conditions are met:

- The standby site cannot reach the VIP address of the active site or the node IP addresses of other load-balancer nodes at the active site.
- The percentage of the arbiter devices managed by the active site is below the failover threshold.

For more information about the device arbitration algorithm, see "Failure Detection by Using the Device Arbitration Algorithm" on page 46.

Prerequisites to Configure Disaster Recovery

You need to ensure that your Junos Space installation meets the following prerequisites before you configure disaster recovery:

- The Junos Space cluster at the primary or active site (which can be a single node or multiple nodes) and the cluster at the remote or standby site (which can be a single node or multiple nodes) must be set up in exactly the same way, with all the same applications, device adapters, same IP family configurations, and so on.
- Both clusters should be configured with SMTP server information from the Junos Space user interface. For more information, see *Managing SMTP Servers*. This configuration enables the clusters at both the active site and the standby site to notify the administrator by e-mail if the replications fail.

NOTE: The number of node(s) in active site and standby site should be the same.

Connectivity Requirements to Configure Disaster Recovery

You need to ensure that the disaster recovery solution meets the following connectivity requirements before you configure disaster recovery:

- Layer 3 connectivity between the Junos Space clusters at the active and standby sites. This means:
  - Every node in a cluster can successfully ping the VIP address of the other cluster
  - Every node in a cluster can use SCP to transfer files between the active and standby sites
  - The bandwidth and latency of the connection between the two clusters are such that real-time database replication is successful. Although the exact bandwidth required depends on the amount of data transferred, we recommend a minimum of a 100-Mbps bandwidth connection with a latency of fewer than 150 milliseconds.
- Independent Layer 3 connectivity between each cluster and managed devices
- Independent Layer 3 connectivity between each cluster and GUI or NBI clients

To set up the disaster recovery process, see "Configuring the Disaster Recovery Process Between an Active and a Standby Site" on page 74.

Disaster Recovery Watchdog

The disaster recovery watchdog, also known as a DR watchdog, is an in-built Junos Space mechanism to monitor the integrity of data replication (MySQL database, configuration files, and RRD files) across sites. The disaster recovery watchdog also monitors the overall health of the disaster recovery setup, initiates a failover from the active to the standby site when the active site fails, and enables the standby site to resume network management services with minimal service disruption. An instance of the disaster recovery watchdog is initiated at the VIP node on both sites when you start the disaster recovery process.

The disaster recovery watchdog provides the following services:

heartbeat

The heartbeat service between the active and standby sites uses ping to check the connectivity between the sites. Both sites send heartbeat messages to each other. The heartbeat service performs the following tasks:

- Detect a failure at the remote site by pinging the remote site at regular intervals.
- When the remote site fails to reply, rule out the possibility of a temporary issue due to a local failover at the remote site.
- Enable or disable automatic failover depending on the disaster recovery configuration settings.
- Avoid split-brain scenarios by running the device arbitration algorithm (default) or the logic configured in the custom script.
- Verify the disaster recovery configuration after a site is rebooted.

mysqlMonitor

The mysqlMonitor service performs the following tasks:

- Monitor the health of MySQL database replication within the site and between the active and standby sites.
- Fix MySQL database replication errors.
- Notify the administrator through e-mail if any of the MySQL database replication errors cannot be fixed automatically.

fileMonitor

The fileMonitor service performs the following tasks:

- Monitor the health of the configuration files and RRD files replicated within the sites and between the active and standby sites.
- Fix errors found during the replication of configuration files and RRD files.
- Notify the administrator through e-mail if any of the replication errors cannot be fixed automatically. You can also view the replication errors in the output of the cron job.

arbiterMonitor

The arbiterMonitor service periodically checks whether the local site can ping all the arbiter devices. If the percentage of arbiter devices that are reachable falls below a configured warning threshold (70%, by default), an e-mail notification is sent to the administrator.

configMonitor

The configMonitor service performs the following tasks:

- Monitor the disaster recovery configuration files at all nodes at both sites.
- Transfer the configuration files across nodes within a site if the files are not in sync.

serviceMonitor

The serviceMonitor service performs the following tasks:

- Monitor the status of selected services (such as jboss, jboss-dc, httpd, and dr-watchdog) within a specific site.
- Start or stop the selected services if they display an incorrect status.

notification

The notification service notifies the administrator about error conditions, warnings, or disaster recovery state changes detected by the disaster recovery watchdog through e-mail. Notification e-mails are sent if:

- Automatic failover is disabled due to connectivity issues between a site and arbiter devices.
- The percentage of arbiter devices that are reachable is lower than the warning threshold.
- A site becomes standby or active.
- The standby site cannot back up files from the active site through SCP.
- A site cannot establish an SSH connection to the remote site.
- The local site cannot fetch the hostname of the MySQL primary node.
- A site cannot fix MySQL database replication errors.

The notification service does not send e-mails to report the same errors within a configurable period of time (by default, 3600 seconds).

Failure Detection by Using the Device Arbitration Algorithm

A device arbitration algorithm is used to detect failure at a site. A list of highly reachable devices running Junos OS and managed by Junos Space Platform are selected as arbiter devices. We recommend that you select arbiter devices based on the following criteria:

- You must be able to reach the arbiter devices through Junos Space–initiated SSH connections from both sites. Do not select devices that use device-initiated connections to Junos Space Platform.
- You must be able to ping arbiter devices from both disaster recovery sites.

- You must choose arbiter devices that stay connected to Junos Space Platform or are less frequently rebooted or shut down because this may impact the device arbitration algorithm results. If you foresee that certain arbiter devices will be offline during some part of their lives, avoid choosing those devices.
- You must choose arbiter devices from different geographical locations to ensure that a problem in the management network at a location does not make all arbiter devices unreachable from your sites.
- You cannot select NAT and ww Junos OS devices as arbiter devices.

The device arbitration algorithm at the active site uses ping to connect to arbiter devices from the active site. The device arbitration algorithm at the standby site logs in to the arbiter devices through SSH connections by using the login credentials from the Junos Space Platform database. Following are the workflows of the device arbitration algorithm at the active and standby sites.

At the active site:

1. Ping all selected arbiter devices.
2. Compute the percentage of arbiter devices that can be pinged.
3. Check whether the percentage of arbiter devices that can be pinged is above or the same as the configured value of the failover threshold.
    - If the percentage of arbiter devices connected is above or the same as the configured value of the failover threshold (failureDetection.threshold.failover parameter in the watchdog section of the disaster recovery API), failover is not initiated because the active site is managing a majority of the arbiter devices.
    - If the percentage of arbiter devices is below the configured value of the failover threshold, failover is initiated (if automatic failover is enabled) and the active site becomes standby. If automatic failover is disabled, the active site remains active.

At the standby site:

1. Log in to arbiter devices through SSH connections.
2. Execute a command on each arbiter device to retrieve the list of SSH connections to any node (managed by the node) at the active site.
3. Calculate the percentage of arbiter devices managed by the active site.
4. Calculate the percentage of arbiter devices that cannot be reached through SSH connections.
    - If the percentage of arbiter devices managed by the active site is above or the same as the configured value of the failover threshold, failover is not required because the active site is still managing a majority of the arbiter devices.

- If the percentage of arbiter devices managed by the active site is below the configured value of the failover threshold, the disaster recovery watchdog concludes that a failover may be required.

5. However, because the devices that cannot be reached from the standby site may be connected and managed by the active site, the standby site assumes that all arbiter devices that cannot be reached are being managed by the active site and calculates the new percentage of devices managed by the active site.

   - If the percentage of devices managed by the active site is below the threshold percentage to adjust managed devices (failureDetection.threshold.adjustManaged parameter in the watchdog section of the disaster recovery API, the default value is 50%), the standby site remains standby. By default, the threshold percentage to adjust managed devices must be below the failover threshold.

   - If the new percentage calculated by adding the devices managed by the active site and arbiter devices that cannot be reached is below the failover threshold, the disaster recovery watchdog concludes that a failover must be initiated.

     If automatic failover is enabled, the standby site initiates the process of becoming active. If automatic failover is disabled, no failover happens.

If you disabled automatic failover or the feature was disabled due to connectivity issues, you must execute `jmp-dr manualFailover` at the standby site to resume network management services from the standby site.

Failure Detection by Using the Custom Failure-Detection Scripts

In addition to using the device arbitration algorithm, you can create custom failure-detection scripts (sh, bash, Perl, or Python) to decide when or whether to fail over to the standby site. Custom failure scripts invoke the `jmp-dr api v1 config --include` command and fetch the disaster recovery configuration and the status of the disaster recovery watchdog services. The disaster recovery configuration and the status of the disaster recovery watchdog services at a site are organized as various sections. Table 1 on page 49 lists these sections.

Use the **-- include <section-name>** option to view the details of a section or use the details of the section in the custom failure-detection script.

Table 1: API Sections

Section	Description	Details Included in the Section	Sample Output
role	Disaster recovery role of the current site	Roles can be active, standby, or standalone.	–
failover	Type of failover that happened last	Value can be active_to_standby, standby_to_active, or empty if failover has not happened yet.	–

core	Core disaster recovery configuration that includes the remote site node details	peerVip–VIP of the load-balancer at the remote site
		
		adminPass–Encrypted administrator passwords of the remote site. Multiple entries are separated by commas.
		
		scpTimeout–Timeout value used to detect SCP transfer failures between sites
		
		peerLoadBalancerNodes–Node IP addresses of the load-balancer nodes at the remote site. Multiple entries are separated by commas.
		
		peerBusinessLogicNodes–Node IP addresses of the JBoss nodes at the remote site. Multiple entries are separated by commas.
		
		peerDeviceMgtIps–Device management IP addresses of the remote site. Multiple entries are separated by commas.
			{
			"core": {
			"peerVip": "10.155.90.210",
			"adminPass": "ABCDE12345",
			"scpTimeout": 120,
			"peerLoadBalancerNodes":
			"10.155.90.211",
			"peerBusinessLogicNodes":
			"10.155.90.211",
			"peerDeviceMgtIps":
			"10.155.90.211"}
			}

Table 1: API Sections *(Continued)*

Section	Description	Details Included in the Section	Sample Output

mysql	Disaster recovery configuration related to the MySQL database at the remote site	hasDedicatedDb–Whether the remote site includes dedicated database nodes
		
		peerVip–VIP of the MySQL nodes at the remote site (either normal node or dedicated database node)
		
		peerNodes–Node IP addresses of the MySQL nodes at the remote site (either normal node or dedicated DB node). Multiple entries are separated by commas.
			<pre>{ "mysql": {
			"hasDedicatedDb": false,
			"peerVip": "10.155.90.210",
			"peerNodes": "10.155.90.211"
			}
			}</pre>

file	Configuration and RRD files–related disaster recovery configuration at the remote site	backup.maxCount–Maximum number of backup files to retain
		
		backup.hoursOfDay–Times of the day to back up files
		
		backup.daysOfWeek–Days of the week to back up files
		
		restore.hoursOfDay–Times of the day to poll files from the active site
		
		restore.daysOfWeek–Days of the week to poll files from the active site
			<pre>{ "file": {
			"backup": {
			"maxCount": 3,
			"hoursOfDay": "*",
			"daysOfWeek": "*" },
			"restore": {
			"hoursOfDay": "*",
			"daysOfWeek": "*" }
			}
			}</pre>

Table 1: API Sections *(Continued)*

Section	Description	Details Included in the Section	Sample Output

watchdog	Disaster recovery configuration related to the disaster recovery watchdog at the current site	heartbeat.retries–Number of times to retry the heartbeat message
		
		heartbeat.timeout–Timeout of each heartbeat message in seconds
		
		heartbeat.interval–Heartbeat message interval between sites in seconds
		
		notification.email–Contact e-mail address to report service issues
		
		notification.interval–Dampening interval between receiving e-mails about affected services
		
		failureDetection.isCustom–Whether the remote site uses custom failure detection
		
		failureDetection.script–Path of the failure-detection script
		
		failureDetection.threshold.failover–Threshold percentage to trigger a failover
		
		failureDetection.threshold.adjustManaged–
		Threshold percentage to adjust the percentage of managed devices
		
		failureDetection.threshold.warning–Threshold percentage to send a warning to ensure that a disaster recovery site can reach more arbiter devices to improve the accuracy of the device arbitration algorithm
		
		failureDetection.waitDuration–Grace period to allow the original active site to become active again when both sites become standby
			{ "watchdog": {
			"heartbeat": {
			"retries": 4,
			"timeout": 5,
			"interval": 30 },
			"notification": {
			"email": "abc@example.com",
			"interval": 3600 },
			"failureDetection": {
			"isCustom": false,
			"script": "/var/cache/jmp-geo/watchdog/bin/arbitration",
			"threshold": {
			"failover": 0.5,
			"adjustManaged": 0.5,
			"warning": 0.7 },
			"waitDuration": "8h",
			"arbiters": [{
			"username": "user1",
			"password": "xxx",
			"host": "10.155.69.114",
			"port": 22,
			"privateKey": ""
			}]
			}
			}
			}

Table 1: API Sections *(Continued)*

Section	Description	Details Included in the Section	Sample Output
		failureDetection.arbiters–List of arbiter devices	
deviceManagement	Device management IP addresses at the remote site	peerNodes–Device management IP addresses of the remote site. Multiple entries are separated by commas. nodes–Device management IP addresses at the current site. Multiple entries are separated by commas. ip–Device management IP address and interface on this node (node on which the `jmp-dr api v1 config --list` command is executed)	`{ "deviceManagement": { "peerNodes": "10.155.90.211", "nodes": "10.155.90.222", "ip": "10.155.90.228,eth0" } }`
states	Runtime information of the disaster recovery watchdog services at the current site. If the disaster recovery watchdog has never run on this site, this section is not available. If the disaster recovery watchdog has stopped, the information in this section is out-of-date.	–	`{ "states": { "arbiterMonitor": { "progress": "idle", "msg": { "service": "arbiterMonitor", "description": "", "state": true, "force": false, "progress": "unknown", "payload": { "code": 0 }, "time": "2015-07-18T22:18:55+00:00" }, "service": {} },`

Table 1: API Sections *(Continued)*

Section	Description	Details Included in the Section	Sample Output
			"configMonitor": { "progress": "idle", "msg": { "service": "configMonitor", "description": "", "state": true, "force": false, "progress": "unknown", "payload": { "code": 0 }, "time": "2015-07-18T22:19:15+00:00" },"service": {} },
			"fileMonitor": { "progress": "idle", "msg": { "service": "fileMonitor", "description": "", "state": true, "force": false, "progress": "unknown", "payload": { "code": 0 }, "time": "2015-07-18T22:18:59+00:00" }, "service": {} },

Table 1: API Sections *(Continued)*

Section	Description	Details Included in the Section	Sample Output
			"heartbeat": { "progress": "unknown", "msg": { "service": "heartbeat", "description": "", "state": true, "force": false, "progress": "unknown", "payload": { "localFailover": false }, "time": "2015-07-18T22:17:49+00:00" }, "service": { "booting": false, "bootEndTime": null, "waitTime": null, "automaticFailover": false, "automaticFailoverEndTime": "2015-07-18T07:41:41+00:00" } },

Table 1: API Sections *(Continued)*

Section	Description	Details Included in the Section	Sample Output

			"mysqlMonitor": {
			"progress": "idle",
			"msg": {
			"service": "mysqlMonitor",
			"description": "",
			"state": true,
			"force": false,
			"progress": "unknown",
			"payload": {
			"code": 0
			},
			"time":
			"2015-07-18T22:19:09+00:00"
			},
			"service": {}
			},

Table 1: API Sections *(Continued)*

Section	Description	Details Included in the Section	Sample Output

			<pre>"serviceMonitor": {
			"progress": "running",
			"msg": {
			"service": "serviceMonitor",
			"description": "",
			"state": true,
			"force": false,
			"progress": "unknown",
			"payload": {
			"code": 0
			},
			"time":
			"2015-07-18T22:19:30+00:00"
			},
			
			"service": {}
			}
			}
			}</pre>

The output from the custom script informs the disaster recovery watchdog whether a failover to the standby site is required. The disaster recovery watchdog interprets the output from the script in the JSON format. The following is an example:

```
{
  "state": "active",
  "action": "nothing",
  "description": "",
  "payload": {
    "waitTime": "",
    "details": {
      "percentages": {
        "connected": 1,
        "arbiters": {
          "10.155.69.114": "reachable"
```

```
        }
      }
    }
  }
}
```

Table 2 on page 57 describes the details of the script output.

