



Gemini Central Administration Guide

Version 3.0

CONTENTS

Prerequisites

Overview & New Features

System Initialization

- Introduction

- Initial setup

 - Creating an Appliance IP address

 - Static IP address assignment

 - DHCP Address assignment

- Gemini Central - Web Interface

 - Localization Settings

 - License Activation

 - Use or request a Licence File

 - Connect to a Licence Server

 - Change Admin Password

 - Initial Login Screen

 - Deployable Roles for Gemini Central

 - Standalone Node

 - The Management Center Node

 - Recommended Actions Feature

 - Bulk Provisioning Feature

 - Bulk Provisioning - Instance Discovery

 - Bulk Provisioning - Network Settings

 - Bulk Provisioning - Hostname assignment

 - Bulk Provisioning - Change Admin Password

 - Bulk Provisioning - Connect to LDAP

 - Bulk Provisioning - SSH Authentication

 - Bulk Provisioning - Summary Screen

 - Adding instances to an existing Manage Cluster

 - Gemini Central OS child instances:

 - Adjust Storage Plan

HOME Dashboard

- Featured Platform

- Integration Center

Splunk Installation

Activate Splunk

Splunk Clustered Environment Installation

Building the Splunk Environment

Step 1: Creating the Splunk Environment

Step 2: Create Splunk Clusters

Step 3: Organize the Nodes

Step 4: Locate Nodes

Initial Splunk Environment

Verification of the Splunk Environment

Post Splunk Cluster Deployment Features and Tasks

Splunk Activation

Splunk Web Port (8000)

Splunk Receiver Port (9997)

Splunk Boot-Start

Splunk Workload Management features

Splunk admin password

SSL Passwords

Heartbeat monitoring feature

Resource activity

Using the Splunk Web interface

Adding more Nodes

Adding a Splunk License Manager

Adding a Splunk Monitoring Console

Standalone Splunk Installation

Gemini Explore

Activating Gemini Central

Licensing Gemini Central (Beta Trial)

Integrating Inspect for Splunk Admin

Tableau Server

Installation of Tableau Server

MinIO Object Storage

LOG Menu

The features of Gemini Log Receiver

Rule Manager

Log Receiver - Rule Manager Dashboard

Log Receiver - Source

Log Receiver - Destination

Destination Log File Splitting - Overview

Destination Log File Splitting - By host

Destination Log File Splitting - By Facility

Destination Log File Splitting - By Level(Severity)

Destination Log File Splitting - By Program

Destination Log File Splitting - User Custom Path

Log Receiver Settings - Filter

Filter by Host

Filter by Netmask

Filter by Match

Forwarding Log Receiver events into Splunk

Enable Splunk Indexers to receive events from the Log Receiver

Verify the Receiver Port at the Indexers

Creating an Index at the Indexers

Testing Log Receiver Rules before enabling forwarding (optional)

Create a Heavy Forwarder to forward Log Receiver data

Creating a High-Availability Syslog environment

Load Balancing a syslog feed

NODE Menu

System Time

Name

Hostname

Local Hosts

Network

NIC Bonding

Port Redirect

OS Users

FTP

FTP Service

FTP User

SSH

SNMP

SNMP Service

SNMP Agent (polling option)

- SNMP Agent version 1

- SNMP Agent version 2c

- SNMP Agent version 3

- SNMP Trap Destinations (trapping option)

- SNMP Trap Thresholds

Failover

- Creating a Failover Group

- Option A - The process to set up a single Failover Group

- Joining a Failover Group

- Option B - Creating a reciprocal failover Group with optional Load Balancing

Storage

- Storage Devices

- Mount disk and mount points

- Encryption and Decryption

- Create Software RAID Disk

- Create a Logical Volume

- Merge Disk

- Add an NFS Mount

- Add a CIFS Mount

- Add an S3 Mount

- Add an iSCSI Target

- Manage Swap space

Log Forwarding

Diagnostics

Rsync Backup

Benchmark

Cluster

- Manage Nodes

- Adding instances to an existing Manage Group

- Manage Groups

- Execute Jobs

- Backup Center

- Backup Job Detail

- Restoring Backups

- Membership Settings

- Bulk Provision

License

- Understanding Gemini Central Licensing
- License Status
- Remote Licenses
- Inventory
- License Server

Splunk

- AWS Provisioning
 - Prerequisites for AWS Provisioning
 - AWS Provisioning Procedure
- Daemon
- Web Interface
- Apps
- Splunk Diag
- Optimizer
- Config Editor
- Versioning
 - Rollback Option
 - Recover option
- Command
 - Installing a Splunk App
- Splunk Environments
 - Adding a Node (Unassigned Nodes)
 - Add a single Node
 - Add a pre-existing Splunk Indexer Cluster
 - Add a pre-existing Splunk Search Head Cluster
 - Add a Group of standalone Splunk Nodes
 - Assigning Unassigned Nodes and Clusters to a Splunk Environment
 - Assign to Environment
 - Create New Cluster
 - Add to Cluster
 - Add an additional Indexer or Search Head to a Splunk Cluster
- New Instance install - Clusters and Standalone Instances
- Adopting remote Splunk Instances and Clusters
 - Creating a 'shell environment'
 - Assigning an external Splunk Cluster to an Environment

- Assigning Splunk Standalone Nodes
 - Assigning a new remote Indexer or Search Head to a Cluster
- Deploy Independent Stream Forwarder
- Operations and Administration of Splunk
 - Splunk Environment Level Options
- Upgrading Splunk using the Rolling Upgrade Feature
 - Manual Upgrade of a Splunk Instance
 - Splunk Cluster Level Options
 - Splunk Standalone Options
 - Unassigned Nodes Panel Options

Gemini Agents

- Known Issues and Restrictions - Read before progressing
- Splunk Interface Detail
- Prerequisites for Installation
 - Supported Operating Systems
 - System Requirements on the Target Host
- Installation of Gemini Agent
 - Step 1: Management Center - Enable Gemini Agent distribution
 - Step 2: Gemini Agent - Prerequisites (Splunk Host)
 - Step 3: Installing the Gemini Agent - Splunk Host
 - Gemini Agent - CLI options at the Splunk host
 - agent status
 - agent --version
 - agent restart
 - agent configure
 - agent uninstall
 - agent stop/start
 - Gemini Agent - Troubleshooting
- Upgrading the Gemini Agent

Settings

- System Admin
 - Admin Web
 - Backup & Restore
 - Diagnose
- System Update
- Information

Software

Hardware

Listen Port

Audit Report

Authentication

Manage User

User Permissions

LDAP

Single Sign-on (SSO)

Password Policy

HTTP Proxy

Login Banner

Reboot

Shutdown

CLI Commands

The help Command

Commands for initial setup

The network Command

The config Command

The admin Command (setup options)

The agent Command - Management Center Node

agent status

agent --version

agent restart

agent configure

agent uninstall

agent stop/start

Commands for Information Gathering

The version Operator

The model Operator

The service-tag Operator

The service Command (status operator)

The service Command (listen-port operator)

The system Command (info operator)

Commands for Troubleshooting

The admin Command (Troubleshooting)

The admin Command (reset-password operator)

The admin Command (set-password operator)

The admin Command (gen-ssl operator)

The network Command (reset operator)

The service Command (restart operator)

The splunk Command

Reset Splunk Environments

The system Command (patch operator)

Commands for System Operations

System Reboot

System Power Off

Default Passwords for CLI Operations

Prerequisites

- Cabling - An ethernet cable and an available ethernet switch connection.
- Networking - One IP address which can be assigned to the Gemini instance ethernet Port 1 (address, netmask and gateway are required for manual configuration. DHCP is also supported)
- Accessories - VGA monitor and USB keyboard
- Client - A client PC on the same network as the Gemini instance running a suitable web browser.
- Splunk: Access to a Splunk Enterprise installation tarball (splunk-*. *-Linux-x86_64.tgz).
- Tableau & Explore: Internet access is required for installation of these Featured Platforms.