Table 2: Details of the Custom Script Output

Property	Description	Data Type	Values or Format	Other Details
state	Current disaster recovery role of this site	String	active standby	Required An empty string is not allowed.
action	Action that the disaster recovery watchdog must perform	String	beActive–Change role to active. beStandby–Change role to standby. nothing–Do not change role. wait–Wait in the current role for the time specified in the payload.waitTime property.	Required An empty string is not allowed.
description	Description of the action field and the message that is sent in the e-mail notification	String	–	Required An empty string is allowed.
payload.waitTime	End time of the grace period when both sites become standby	String (Date)	YYYY-MM-DD, UTC time in HH:MM+00:00 format	Required An empty string is allowed. This property is used when you specify the value of action as wait.

Table 2: Details of the Custom Script Output *(Continued)*

Property	Description	Data Type	Values or Format	Other Details
payload.details	User- customized information that can be used to debug the script	–	JSON object	Optional An empty string is not allowed.

Steps to Configure Disaster Recovery

To configure disaster recovery between an active site and a standby site:

1. Stop the disaster recovery process configured during earlier releases before upgrading to Junos Space Network Management Platform Release 15.2R1. For more information on the upgrade process, see the Upgrade Instructions section in the Junos Space Network Management Platform Release Notes 15.2R1.

   For more information about stopping the disaster recovery process configured during earlier releases, see "Stopping the Disaster Recovery Process on Junos Space Network Management Platform Release 14.1R3 and Earlier" on page 86.

   You do not require to perform this step for a clean installation of Junos Space Network Management Platform Release 15.2R1.

2. Set up SMTP servers at both sites from the Junos Space user interface to receive notifications. For more information, see *Managing SMTP Servers* in the *Junos Space Network Management Platform* Workspaces User Guide.

3. Copy the file with the list of arbiter devices (if you are using the device arbitration algorithm) or the custom failure-detection script to the appropriate location at the active site. Ensure that all arbiter devices are discovered at the active site. For more information, see *Device Discovery Profiles Overview* in the *Junos Space Network Management Platform* Workspaces User Guide.

4. Configure the disaster recovery configuration file at the active site. The disaster recovery configuration includes SCP settings to synchronize configuration and RRD files, heartbeat settings, notifications settings, and the failure-detection mechanism.

5. Configure the disaster recovery configuration file at the standby site. The disaster recovery configuration includes SCP settings to synchronize configuration and RRD files, heartbeat settings, and notification settings.

6. Start the disaster recovery process from the active site.

For more information, see "Configuring the Disaster Recovery Process Between an Active and a Standby Site" on page 74.

Disaster Recovery Commands

You use the disaster recovery commands listed in Table 3 on page 59 to configure and manage disaster recovery sites. You must execute these commands at the VIP node of the site. You can use the `--help` option with these commands to view more information.

Table 3: Disaster Recovery Commands

Command	Description	Options
`jmp-dr init`	Initialize the disaster recovery configuration files at both sites. You need to enter values for the parameters prompted by the command. Create MySQL users and passwords required to replicate data and monitor the replication across disaster recovery sites. The following users are created: • User with a default username repUser and password repPass for MySQL database replication. • User with a default username repAdmin and password repAdminPass to monitor the MySQL database replication health and failover.	`-a`–Initialize the disaster recovery configuration file only at the active site. `-s`–Initialize the disaster recovery configuration file only at the standby site.
`jmp-dr start`	Start the disaster recovery process at both sites. You must execute this command at the VIP node of the active site. The active site establishes an SSH connection to the standby site and executes the `jmp-dr start` command at the standby site. When you execute this command, MySQL database replication and configuration and RRD files backup to the standby site are initiated. You execute this command: • To initially start the disaster recovery process	`-a`–Start the disaster recovery process only at the active site.

Table 3: Disaster Recovery Commands *(Continued)*

Command	Description	Options
	• To restart the disaster recovery process after you stopped the process to upgrade your Junos Space setup.	`-s`–Start the disaster recovery process only at the standby site.
`jmp-dr toolkit config update`	When the command is executed without options, the command: • Displays the modified cluster configuration at a site and updates this at the local site. • Accepts and updates the modified cluster configuration at the remote site. You must execute the command in the following order: 1. Accept and update the cluster configuration changes at both sites. 2. Update load-balancer changes, and modify and update SCP timeout settings at both sites. 3. Modify and update other disaster recovery configuration parameters. You must execute this command at the VIP node of the local site to modify the configuration and the VIP node of the remote site to accept the modified configuration.	Use these options to modify the disaster recovery configuration at a site and update the change at the peer site: `-user-core`–Modify the VIP address, password, and SCP timeout settings. `-user-file-backup`–Modify configuration and RRD files backup settings. `-user-file-restore`–Modify configuration and RRD files replication to standby site settings. `-user-watchdog-heartbeat`–Modify disaster recovery watchdog heartbeat settings. `-user-watchdog-notification`–Modify e-mail notification settings. `-user-watchdog-failureDetection`–Modify failure-detection settings.

Table 3: Disaster Recovery Commands *(Continued)*

Command	Description	Options
`jmp-dr health`	Check the status of the disaster recovery process. The command checks whether MySQL databases are replicated and configuration and RRD files are backed up, and verifies the status of the disaster recovery watchdog and reports errors.	–
`jmp-dr stop`	Stop the disaster recovery process between sites. When you execute this command, MySQL database replication and configuration and RRD files backup between sites are stopped. The disaster recovery data files are not deleted. The status of services such as JBoss, Apache remains unchanged.	–
`jmp-dr reset`	Stop the disaster recovery process and delete the disaster recovery data files from a site. The site initiates services as a standalone cluster. You must execute this command at the VIP node of both sites to stop the disaster recovery process completely and delete the disaster recovery data files from both sites.	–
`jmp-dr manualFailover`	Manually fail over to the standby site. When you execute this command, the standby site becomes the new active site and the active site becomes the new standby site.	`-a`–Manually change the role of the site to active. `-s`–Manually change the role of the site to standby.
`jmp-dr toolkit watchdog status [options]`	Enable automatic failover to the standby site or disable automatic failover to the standby site for a specified duration. **NOTE**: You can execute this command only if the disaster recovery watchdog is active at the site.	`-enable-automatic-failover`–Enable automatic failover to the standby site.

Table 3: Disaster Recovery Commands *(Continued)*

Command	Description	Options
		`-disable-automatic-failover` *duration*–Disable automatic failover to the standby site for a specified time duration. Enter the time duration in hours or minutes. For example, 1h or 30m. If you do not enter “h” or “m” along with the value—for example, 2—the default duration is calculated in hours. If you enter zero, automatic failover is disabled permanently.
`jmp-dr api v1 config`	View the disaster recovery configuration and runtime information in the JSON format.	`--list`–View specific sections of the disaster recovery configuration and status of the disaster recovery watchdog services. Table 1 on page 49 lists the section names.

Table 3: Disaster Recovery Commands *(Continued)*

Command	Description	Options
		`--include <sections>`—Include specific sections of the disaster recovery configuration and status of the disaster recovery watchdog services in the custom failure-detection script. Separate multiple section names with commas. When you include this command in a custom failure-detection script, the command fetches the disaster recovery configuration and status of the disaster recovery watchdog services and executes the logic in the script.

RELATED DOCUMENTATION

- Understanding High Availability Nodes in a Cluster | 20
- Configuring the Junos Space Cluster for High Availability Overview | 23
- Configuring the Disaster Recovery Process Between an Active and a Standby Site | 74
- Modifying Applications and Nodes on a Disaster Recovery Setup | 117

Understanding the Normal Operation of Active and Standby Sites

During the normal operation of the active and standby sites, you use the virtual IP (VIP) address of the active site to access its GUI and API for all network management services. On the active site, a cron job is run based on the disaster recovery configuration. MySQL databases at the active site are asynchronously replicated at the standby site. This ensures that if the active site fails due to a disaster, the databases at the standby site contain the most recent data from the active site. Performance

monitoring data in the RRD files and certain configuration files are periodically backed up at the active site and transferred to the standby site by using scripts that are configured to run as cron jobs.

To view the cron job to back up files at the active site, execute the `crontab -l` command at the active site. The following is a sample output:

The output shows the time you scheduled to run backups at the active site.

The backup is archived into a **tgz** file in the **/var/cache/jmp-geo/backup/data** directory. Only the most recent three backups (default value) or as configured in the disaster recovery configuration are retained in this directory. The older backups are purged. To view a log of all backups by using the `backupReal.sh` script, see the **backup.log** file located at **/var/cache/jmp-geo/backup**.

To view the cron job to fetch files from the active site, execute the `crontab -l` command at the standby site. The following is a sample output:

The output shows the time you scheduled to restore the backups from the active site.

The **poll.sh** script transfers the most recent backup file from the active site using SCP. The backup files are stored in the **/var/cache/jmp-geo/restore/data** directory. The script ensures that only the most recent three backups (default value) or as configured in the disaster recovery configuration are retained in this directory and older files are purged. To view a log of all backups from the active site by using the `poll.sh` script, see the **restore.log** file located at **/var/cache/jmp-geo/restore**.

You cannot discover or manage any devices at the standby site during the normal operation of a disaster recovery setup.

RELATED DOCUMENTATION

- Disaster Recovery Overview | 40
- Understanding How the Standby Site Becomes Operational When the Active Site Goes Down | 72
- Configuring the Disaster Recovery Process Between an Active and a Standby Site | 74

Understanding Disaster Recovery Failure Scenarios

IN THIS SECTION

The following sections explain failure scenarios such as the active and standby sites (with automatic failover enabled) going down due to a disaster, losing connectivity between sites, and losing connectivity with arbiter devices. The device arbitration algorithm is used for failure detection.

For the scenarios, assume that the active site is site1 and standby site is site2.

Active Site (site1) Goes Down Due to a Disaster or Is Powered Down

Detection

The disaster recovery watchdog at site2 does not receive replies to successive ping retries to site1. The disaster recovery watchdog at site2 initiates the device arbitration algorithm and finds that arbiter devices (all or most) are not managed by site1.

An e-mail is sent to the administrator with this information.

Impact

MySQL database replication to site2 is stopped. If you configured any file transfers through SCP during downtime, site2 may lose that version of configuration and RRD files.

MySQL databases at site2 now contain the latest data that was replicated in real time from site1 before it went down. This includes configuration, inventory, alarm-related data of all managed devices, and data maintained by Junos Space Platform and Junos Space applications. The latest version of configuration and RRD files available at site2 are from the most recent file transfer through SCP.

Junos Space users and NBI clients need to wait until site2 becomes active and use the VIP address of site2 to access all network management services.

Recovery

The disaster recovery watchdog at site2 initiates the process to become active. The complete process may take around 15 to 20 minutes. This can vary depending on the number of devices that are managed on your Junos Space setup.

When the failover is complete, site2 establishes connections with all devices and resynchronizes configuration and inventory data if required. site2 starts receiving alarms and performance management data from managed devices.

NOTE: When you rebuild or power on site1, if the disaster recovery configuration is deleted, you must reconfigure disaster recovery between the sites.

No Connectivity Between the Active and Standby Sites and Both Sites Lose Connectivity with Arbiter Devices

Detection

The disaster recovery watchdog at both sites do not receive replies to successive ping retries. The disaster recovery watchdog at both sites initiates the device arbitration algorithm.

An e-mail is sent to the administrator regarding the failure of MySQL replication and file transfer through SCP between sites.

Impact

MySQL database replication to site2 is stopped. If you configured any file transfers through SCP during downtime, site2 may lose that version of configuration and RRD files.

Because both sites cannot connect to arbiter devices (all or most), both sites cannot determine the status of the other site. site1 starts to become standby and site2 remains standby to avoid a split-brain situation.

Even if connectivity between the two sites is restored, both sites remain standby because the sites cannot connect to arbiter devices.

The network management services are stopped at both sites until one of the sites becomes active.

If connectivity to arbiter devices is not restored within the grace period (by default, eight hours), automatic failover functionality is disabled at both sites. An e-mail is sent every hour to the administrator with this information.

Recovery

If connectivity to arbiter devices is restored within the grace period (by default, eight hours), site1 becomes active again. site2 remains standby.

If both sites are standby, enable disaster recovery by executing the `jmp-dr manualFailover -a` command at the VIP node of site1. To enable automatic failover at the sites, execute the `jmp-dr toolkit watchdog status --enable-automatic-failover` command at the VIP node of site1 and site2.

Fix connectivity issues between site1 and site2 to resume MySQL replication and file transfer through SCP.

No Connectivity Between the Active and Standby Sites

Detection

The disaster recovery watchdog at both sites do not receive replies to successive ping retries. The disaster recovery watchdog at both sites initiates the device arbitration algorithm and finds that arbiter devices (all or most) are managed by site1.

An e-mail is sent to the administrator regarding the failure of MySQL database replication and file transfer through SCP between sites.

Impact

MySQL database replication to site2 is stopped. If you configured any file transfers through SCP during downtime, site2 may lose that version of configuration and RRD files.

Recovery

site1 remains active and site2 remains standby. Fix connectivity issues between site1 and site2 to resume MySQL database replication and file transfer through SCP.

No Connectivity Between the Active and Standby Sites and the Active Site (site1) Loses Connectivity with Arbiter Devices

Detection

The disaster recovery watchdog at both sites do not receive replies to successive ping retries. The disaster recovery watchdog at both sites initiates the device arbitration algorithm.

An e-mail is sent to the administrator regarding the failure of MySQL database replication and file transfer through SCP between sites.

Impact

MySQL database replication to site2 is stopped. If you configured any file transfers through SCP during downtime, site2 may lose that version of configuration and RRD files.

Because site1 cannot connect to arbiter devices, site1 starts to become standby. Because site2 finds that arbiter devices (all or most) are not managed by site1, a failover is initiated. As part of becoming standby, all network management services are stopped at site1.

site2 now contains the latest MySQL data that was replicated in real time from site1. The latest version of configuration and RRD files available at site2 are from the most recent file transfer through SCP.

Junos Space users and NBI clients need to wait until site2 becomes active and use the VIP address of site2 to access all network management services.

Recovery

The disaster recovery watchdog at site2 initiates the process to become active. The complete process may take around 15 to 20 minutes. This can vary depending on the number of devices that are managed on your Junos Space setup.

When the failover is complete, site2 establishes connections with all devices and resynchronizes configuration and inventory data if required. site2 starts receiving alarms and performance management data from managed devices.

Fix connectivity issues between site1 and site2 to resume MySQL database replication and file transfer through SCP.

No Connectivity Between the Active and Standby Sites and the Standby Site (site2) Loses Connectivity With Arbiter Devices

Detection

The disaster recovery watchdog at both sites do not receive replies to successive ping retries. The disaster recovery watchdog at site1 initiates the device arbitration algorithm and finds that arbiter devices (all or most) are managed by site1. The disaster recovery watchdog at site2 initiates the device arbitration algorithm.

An e-mail is sent to the administrator regarding the failure of MySQL replication and file transfer through SCP between sites.

Impact

MySQL database replication to site2 is stopped. If you configured any file transfers through SCP during downtime, site2 may lose that version of configuration and RRD files.

Because site2 cannot connect to arbiter devices (all or most), site2 remains standby.

site2 retries to connect to arbiter devices, but does not become active again even if it can connect to enough arbiter devices within eight hours. During these eight hours, site2 requests disaster recovery runtime information of the remote site to ensure that the remote site is active and not in the process of a failover. If site2 cannot connect to enough arbiter devices within eight hours, site2 disables automatic failover permanently until you manually enable automatic failover. An e-mail is sent every hour to the administrator with this information.

Recovery

Fix connectivity issues between site1 and site2 to resume MySQL database replication and file transfer through SCP.

To enable automatic failover at the standby site, execute the `jmp-dr toolkit watchdog status --enable-automatic-failover` command at the VIP node of site2.

Standby Site (site2) Goes Down Due to Disaster or Is Powered Down

Detection

The disaster recovery watchdog at site1 does not receive replies to successive ping retries to site2. The disaster recovery watchdog at site1 initiates the device arbitration algorithm and finds that arbiter devices (all or most) are managed by site1.

An e-mail is sent to the administrator regarding the failure of MySQL replication and file transfer through SCP between sites.

Impact

MySQL database replication to site2 is stopped. If you configured any file transfers through SCP during downtime, site2 may lose that version of configuration and RRD files.

Recovery

site1 remains active. When you power on site2, site2 becomes standby. If you powered down or if the disaster recovery configuration is not deleted from site2, MySQL database replication and file transfer through SCP are initiated.