Important Note on Network Access Control

To administer and run Gemini instances and its services, certain communication channels between clients and nodes are required. As a minimum, the following ports are required:

Port	Reason
443/TCP	Https access
22/TCP	SSH access
4444/TCP	Cluster communication

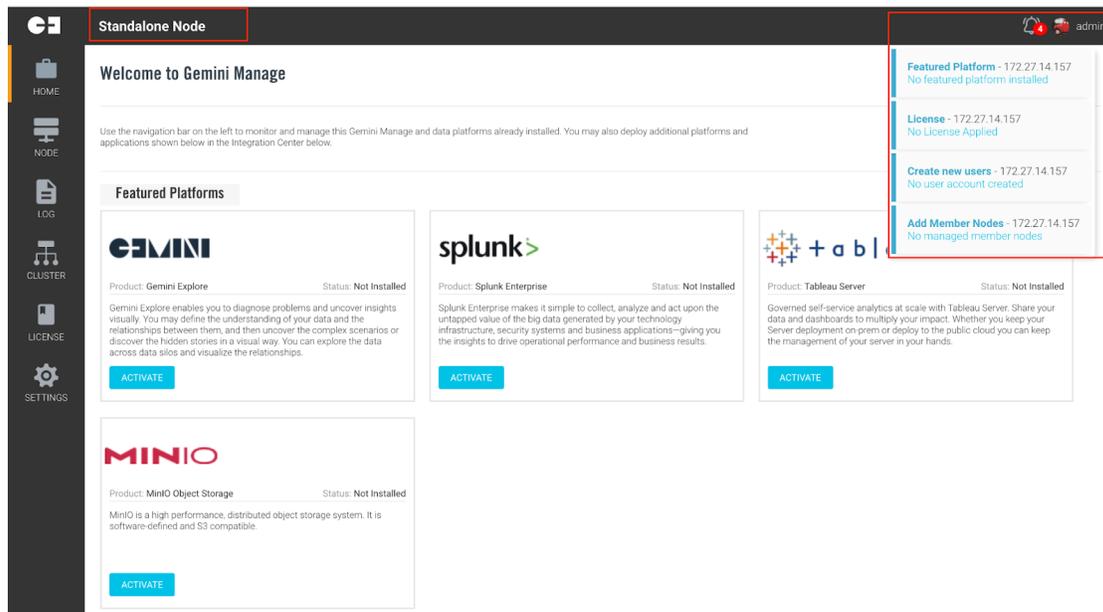
As the Web Interface and SSH console offer low-level system access, try to ensure that network settings are biased towards a **'host-only'** approach and are not exposed to public access (ie. Anywhere, 0.0.0.0/0). Depending on the deployment, add inbound/outbound rules as needed.

Overview & New Features

Welcome to **Gemini Central** - the new name for Gemini Enterprise Manage!

Recent iterations of Gemini software have been demonstrating the advantages of having a **Management Center** node at its core that can oversee an entire Splunk environment, or even multiple Splunk environments.

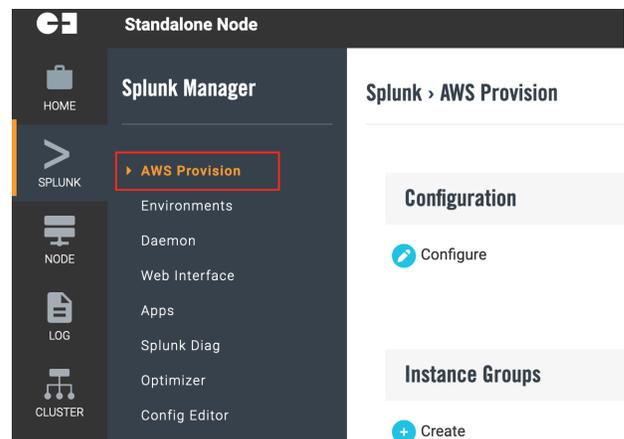
Recently introduced features such as the central collation of Splunk diags and a central backup solution for all Splunk/Gemini config files have been very popular with customers. To complement this, we are introducing another new feature called, '**Recommended Actions**' that advises on outstanding issues that need to be addressed on a Gemini instance. From the **Management Center** node, selecting a 'Recommended Action' relevant to any node will automatically open it in a separate browser tab for immediate addressing.



Another exciting new feature, **AWS Provisioning**, enables central provisioning of complete Splunk Indexer and Search Head clusters based on AWS EC2 instances using **Splunk AMI's** direct from a Gemini Central node, typically the **Management Center**.

This is a big departure to the use of our hardened OS built-in to the **Gemini AMI**, but will enable customers to get up and running quickly with complex Splunk environments on **AWS** in just a few minutes, provided the customer has their AWS secret keys and passwords available.

A **Gemini Agent** will be installed to each Splunk instance during the install, allowing the **Management Center** node central observation and a fast Splunk upgrade facility.



System Initialization

Introduction

If this is the first, or primary Gemini Central instance to be installed, use the following section to guide you through creating a **Management Center** that can be used to import existing, remote, or newly created Gemini instances for centralized Splunk and operational management purposes.

In order to import existing remote Splunk environments (non-Gemini instances) see the [Gemini Agent Installation](#) section of this manual.

Gemini Central software should be installed to a suitable host platform using an ISO file. Gemini appliances will have this software built-in, offering a security-hardened environment out of the box.

The naming convention for Gemini Central software uses the following configuration:

gemini-appliance-<major_release>.<minor_release>.<bug_fix>.iso

An example would therefore be:

gemini-appliance-2.7-251.iso

Please contact Gemini Support (support@gemini-data.com) if you require the very latest, or indeed an earlier working version of the ISO file.

Use the chosen ISO file to bring up a suitable instance. This instance could be a physical appliance, a VMware server or a Cloud-based machine within AWS for example.

Recommended hardware, specifications and quickstart instructions for these instances can be found using the following links:

- Gemini Central Quickstart Guide 3.0 (Physical Hardware)
- Gemini Central Quickstart Guide 3.0 (VMware)
- Gemini Central Quickstart Guide 3.0 (AWS)
- Gemini Central Quickstart Guide 3.0 (Azure)

When starting an instance for the very first time, ensure it has an IP address. Once an IP address has been assigned, all subsequent configuration, amendments, installations and future upgrades for this and other instances can all be achieved using the **Management Center**.

Initial setup

When the host instance 'boots' to our ISO file on a physical appliance, the following screen should be observed.

If the 'login' prompt is not visible on the connected monitor, press 'enter' on the keyboard a few times. The login prompt should then appear.



Due to the automatic selection of option 1; 'Install Gemini Central', the above screen may not actually be observed, but following installation of Gemini software, which can take 10 - 15 minutes, the screen will return a login prompt, similar to this one below:

```
gemini-1cece3 login:
```

You now have two options; If this appliance is to be included in an existing Gemini Central environment, and it has been issued an IP address via DHCP, simply login to your Gemini **Management Center** node and continue with instructions contained in the [Provisioning Appliances](#) section of this manual.

Alternatively, continue with this section if you need to achieve any of the following:

- If this is to become the Gemini Management Center or a standalone instance
- If you need to provision this instance with an IP address or DNS hostname
- If you need to change an IP address already assigned to this instance

At the terminal prompt, login to the interface using the following credentials:

username: **sbox**

password: **facing jet function drive** (note the spaces are important!)

You will be prompted to change the default password. Please complete this exercise and ensure that you record the new password. Note that there is a default expiry policy of 60 days on this account. If you wish to freeze this for the foreseeable future, navigate immediately to the **Settings / Password Policy** dashboard and remove the checkmark from the relevant box.

Note: Contact support@gemini-data.com if you have any issues or questions with the initial setup process.

Creating an Appliance IP address

If you need to create or change the IP address assigned to the instance, remain at the terminal prompt and use the '**sbox network**' command to create/change an IP address that will be used for this Gemini instance.

Type the following at the prompt to reveal all network options available.

```
sbox network
```

```
[sbox@gemini ~]$ sbox network
Usage: sbox network [OPTIONS]

[OPTIONS]
  --reset      Reset all network interface to default value.
  -nic         Setup specific network interface, required for
              below options.
  --disable    Disable specific NIC, when given it ignores
              --dhcp, -ip, -netmask and -gateway options.
  --dhcp       Config the specific NIC as DHCP.
  -ip          Set IP address for specific NIC.
  -netmask     Set subnet mask for specific NIC, required when
              set IP address.
  -gateway     Set gateway on specific NIC. Optional.

[sbox@gemini ~]$
```

Using this command it is required that we create an IP address. Two methods are available to achieve this.