> **NOTE**: When you rebuild or power on site2, if the disaster recovery configuration is deleted, you must reconfigure disaster recovery between both sites.

No Connectivity Between the Active Site (site1) and Arbiter Devices

Detection

The arbiterMonitor service of the disaster recovery watchdog at site1 detects that the percentage of reachable arbiter devices is below the configured warning threshold. An e-mail is sent to the administrator with this information.

Impact

There is no impact on the disaster recovery solution until the percentage of reachable arbiter devices goes below the failover threshold.

Recovery

No recovery is required because network management services are available from site1.

No Connectivity Between the Standby Site (site2) and Arbiter Devices

Detection

The arbiterMonitor service of the disaster recovery watchdog at site2 detects that the percentage of reachable arbiter devices is below the configured warning threshold. An e-mail is sent to the administrator with this information.

Impact

There is no impact on the disaster recovery solution.

Recovery

No recovery is required because network management services are available from site1.

Both active and standby sites are active and do not lose connectivity with arbiter devices

Detection

The disaster recovery watchdog receives replies to successive pings at both active and standby sites and retries after recovery of automatic disaster recovery failover.

Impact

- Both active and standby sites connect to all or most of the arbiter devices.
- Both active and standby sites fails to determine the status of the other site. Site1 starts as an active site. where as site2 remains active.
- Even when connectivity between active and standby sites is restored, both sites remain active.
- The network management services starts at both active and standby sites until one of the sites becomes a standby site.

Recovery

If connectivity to arbiter devices is restored within the grace period (by default, eight hours), site1 becomes active and site2 remains as a standby site.

If both active and standby sites are active, enable disaster recovery by executing the `jmp-dr manualFailover -s` command at the VIP node of site1.

To enable automatic failover at the sites, execute the `jmp-dr toolkit watchdog status --enable-automatic-failover` command at the VIP node of site1 and site2.

RELATED DOCUMENTATION

- Disaster Recovery Overview | 40
- Modifying Applications and Nodes on a Disaster Recovery Setup | 117
- Manually Failing Over the Network Management Services to the Standby Site | 130
- Understanding High-Availability Failover Scenarios | 28

Understanding How the Standby Site Becomes Operational When the Active Site Goes Down

When a disaster causes the active site to go down, if automatic failover is enabled and the standby site can exceed the failure threshold, the standby site becomes operational. Otherwise, you may need to execute the `jmp-dr manualFailover` or `jmp-dr manualFailover -a` command at the standby site to resume network management services.

The disaster recovery watchdog at the standby site performs the following failover operations to become an active site:

- Verify that the VIP address at the active site is not reachable.
- Stop database replication and SCP file transfer between the two sites.
- Remove the cron job from the standby site for fetching backup files from the active site.
- Add a cron job at the standby site to back up configuration and RRD files.
- Modify the role of the standby site to active.
- Open port 7804 on all nodes at the standby site.
- Start all services at the standby site.
- Copy system configuration files contained in the backup to appropriate locations.

- Configure all devices to send SNMP traps to the VIP address of the standby site. If eth3 is used for device management at the standby site, the eth3 IP address of the active-VIP node at the standby site is configured as the trap destination, instead of the VIP address.

After the failover is complete, the disaster recovery role of the site is set to Active and the state of the cluster is set to active (1). You can access the GUI and API of the standby site from its VIP to perform all network management tasks. In most cases, the failover should happen within 20 to 30 minutes. When the active site becomes operational again, it becomes the standby site. You can either retain the failed state or choose to revert to the original state.

RELATED DOCUMENTATION

Disaster Recovery Overview | **40**

Understanding the Normal Operation of Active and Standby Sites | **63**

Configuring the Disaster Recovery Process Between an Active and a Standby Site | **74**

Understanding High-Availability Failover Scenarios | **28**

CHAPTER 7

Configuring the Disaster Recovery Process

IN THIS CHAPTER

Configuring the Disaster Recovery Process Between an Active and a Standby Site

IN THIS SECTION

You configure disaster recovery between an active site and a standby site to ensure geographical redundancy of network management services.

Before you initiate the disaster recovery process between both sites, perform the following tasks:

- Ensure that the connectivity requirements as described in the "Disaster Recovery Overview" on page 40 topic are met.
- Check whether identical cluster configurations exist on both sites. We recommend that both clusters have the same number of nodes so that, even in the case of a disaster, the standby site can operate with the same capacity as the active site.

- Ensure that the same versions of Junos Space Network Management Platform, high-level Junos Space applications, and device adapters are installed at both sites.
- Shut down the disaster recovery process configured on Junos Space Network Management Platform Release 14.1R3 and earlier before upgrading to Junos Space Network Management Platform Release 15.2R1 and configuring the new disaster recovery process. For more information, see "Stopping the Disaster Recovery Process on Junos Space Network Management Platform Release 14.1R3 and Earlier" on page 86.

  You cannot configure the new disaster recovery process if you do not stop the disaster recovery you set up on 14.1R3 and earlier releases. You do not need to perform this step on a clean installation of Junos Space Network Management Platform Release 15.2R1.
- Ensure that the same SMTP server configuration exists on both sites to receive e-mail alerts related to the disaster recovery process. You can add SMTP servers from the SMTP Servers task group in the Administration workspace. For more information about adding SMTP servers, see *Adding an SMTP Server* in the *Junos Space Network Management Platform Workspaces Feature Guide*.
- Copy a file with the list of arbitrator devices (one IP address per row) in the CSV format or the custom failure-detection scripts on the VIP node at the active site. You can refer to the sample files at **/var/cache/jmp-geo/doc/samples/**.
- Decide on the values for the following parameters depending on your network connectivity and disaster recovery requirements:
  - VIP address and password of both the active and standby sites
  - Backup, restoration, and Secure Copy Protocol (SCP) synchronization settings
  - Heartbeat time intervals
  - E-mail address of the administrator and the dampening interval in seconds to avoid reporting the same errors to avoid an e-mail flood
  - Failure-detection settings such as the failover threshold and the time during which the standby site stays standby if the arbiter devices are unreachable

The following sections explain how to configure disaster recovery at the active and standby sites and initiate the disaster recovery between both sites.

Configuring Disaster Recovery at the Active Site

You use the `jmp-dr init -a` command to configure disaster recovery at the active site. You need to enter values for the parameters that are displayed. The values you enter here are saved in a configuration file.

To configure disaster recovery at the active site:

1. Log in to the CLI of the Junos Space node at the active site on which the VIP or the eth0:0 interface is configured.

   The Junos Space Settings Menu is displayed.

2. Enter **7** (if you are using a virtual appliance) at the Junos Space Settings Menu prompt to run shell commands.

   The following is a sample output from a virtual appliance:

```
admin@10.206.41.183's password:
Last login: Mon Aug 17 06:17:58 2015 from 10.206.41.42

Welcome to the Junos Space network settings utility.

Initializing, please wait



Junos Space Settings Menu

1> Change Password
2> Change Network Settings
3> Change Time Options
4> Retrieve Logs
5> Security
6> Expand VM Drive Size
7> (Debug) run shell

A> Apply changes
Q> Quit
R> Redraw Menu

Choice [1-7,AQR]: 7
```

   You are prompted to enter the administrator password.

3. Enter the administrator password.
4. Enter `jmp-dr init -a` at the shell prompt.

   The values you need to input to configure disaster recovery at the active site are displayed.

   The Load Balancers part of the disaster recovery configuration file is displayed.

5. Enter the values for the parameters displayed:

   a. Enter the VIP address of the standby site and press Enter.

   b. Enter the administrator passwords of the load-balancer nodes at the standby site and press Enter.

You can enter multiple passwords separated with commas.

If multiple nodes use a common password, you need to enter the password only once.

c. Enter the timeout value to detect a failure in transferring files through SCP from the active site to the standby site, in seconds, and press Enter.

   The minimum and default value is 120.

d. Enter the maximum number of backups to retain at the active site and press Enter.

   The minimum and default value is 3.

e. Enter the times of the day to back up files (in hours) at the active site, separated with commas, and press Enter.

   You can enter any value from 0 through 23. You can also enter * to back up files every hour.

f. Enter the days of the week to back up files at the active site, separated with commas, and press Enter.

   You can enter any value from 0 through 6, where Sunday equals zero. You can also enter * to back up files every day.

g. Enter the times of the day to copy files (in hours) from the active site to the standby site, separated with commas, and press Enter.

   You can enter any value from 0 through 23. You can also enter * to poll files every hour.

h. Enter the days of the week to copy files from the active site to the standby site, separated with commas, and press Enter.

   You can enter any value from 0 through 6, where Sunday equals zero. You can also enter * to poll files every day.

The following is a sample output:

```
########################
#
# Load Balancers
#
########################

What's the vip for load balancers at the standby site? 10.206.41.225
What are the unique admin passwords for load balancer nodes at the standby site (separated by
comma, no space)? $ABC123
What's the scp timeout value (seconds)? 120

# backup for data in file system instead of DB
```

```
What's the max number of backup files to keep? 3
What are the times of the day to run file backup (0-23)?
0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17,18,19,20,21,22,23
What are the days of the week to run file backup (0-6)? 0,1,2,3,4,5,6

# restore for data in file system instead of DB

What are the times of the day to poll files from the active site (0-23)?
0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17,18,19,20,21,22,23
What are the days of the week to poll files from the active site (0-6)? 0,1,2,3,4,5,6
```

When you enter the values for all parameters, the DR Watchdog part of the disaster recovery configuration file is displayed.

6. Enter values for the parameters displayed:

   a. Enter the number of times the active site should send heartbeat messages to the standby site through ping after a heartbeat message times out and press Enter.

      The minimum and default value is 4.

   b. Enter the timeout value of each heartbeat message, in seconds, and press Enter.

      The minimum and default value is 5.

   c. Enter the time interval between two consecutive heartbeat messages to the standby site, in seconds, and press Enter.

      The minimum and default value is 30.

   d. Enter the e-mail address of the administrator to whom e-mail messages about disaster recovery service issues must be sent and press Enter.

   e. Enter the time interval during which the same issues are not reported through e-mail (dampening interval), in seconds, and press Enter.

      The default value is 3,600. The minimum value is 300.

   f. Specify the failure-detection mechanism.

      If you intend to use a custom failure-detection script:

      - Enter **Yes** in the failureDetection section and press Enter.

      If you intend to use the device arbitration algorithm:

      i. Enter **No** in the failureDetection section and press Enter.

      ii. Enter the threshold percentage to trigger a failover to the standby site by using the device arbitration algorithm and press Enter.

You can enter any value from 0 to 1. The default value is 0.5.

g. Enter the path of the file containing the arbiter devices or the custom failure-detection scripts and press Enter.

The following is a sample output:

```
########################
#
# DR Watchdog
#
########################


# heartbeat

What's the number of times to retry heartbeat message? 4
What's the timeout of each heartbeat message (seconds)? 5
What's the heartbeat message interval between sites (seconds)? 30

# notification

What's the contact email address of service issues? user1@example.com
What's the dampening interval between emails of affected services (seconds)? 300

# failureDetection

Do you want to use custom failure detection? No
What's the threshold percentage to trigger failover? 0.5
What's the arbiters list file (note: please refer to example in /var/cache/jmp-geo/doc/
samples/arbiters.list)? /home/admin/user1
Check status of DR remote site: up
Prepare /var/cache/jmp-geo/
incoming                                                              [ OK ]
Configure contact
email                                                                      [ OK ]
Modify firewall for DR remote
IPs                                                                [ OK ]
Configure
NTP
[ OK ]
Configure MySQL
database                                                                    [ OK ]
```

```
Configure PostgreSQL
database                                                              [ OK ]
Copy files to DR
slave                                                                    [ OK ]
Command completed.
```

When you have entered values for all parameters, disaster recovery is initialized at the active site.

Configuring Disaster Recovery at the Standby Site

You use the `jmp-dr init -s` command to configure disaster recovery at the standby site. You need to enter values for the parameters that are displayed. The values you enter here are saved in a configuration file. By default, the standby site uses the failure-detection mechanism you configured at the active site, values you entered for file backup and restoration, heartbeat, and notifications if the standby site becomes an active site.

To configure disaster recovery at the standby site:

1. Log in to the CLI of the Junos Space node at the standby site on which the VIP or the eth0:0 interface is configured.

   The Junos Space Settings Menu is displayed.
2. Enter **7** while using a virtual appliance at the Junos Space Settings Menu prompt to run shell commands.

   You are prompted to enter the administrator password.
3. Enter the administrator password.
4. Enter `jmp-dr init -s` at the shell prompt.

   The values you need to input to configure disaster recovery at the standby site are displayed.

   The Load Balancers part of the disaster recovery configuration file is displayed.
5. The script asks you if you have re-initialised the DR active site, that is run `jmp-dr init -a --skip-user-config` at the DR active site. Select Yes or No accordingly.
6. Enter the values for the parameters displayed:
   a. Enter the VIP address of the active site and press Enter.
   b. Enter the administrator passwords of the load-balancer nodes at the active site and press Enter.

      You can enter multiple passwords separated with commas.

      If multiple nodes use a common password, you need to enter the password only once.
   c. Enter the timeout value to detect a failure in transferring files through SCP from the standby site to the active site, in seconds, and press Enter.

      The minimum and default value is 120.

d. Enter the maximum number of backups to retain at the standby site and press Enter.

   The minimum and default value is 3.

e. Enter the times of the day to back up files (in hours) at the standby site, separated with commas, and press Enter.

   You can enter any value from 0 through 23. You can also enter * to back up files every hour.

f. Enter the days of the week to back up files at the standby site, separated with commas, and press Enter.

   You can enter any value from 0 through 6, where Sunday equals zero. You can also enter * to back up files every day.

g. Enter the times of the day to copy files (in hours) from the standby site to the active site (when failed over to the standby site), separated with commas, and press Enter.

   You can enter any value from 0 through 23. You can also enter * to restore files every hour.

h. Enter the days of the week to copy files from the standby site to the active site (when failed over to the standby site), separated with commas, and press Enter.

   You can enter any value from 0 through 6, where Sunday equals zero. You can also enter * to restore files every day.

The following is a sample output:

```
########################
#
# Load Balancers
#
########################

What's the vip for load balancers at the active site? 10.206.41.220
What are the unique admin passwords for load balancer nodes at the active site (separated by
comma, no space)? $ABC123
What's the scp timeout value (seconds)? 120

# backup for data in file system instead of DB

What's the max number of backup files to keep? 3
What are the times of the day to run file backup (0-23)?
0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17,18,19,20,21,22,23
What are the days of the week to run file backup (0-6)? 0,1,2,3,4,5,6

# restore for data in file system instead of DB
```

```
What are the times of the day to poll files from the active site (0-23)?
0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17,18,19,20,21,22,23
What are the days of the week to poll files from the active site (0-6)? 0,1,2,3,4,5,6
```

When you enter the values for all parameters, the DR Watchdog part of the disaster recovery configuration file is displayed.

7. Enter the values for the parameters displayed.

   a. Enter the number of times the standby site should send heartbeat messages to the active site through ping after a heartbeat message times out and press Enter.

      The minimum and default value is 4.

   b. Enter the timeout value for each heartbeat message, in seconds, and press Enter.

      The minimum and default value is 5.

   c. Enter the time interval between two consecutive heartbeat messages to the active site, in seconds, and press Enter.

      The minimum and default value is 30.

   d. Enter the e-mail address of the administrator to whom e-mail messages about disaster recovery service issues must be sent and press Enter.

   e. Enter the time during which the same issues are not reported through e-mail (dampening interval), in seconds, and press Enter.

      The default value is 3,600. The minimum value is 300.

   The following is a sample output:

```
########################
#
# DR Watchdog
#
########################


# heartbeat

What's the number of times to retry heartbeat message? 4
What's the timeout of each heartbeat message (seconds)? 5
What's the heartbeat message interval between sites (seconds)? 30

# notification
```

```
What's the contact email address of service issues? user1@example.com
What's the dampening interval between emails of affected services (seconds)? 300
Check status of DR remote site: up
Load /var/cache/jmp-geo/incoming/
init.properties                                                  [ OK ]
Configure contact
email                                                                  [ OK ]
Modify firewall for DR remote
IPs                                                                 [ OK ]
Configure
NTP
[ OK ]
Sync jmp-geo
group
[ OK ]
Configure MySQL
database                                                                 [ OK ]
Configure PostgreSQL
database                                                               [ OK ]
Command completed.
```

When you have entered values for all parameters, disaster recovery is initialized at the standby site.