- Assign a static IP address (preferred option)
- Assign an IP address using DHCP

Static IP address assignment

We recommend that you create a permanent IP address for Gemini Manage.

Identify the name of the device network interface using the following command at the terminal:

```
ip a
```

The example output shown below reveals that the interface name is 'nic0', and the current ip address is 192.168.1.100.

```
[sbox@sboxnode1 ~]$ ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
```

```

inet6 ::1/128 scope host
    valid_lft forever preferred_lft forever
2: nic0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP qlen 1000
    link/ether 08:00:27:9e:96:c0 brd ff:ff:ff:ff:ff:ff
    inet 192.168.1.100/24 brd 192.168.1.255 scope global dynamic nic0
        valid_lft 68091sec preferred_lft 68091sec

```

To create or change the current IP Address, at the terminal prompt use the following command to apply your chosen IP address.

```

sbox network -nic <network_interface> -ip <chosen_IP_address>
-netmask <Netmask> -gateway <Gateway_IP_address>

```

Example:

```

sbox network -nic nic0 -ip 192.168.1.100 -netmask 255.255.255.0
-gateway 192.168.1.1

```

Note: Network interface values can be verified using the 'ifconfig' or 'ip a' commands. Additional IP addresses can be assigned using the Management Center, if required.

Note

To create a **'Host only'** instance to operate solely within the local network, omit the **-gateway** option when using the above command.

DHCP Address assignment

If you wish to configure the network interface using DHCP use the following alternative command:

```

sbox network -nic <Network interface name> --dhcp

ie. sbox -nic nic0 --dhcp

```

Gemini Central - Web Interface

Further configuration of the instance can be completed using the web interface.

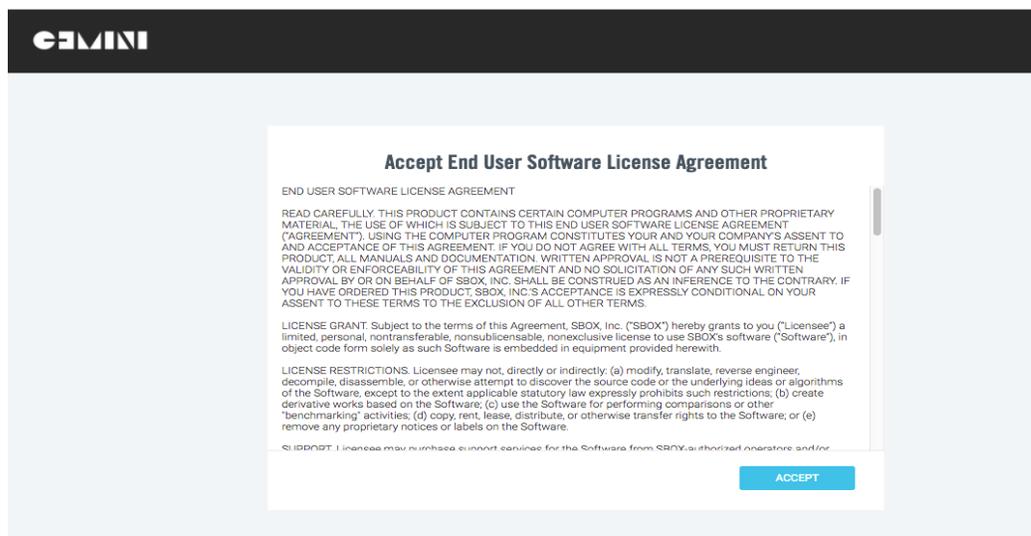
Using a supported web browser, navigate to:

https://<ip_address or FQDN of instance>

If you have an instance already dedicated to the role of **Management Center**, simply use this interface to add and manage all other Gemini instances which can be added as **'Unassigned Nodes'**.

If this instance is to become the **Management Center**, or if you just want to login to the web interface, use a suitable browser to access the IP address assigned, ensuring that you use the **'https:'** prefix.

Accept the browser certificate exception on initial opening, to reveal the following **End User Software License Agreement** license screen.



After reading the terms and conditions, select **'Accept'** to advance to the next screen which will vary depending on the operating platform. This section specifically references our hardware **Appliance** platform. For Software platforms, please refer to the individual AWS or Azure Quickstart guides for details of the start-up procedure.

Localization Settings

This screen allows you to set **locale** information regarding **language**, **timezone** and an appropriate DNS recognized **Hostname** for your instance.

Gemini Central natively supports four languages;

- English (American)
- Traditional Chinese
- German
- Japanese

Select your preferred language to adjust the user experience accordingly.

Use the **'Node'** menu of the web interface at any time to modify these settings in the future.

License Activation

This step allows you to activate the appropriate license for your intended use.

Enterprise Edition (Purchased Licence)

The most common option, select if you have an appropriate **Gemini Licence Server** in place.

Enterprise Edition (30 days Trial)

Select if you are entering into a trial of our product and simply want to test out the features, or you have not yet obtained a license from Gemini Data.

The License status can be upgraded to **Enterprise Edition** at any time during the 30 day trial period.

Free Edition

The Free Edition may be used indefinitely, but is restricted to 3 instances and some features have been disabled (see Note below)

Select License Option

Enterprise Edition (Purchased License)

Select this option to use a previously-purchased paid license. You will be provided the option of applying a newly generated license or connecting to a license server.

Enterprise Edition (30 Days Trial)

The full-featured Enterprise Edition can be used for trial purposes for 30 days with no limitations. If not purchased within the trial period, the appliance reverts to the Free Edition.

Free Edition

The Free Edition can be used indefinitely but is restricted to maximum of 3 servers and also restricted in terms of features. Click [here](#) for a comparison of features.

Note

In the Free Edition the following features are restricted;

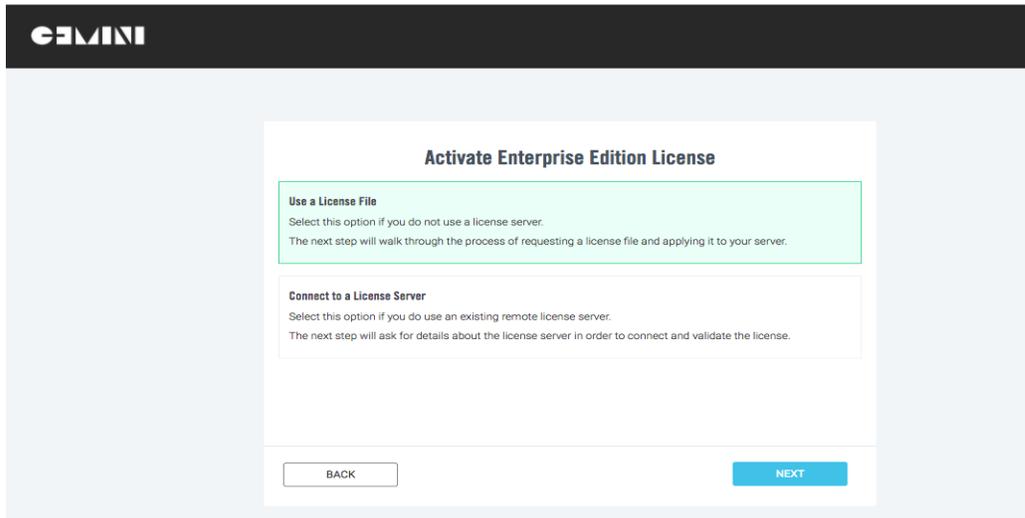
- No Failover function
- No LDAP Authentication
- No support of external storage, including NFS, CIFS, and S3.
- No remote license server.
- Limited Splunk configuration Versioning, restricted to roll back to last 3 versions.

Gemini Cluster features affected;

- Up to 4 nodes in a cluster in maximum.
- No scheduled jobs.
- Jobs for Splunk upgrade are restricted.
- Jobs for Gemini instance boot control are restricted.
- Jobs for Splunk service control are restricted.

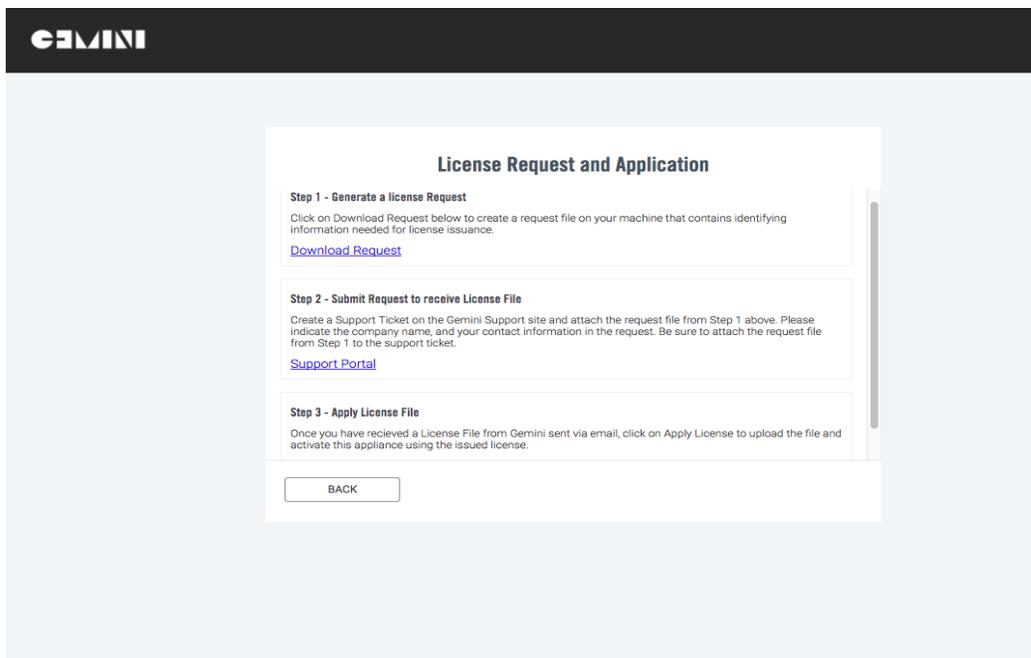
Use or request a Licence File

If you want to activate a pre-purchased Enterprise license, or initiate the request of an Enterprise License, and you do not have an existing Gemini License Server - select the first option, '**Use a License file**'.



This reveals a three-step process, as outlined in the screen below, beginning with a License request process.

Use the 'Licences' menu of the web interface at any time to modify or view License settings.

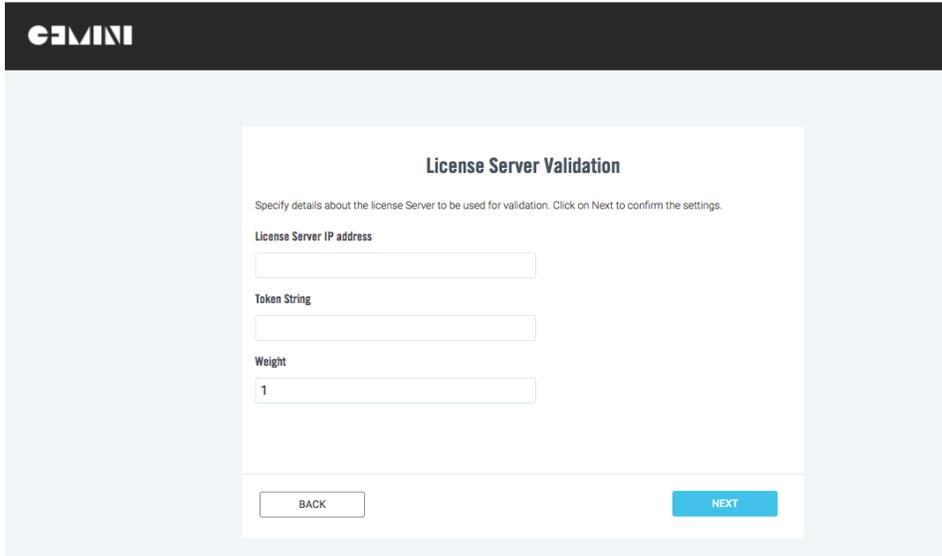


If you are applying for a License file during the **Bulk Provision** process, there can be a delay of up to 72 hours before your License file is made available. In order to proceed with the Bulk Provision process, use the '**Back**' button after completing **Step 2** of the License application wizard, and select the '**Trial License**' option.

On receipt of your **License File**, complete **Step 3** of the License application wizard using the '**License**' menu of the web interface at any time within the 30 day trial period.

Connect to a Licence Server

Alternatively, if you are using a Gemini License Server to manage all available licenses, selecting this option will allow you to specify license server information, including IP Address and token, in order to perform the validation.

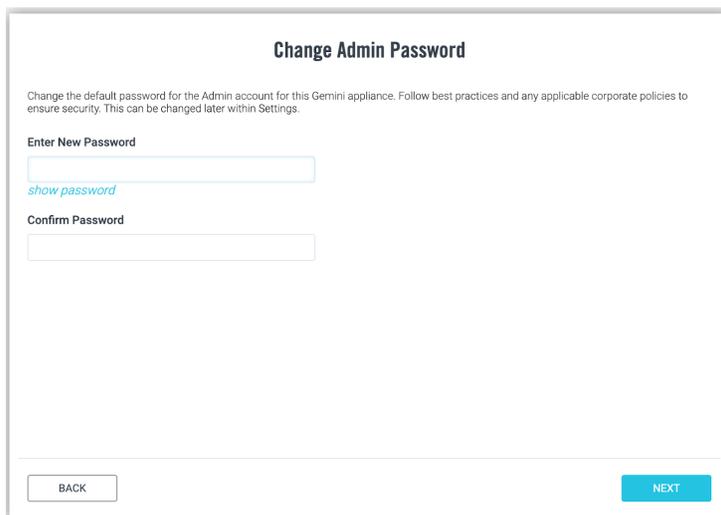


The screenshot shows the Gemini logo in the top left corner. The main content area is titled "License Server Validation". Below the title, there is a sub-header "License Server Validation" and a paragraph of instructions: "Specify details about the license Server to be used for validation. Click on Next to confirm the settings." There are three input fields: "License Server IP address", "Token String", and "Weight". The "Weight" field contains the number "1". At the bottom of the form, there are two buttons: "BACK" and "NEXT".

Change Admin Password

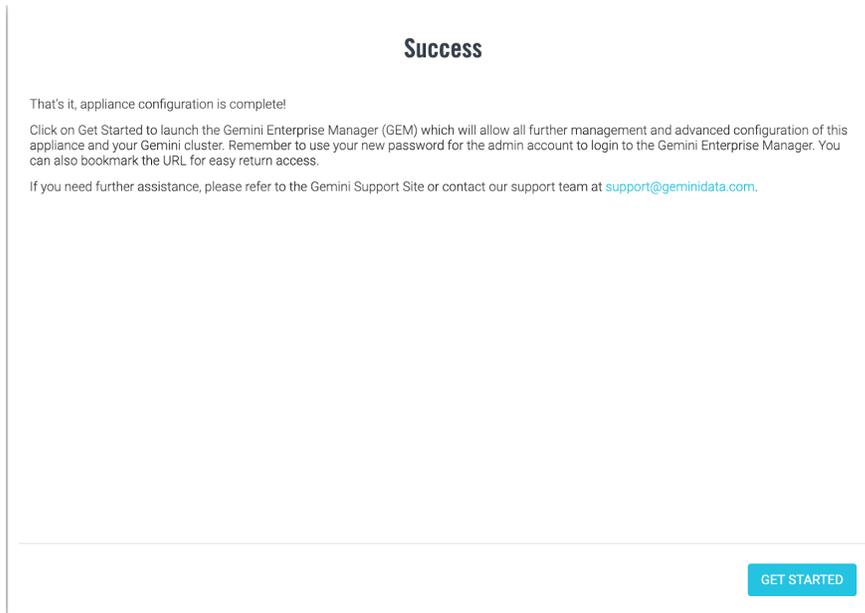
The final step of the Gemini instance initialization wizard is to create an administration account to enable access to the web interface. This is a mandatory requirement.

***** It is essential to record the admin account and password in a safe place *****



The screenshot shows the Gemini logo in the top left corner. The main content area is titled "Change Admin Password". Below the title, there is a paragraph of instructions: "Change the default password for the Admin account for this Gemini appliance. Follow best practices and any applicable corporate policies to ensure security. This can be changed later within Settings." There are two input fields: "Enter New Password" and "Confirm Password". Below the "Enter New Password" field, there is a link "show password". At the bottom of the form, there are two buttons: "BACK" and "NEXT".