Starting the Disaster Recovery Process

You use the `jmp-dr start` command to start the disaster recovery process at both sites. You can also use the `jmp-dr start-a` command to start the disaster recovery process on the active site and the `jmp-dr start-s` command to start the disaster recovery process on the standby site.

NOTE: When you trigger DR start operation from the active site by choosing option both through the GUI and the operation completes without triggering start on standby site due to any network or environmental issue, apply the following workaround:

`Workaround:Login to the CLI on target Standby site VIP node, and use the jmp-dr start -s command.`

To start the disaster recovery process:

1. Log in to the CLI of the Junos Space node at the active site on which the VIP or the eth0:0 interface is configured.

   The Junos Space Settings Menu is displayed.

2. Enter **7**while using a virtual appliance at the Junos Space Settings Menu prompt to run shell commands.

   You are prompted to enter the administrator password.

3. Enter the administrator password.

4. Enter `jmp-dr start` at the shell prompt.

   The disaster recovery process is initiated on both sites.

   The following is a sample output at the active site:

```
[user1@host]# jmp-dr start
Stop dr-watchdog if it's
running                                                          [ OK ]
Check status of DR remote site: up
Check current DR role: active

INFO: => start DR at current site: active

Add device management IPs of DR remote site to up
devices                                          [ OK ]
Setup MySQL replication: master-
master                                                       [ OK ]
Start MySQL
dump
[ OK ]
Setup PostgreSQL
replication                                                            [ OK ]
Start file & RRD
replication                                                            [ OK ]
Open firewall for device
traffic                                                          [ OK ]
Start services(jboss,jboss-
dc,etc.)                                                     [ OK ]
Start dr-
watchdog
[ OK ]
Copy files to DR slave
site                                                               [ OK ]
Update DR role of current site:
active                                                       [ OK ]

INFO: => start DR at DR remote site: standby
```

```
Stop dr-watchdog if it's
running                                                          [ OK ]
Check status of DR remote site: up
Check current DR role: standby
Load /var/cache/jmp-geo/incoming/
start.properties                                               [ OK ]
Stop services(jboss,jboss-
dc,etc.)                                                       [ OK ]
Block firewall for device
traffic                                                        [ OK ]
Reset MySQL init script and stop
replication                                                    [ OK ]
Scp backup file from peer site: /var/cache/jmp-geo/data/
db.gz                                          [ OK ]
Start MySQL
restore
[ OK ]
Setup MySQL replication and start
replication                                                    [ OK ]
Setup PostgreSQL
replication                                                                [ OK ]
Start files & RRD
replication                                                                [ OK ]
Start dr-
watchdog
[ OK ]
Clean up /var/cache/jmp-geo/
incoming                                                           [ OK ]
Update DR role of current site:
standby                                                          [ OK ]
Command completed.
Command completed.
```

The disaster recovery process is initialized on the active site and the standby site.

Verifying the Status of the Disaster Recovery Process

We recommend that you execute the `jmp-dr health` command to verify the status (overall health) of the disaster recovery process at both the active and standby sites when you start the disaster recovery process on both sites. For more information about executing the `jmp-dr health` command, see "Checking the Status of the Disaster Recovery Configuration" on page 100.

RELATED DOCUMENTATION

Disaster Recovery Overview | 40

Modifying the Disaster Recovery Configuration | 107

Modifying Applications and Nodes on a Disaster Recovery Setup | 117

Stopping the Disaster Recovery Process on Junos Space Network Management Platform Release 14.1R3 and Earlier

IN THIS SECTION

To configure the disaster recovery enhancements added as part of the Junos Space Network Management Platform Release 15.2R1, you must first disable the disaster recovery feature on your current Junos Space setup (Junos Space Network Management Platform Release 14.1R3 and earlier) before upgrading to Junos Space Network Management Platform Release 15.2R1. You must stop backups at the active site and the standby site must stop collecting backups from the active site. The scripts to stop the backup and restoration process configured during earlier releases are stored at **/opt/jmp-geo/backup/script/** and **/opt/jmp-geo/restore/script/**.

Stopping the Backup Process at the Active Site

Stopping the backup process at the active site removes the cron job and stops the backup operation from being performed.

To stop the backup process at the active site:

1. Log in to the CLI of the Junos Space node at the active site on which the VIP or the eth0:0 interface is configured.

   The Junos Space Settings Menu is displayed.
2. Enter **7** while using a virtual appliance at the Junos Space Settings Menu prompt to run shell commands.

The following is a sample output from a virtual appliance:

```
admin@10.206.41.183's password:
Last login: Mon Aug 17 06:17:58 2015 from 10.206.41.42

Welcome to the Junos Space network settings utility.

Initializing, please wait


Junos Space Settings Menu

1> Change Password
2> Change Network Settings
3> Change Time Options
4> Retrieve Logs
5> Security
6> Expand VM Drive Size
7> (Debug) run shell

A> Apply changes
Q> Quit
R> Redraw Menu

Choice [1-7,AQR]: 7
```

You are prompted to enter the administrator password.

3. Enter the administrator password.
4. Type `/opt/jmp-geo/backup/script/./backup.sh stop` at the shell prompt and press Enter.

   The following is a sample output:

```
Demoting this cluster from the DR Master Cluster Role ...
update cluster state successful
Stoping backup cron job...
Stopping crond: [ OK ]
Starting crond: [ OK ]
```

The backup process at the active site is stopped.

Stopping Collecting Backups from the Active Site

Stopping the restoration process at the standby site removes the cron job and stops collecting backups from the active site.

To stop the restoration process at the standby site:

1. Log in to the CLI of the Junos Space node at the active site on which the VIP or the eth0:0 interface is configured.

   The Junos Space Settings Menu is displayed.
2. Enter **7** while using a virtual appliance at the Junos Space Settings Menu prompt to run shell commands.

   The following is a sample output from a virtual appliance:

```
admin@10.206.41.183's password:
Last login: Mon Aug 17 06:17:58 2015 from 10.206.41.42

Welcome to the Junos Space network settings utility.

Initializing, please wait


Junos Space Settings Menu

1> Change Password
2> Change Network Settings
3> Change Time Options
4> Retrieve Logs
5> Security
6> Expand VM Drive Size
7> (Debug) run shell

A> Apply changes
Q> Quit
R> Redraw Menu

Choice [1-7,AQR]: 7
```

   You are prompted to enter the administrator password.
3. Enter the administrator password.
4. Type `/opt/jmp-geo/restore/script/./restore.sh stopPoll` at the shell prompt and press Enter.

The following is a sample output:

```
Stoping restore cron job...
Stopping crond: [ OK ]
Starting crond: [ OK ]
Demoting this cluster from the DR Slave Cluster Role ...
update cluster state successful
opening port 7804 on user1@host....
jmp-firewall is stopped. Skip reloading
<response>
<message
</message>
<status>SUCCESS</success>
</response>
```

The standby site stops collecting backups from the active site.

RELATED DOCUMENTATION

Disaster Recovery Overview | 40

Configuring the Disaster Recovery Process Between an Active and a Standby Site | 74

CHAPTER 8

Configuring the Disaster Recovery Process in the GUI

IN THIS CHAPTER

Validate Peer Site

Use the Validate Peer Site page to check the reachability of the peer site, before you add it to the Disaster Recovery (DR) environment.

Before you configure the DR, ensure that your Junos Space installation meets the following prerequisites:

- The Junos Space cluster at the primary or active site (single node or multiple nodes) and the cluster at the remote or standby site (single node or multiple nodes) must have the same configuration, with the same applications, device adapters, same IP family configurations, and so on.
- Passwords used must be valid.
- Both clusters must be configured with SMTP server information from the Junos Space GUI. For more information, see *Managing SMTP Servers*. This configuration enables the clusters at both the active site and the standby site to notify the administrator by e-mail if the replications fail.
- The arbitrary devices used must be reachable.

To validate peer site in active and standby site:

1. Select **Administration** > **Disaster Recovery** > **Validate Peer Site**.

   The Validate Peer Site page appears.
2. Enter the required parameters and select one or more devices from the list that you want to validate. See Table 4 on page 91 for more details on the Validate Peer Site page.

Table 4: Fields on Validate Peer Site page

Field	Description
Peer Site VIP Address	Enter a valid peer site VIP address.
Load Balancer's CLI Admin Password	Enter the correct load balancer password.
Confirm Password	Re-enter the above password.
Arbitrary Devices	Select one or more devices from the list of devices used during the DR auto failover. You can also search and filter the devices.
Device Name	Displays the name of the device.
Device Alias	Displays the alias for the device.
IP Address	Shows the IP addresses for the devices.
Platform	Displays the platform for the devices.
OS Version	Displays the OS version of devices.
Connection Status	Displays the connection status of the devices.
Validate Peer Site	Select to validate the selections and perform the validation. This is enabled when the mandatory fields are filled.
Cancel	Select to cancel the selections and go back to the landing page of DR.

RELATED DOCUMENTATION

Disaster Recovery Overview

Manage Disaster Recovery | 92

Manage Disaster Recovery

IN THIS SECTION

Configuration of Disaster Recovery (DR) between an active site and a standby site ensures geographical redundancy of network management services.

Before you initiate the DR process between both sites, perform the following tasks:

- Ensure that the connectivity requirements as described in the *Disaster Recovery Overview* topic are met.
- Check whether identical cluster configurations exist on both sites. We recommend that both clusters have the same number of nodes so that, even in the case of a disaster, the standby site can operate with the same capacity as the active site.
- Ensure that the same versions of Junos Space Network Management Platform, high-level Junos Space applications, and device adapters are installed at both sites.
- Shut down the DR process configured on Junos Space Network Management Platform Release 14.1R3 and earlier before upgrading to Junos Space Network Management Platform Release 15.2R1 and configuring the new DR process. For more information, see "Stopping the Disaster Recovery Process on Junos Space Network Management Platform Release 14.1R3 and Earlier" on page 86.

  You cannot configure the new DR process if you do not stop the DR you set up on 14.1R3 and earlier releases. You do not need to perform this step on a clean installation of Junos Space Network Management Platform Release 15.2R1.

- Ensure that the same SMTP server configuration exists on both sites to receive e-mail alerts related to the DR process. You can add SMTP servers from the SMTP Servers task group in the Administration workspace. For more information about adding SMTP servers, see *Adding an SMTP Server*.

To configure Disaster Recovery:

1. Select **Administration** > **Disaster Recovery** > **Manage Disaster Recovery**.

   The Configure Disaster Recovery Wizard page opens.

2. Enter the required parameters and select one or more devices from the list that you want to validate. See Table 5 on page 93 for more details on the Configure Disaster Recovery Wizard page.

Table 5: Fields on the Configure Disaster Recovery Wizard Page

Field	Description
Site Role	Select an option for which you want to configure the DR. The available options are Active and Standby Site. **NOTE**: Its is mandatory to initiate the DR on the Active Site first followed by Standby Site or else system prompts you to do so.
Peer Site VIP Address	Enter a valid IP address for configuration. **NOTE**: You cannot edit this information if the DR is not in the Initialized state.
Load Balancer's CLI Admin Password	Enter a valid admin CLI password. **NOTE**: If you have more than one password, you can enter both separated by a comma. You cannot edit this information if the DR is not in the Initialized state.
Confirm Password	Re-enter the previously entered password to configure the DR Wizard.
Arbitrary Devices	Select one or more devices from the list of devices used during DR auto failover. You can also search and filter the devices.
Next	Select Next to configure Disaster Recovery at the Active Site followed by Standby Site. See "Configuring Disaster Recovery at the Active Site" on page 94 and "Configuring Disaster Recovery at the Standby Site" on page 96. It is enabled only when all the parameters are fulfilled.

Next, the window to configure Disaster Recovery at the Active Site followed by Standby Site gets displayed. For more details, see "Configuring Disaster Recovery at the Active Site" on page 94 and "Configuring Disaster Recovery at the Standby Site" on page 96.

The following sections explains the procedure to configure DR at the Active and Standby Sites and initiate the disaster recovery between both sites.

Configuring Disaster Recovery at the Active Site

To configure the Disaster Recovery at the Active Site:

1. Select **Next** after you have filled all the parameters in the Configure Disaster Recovery Wizard page.

   The Configure Disaster Recovery Wizard for Active Site opens.
2. Enter all the required details for the parameters that are displayed on the page. For more details on the fields, see Table 6 on page 94.

Table 6: Fields on the Configure Disaster Recovery Wizard page at the Active Site

Field	Description
Peer Site VIP	Displays the IP address entered in the Configure Disaster Recovery Wizard page.
Arbitrary Devices	Displays all the devices that are selected in the Configure Disaster Recovery Wizard page.
SCP Timeout	Displays the timeout value to detect a failure in transferring files from standby to active site through Secure Copy Protocol (SCP). The time is dispayed in seconds. **NOTE**: You cannot edit the value if DR is not in the Initialized state.
Maximum number of backup	Displays the numbers of files that you want to retain. **NOTE**: You cannot edit the value if DR is not in the Initialized state.
Backup Schedule **NOTE**: You cannot edit the parameters if DR is not in the Initialized state.	
Time of the day (in Hrs)	The time of the day when you want to schedule the backup. Time is in 24 hours format.

Table 6: Fields on the Configure Disaster Recovery Wizard page at the Active Site *(Continued)*

Field	Description
Days of the week	The days when you want to schedule the backup.
Restore Schedule **NOTE**: You cannot edit the parameters if DR is not in the Initialized state.	
Time of the day (in Hrs)	The time of the day to copy files from active site to standby site. Time is in 24 hours format.
Days of the week	The days to copy files from active site to standby site.
Watchdog **NOTE**: You cannot edit the parameters if DR is not in the Initialized state.	
Heartbeat retry times	The number of times the active site should send heartbeat messages to the standby site. It ranges from 4 to 15.
Heartbeat message timeout	The timeout value of each heartbeat message in seconds. The maximum and default value is 5.
Heartbeat message interval	Displays the time interval between two consecutive heartbeat messages to the standby site in seconds, ranging from 30 seconds to 120 seconds.
Notification email	The e-mail address of the administrator to whom e-mail messages about disaster recovery service issues must be sent.
Notification interval	The time interval during which the same issues are not reported through e-mail (dampening interval) in seconds. It ranges from 300 to 1800 seconds.
Failure Detection	

Table 6: Fields on the Configure Disaster Recovery Wizard page at the Active Site *(Continued)*

Field	Description
Failure detection method	Displays the method of failure detection. **NOTE**: In Junos Space Network Management Platform 20.3R1, only default option is allowed through GUI.
Failure detection threshold percentage	Displays the threshold percentage for failure detection.

When you have entered values for all parameters, disaster recovery is initialized at the active site.

Configuring Disaster Recovery at the Standby Site

To configure the Disaster Recovery at the Standby Site:

1. Select **Next** after you have filled all the parameters in the Configure Disaster Recovery Wizard page.
   The Configure Disaster Recovery Wizard for Standby Site opens.
2. Enter all the required details for the parameters that are displayed on the page. For more details on the fields, see Table 7 on page 96.

NOTE: Its mandatory to initialize the Active Site before initializing the Standby Site. Arbitrary devices can be selected only in the Active Site.

Table 7: Fields on the Configure Disaster Recovery Wizard page at the Standby Site

Field	Description
Peer Site VIP	Displays the IP address entered in the Configure Disaster Recovery Wizard page.
Arbitrary Devices	Displays all the devices that are selected in the Configure Disaster Recovery Wizard page.
SCP Timeout	Displays the timeout value to detect a failure in transferring files from standby to active site through Secure Copy Protocol (SCP). The time is dispayed in seconds. **NOTE**: You cannot edit the value if DR is not in the Initialized state.

Table 7: Fields on the Configure Disaster Recovery Wizard page at the Standby Site *(Continued)*

Field	Description
Maximum number of backup	Displays the maximum number of backups to retain at the standby site. **NOTE**: You cannot edit the value if DR is not in the Initialized state.
Backup Schedule **NOTE**: You cannot edit the parameters if DR is not in the Initialized state.	
Time of the day (in Hrs)	The time of the day when you want to schedule the backup. Time is in 24 hours format.
Days of the week	The days when you want to schedule the backup.
Restore Schedule **NOTE**: You cannot edit the parameters if DR is not in the Initialized state.	
Time of the day (in Hrs)	The time of the day to copy files from active site to standby site. Time is in 24 hours format.
Days of the week	The days to copy files from active site to standby site.

When you have entered values for all parameters, disaster recovery is initialized at the standby site.