This completes all options regarding the initial setup process of Gemini Central software, and the following screen should confirm this;



Initial Login Screen

Upon completion of the setup process and after selecting the **Get Started** button from the Success screen you will be presented with the login screen.

The image shows the 'Welcome to Gemini Enterprise: Manager' login screen. It features the following elements:

- Welcome to Gemini Enterprise: Manager** (Title)
- Username** label above a text input field containing 'admin'.
- Password** label above a text input field containing 'Password'.
- A blue button labeled **LOGIN** at the bottom right.

Login to Gemini Manage with the username '**admin**' and the password created in the previous step.

Return visits to this interface will proceed directly to the login screen. Configured settings may still be changed within the corresponding areas within Manage.

Further options can be achieved from the Gemini web interface, and some useful examples are given below;

- Create a **Management Center** from this instance used to control multiple Gemini nodes
- Activate **Splunk** on this instance for use as a Splunk server
- Activate **Gemini Explore** on this instance to make use of **Inspect for Splunk Admin** or create analysis models
- Activate **Tableau** on this instance for Business Intelligence purposes
- Activate **MinIO** for additional storage purposes like **Splunk Smartstore**

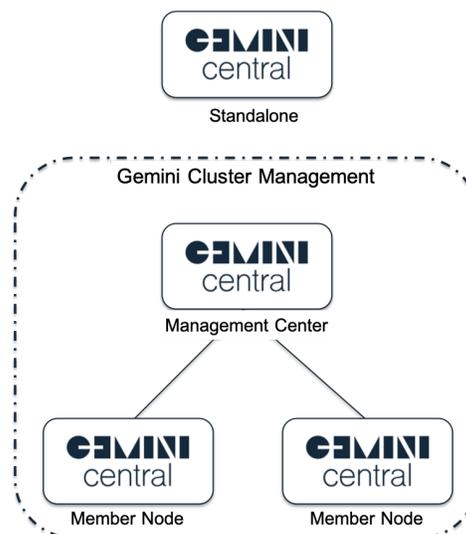
Deployable Roles for Gemini Central

Gemini Central instances will fall into three categories that will determine its deployment role: **Standalone**, **Management Center**, or a **Member** node.

This category will be displayed at the top of the web interface, and all nodes will start out as a Standalone node.

When nodes are grouped and in control of a Management Center, such as when the **Bulk Provisioning** feature has been used, nodes will naturally become Member nodes controlled by a Management Center node. All other nodes will remain as Standalone.

Irrelevant functions for specific roles will be restricted, thus preventing misuse and confusion from inappropriate functions.



Standalone Node

If you wish to use this instance simply as a standalone node, for instance as a Log Receiver, or Splunk Deployment server, use the web interface menus to configure the features required.

If you wish to use Splunk on the instance, you will first have to **'Activate'** it from the **Home** dashboard of the Gemini web interface.

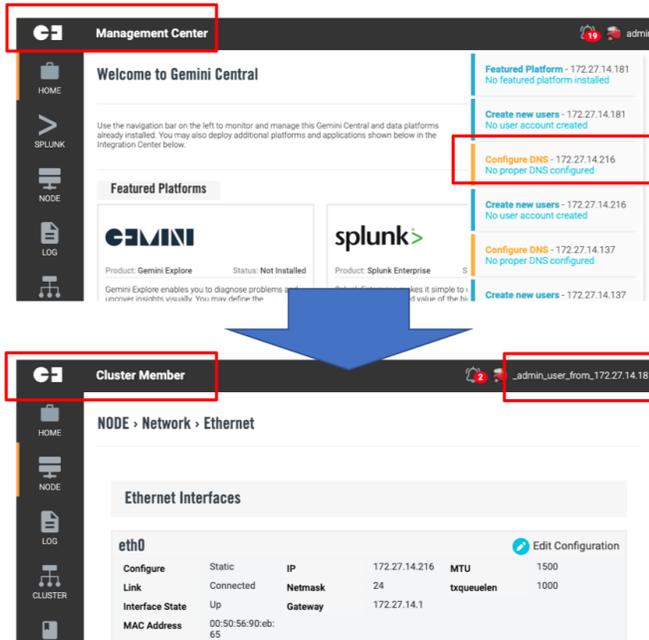
The Management Center Node

The Gemini **Management Center** should be considered a special standalone instance. It is usually the first Gemini Central instance created and can provide a central location for the control and management of all your Gemini instances, with features that include;

- Centralized control and monitoring of Gemini instances

- Fast and efficient creation of Splunk clusters (Bulk Provisioning)
- Management, scaling and upgrading of Splunk instances and clusters(Splunk Environments)
- Provides a backup repository for all Gemini and Splunk config files(Centralized Backup Center)
- Provides a central 'jump server' for outstanding issues discovered by the Recommended Actions feature(see below)
- Provides a central repository for Splunk Diag files.

Using the Management Center as a jump server, most of the administration in a Gemini Cluster can be achieved centrally.



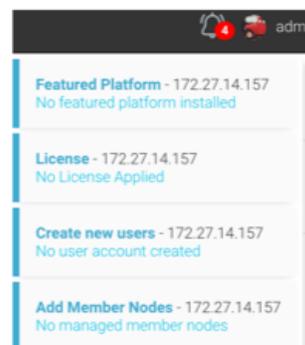
- The SSH key of the Management Center is authorized on each Member Node
- Recommended Action alerts are consolidated on the Management Center
- Jump straight to another Member Node for administration without additional authentication

Recommended Actions Feature

In the top right-hand corner of Gemini Central's **Home** dashboard, a small bell icon can often be seen together with a number.

This number represents an unusual system status or incomplete setting within a specific Gemini node.

Selecting the issue itself will provide recommendations based on best practices to prevent misuse and maximize the capabilities.



When used from the Management Center node, this feature will list issues from **all** Gemini Nodes within its Managed Group. Selecting an issue from another instance will open the instance concerned in another browser tab so that it can be addressed locally.

The main advantages of this feature are listed below;

- The ability to both detect and report various types of problems prompting the administrator to correct them
- Alert messages include INFO, WARNING, and DANGER are color-coded for ease of reference

- Alerts for all Member Nodes can be viewed and addressed directly from the Management Center node acting as a 'jump server'

The entire list of potential Alerts can be seen in the table below;

Issue/Problem/Failure Detected	Importance	Advisory Message
Detects if the ethernet cable is connected but not yet configured.	WARNING	Configure Ethernet
Detects if the DNS has been configured or if the instance is unable to resolve a domain name.	WARNING	Configure DNS
Detects if there are block devices unassigned/undefined.	INFO	Configure Storage
Detects if it is in Free tier.	INFO	Apply License
Detects if there is a valid licence.	INFO	Apply License
Management Center only. Detects if there are member nodes added.	INFO	Add member nodes
Detects if there is a featured platform installed or if the log receiver has not been configured.	INFO	Featured Platform
Detects if any accounts have been created other than the defaults.	INFO	Create new users
Detects if any one of the block devices has reached its high watermark(90%).	DANGER	Expand disk space

Bulk Provisioning Feature

The Bulk Provisioning process is a step-by-step workflow that allows the configuration of multiple instances. This is usually achieved by the nominated **Management Center** instance.

If you are using **Gemini Central** to build Splunk clusters, whatever the platform; hardware appliances, cloud instances, or virtual nodes, you will almost certainly want to make use of the **Bulk Provisioning** method for the deployment of multiple instances. For your convenience, we have made this one of the first sections of this Administration Guide - your **Gemini Central** journey therefore, starts here!

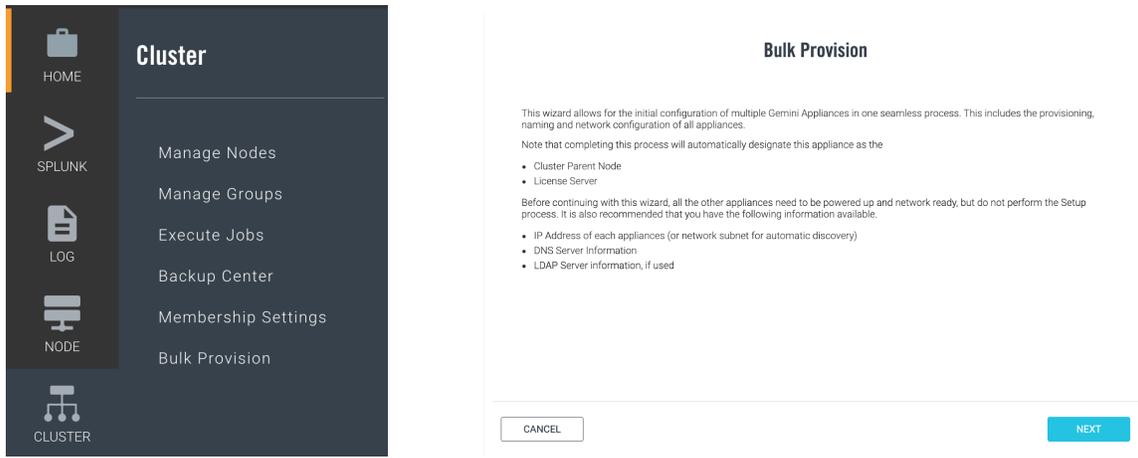
Ensure that the following are true before continuing;

- Each instance contains the installed Gemini Central software
- Each instance has been powered on

- Each instance is accessible on the network

To begin the **Bulk Provision** process of your Gemini instances, login to the Gemini web interface of your chosen **Management Center** as the **'admin'** user.

Select the **'Bulk Provision'** option from the **Cluster** menu.



This will invoke the **Bulk Provisioning wizard** allowing for the simultaneous configuration of multiple instances. This option covers all aspects of configuration including; server naming, network configuration, LDAP access, and software installation.

Select the **'Next'** button to navigate through the wizard.

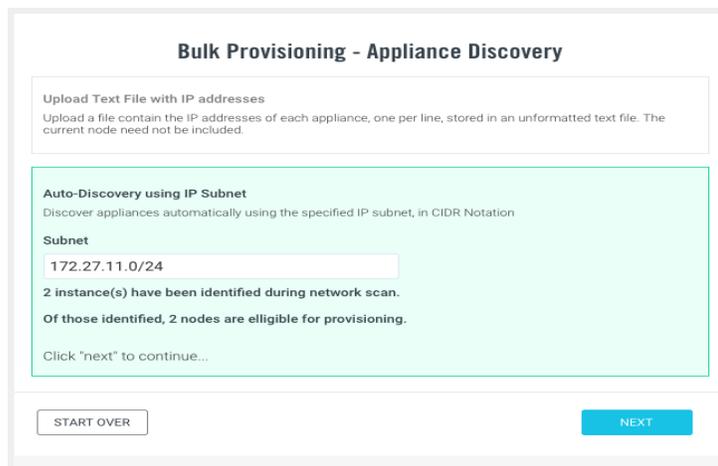
Bulk Provisioning - Instance Discovery

The first step will discover and confirm instances that are to be configured. The discovery process can be achieved in one of two ways:

Use of a text file: If all the IP addresses of instances are known, you may add them to a text file containing one IP/DNS name per line. This can be uploaded during the discovery process and has the advantage of naming the instances with logical Splunk instance names, ie. idx_01, idx_02

Subnet search: Alternatively, the instances can be discovered on the network by performing an IP subnet scan, using CIDR notation to specify (ie, 192.168.156.0/24).

Note that the scope of the subnet will determine the length of time taken for the scan to complete.



Bulk Provisioning - Network Settings

This step determines the mechanism used for assigning IP addresses to the instances. Two options are available here;

- **Network settings - DHCP assigned**

The default option uses a dynamic assignment of IP addresses using a DHCP Server. Note: This could mean a change to an instance IP address under some circumstances, but it is a fully automated solution.

- **Network settings - Static assigned (preferred option)**

This option will also utilize DHCP. The difference here is that DHCP is used for the initial IP address assignment only. From then on, the IP address becomes 'static'. This is useful when using a temporary DHCP server during deployment.

The screenshot displays the 'Bulk Provisioning - Network Settings' configuration page. At the top, the Gemini logo is visible. The main content area contains a form with the following elements:

- DNS Server:** A text input field containing the value '8.8.8.8'.
- Network settings - DHCP assigned:** A radio button option that is currently selected. Below it, a small note reads: 'Use dynamically assigned network settings (by DHCP server)'.
- Network settings - Static assigned:** A radio button option that is currently unselected. Below it, a small note reads: 'Use static assigned network settings. All the settings will use currently assigned configurations.'
- Navigation:** 'BACK' and 'NEXT' buttons are located at the bottom of the form.
- Notification:** A green banner at the bottom of the page indicates 'Provisioning Configuration updated'.

Bulk Provisioning - Hostname assignment

If you have chosen the '**Static assigned**' option you have two options to consider regarding **hostname** assignment;

Using a Reverse DNS Lookup

If DNS records have been assigned for each instance, this option will use the hostname discovered through the DNS server. This requires an 'A record' entry for each device, and the inclusion of a valid **DNS server** address in the entry box provided.

Specify Custom Hostname Pattern

This option allows the specification of a custom string to automatically compose each hostname. The following tokens can be used dynamically in this process:

- **`$$service_tag$`** - will use the 'service tag' of the instance, as indicated on the box or available from the Gemini support website.

- **\$increment\$** - provides an automatically incrementing number usually used with a text prefix.

The example below will create instances named; **gemini-1, gemini-2, gemini-3**, etc

Bulk Provisioning - Change Admin Password

This step is used to specify the password for the **'admin'** account on **each** instance to be configured, allowing access to the Gemini web interface. This can be changed at a later date if required.

It is recommended that you use a strong password and follow the recommended appropriate password policy guidelines as required. Note that when using bulk provisioning, all instances will be updated with the same admin password.

Bulk Provisioning - Connect to LDAP

If you wish to use an LDAP server as a resource to access each instance, carefully enter the **Base DN** and other details required in the following screen.

This can be set at a later date, so feel free to 'skip' this configuration.

Please refer to the [LDAP](#) section, located within **Settings / Authentication** of this guide for more details.

Please note that LDAP authentication is optional, and by default the instance resorts to local account-based authentication.

If you do not require to set up an LDAP connection, select the **'skip'** button to move on.

Bulk Provisioning - Connect to LDAP

Resource Name
Admin from IT Group

Base DN
OU=IT, DC=example, DC=com

Login Attribute
uid

Host
ad1.example.com

Port
389

SSL

BACK SKIP NEXT

Bulk Provisioning - SSH Authentication

This step allows for the control of SSH authentication.

Use the first panel to set SSH passwords for the built-in **'sbox'** and **'splunk'** user accounts. Be sure to carefully record both of these assigned passwords.

Alternatively, use the second panel here to upload the **SSH private key** to complete the connection

Bulk Provisioning - SSH Authentication

Use a password for SSH authentication
Set a password for the 'sbox' and 'splunk' user which are available as SSH credential. Minimum length 6 characters.

Password for 'sbox'
.....
show password

Password for 'splunk'
.....
show password

Upload SSH key
Use public-key authentication to connect to the CLI using SSH.

Drop file here or click to select...

BACK SKIP NEXT

Bulk Provisioning - Summary Screen

The **summary** screen lists all the instances about to be provisioned along with their configuration settings. It should be used as a final review and confirmation step before starting an automated configuration.

It is also strongly recommended that you use the **Download CSV** option here, as it contains all the instance details provisioned during this process for future reference.

Select the **'Start'** button to initiate the automated provisioning process, which can take several minutes or longer depending on the number of instances being provisioned.

The status of each instance is updated in real-time. After all instances have been provisioned, select **Finished** to complete the process.