Actions common for both Active and Standby Site

Table 8 on page 97 shows the actions common for configuring both Active and Standby Sites.

Table 8: Actions common for both Active and Standby Site configuration

Field	Action
Initialize	Starts the initialization of DR with the given values. This is enabled only when all the parameters are provided with correct vales on both the sites.

Table 8: Actions common for both Active and Standby Site configuration *(Continued)*

Field	Action
Reset	Resets the DR configuration. This is enabled only when the DR is already initialized or else stopped.
Start	Starts the DR process. This is enabled when the DR is already initialized. **NOTE**: When you trigger DR start operation from the active site by choosing option **both** through the GUI and the operation completes without triggering start on standby site due to any network or environmental issue, apply the following workaround: Workaround:Login to the CLI on target Standby site VIP node, and use the jmp-dr start -s command
Stop	Allows you to stop the configuration on either of the sites or both the sites.
Manual Failover	This performs manual fail over on the standby site. This parameter is available only when the DR has started or is stopped.

Disaster Recovery Health

To check the Disaster Recovery health status:

1. Select **Administration** > **Disaster Recovery**.

   The landing page opens with a graphical representation of both the Active and Standby Site.
2. Right click on the site you want to check the health status.

   The options available are Current Configuration, Health and Start.
3. Select **Health**.

   The health report status for the selected site is displayed. The report shows the last verified status for a particular site with the date and time of generation of the report.
4. Select **Trigger Health Report** to check the current health report status for the selected site.

   The Health Command starts and after completion, it shows all the relevant messages with their status.

RELATED DOCUMENTATION

Disaster Recovery Overview

Validate Peer Site | 90

CHAPTER 9

Managing the Disaster Recovery Solution

IN THIS CHAPTER

Checking the Status of the Disaster Recovery Configuration

You check the status of the disaster recovery configuration:

- After starting the disaster recovery process to ensure that the disaster recovery configuration is accurate, files are being replicated, and the disaster recovery watchdog is monitoring the disaster recovery setup
- After stopping the disaster recovery process, to ensure that file replication and disaster recovery watchdog process have stopped

You execute the `jmp-dr health` command to check the status of the disaster recovery configuration. This command checks the status of asynchronous data replication, file transfer, and disaster recovery watchdog, and the role of clusters in the disaster recovery setup. Errors found during command execution are listed in the command output.

NOTE: If you have already executed the `jmp-dr health` command and the execution is in progress, executing another `jmp-dr health` command can display incorrect output. The output from the `jmp-dr health` command also lists whether another instance of the command is being executed.

To check the status of the disaster recovery configuration at a site:

1. Log in to the CLI of the Junos Space node at the site on which the VIP or the eth0:0 interface is configured.

   The Junos Space Settings Menu is displayed.

2. Enter **7**while using a virtual appliance at the Junos Space Settings Menu prompt to run shell commands.

   The following is a sample output from a virtual appliance:

```
admin@10.206.41.183's password:
Last login: Mon Aug 17 06:17:58 2015 from 10.206.41.42

Welcome to the Junos Space network settings utility.

Initializing, please wait


Junos Space Settings Menu

1> Change Password
2> Change Network Settings
3> Change Time Options
4> Retrieve Logs
5> Security
6> Expand VM Drive Size
7> (Debug) run shell

A> Apply changes
Q> Quit
R> Redraw Menu

Choice [1-7,AQR]: 7
```

   You are prompted to enter the administrator password.

3. Enter the administrator password.

4. Enter `jmp-dr health` at the shell prompt.

   Junos Space Platform checks the status (overall health) of the disaster recovery configuration at the site.

   - The following is a sample output of the `jmp-dr health` command after you start the disaster recovery process and execute the command at the active site:

```
[user1@host]# jmp-dr health
The DR role of this site: active
The DR state of this site: started
The status of DR watchdog: ready
The status of DR remote site: up
Check admin password of DR remote
site                                                                [ OK ]
Check mysql ca and server
certificates                                                               [ OK ]
Check mysql replication
users                                                                        [ OK ]
Check file replication: backup cron job should be
added                                                  [ OK ]
Check mysql
replication:
   node (10.206.41.221, user1@host) and peer node of same
site
   (10.206.41.222, user2@host) should be master-
master                                              [ OK ]
                                           [ OK ]
Services (jboss, jboss-dc, etc.) should be
up                                                          [ OK ]
DR watchdog should be
up                                                                           [ OK ]
Command completed.
```

   - The following is a sample output of the `jmp-dr health` command after you start the disaster recovery process and execute the command at the standby site:

```
[user3@host]# jmp-dr health
The DR role of this site: standby
The DR state of this site: started
The status of DR watchdog: ready
The status of DR remote site: up
```

```
Check admin password of DR remote
site                                                                [ OK ]
Check mysql ca and server
certificates                                                              [ OK ]
Check mysql replication
users                                                                       [ OK ]
Check file replication: poll cron job should be
added                                                  [ OK ]
Check mysql
replication:
   node (10.206.41.226, user3@host) and peer node of same
site
   (10.206.41.227, user4@host) should be master-
slave,
   and node (10.206.41.226, user3@host) and mysql VIP node of remote
site
   (10.206.41.220) should be slave-
master                                                              [ OK ]

                                                                 [ OK ]
Services (jboss, jboss-dc, etc.) should be
down                                                          [ OK ]
DR watchdog should be
up                                                                          [-]
 [ OK ]
Command completed.
```

- The following is a sample output of the `jmp-dr health` command after you stop the disaster recovery process and execute the command at the active site:

```
[user2@host]# jmp-dr health
The DR role of this site: active
The DR state of this site: stopped
The status of DR watchdog: ready
The status of DR remote site: up
Check admin password of DR remote
site                                                                [ OK ]
Check mysql ca and server
certificates                                                              [ OK ]
Check mysql replication
users                                                                       [ OK ]
Check file replication: cron jobs should be
```

```
removed                                                    [ OK ]
Check mysql
replication:
   node (10.206.41.222, user2@host) and peer node of same
site
   (10.206.41.221, user1@host) should be master-
master                                          [ OK ]
Services (jboss, jboss-dc, etc.) should be
up                                                          [ OK ]
DR watchdog should be
down                                                                  [ OK ]
Command completed.
```

- The following is a sample output of the `jmp-dr health` command after you stop the disaster recovery process and execute the command at the standby site:

```
[user3@host]# jmp-dr health
The DR role of this site: standby
The DR state of this site: stopped
The status of DR watchdog: ready
The status of DR remote site: up
Check admin password of DR remote
site                                                          [ OK ]
Check mysql ca and server
certificates                                                         [ OK ]
Check mysql replication
users                                                                  [ OK ]
Check file replication: cron jobs should be
removed                                                    [ OK ]
Check mysql
replication:
   node (10.206.41.226, user3@host) and peer node of same
site
   (10.206.41.227, user4@host) should be master-
master                                          [ OK ]
Services (jboss, jboss-dc, etc.) should be
down                                                        [ OK ]
DR watchdog should be
down                                                                  [ OK ]
Command completed.
```

RELATED DOCUMENTATION

Configuring the Disaster Recovery Process Between an Active and a Standby Site | 74

Resetting the Disaster Recovery Configuration | 136

Stopping the Disaster Recovery Process | 133

Viewing the Disaster Recovery Configuration and Status of Watchdog Services

You execute the `jmp-dr api v1 config` command to view the disaster recovery configuration and the status of the disaster recovery watchdog services at the local site. You can use this command to create custom failure-detection scripts. For more information about using custom failure-detection scripts, see the *Failure Detection by Using Custom Failure-Detection Scripts* section in the "Disaster Recovery Overview" on page 40 topic. You can also refer to the sample scripts located at **var/cache/jmp-geo/doc/samples/**.

To view the disaster recovery configuration and the status of the disaster recovery watchdog services:

1. Log in to the CLI of the Junos Space node at the site on which the VIP or the eth0:0 interface is configured.

   The Junos Space Settings Menu is displayed.

2. Enter **7** while using a virtual appliance at the Junos Space Settings Menu prompt to run shell commands.

   The following is a sample output from a virtual appliance:

```
admin@10.206.41.183's password:
Last login: Mon Aug 17 06:17:58 2015 from 10.206.41.42

Welcome to the Junos Space network settings utility.

Initializing, please wait


Junos Space Settings Menu

1> Change Password
2> Change Network Settings
3> Change Time Options
4> Retrieve Logs
5> Security
```

```
6> Expand VM Drive Size
7> (Debug) run shell

A> Apply changes
Q> Quit
R> Redraw Menu

Choice [1-7,AQR]: 7
```

You are prompted to enter the administrator password.

3. Enter the administrator password.
4. Enter `jmp-dr api v1 config` at the shell prompt.

   You can view the disaster recovery configuration and the status of all the disaster recovery watchdog services.
5. (Optional) To view the sections of the disaster recovery configuration, enter `jmp-dr api v1 config --list` at the shell prompt.

   All available sections of the disaster recovery configuration from the remote site are displayed.

   The following is a sample output:

```
[user1@host]# jmp-dr api v1 config --list
{
  "sections": "role,failover,states,core,mysql,psql,file,watchdog,deviceManagement"
}
```

6. (Optional) To view selected sections of the disaster recovery configuration, enter `jmp-dr api v1 config --include <section1>,<section2>` at the shell prompt.

   The following is a sample output of the core and deviceManagement sections:

```
[user1@host]# jmp-dr api v1 config --include core,deviceManagement
{
  "core": {
    "peerVip": "10.206.41.41",
    "adminPass": "53616c7465645f5f737061636513233126c3f3e6fd6257a81cded28f55d465c",
    "scpTimeout": 120,
    "peerLoadBalancerNodes": "10.206.41.42,10.206.41.44",
    "peerBusinessLogicNodes": "10.206.41.42,10.206.41.44,10.206.41.50",
    "peerDeviceMgtIps": "10.206.41.42,10.206.41.44,10.206.41.50"
  },
  "deviceManagement": {
    "peerNodes": "10.206.41.42,10.206.41.44,10.206.41.50",
```

```
    "nodes": "10.206.41.182,10.206.41.183,10.206.41.184",
    "ip": "10.206.41.183,eth0"
  }
}
```

RELATED DOCUMENTATION

Configuring the Disaster Recovery Process Between an Active and a Standby Site | 74

Checking the Status of the Disaster Recovery Configuration | 100

Stopping the Disaster Recovery Process | 133

Modifying the Disaster Recovery Configuration

After you have initially configured the disaster recovery setup, you may need to modify the Junos Space cluster at either sites such as addition or removal of nodes from your Junos Space setup, change IP addresses of interfaces, change device management interfaces, load balancer details (VIP address and password), and change VIP address of dedicated services such as FMPM or database. You may also need to modify disaster recovery parameters such as backup and restore settings, heartbeat settings, and failure detection settings.

You use the `jmp-dr toolkit config update` command to:

- Reinitiate MySQL replication at a site and to the standby site if you added or removed dedicated nodes and update these changes at the standby site.
- Update the changes in cluster configuration (addition or removal of load balancer nodes), at both sites.

  Refer to "Modifying Applications and Nodes on a Disaster Recovery Setup" on page 117 for more information about modifying nodes.

You use the options along with `jmp-dr toolkit config update` command to modify backup and restore settings, heartbeat settings, failure detection settings, update load balancer details (VIP address and password), and SCP timeout settings, and update these changes at the peer site.

NOTE: You must update the changes to the load balancers at both sites by using the `--user-core` option before modifying and updating the other sections of disaster recovery such as heartbeat settings, notification settings, failure detection settings, file backup and restore settings.

Table 9 on page 108 lists the options and the group of disaster recovery configuration parameters included with the option.

Table 9: jmp-dr toolkit config update command options

Configuration Update Option	Description
`--user-core`	Load balancer VIP and password, and SCP timeout settings
`--user-file-backup`	Configuration and RRD files backup settings
`--user-file-restore`	Configuration and RRD files replication to standby site settings
`--user-watchdog-heartbeat`	Watchdog heartbeat settings
`--user-watchdog-notification`	Email notification settings
`--user-watchdog-failureDetection`	Failure detection settings

To modify the disaster recovery configuration:

1. Log in to the CLI of the Junos Space node at the active site on which the VIP or the eth0:0 interface is configured.

   The Junos Space Settings Menu is displayed.
2. Enter **7** while using a virtual appliance at the Junos Space Settings Menu prompt to run shell commands.

   The following is a sample output from a virtual appliance:

```
admin@10.206.41.183's password:
Last login: Mon Aug 17 06:17:58 2015 from 10.206.41.42

Welcome to the Junos Space network settings utility.

Initializing, please wait


Junos Space Settings Menu

1> Change Password
```

```
2> Change Network Settings
3> Change Time Options
4> Retrieve Logs
5> Security
6> Expand VM Drive Size
7> (Debug) run shell

A> Apply changes
Q> Quit
R> Redraw Menu

Choice [1-7,AQR]: 7
```

3. Enter the administrator password.
4. Stop the disaster recovery process on both sites. To do so, type `jmp-dr stop` at the shell prompt and press Enter.
5. Perform the following steps to modify the disaster recovery configuration at sites.
   - Use the `jmp-dr toolkit config update` command to:
     - Update the addition or removal of load balancer nodes.
     - Reinitiate MySQL replication at a site (after adding or removing dedicated nodes).
     - Update changes to VIP address of FMPM node or database node of a site, at both sites.
     - Update the changes to IP address of the nodes of a site, at both sites.
     - Update the addition or removal of eth3 interface for device management or change in the IP address of the eth3 interface, at both sites.
     - Update the change of IP address version for device management (IPv4 to IPv6 or IPv6 to IPv4) of a site, at both sites.

   a. Log in to the CLI of the Junos Space VIP node at the site where you made the modifications.

   b. Type `jmp-dr toolkit config update` at the shell prompt and press Enter.

   c. Type No and press Enter to view the cluster modifications.

      The modified cluster configuration is displayed in JSON format.

   d. Press Enter to accept the changes.

The following is a sample output when the disaster recovery configuration is updated after adding a dedicated MySQL node.

```
[user1@host]# jmp-dr toolkit config update
If admin password of any node belonging to remote site is changed or if a new node
with different admin password is added then please use option --user-core. Continue?
Yes
The modified user configuration in JSON format is as follows:
{
  "user_mysql_hasDedicatedDb": {
    "lhs": false,
    "rhs": true
  },
  "user_mysql_peerVip": {
    "lhs": "10.206.41.225",
    "rhs": "10.206.41.84"
  },
  "user_mysql_peerNodes": {
    "lhs": "10.206.41.226,10.206.41.227",
    "rhs": "10.206.41.85,10.206.41.86"
  }
}

Do you want to apply these changes? Yes
Check status of DR remote site: up
Stop mysql replication if
applicable                                                              [ OK ]
Update mysql repUser &
repAdmin                                                                   [ OK ]
Update
firewall
        [ OK ]
Update
ntp
        [ OK ]
Update mysql configuration if
applicable                                                            [ OK ]
Update services (such as jboss-dc, httpd,
etc.)                                                        [ OK ]
The configuration change is updated only at current site, please ensure to update at
```

```
the remote site accordingly.
The `toolkit config` command is done
```

The disaster recovery configuration file at the local site is updated with the modified configuration of the cluster.

e. Log in to the CLI of the Junos Space VIP node at the peer site to update the modifications made at the local site.

f. Type `jmp-dr toolkit config update` at the shell prompt and press Enter.

g. Type No and press Enter to view the modified cluster configuration at the remote site.

The modified configuration is displayed in JSON format.

h. Press Enter to accept the changes.

The disaster recovery configuration file at the peer site is updated with the modified configuration of the cluster at the local site.

- To modify the heartbeat settings at a site:

a. Log in to the CLI of the Junos Space VIP node at the site where the heartbeat settings must be modified.

b. Type `jmp-dr toolkit config update --user-watchdog-heartbeat` at the shell prompt and press Enter.

c. Type No and press Enter to modify the heartbeat settings.

The disaster recovery configuration parameters to modify heartbeat settings are displayed.

The following is a sample screen output.

```
? If admin password of any node belonging to remote site is changed or if a new node
with different admin password is added then please use option --user-core. Continue?
Yes

########################
#
# DR Watchdog
#
########################


# heartbeat
```

```
? What's the number of times to retry heartbeat message? 4
? What's the timeout of each heartbeat message (seconds)? 5
? What's the heartbeat message interval between sites (seconds)? 30
```

d. Modify the heartbeat settings.