Bulk Provisioning - Summary

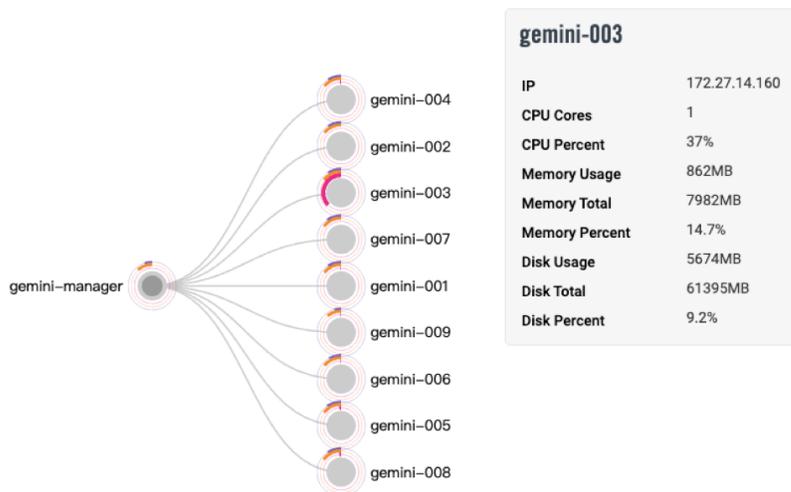
Provisioning Nodes				
IP Address	Hostname	Use LDAP	SSH Authentication	Status
172.27.11.109	gemini-5004-57	No	Password	✓
172.27.11.110	gemini-001	No	Password	✓
172.27.11.111	gemini-002	No	Password	✓
172.27.11.112	gemini-003	No	Password	✓
172.27.11.113	gemini-004	No	Password	✓
172.27.11.118	gemini-005	No	Password	✓
172.27.11.115	gemini-006	No	Password	✓
172.27.11.116	gemini-007	No	Password	✓
172.27.11.117	gemini-008	No	Password	✓

Success!
 You are all set. All appliances have been setup successfully.
 To configure additional settings such as Timezone and install Splunk, continue by:

- View Cluster Topology
- Create Node Groups
- Execute Cluster Jobs
- Activate Splunk Manager

BACK
START OVER
DOWNLOAD CSV
FINISHED

Use the **Cluster / Manage Nodes** dashboard of the **Management Center** node to view the Manage Cluster of Gemini instances.



At this stage, it is likely that you will want to install and configure Splunk clusters using the provisioned Nodes. For details regarding this process, please refer to the [Splunk Installation](#) section for details.

Adding instances to an existing Manage Cluster

The addition of other Gemini nodes into an existing **Manage Cluster**, must be achieved from the **'Parent'** node. A parent node is the main control node for each unique Gemini Cluster. The most common parent node is of course the **Management Center**, and the addition of other instances to the parent cluster can be achieved using its **Cluster / Manage Nodes** dashboard.

Note

In Version 2.8 and above, the addition of instances can be completed from the Management Center itself. Previous versions of Manage required the Parent Token String to be added to each child instance individually.

Gemini Central OS child instances:

Additional child nodes that include Gemini Central software could be; Gemini appliances, cloud instances or virtual nodes. In order to add these as child nodes, navigate to the **Cluster / Manage Nodes** dashboard, and select the '+ Add Node' button to enter its hostname or IP address.

Adjust Storage Plan

Beware of disk partitions and mount points. Not all disks are mounted on the system with default partitions, especially the Gemini Appliance models. There are different storage configurations on each appliance model;

- Onboard flash drives
- All hard disk drive(HDD)
- All solid-state disks(SSD)
- Hybrid configurations of the above

If the Appliance is planned to store a large amount of data, ie. a Splunk indexer, complete the following checks before you start to deploy any applications:

1. Understand the storage devices and mount points on the Appliance. Navigate to **NODE / Storage / Storage** and you will see a list view of storage devices and mount points.
2. HDD and SSD are mounted on the following mount points by default:
 - HDD disk will be mounted on **/opt/mnt/hdd01**.
 - SSD disk will be mounted on **/opt/mnt/ssd01**.
 - Not applicable to the following models: G1000, IB-1050D.
3. Design your storage plan and adjust the logical volumes and mount points. The default storage plan might not satisfy your needs. You may adjust them to the new logical volumes and new mount points. For example:
 - Unmount SSD and remove it from a logical volume, and merge it into **/opt** to extend it's capacity.
 - Unmount HDD and mount it on **/opt/splunk** so that the whole splunk including binary, configs, and data are stored on the same disk.

4. Mount points are available for the various **Featured Platforms**, as shown opposite.

Alternatively, create and include your own, including subdirectory. Using an individual disk volume for specific featured platforms will help prevent the application from impacting the system and data store.

Easy to expand or migrate data especially when the instance is running on public cloud platforms like AWS or Azure.

See also the section on [Managing Swap space](#), for guidance on whether this should be enabled.

The screenshot shows a web interface for configuring file system mounts. The title is 'Mount to File System - sdb'. Below the title is a section labeled 'Mount Point' with the instruction 'Choose mount points below and fill subdirectory if required.' There are six input fields, each containing a default path: /opt/splunk, /opt/explore, /opt/tableau, /opt/minio, /opt/parseme, and /opt/mnt/. The /opt/mnt/ field is currently empty.

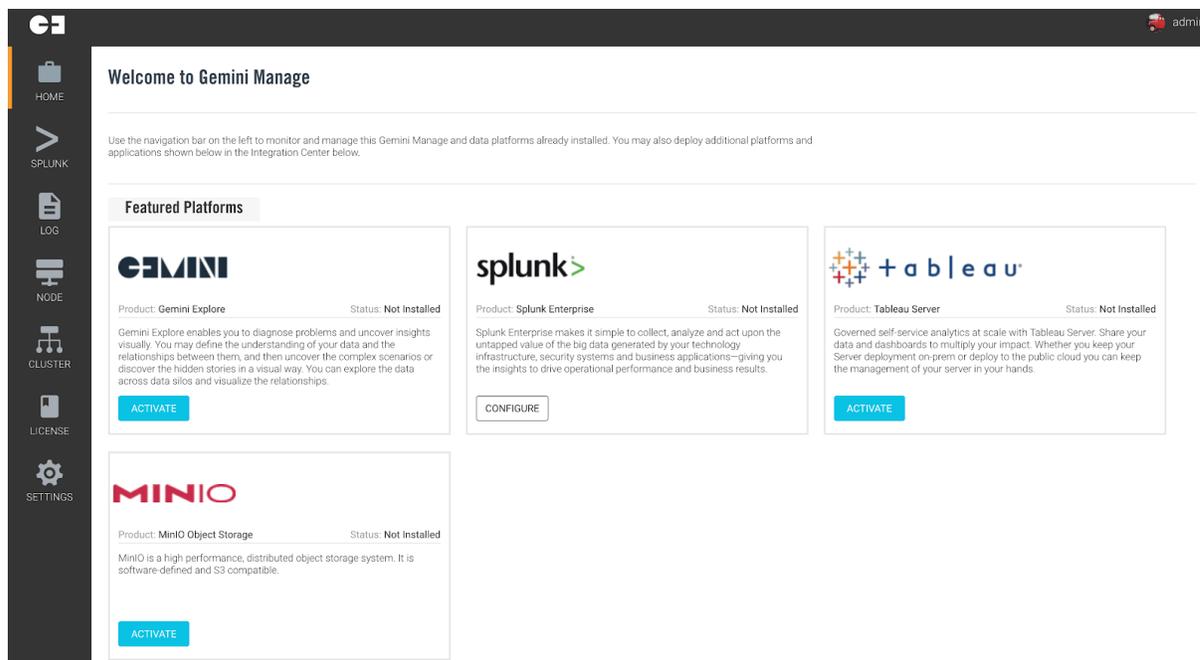
Note

A mix of variant storage types and speed, e.g. SSD, HDD and iSCSI connected disks in one RAID disk or in one logical volume is not recommended. It will slow down disk performance and make it unstable.

HOME Dashboard

The **Home** dashboard is both the first and default screen following a successful login.