   The modified heartbeat settings are displayed in JSON format.

e. Press Enter to accept the changes.

   The heartbeat settings are modified when the command is executed.

   The following is a sample screen output.

```
The configuration change is updated only at current site, please ensure to update at
the remote site accordingly.
Command completed.
```

f. Log in to the CLI of the Junos Space VIP node at the peer site to update the modifications made at the local site.

g. Type `jmp-dr toolkit config update --user-watchdog-heartbeat` at the shell prompt and press Enter.

h. Type No and press Enter to view the modified heartbeat settings at the local site.

   The modified heartbeat settings are displayed in JSON format.

i. Press Enter to accept the changes.

   The heartbeat settings are updated when the command is executed.

- To modify the notification settings at a site:

  a. Log in to the CLI of the Junos Space VIP node at the site where the notification settings must be modified.

  b. Type `jmp-dr toolkit config update --user-watchdog-notification` at the shell prompt and press Enter.

  c. Type No and press Enter to modify the notification settings.

     The disaster recovery configuration parameters to modify notification settings are displayed.

  d. Modify the notification settings.

     The modified notification settings are displayed in JSON format.

  e. Press Enter to accept the changes.

The notification settings are modified when the command is executed.

f. Log in to the CLI of the Junos Space VIP node at the peer site to update the modifications made at the local site.

g. Type `jmp-dr toolkit config update --user-watchdog-notification` at the shell prompt and press Enter.

h. Type No and press Enter to view the modified notification settings at the local site.

The modified notification settings are displayed in JSON format.

i. Press Enter to accept the changes.

The notification settings are updated when the command is executed.

- To modify the failure detection settings at the active site:

a. Log in to the CLI of the Junos Space VIP node at the active site.

b. Type `jmp-dr toolkit config update --user-watchdog-failureDetection` at the shell prompt and press Enter.

c. Type No and press Enter to modify the failure detection settings.

The disaster recovery configuration parameters to modify failure detection settings are displayed.

d. Modify the failure detection settings.

The modified failure detection settings are displayed in JSON format.

e. Press Enter to accept the changes.

The failure detection settings are modified when the command is executed.

The following is a sample screen output.

```
? If admin password of any node belonging to remote site is changed or if a new node
with different admin password is added then please use option --user-core. Continue?
Yes

########################
#
# DR Watchdog
#
########################
```

```
# failureDetection

? Do you want to use custom failure detection? No
? What's the threshold percentage to trigger failover? 50
? What's the arbiters list file (note: please refer to example in /var/cache/jmp-
geo/doc/samples/arbiters.list)? /var/cache/jmp-geo/doc/arbiters.list
The modified user configuration in JSON format is as follows:
{
  "user_watchdog_failureDetection_arbiters": {
    "lhs": "/var/cache/jmp-geo/config/arbiters.list",
    "rhs": "/var/cache/jmp-geo/doc/arbiters.list"
  }
}

? Do you want to apply these changes? Yes
Check status of DR remote site: up
Update MySQL configuration if
applicable                                                          [ OK ]
Update services (such as jboss-dc, httpd,
etc.)                                                    [ OK ]
The configuration change is updated only at current site, please ensure to update at
the remote site accordingly.
Command completed.
```

f. Log in to the CLI of the Junos Space VIP node at the peer site to update the modifications made at the active site.

g. Type `jmp-dr toolkit config update --user-watchdog-failureDetection` at the shell prompt and press Enter.

h. Type No and press Enter to view the modified failure detection settings at the local site.

   The modified failure detection settings are displayed in JSON format.

i. Press Enter to accept the changes.

   The failure detection settings are updated when the command is executed.

- To modify the file backup settings at a site:

  a. Log in to the CLI of the Junos Space VIP node at the site where the file backup settings must be modified.

  b. Type `jmp-dr toolkit config update --user-file-backup` at the shell prompt and press Enter.

  c. Type No and press Enter to modify the file backup settings.

The disaster recovery configuration parameters to modify file backup settings are displayed.

d. Modify the file backup settings.

The modified file backup settings are displayed in JSON format.

e. Press Enter to accept the changes.

The file backup settings are modified when the command is executed.

f. Log in to the CLI of the Junos Space VIP node at the peer site to update the modifications made at the local site.

g. Type `jmp-dr toolkit config update --user-file-backup` at the shell prompt and press Enter.

h. Type No and press Enter to view the modified file backup settings at the local site.

The modified file backup settings are displayed in JSON format.

i. Press Enter to accept the changes.

The file backup settings are updated when the command is executed.

- To modify the file restore settings at a site:

a. Log in to the CLI of the Junos Space VIP node at the site where the file restore settings must be modified.

b. Type `jmp-dr toolkit config update --user-file-restore` at the shell prompt and press Enter.

c. Type No and press Enter to modify the file restore settings.

The disaster recovery configuration parameters to modify file restore settings are displayed.

d. Modify the file restore settings.

The modified file restore settings are displayed in JSON format.

e. Press Enter to accept the changes.

The file restore settings are modified when the command is executed.

f. Log in to the CLI of the Junos Space VIP node at the peer site to update the modifications made at the local site.

g. Type `jmp-dr toolkit config update --user-file-restore` at the shell prompt and press Enter.

h. Type No and press Enter to view the modified file restore settings at the local site.

The modified file restore settings are displayed in JSON format.

i. Press Enter to accept the changes.

   The file restore settings are updated when the command is executed.

- Use the `--user-core` option to:
  - Update the modified VIP address of the load balancers of a site at the peer site.
  - Update the modified password of the load balancers of the standby a site at the active site.

    NOTE: You can use the `--user-core` option to update the modified password of the active site at the standby site.

    Refer to *Modifying the Network Settings of a Node in the Junos Space Fabric* in the *Junos Space Network Management Platform Workspaces Feature Guide* for more information about modifying the VIP address of the load balancers.

  - Modify the SCP timeout settings at a site.

a. Log in to the CLI of the Junos Space VIP node at the site where the load balancer details are modified or where the SCP timeout settings must be modified.

b. Type `jmp-dr toolkit config update --user-core` at the shell prompt and press Enter.

c. Press Enter to update the load balancer modifications or modify the SCP timeout settings.

   The disaster recovery configuration parameters to modify load balancer settings are displayed.

d. Modify the load balancer settings.

   The modified load balancer settings are displayed in JSON format.

e. Press Enter to apply the changes.

   The load balancer settings are modified when the command is executed.

   The following is a sample screen output.

```
[user1@host]# jmp-dr toolkit config update --user-core

#########################
#
# Load Balancers
#
#########################
```

```
? What's the vip for load balancers at the standby site? 10.206.41.101
? What are the unique admin passwords for load balancer nodes at the standby site
(separated by comma, no space)? $ABC123
? What's the scp timeout value (seconds)? 120
Check status of DR remote site: up
Update MySQL configuration if
applicable                                                          [ OK ]
Update services (such as jboss-dc, httpd,
etc.)                                                    [ OK ]
The configuration change is updated only at current site, please ensure to update at
the remote site accordingly.
Command completed.
```

f. Log in to the CLI of the Junos Space VIP node at the peer site.

g. Type `jmp-dr toolkit config update --user-core` at the shell prompt and press Enter.

   The modified load balancer settings are displayed in JSON format.

h. Press Enter to accept the changes.

   The load balancer settings are modified when the command is executed.

6. Start the disaster recovery process on both sites from the active site. To do so, type `jmp-dr start` at the shell prompt and press Enter.

RELATED DOCUMENTATION

- Configuring the Disaster Recovery Process Between an Active and a Standby Site | 74
- Resetting the Disaster Recovery Configuration | 136
- Stopping the Disaster Recovery Process | 133

Modifying Applications and Nodes on a Disaster Recovery Setup

IN THIS SECTION

You need to log in to the Junos Space user interface and initiate Junos Space Platform workflows to modify the applications and nodes.

NOTE: You must stop the disaster recovery process when you make changes to the disaster recovery setup. Missing transactions at the standby site are collected from the active site when you restart the disaster recovery process after modifying the setup.

NOTE: We recommend that you install the same set of applications on both sites to ensure that you can use all the applications from the standby site in case of a failover.

NOTE: When you execute the scripts to install and upgrade Junos Space Platform and Junos Space applications, you must enter only the release version. For example, `/var/www/cgi-bin/executeUpgradeOnDr.pl` *`16.1R1.XX`* and not `/var/www/cgi-bin/executeUpgradeOnDr.pl` *`16.1R1.XX.img`*.

When you execute disaster recovery scripts, ensure that you use only the following special characters to create user names and passwords:

Table 10: Supported Special Characters

Supported Special Characters
!
#

%
*
–
-
=
+
[
{
]
}
;
,
.
/

The following sections contain steps to modify the applications or nodes on a disaster recovery setup.

Upgrading the Junos Space Network Management Platform Software

You upgrade Junos Space Platform at both the active and standby sites to upgrade the version of Junos Space Platform on your Junos Space deployment. You can upgrade Junos Space Platform on both sites as follows:

- Upgrade Junos Space Platform at the standby site before you upgrade Junos Space Platform at the active site. By upgrading the software image on the standby site first, you can verify the software upgrade process without impacting normal operations at the active site. Although you can upgrade the software on the standby site by using scripts, you must manually failover to the standby site to verify the functionality and features of Junos Space Platform from the user interface. By upgrading the software image on the standby site first, you also ensure that the new software and new database schema are first made available on the standby site to enable it to receive new backup files from the active site after upgrading the software on the active site and restarting disaster recovery.
- Upgrade Junos Space Platform at the active site before you upgrade Junos Space Platform at the standby site. By upgrading and testing Junos Space Platform for a duration that allows no disaster recovery functionality on your Junos Space setup, and using the newer version of Junos Space Platform on the active site first, you ensure that all functionality and features accessible through the user interface work as expected. You can then upgrade the software on the standby site by using scripts or by manually failing over to the standby site and upgrading from the user interface.

You execute the `./executeScpImageOnDr.pl` and `./executeUpgradeOnDr.pl` scripts to upgrade Junos Space Platform Release to later releases. You need to stop the disaster recovery process on both sites before uploading and upgrading the software on both sites, reboot all nodes at both sites, and start the disaster recovery process from the active site.

NOTE: See Table 10 on page 118 for information about the usage of supported special characters to create user name and passwords, while executing disaster recovery scripts.

To upgrade Junos Space Platform to a later release:

NOTE: If you are upgrading Junos Space Platform to Release 18.1 from a version earlier than Release 16.1, you must first upgrade Junos Space Platform to Release 16.1, and then upgrade Junos Space Platform Release 16.1 to Release 17.1 or Release 17.2.

If you are upgrading to Junos Space Platform Release 16.1 from an earlier version, follow the steps listed in the "Upgrading to Junos Space Network Management Platform Release 16.1R1" on page 124 section.

NOTE: Before you upgrade Junos Space Platform to Release 18.1, ensure that the time on all Junos Space nodes is synchronized. For information about synchronizing time on Junos Space nodes, see Synchronizing Time Across Junos Space Nodes

1. Log in to the CLI of the Junos Space node at the active site on which the VIP or the eth0:0 interface is configured.

The Junos Space Settings Menu is displayed.

2. Enter **7** while using a virtual appliance at the Junos Space Settings Menu prompt to run shell commands.

   The following is a sample output from a virtual appliance:

```
admin@10.206.41.183's password:
Last login: Mon Aug 17 06:17:58 2015 from 10.206.41.42

Welcome to the Junos Space network settings utility.

Initializing, please wait


Junos Space Settings Menu

1> Change Password
2> Change Network Settings
3> Change Time Options
4> Retrieve Logs
5> Security
6> Expand VM Drive Size
7> (Debug) run shell

A> Apply changes
Q> Quit
R> Redraw Menu

Choice [1-7,AQR]: 7
```

   You are prompted to enter the administrator password.

3. Enter the administrator password.
4. You can start upgrading the software at the active site or the standby site.

   To upgrade the active site first:

   a. Stop the disaster recovery process on both sites. To do so, type `jmp-dr stop` at the shell prompt of the active site and press Enter.

   b. Go to the Junos Space user interface > Administration workspace > Applications page and upload the software image to the active site. The software image file should be listed on the Upgrade Platform page. Refer to Upgrading Junos Space Network Management Platform in the *Junos Space Network Management Platform Workspaces Feature Guide*.

c. Use SCP to copy the software image from the active site to the standby site. To do so, type `/var/www/cgi-bin/executeScpImageOnDr.pl` *`software-image-name`* at the shell prompt at the active site and press Enter.

   The software image is copied from the **/var/cache/jboss/jmp/** directory at the active site to the **/var/cache/jboss/jmp/payloads/** directory at the standby site.

d. Go to the Junos Space user interface > Administration workspace > Applications page and upgrade the software at the active site. Refer to *Upgrading Junos Space Network Management Platform* in the *Junos Space Network Management Platform Workspaces Feature Guide*.

e. When the upgrade is complete and all nodes reboot, check the functionality of Junos Space Platform from the user interface.

f. Upgrade the software on the standby site. To do so, type `/var/www/cgi-bin/executeUpgradeOnDr.pl` *`software-image-name`* at the shell prompt at the active site and press Enter.

g. Verify that the software is upgraded at the standby site as follows:

   - Verify from the log entry in the **install.log** file located at **/var/log/**.
   - Execute the `rpm -qa | grep jmp-` command and verify that the following RPMs are upgraded: **Jmp-nma**, **Jmp-cmp**, **jmp-ems**, and other **jmp**-related RPMs.

h. Reboot all nodes at the standby site from the CLI. To do so, type `reboot` at the shell prompt of each node and press Enter.

i. Since the standby site cannot be accessed through the user interface, you must manually failover to the standby site to access the user interface. To do so, type `jmp-dr manualFailover` at the shell prompt of the standby site and press Enter.

j. Verify the functionality of Junos Space Platform on the standby site.

k. Manually failover to the original active site. To do so, type `jmp-dr manualFailover` at the shell prompt of the current standby site and press Enter.

l. Start the disaster recovery process on both sites from the active site. To do so, type `jmp-dr start` at the shell prompt and press Enter.

To upgrade the standby site first:

a. Since the standby site cannot be accessed through the user interface, you must manually failover to the standby site to access the user interface on the standby site and upgrade the software. To do so, type `jmp-dr manualFailover` at the shell prompt of the standby site and press Enter.

   The standby site is the new active site. From steps 4.b through 4.j the original active site is referred as the standby site and the original standby site is referred as the active site.

b. Stop the disaster recovery process on both sites. To do so, type `jmp-dr stop` at the shell prompt of the active site and press Enter.

c. Go to the Junos Space user interface > Administration workspace > Applications page and upload the software image to the active site. The software image file should be listed on the Upgrade Platform page. Refer to *Upgrading Junos Space Network Management Platform* in the *Junos Space Network Management Platform Workspaces Feature Guide*.

d. Use SCP to copy the software image from the active site to the standby site. To do so, type `/var/www/cgi-bin/executeScpImageOnDr.pl` *`software-image-name`* at the shell prompt at the active site and press Enter.

   The software image is copied from the **/var/cache/jboss/jmp/** directory at the active site to the **/var/cache/jboss/jmp/payloads/** directory at the standby site.

e. Go to the Junos Space user interface > Administration workspace > Applications page and upgrade the software at the active site. Refer to *Upgrading Junos Space Network Management Platform* in the *Junos Space Network Management Platform Workspaces Feature Guide*.

f. When the upgrade is complete and all nodes reboot, check the functionality of Junos Space Platform from the user interface.

g. Upgrade the software on the standby site. To do so, type `/var/www/cgi-bin/executeUpgradeOnDr.pl` *`software-image-name`* at the shell prompt at the active site and press Enter.

h. Verify that the software is upgraded at the standby site as follows:

   - Verify from the log entry in the **install.log** file located at **/var/log/**.
   - Execute the `rpm -qa | grep jmp-` command and verify that the following RPMs are upgraded: **Jmp-nma**, **Jmp-cmp**, **jmp-ems**, and other **jmp**-related RPMs.

i. Reboot all nodes at the standby site from the CLI. To do so, type `reboot` at the shell prompt of each node and press Enter.

j. Since the standby site cannot be accessed through the user interface, you must manually failover to the standby site (the original active site at the start of the upgrade process) to access the user interface. To do so, type `jmp-dr manualFailover` at the shell prompt of the original active site and press Enter.

   The

k. Verify the functionality of Junos Space Platform on the active site.

l. Start the disaster recovery process on both sites from the active site. To do so, type `jmp-dr start` at the shell prompt and press Enter.

Junos Space Platform is upgraded on the active and standby sites.

NOTE: We recommend that you execute the `jmp-dr health` command at both sites and verify the output after starting disaster recovery on the upgraded setup.

Upgrading to Junos Space Network Management Platform Release 16.1R1

You can upgrade to Junos Space Network Management Platform Release 16.1R1 only from Junos Space Platform Release 15.2R2. To upgrade to Junos Space Platform Release 16.1R1 from releases prior to Junos Space Platform Release 15.2R2, you must first upgrade Junos Space Platform to Junos Space Platform Release 15.2R2. For more information about upgrading to Junos Space Platform Release 15.2R2, refer to the Junos Space Network Management Platform Release 15.2R2 Release Notes.

In Junos Space Platform Release 16.1R1, CentOS 6.8 is used as the underlying OS. A direct upgrade of the OS from CentOS 5.9 to CentOS 6.8 is not recommended, therefore, a direct upgrade to Junos Space Platform Release 16.1R1 by using the Junos Space Platform UI is not supported. You must follow a multi-step procedure to upgrade to Junos Space Platform Release 16.1R1.

To upgrade to Junos Space Platform Release 16.1R1 on a setup that has disaster recovery configured, you must upgrade both the active and standby sites by following the procedure outlined in Upgrading to Junos Space Network Management Platform Release 16.1R1 and then reconfigure disaster recovery. For more information about configuring disaster recovery, see "Configuring the Disaster Recovery Process Between an Active and a Standby Site" on page 74.

Installing a Junos Space Application

You install a Junos Space application on both sites to add the application to your Junos Space deployment. You execute the `./executeScpImageOnDr.pl` and `./executeInstallAppOnDr.pl` scripts to install an application.

NOTE: See Table 10 on page 118 for information about the usage of supported special characters to create user name and passwords, while executing disaster recovery scripts.

To install a Junos Space application on both sites:

1. Log in to the CLI of the Junos Space node at the active site on which the VIP or the eth0:0 interface is configured.

   The Junos Space Settings Menu is displayed.
2. Enter **7** while using a virtual appliance at the Junos Space Settings Menu prompt to run shell commands.

   You are prompted to enter the administrator password.
3. Enter the administrator password.

4. Stop the disaster recovery process on both sites. To do so, type `jmp-dr stop` at the shell prompt and press Enter.
5. Go to the Junos Space user interface > Administration workspace > Applications page to upload the application image to the active site. Refer to the Adding a Junos Space Application workflow in the *Junos Space Network Management Platform Workspaces Feature Guide*.
6. Use SCP to copy the application image to the standby site from the active site. To do so, type `/var/www/cgi-bin/./executeScpImageOnDr.pl` *`application-image-name`* at the shell prompt at the active site and press Enter.

   The application image is copied from the **/var/cache/jboss/jmp/** directory at the active site to the **/var/cache/jboss/jmp/payloads/** directory at the standby site.
7. Go to the Junos Space user interface > Administration workspace > Applications page to install the application on the active site. Refer to *Adding a Junos Space Application* in the *Junos Space Network Management Platform Workspaces Feature Guide*.
8. Go to the Junos Space user interface > Job Management page to verify that the application is installed on the active site.
9. Install the application image on the standby site. To do so, type `/var/www/cgi-bin/./executeInstallAppOnDr.pl` *`application-image-name`* at the shell prompt at the active site and press Enter.
10. Verify the following on the standby site:
    - RPMs of the application are installed. To verify, execute the following command: `rpm -qa | grep <application-rpm-name>`.
    - **.ear** files of the application are available at **/usr/local/jboss/standalone/deployments/**.
11. Start the disaster recovery process on both sites from the active site. To do so, type `jmp-dr start` at the shell prompt of the VIP node at the active site and press Enter.

The Junos Space application is installed on the active and standby sites.

Upgrading a Junos Space Application

You upgrade a Junos Space application on both sites to upgrade the application on your Junos Space deployment. You execute the `./executeScpImageOnDr.pl` and `./executeInstallAppOnDr.pl` scripts to upgrade a Junos Space application.

To upgrade a Junos Space application on both sites:

1. Log in to the CLI of the Junos Space node at the active site on which the VIP or the eth0:0 interface is configured.

   The Junos Space Settings Menu is displayed.
2. Enter **7** while using a virtual appliance) at the Junos Space Settings Menu prompt to run shell commands.

   You are prompted to enter the administrator password.
3. Enter the administrator password.

4. Stop the disaster recovery process on both sites. To do so, type `jmp-dr stop` at the shell prompt and press Enter.
5. Go to the Junos Space user interface > Administration workspace > Applications page to upload the application image to the active site. Refer to the Upgrading a Junos Space Application workflow in the *Junos Space Network Management Platform Workspaces Feature Guide*.
6. Use SCP to copy the application image to the standby site from the active site. To do so, type `/var/www/cgi-bin/./executeScpImageOnDr.pl` *`application-image-name`* at the shell prompt at the active site and press Enter.
7. Go to the Junos Space user interface > Administration workspace > Applications page to upgrade the application on the active site. Refer to *Upgrading a Junos Space Application* in the *Junos Space Network Management Platform Workspaces Feature Guide*.
8. Go to the Junos Space user interface > Job Management page to verify that the application is upgraded on the active site.
9. Upgrade the application on the standby site. To do so, type `/var/www/cgi-bin/./executeInstallAppOnDr.pl` *`application-image-name`* at the shell prompt of the active site and press Enter.
10. Verify the following on the standby site:
    - RPMs of the application are installed. To verify, execute the following command: `rpm -qa | grep <application-rpm-name>`.
    - **.ear** files of the application are available at **/usr/local/jboss/standalone/deployments/**.
11. Start the disaster recovery process on both sites from the active site. To do so, type `jmp-dr start` at the shell prompt of the VIP node at the active site and press Enter.

The Junos Space application is upgraded on the active and standby sites.

Uninstalling a Junos Space Application

You uninstall a Junos Space application from both sites to remove the application from your Junos Space deployment. You execute the `./executeUninstallAppOnDr.pl` script to uninstall an application.

To uninstall a Junos Space application from both sites:

1. Log in to the CLI of the Junos Space node at the active site on which the VIP or the eth0:0 interface is configured.

   The Junos Space Settings Menu is displayed.
2. Enter **7**while using a virtual appliance at the Junos Space Settings Menu prompt to run shell commands.

   You are prompted to enter the administrator password.
3. Enter the administrator password.
4. Stop the disaster recovery process on both sites. To do so, type `jmp-dr stop` at the shell prompt and press Enter.

5. Go to the Junos Space user interface > Administration workspace > Applications page to uninstall the application from the active site. Refer to *Uninstalling a Junos Space Application* in the *Junos Space Network Management Platform Workspaces Feature Guide*.
6. Go to the Junos Space user interface > Job Management page to verify that the application is completely removed from the active site.
7. Uninstall the application from the standby site. To do so, type `/var/www/cgi-bin/./executeUninstallAppOnDr.pl` *`ear-filename`* at the shell prompt at the standby site and press Enter.

   NOTE: You must add the filename without the extension (.ear) as follows:

   `/var/www/cgi-bin/executeUninstallAppOnDr.pl aim`

8. Verify the following on the standby site:
   - All database-related application data is removed.
   - RPMs of the application are removed. To verify, execute the following command: `rpm -qa | grep <application-rpm-name>`.
   - **.ear** files related to the application under **/usr/local/jboss/standalone/deployments/** are removed.
9. Start the disaster recovery process on both sites from the active site. To do so, type `jmp-dr start` at the shell prompt of the VIP node at the active site and press Enter.

The Junos Space application is uninstalled from the active and standby sites.

Adding or Removing a JBoss Node

We recommend that you meet the prerequisites to add a JBoss node or to know the impact of removing a JBoss node from the Junos Space setup. Refer to the *Adding a Node to an Existing Junos Space Fabric* and *Deleting a Node from the Junos Space Fabric* topics in the *Junos Space Network Management Platform Workspaces Feature Guide*. We also recommend that the active site and standby site be symmetric to ensure that the performance of the disaster recovery solution is effective and stable.

You add a JBoss node to improve the performance of your Junos Space setup. You remove a JBoss node if it is faulty and needs to be replaced. You may need to modify the disaster recovery configuration depending on why you added or removed the JBoss node to or from the site.

To add or remove a JBoss node to or from both sites:

1. Log in to the CLI of the Junos Space node at the active site on which the VIP or the eth0:0 interface is configured.

   The Junos Space Settings Menu is displayed.
2. Enter **7** while using a virtual appliance at the Junos Space Settings Menu prompt to run shell commands.

You are prompted to enter the administrator password.

3. Enter the administrator password.
4. Stop the disaster recovery process on both sites. Type `jmp-dr stop` at the active site shell prompt and press Enter.
5. Go to the Junos Space user interface > Administration workspace > Fabric page to add or remove the JBoss node to or from the active site. Refer to the *Adding a Node to an Existing Junos Space Fabric* and *Deleting a Node from the Junos Space Fabric* topics in the *Junos Space Network Management Platform Workspaces Feature Guide*.
6. Log in to the CLI of the Junos Space node at the standby site on which the VIP or the eth0:0 interface is configured.

   The Junos Space Settings Menu is displayed.
7. Enter 7 while using a virtual appliance at the Junos Space Settings Menu prompt to run shell commands.

   You are prompted to enter the administrator password.
8. Enter the administrator password.
9. Update the disaster recovery configuration on the standby site. Type `jmp-dr toolkit config update` at the shell prompt of the VIP node at the standby site and press Enter.
10. Configure the current standby site as the active site. Type `jmp-dr manualFailover` at the standby site shell prompt and press Enter. For more information, see "Manually Failing Over the Network Management Services to the Standby Site" on page 130.
11. Go to the **Junos Space user interface** > **Administration workspace** > **Fabric page** to add or remove the JBOSS node to or from the current active site.
12. Update the disaster recovery configuration on the current standby site. Type `jmp-dr toolkit config update` at the shell prompt of the VIP node at the current standby site and press Enter.
13. Configure the original active site back as the active site again. Type `jmp-dr manualFailover` at the current standby site shell prompt and press Enter.
14. Start the disaster recovery process on both sites from the active site. Type `jmp-dr start` at the current active site shell prompt and press Enter.

The JBoss node is added to or removed from the active and standby sites.

Adding or Removing a Dedicated Junos Space Node

We recommend that you meet the prerequisites to add a Junos Space node or to know the impact of removing a Junos Space node from a Junos Space setup. Refer to the *Adding a Node to an Existing Junos Space Fabric* and *Deleting a Node from the Junos Space Fabric* topics in the *Junos Space Network Management Platform Workspaces Feature Guide*. We also recommend that the active site and standby site be symmetric to ensure that the performance of the disaster recovery solution is efficient and stable.

You add a dedicated Junos Space node to improve the performance of your Junos Space setup. You remove a dedicated Junos Space node if it is faulty and needs to be replaced. You may need to modify the disaster recovery configuration depending on why you added or removed the dedicated Junos Space node to or from the site.

To add or remove a dedicated Junos Space node to or from both sites:

1. Log in to the CLI of the Junos Space node at the active site on which the VIP or the eth0:0 interface is configured.

   The Junos Space Settings Menu is displayed.
2. Enter **7** while using a virtual appliance) at the Junos Space Settings Menu prompt to run shell commands.

   You are prompted to enter the administrator password.
3. Enter the administrator password.
4. Stop the disaster recovery process on both sites. To do so, type `jmp-dr stop` at the active site shell prompt and press Enter.
5. Go to the Junos Space user interface > Administration workspace > Fabric page to add or remove the dedicated Junos Space node to or from the active site. Refer to the *Adding a Node to an Existing Junos Space Fabric* and *Deleting a Node from the Junos Space Fabric* topics in the *Junos Space Network Management Platform Workspaces Feature Guide*.
6. Log in to the CLI of the Junos Space node at the standby site on which the VIP or the eth0:0 interface is configured.

   The Junos Space Settings Menu is displayed.
7. Enter **7** while using a virtual appliance at the Junos Space Settings Menu prompt to run shell commands.

   You are prompted to enter the administrator password.
8. Enter the administrator password.
9. Update the disaster recovery configuration on the standby site. To do so, type `jmp-dr toolkit config update` at the shell prompt of the VIP node at the standby site and press Enter.
10. Configure the current standby site as the active site. To do so, type `jmp-dr manualFailover` at the shell prompt and press Enter. For more information, see "Manually Failing Over the Network Management Services to the Standby Site" on page 130.
11. Go to the Junos Space user interface > Administration workspace > Fabric page to add or remove the dedicated Junos Space node to or from the standby site.
12. Update the disaster recovery configuration on the active site. To do so, type `jmp-dr toolkit config update` at the shell prompt of the VIP node at the active site and press Enter.
13. Configure the original active site back as the active site. To do so, type `jmp-dr manualFailover` at the shell prompt and press Enter.
14. Start the disaster recovery process on both sites from the active site. To do so, type `jmp-dr start` at the shell prompt and press Enter.

The dedicated Junos Space node is added to or removed from the active and standby sites.

RELATED DOCUMENTATION

Disaster Recovery Overview | 40

Configuring the Disaster Recovery Process Between an Active and a Standby Site | 74

Viewing the Disaster Recovery Configuration and Status of Watchdog Services | 105

Modifying the Disaster Recovery Configuration | 107

Manually Failing Over the Network Management Services to the Standby Site

You may need to fail over the network management services to the standby site even when the active site is fully operational. You execute the `jmp-dr manualFailover` command at the standby site to fail over the network management services to the standby site. When the failover is complete, the standby site becomes the new active site.

NOTE: We recommend that you check the status of the disaster recovery configuration before and after executing the `jmp-dr manualFailover` command. To do so, execute the `jmp-dr health` command at both sites.

To manually fail over the network management services to the standby site:

1. Log in to the CLI of the Junos Space node at the standby site on which the VIP or the eth0:0 interface is configured.

   The Junos Space Settings Menu is displayed.

2. Enter **7**while using a virtual appliance at the Junos Space Settings Menu prompt to run shell commands.

   The following is a sample output from a virtual appliance:

```
admin@10.206.41.183's password:
Last login: Mon Aug 17 06:17:58 2015 from 10.206.41.42

Welcome to the Junos Space network settings utility.

Initializing, please wait
```

```
Junos Space Settings Menu

1> Change Password
2> Change Network Settings
3> Change Time Options
4> Retrieve Logs
5> Security
6> Expand VM Drive Size
7> (Debug) run shell

A> Apply changes
Q> Quit
R> Redraw Menu

Choice [1-7,AQR]: 7
```

You are prompted to enter the administrator password.

3. Enter the administrator password.
4. Type `jmp-dr manualFailover` at the shell prompt and press Enter.
5. Enter **Yes**.

   The following is a sample output:

```
[user1@host]# jmp-dr manualFailover
Do you really want to start manual failover: Yes
Check DR state of this site: started

INFO: => switchover DR at current site: active

Stop dr-watchdog if it's
running                                                              [ OK ]
Check status of DR remote site: up
Check current DR role: standby
Restore configuration
files                                                                  [ OK ]
Setup MySQL replication: master-
master                                                            [ OK ]
Skip MySQL data
backup                                                                      [ OK ]
Setup PostgreSQL
```

```
replication                                                    [ OK ]
Start file & RRD
replication                                                    [ OK ]
Open firewall for device
traffic                                                        [ OK ]
Start services(jboss,jboss-
dc,etc.)                                                       [ OK ]
Start dr-
watchdog
[ OK ]
Copy files to DR slave
site                                                           [ OK ]
Update DR role of current site:
active                                                         [ OK ]

INFO: => switchover DR at DR remote site: standby

Check DR state of this site: started
Stop dr-watchdog if it's
running                                                        [ OK ]
Check status of DR remote site: up
Check current DR role: active
Stop services(jboss,jboss-
dc,etc.)                                                       [ OK ]
Block firewall for device
traffic                                                        [ OK ]
Reset MySQL init script and stop
replication                                                    [ OK ]
Skip MySQL data
restore                                                        [ OK ]
Setup MySQL replication and start
replication                                                    [ OK ]
Setup PostgreSQL
replication                                                    [ OK ]
Start files & RRD
replication                                                    [ OK ]
Start dr-
watchdog
[ OK ]
Clean up /var/cache/jmp-geo/
incoming                                                       [ OK ]
Update DR role of current site:
standby                                                        [ OK ]
```

```
The manualFailover command is done.
The manualFailover command is done.
```

The standby site becomes the new active site.

NOTE: If you have made any NAT-related updates in any of the disaster recovery sites, after a manual failover, run the following commands to ensure that NAT devices work seamlessly with the new active site:

1. Move the backup file from **/var/cache/jmp-geo/config/diff.properties_backup** to **/var/cache/jmp-geo/config/diff.properties**.
2. Run the following command on the VIP node to update the changed standby cluster device management, NAT, and IP configuration on the current active site:

   ```
   /var/cache/jmp-geo/script/toolkit-config-update.pl
   ```

RELATED DOCUMENTATION

Disaster Recovery Overview | 40

Configuring the Disaster Recovery Process Between an Active and a Standby Site | 74

Stopping the Disaster Recovery Process | 133

Stopping the Disaster Recovery Process

You stop the disaster recovery process from the active site or the standby site when you need to update the disaster recovery configuration or add nodes or applications to the disaster recovery setup. You use the `jmp-dr stop` command to stop the disaster recovery process on both sites. Stopping the disaster recovery process does not clean up the disaster recovery configuration from the sites.

The `jmp-dr stop` command does the following:

- Stops the disaster recovery watchdog at the sites
- Stops the replication of MySQL data, configuration files, and round-robin database (RRD) files between sites

We recommend that you execute the `jmp-dr health` command at both sites after you stop the disaster recovery process. This is to ensure that file replication, disaster recovery watchdog services, and other

services are stopped. For more information, see "Checking the Status of the Disaster Recovery Configuration" on page 100.

To stop the disaster recovery process at both sites:

1. Log in to the CLI of the Junos Space node at the active site on which the VIP or the eth0:0 interface is configured.

   The Junos Space Settings Menu is displayed.
2. Enter **7** while using a virtual appliance at the Junos Space Settings Menu prompt to run shell commands.

   The following is a sample output from a virtual appliance:

```
admin@10.206.41.183's password:
Last login: Mon Aug 17 06:17:58 2015 from 10.206.41.42

Welcome to the Junos Space network settings utility.

Initializing, please wait


Junos Space Settings Menu

1> Change Password
2> Change Network Settings
3> Change Time Options
4> Retrieve Logs
5> Security
6> Expand VM Drive Size
7> (Debug) run shell

A> Apply changes
Q> Quit
R> Redraw Menu

Choice [1-7,AQR]: 7
```

   You are prompted to enter the administrator password.
3. Enter the administrator password.
4. Type `jmp-dr stop` at the shell prompt and press Enter.

The following is a sample output:

```
[user1@host]# jmp-dr stop
Check status of DR remote site: up
Check DR stop mode: both
Check current DR role: active
Stop order: DR remote site and then current

INFO: => stop DR at remote site

Check status of DR remote site: up
Check DR stop mode: solo
Check current DR role: standby
Stop dr-
watchdog
[ OK ]
Stop mysql replication between
sites                                                          [ OK ]
                                                             [ OK ]
Stop files & RRD
replication                                                                [ OK ]
The stop command is done.

INFO: => stop DR at current site: active

Stop dr-
watchdog
[ OK ]
Stop files & RRD
replication                                                                [ OK ]
The stop command is done.
```

The disaster recovery process is stopped.

RELATED DOCUMENTATION

- Resetting the Disaster Recovery Configuration | **136**
- Modifying the Disaster Recovery Configuration | **107**
- Modifying Applications and Nodes on a Disaster Recovery Setup | **117**

Resetting the Disaster Recovery Configuration

You reset the disaster recovery configuration on both the active and the standby sites to stop the disaster recovery process and clean up the disaster recovery configuration from both sites. To reset the disaster recovery configuration, you execute the `jmp-dr reset` command.

The `jmp-dr reset` command does the following:

- Stops the disaster recovery watchdog at the sites
- Stops the replication of MySQL data, configuration files, and round-robin database (RRD) files between sites
- Starts services such as JBoss, Apache, and so on at the standby site
- Modifies the role of the cluster at the site (from active or standby to standalone)

To reset the disaster recovery configuration:

1. Log in to the CLI of the Junos Space node at the active site on which the VIP or the eth0:0 interface is configured.

   The Junos Space Settings Menu is displayed.
2. Enter **7** while using a virtual appliance at the Junos Space Settings Menu prompt to run shell commands.

   The following is a sample output from a virtual appliance:

```
admin@10.206.41.183's password:
Last login: Mon Aug 17 06:17:58 2015 from 10.206.41.42

Welcome to the Junos Space network settings utility.

Initializing, please wait


Junos Space Settings Menu

1> Change Password
2> Change Network Settings
3> Change Time Options
4> Retrieve Logs
5> Security
6> Expand VM Drive Size
7> (Debug) run shell
```

```
A> Apply changes
Q> Quit
R> Redraw Menu

Choice [1-7,AQR]: 7
```

You are prompted to enter the administrator password.

3. Enter the administrator password.
4. Type `jmp-dr reset` at the shell prompt and press Enter.

   The following is a sample output:

```
[user1@host]# jmp-dr reset
Check status of DR remote site: up
Check current DR role: active
Stop DR at both sites if it's
running                                                    [ OK ]
Clean up DR related tables in
DB                                                         [ OK ]
Clean up mysql repUser and
repAdmin                                                      [ OK ]
Clean up
NTP
[ OK ]
Remove DR
data
[ OK ]
Start services in standalone
mode                                                        [ OK ]
Remove DR
configuration
[ OK ]
Clean up
firewall
[ OK ]
Command completed.
```

5. To reset the disaster recovery configuration at the standby site, repeat steps 1 through 4 at the standby site.

The disaster recovery configuration is reset.

NOTE: We recommend that you execute the `jmp-dr health` command on both sites after resetting the disaster recovery configuration to check the status of the role, disaster recovery process, services, replication process, and disaster recovery watchdog.

RELATED DOCUMENTATION

Configuring the Disaster Recovery Process Between an Active and a Standby Site | 74

Modifying the Disaster Recovery Configuration | 107

Reimage a Node and Add the Node Back with the Same IP Address

Perform the following procedure to reimage a node and add the node back with the same IP address when the primary database node in the Disaster Recovery (DR) cluster goes down:

1. Reset the DR when the primary database node goes down on the active site.
2. Delete the node that is down by selecting to transfer its role to other JBoss node in the cluster.
3. Reimage the new Virtual Machine (VM) with the IP address of the deleted node and add it back to the cluster.
4. When the setup goes into maintenance mode, cluster formation stops. To create the cluster again, perform the following:

   NOTE: You must create the cluster formation between the JBoss VIP node and the node that took the primary database role in **Step 2**.

   a. Run the following command on the initial Virtual IP address (VIP) node before addition of nodes:

   ```
   pcs cluster stop --force
   pcs cluster destroy
   ```

   b. Run the following command on the VIP node:

   ```
   /usr/bin/systemctl start corosync
   pcs cluster setup jmp-CLUSTER <node-hostname>
   ```

c. Run the following command to authorize the host on both VIP node and the node that took the primary database role in **step 2**:

```
pcs host auth <other node's hostname>.
Username: hacluster
Password:
jnpr123! pcs cluster start --all
pcs cluster enable
```

d. Run the following command to check pcs cluster status:

```
pcs cluster status
Now On other node:
pcs cluster start
```

Check the `pcs cluster status` on both VIP node and the node that took the primary database role in **step 2**. Both the nodes must be in active status.

e. To add a Virtual IP address, run the following command:

```
pcs resource create virtual_ip ocf:heartbeat:IPaddr2 ip=<ip address> cidr_netmask=<subnet>
nic=eth0:0 op monitor interval=30s
pcs resource create before_vip lsb:beforeSwitchingVIP
```

f. To start the pacemaker, run the following command:

```
/usr/bin/systemctl start pacemaker
```

g. Run the following command when the pcs resource status for the VIP stops:

```
pcs property set stonith-enabled=false
Chk ifconfig for VIP association on all nodes
```

h. To remove a node from the cluster, run the following command:

```
pcs cluster node remove <node-hostname>
```

i. You may need to restart the services. Use the following command to restart services on all the nodes:

```
systemctl stop jmp-watchdog
systemctl stop jboss
systemctl stop jboss-dc
```

j. Check for the https services on the VIP node. If the service is down, bring it up.

   The cluster is now up and running.

k. Configure and start the DR.

CHAPTER 10

Upgrading Junos Space Network Management Platform with Disaster Recovery Enabled

IN THIS CHAPTER

Upgrade Procedure

IN THIS SECTION

You must reset the disaster recovery configuration on both the active and the standby sites to stop the disaster recovery process and clean up the disaster recovery configuration from both the sites. You can upgrade the Junos Space Platform software at both the active and the standby sites at the time of deployment. Since the disaster recovery upgrade procedure is time consuming, we recommend you to reset the disaster recovery configuration and upgrade the sites to the required version.

The upgrade procedure comprises of the following tasks:

1. Back up the Current Disaster Recovery Configuration
2. Reset the Disaster Recovery Configuration
3. Upgrade the Junos Space Network Management Platform and Application
4. Configure and Perform Disaster Recovery

1. Back up the Current Disaster Recovery Configuration

Backup the current disaster recovery configuration on both the sites to refer the parameters when configuring the disaster recovery again after upgrade. Run the following command to create a backup for current configuration on both sites:

```
$ jmp-dr api v1 config --include role,failover,states,core,deviceManagement,mysql,file,watchdog > /home/admin/dr-config.txt
```

NOTE: Enter the following command in a separate shell prompt on the VIP nodes of both the sites to check log errors:

```
$tail-f /var/log/jmp-geo/dr-cli-reset.log
```

2. Reset the Disaster Recovery Configuration

Reset the current disaster recovery configuration to ensure that both the active and the passive sites refer to the same parameters, when you configure the disaster recovery, post the upgrade.

Execute the following command to create a backup for the current configuration.

```
$ jmp-dr api v1 config --include role,failover,states,core,deviceManagement,mysql,file,watchdog > /home/admin/dr-config.txt
```

Execute the following commands to reset disaster recovery on both sites running the `jmp-dr reset` commands on VIP nodes of both sites.

```
$jmp-dr reset
```

After executing this commands, the sites become standalone sites. Junos Space Platform comes up and you can access the UI from both the sites.

3. Upgrade the Junos Space Network Management Platform and Application

You can upgrade the Junos Space Platform and the application on both the active and the standby sites simultaneously, as both the sites are in standalone mode.

To upgrade the Junos Space Network Management Platform:

1. Login to the Junos Space Network Management Platform.

   The dashboard appears.
2. Upload the software image to the active and standby sites.
   a. Select **Network Management Platform** from the **Applications** drop-down menu.

b. Click **Upgrade**.

The upgrade window appears.

c. Upload the image.

> **NOTE**: The system takes approximately 5 minutes to upload the new image. Enter the following command in a separate shell prompt at each node of current active and standby sites to check the logs for errors.
>
> `$tail-f /var/log/install.log`

3. Upgrade the Junos Space software at the active and standby sites.

Reboot occurs as part of post upgrade.

a. Select **Network Management Platform > Administration > Application > Network Management Platform**.

b. Click **Upgrade**.

The upgrade window appears.

c. In the **Upload Platform** menu, select the specific row and click **Upgrade**.

The maintenance mode screen displays the upgrade progress. The reboot screen appears, once the upgrade is completed.

> **NOTE**: Type the following command in the separate shell prompt on each node of current active and standby sites to check the error logs.
>
> `$tail-f /var/log/install.log .`

> **NOTE**: Type the following command in a separate shell prompt at each node on the active and standby sites to check the upgrade progress.
>
> `$tail-f /var/jmp_upgrade/slave/log/<oldRelease_newRelease>`

d. Click **Reboot**.

The nodes starts to reboot.

> **NOTE**: Execute the following command to check for boot to complete.
>
> `$tail -f /tmp/systemStartup.log`

Log Example:

```
2018/03/07 10:05:51.207 Appmgt is now deploying ... [ 9 of 9 ] 2018/03/07 10:05:57.607 Appmgt
```

4. Upload the Junos Space application in active and the standby site.

   a. Select **Network Management Platform > Administration > Application > *<application-name>*> Upgrade > Action**.

   b. Click **Upload**.

      The upload window appears.

   c. Upload the selected Junos Space application.

      NOTE: Enter the following command in the separate shell prompt on each node of current active and standby sites to check the logs for errors.

      ```
      $tail-f /var/log/install.log
      ```

5. Upgrade the Junos Space application in the active and standby site.

   a. Select **Network Management Platform > Administration > Application > *<application-name>*> Upgrade > Action**.

   b. Click **Upgrade**.

      The Upgrade window appears.

   c. In the **Upgrade Application** menu, select the specific row and click **Upgrade**.

      NOTE: Enter the following command in a separate shell prompt on each node of current active and standby sites to check the logs for error.

      ```
      $tail-f /var/log/install.log
      ```

   d. Select **Network Management Platform > Administration > Application**.

      The Application window displays the expected upgrade version.

4. Configure and Perform Disaster Recovery

NOTE: This is an optional procedure. If you plan to configure multiple upgrade paths: upgrade the Junos Space Network Management Platform and the application to the required version, before you configure disaster recovery.

To configure and start the disaster recovery process on both the active and the passive sites:

1. Configure the disaster recovery process.
   a. Execute the following command in the VIP node of active site to configure disaster recovery.

```
$ jmp-dr init -a
```

   b. Execute the following command in the VIP node of standby site to configure disaster recovery.

```
$ jmp-dr init -s
```

NOTE: Ensure that you initialize the disaster recovery process on the standby site, only after the disaster recovery initialization is complete on the active site. For more information about configuring the disaster recovery process, see Configuring the Disaster Recovery Process Between an Active and a Standby Site.

2. Execute the following command to start the disaster recovery process from the active site.

```
jmp-dr start
```

NOTE: Execute the following command in the separate shell prompt on each of the VIP nodes of the site to check the logs for errors.

```
$tail-f /var/log/jmp-geo/dr-cli-start.log
```

If the standby node is not updated, stop and the start the disaster recovery process over again.

Execute the command again for verify the update.

3. Execute the following command to check the disaster recovery process health:

```
$jmp-dr health
```

NOTE: Execute the following command in a separate shell prompt on each site VIP nodes to check the logs for errors:

```
$tail-f /var/log/jmp-geo/dr-cli-health.log
```

4. Execute the following command to stop the disaster recovery process on both the sites from the active site:

```
$jmp-dr stop
```

NOTE: Execute the following command in the separate shell prompt on each VIP nodes of the site to check the logs for errors:

```
$tail-f /var/log/jmp-geo/dr-cli-stop.log
```

5. Execute the following command on the current standby site, to verify the functionality on standby site by manually fail-over from active site:

```
$jmp-dr manualFailover
```

NOTE: Execute the following command in the separate shell prompt on VIP node of the standby site to check the logs for error:

```
$tail-f /var/log/jmp-geo/dr-cli-manual-failover.log
```

NOTE: Execute the following command to verify is the application starts normally, by checking the JBoss log:

```
$tail-f /var/log/jboss/server/server1/server.log
```

The Application page shows the upgrade version.

6. Execute the following command on the current standby site VIP node to apply manual fail-over to the active site from the standby site:

```
$jmp-dr manualFailover
```

NOTE: Execute the following command in the separate shell prompt on VIP node of the current standby site to check the logs for errors.

```
$tail-f /var/log/jmp-geo/dr-cli-manual-failover.log
```

NOTE: Execute the following command to check for the application start by checking the JBoss log:

```
$tail-f /var/log/jboss/server/server1/server.log
```

7. Execute the following command to check the disaster recovery health on both the sites.

```
$jmp-dr health
```

NOTE: Execute the following command in the separate shell prompt on each VIP nodes of the site to check the logs for errors.

```
$tail-f /var/log/jmp-geo/dr-cli-health.log
```

8. Execute the following command to start the disaster recovery process on both the sites from active site.

```
$jmp-dr start
```

NOTE: Execute the following command in the separate shell prompt on each VIP nodes of the site to check the logs for errors.

```
$tail-f /var/log/jmp-geo/dr-cli-start.log
```

9. Execute the following command to check the disaster recovery process health on both sites.

```
$jmp-dr health
```

NOTE: Execute the following command in the separate shell prompt on each VIP nodes of the site to check the logs for errors.

```
$tail-f /var/log/jmp-geo/dr-cli-health.log
```

RELATED DOCUMENTATION

Configuring the Disaster Recovery Process Between an Active and a Standby Site | 74

Modifying the Disaster Recovery Configuration | 107