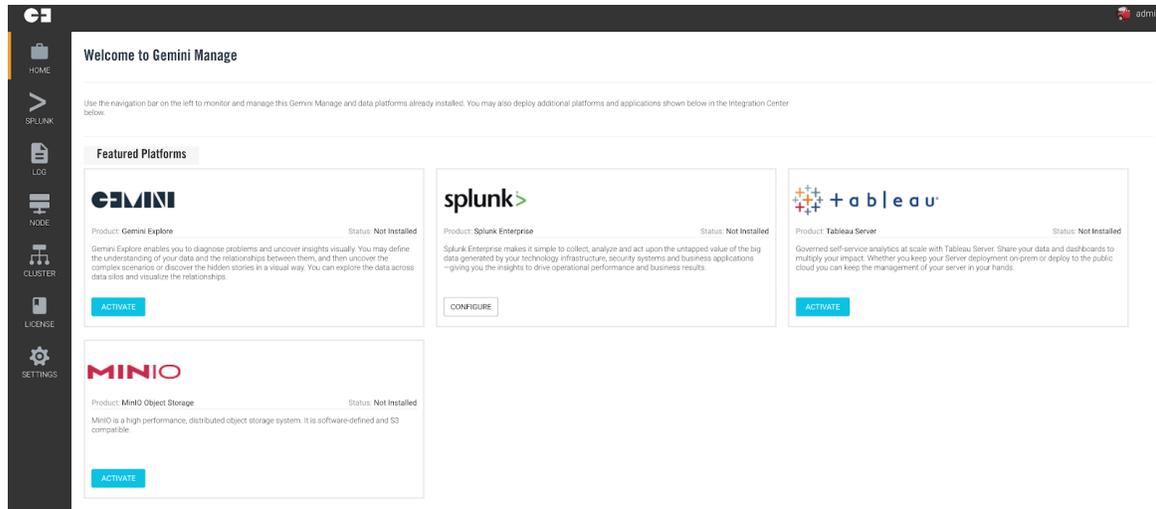
The entire **Gemini Central** experience is organized into different areas accessed from the vertical navigation menu bar on the left of the screen.



HOME	This is the default home page upon login. It provides access to the setup environment and offers the ability to operate as a standalone instance or Management Center.
SPLUNK	On completion of Splunk Enterprise ' Activation ' on the instance, this area will offer configuration parameters for the Splunk environment.
LOG	Provides configuration for Gemini Log Receiver (Syslog-NG server)
NODE	Provides configuration settings for network, identity, and instance preferences.
CLUSTER	Allows for the management of instances as part of a multi-node environment.
LICENSE	Allows for the centralized management of Manage licenses.
EXPLORE	On completion of an ' Activation ' of Explore on the instance, our intuitive Gemini Explore investigation product will be made available.
TABLEAU	On completion of an ' Activation ' of Tableau on the instance, Tableau server is made available for all your business analysis purposes.
MINIO	On completion of an ' Activation ' of MinIO, this instance can be used as part of a MinIO SmartStore S3 compatible cluster.
SETTINGS	This section provides configuration settings for Manage, such as; system update, authentication, security settings, and backup/restore.

Featured Platform

Featured platforms consist of processing platforms available for activation on each instance. Select the 'Activate' button for easy deployment of a **feature platform** with just a few clicks of the mouse.



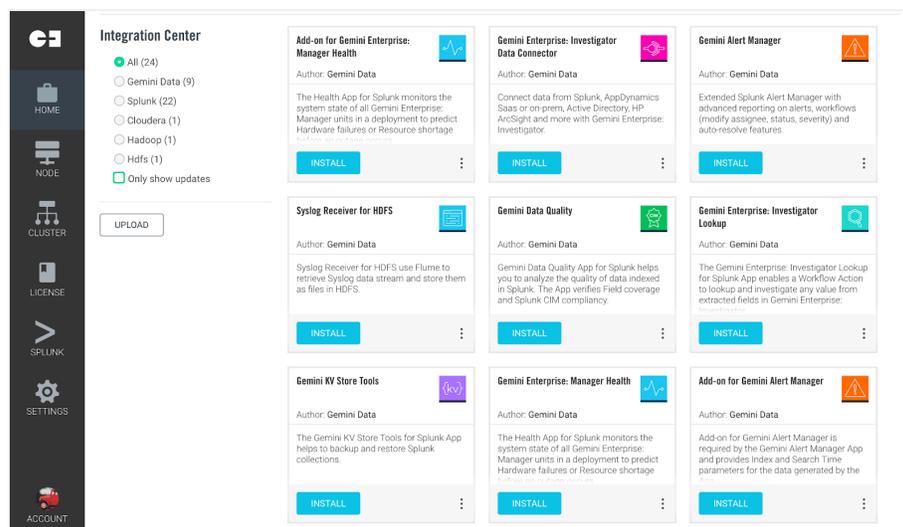
Please note that each platform will consume resources. Installing more than one platform on a Gemini instance will naturally result in competition for available resources.

Ensure your instance meets the hardware specifications suited to the workload required by each platform. Please consult with **Gemini Support** (support@gemini.com) on sizing considerations.

Integration Center

The **Integration Center** contains a repository of applications and apps that provide insight tools, management tools, and connectivity options to sources of data.

Some apps will mirror those provided from other platform ecosystems such as Splunkbase, whilst others are specific to Gemini. Regardless of origin, all apps are fully supported by Gemini Data.



Each solution or application is represented by a descriptive card that allows for easy installation, configuration or removal. Use the **vertical ellipsis** configuration menu for specific screens pertaining to each application.

Splunk Installation

The following process ensures the proper installation and configuration of Splunk inside Gemini Central.

Note

The instance does not ship with the Splunk binaries. Download the required version from the Splunk website. You will need a Splunk account to achieve this.

Before proceeding, ensure that you have the following prerequisites:

- An active Splunk account to access Splunk installation binaries.
- A downloaded Splunk Enterprise 64-bit Linux binary (.tgz)
- Splunk Enterprise Licence (Note, this can be added later)

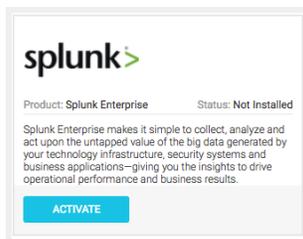
If you do not have a Splunk account or have not downloaded the latest installation file, instructions and links are provided within the **Splunk / Daemon / Install Splunk Enterprise** dashboard within the Management Center, following **Activation**.

Note

Apple/Safari users, you will want to ensure you have disabled 'Open Safe Files after downloading' to retain the .tgz extension.

Activate Splunk

From the **Gemini Central Home** dashboard, enable the **Splunk** menu by selecting the '**Activate**' button located within the **Splunk Enterprise** panel.



Splunk Clustered Environment Installation

A '**Splunk Environment**' can contain **Indexer Clusters** and **Search Head Clusters** created in one or more locations, together with one or more Splunk standalone instances.

If you require one or more Splunk environments to be created from existing Gemini instances, or if you want to add more instances to grow your existing Splunk environment, use the **Splunk Environments** dashboard at the Gemini **Management Center** node.

If you need to interact with **Splunk** on any Gemini Central instance, it will first need to be **'Activated'**.

Note: The Management Center node does not have to have Splunk installed. It is simply used as a control node for your Splunk instances.

From the **HOME** page of the **Gemini web interface**, locate the **Splunk** panel, and from this select the **'Activate'** button. This will create a **'Splunk icon'** at the vertical menu bar if it does not already exist.

To manage existing Splunk instances, login to the dedicated Gemini **Management Center** and open the **Splunk Environments** dashboard using the **Splunk icon**.

To create a new **Splunk Environment**, refer to the [Bulk Provisioning](#) process.

Nodes already created using the **Bulk Provisioning** process should now be visible in the **'Unassigned Nodes'** panel (see below).

SPLUNK > Splunk Environments

Environments
 Clusters
 Nodes

No Environments

Unassigned Nodes (7)

<input type="checkbox"/> Name	IP	Type	Splunk Software	Site	Deployment Type	Status
<input type="checkbox"/> gemini-005	10.1.5.45	Software Appliance	Not Installed	-		<input checked="" type="checkbox"/>
<input type="checkbox"/> gemini-007	10.1.5.47	Software Appliance	Not Installed	-		<input checked="" type="checkbox"/>
<input type="checkbox"/> gemini-004	10.1.5.44	Software Appliance	Not Installed	-		<input checked="" type="checkbox"/>
<input type="checkbox"/> gemini-003	10.1.5.43	Software Appliance	Not Installed	-		<input checked="" type="checkbox"/>
<input type="checkbox"/> gemini-006	10.1.5.46	Software Appliance	Not Installed	-		<input checked="" type="checkbox"/>

If there are missing nodes following the Bulk Provisioning process, or if you want to bring in remote Splunk clusters to a Splunk Environment, these will need to be added as **'Unassigned Nodes'** before they can be re-assigned to a Splunk Environment.

Details on adding **Unassigned Nodes** as either standalone instances or complete Splunk Clusters, can be obtained from the **'Adding a Node'** section located in the [Splunk Environments](#) section of this Administration Guide.

Building the Splunk Environment

When **Unassigned Nodes** have been made available and a **Splunk Environment** already exists, simply add them to your existing environment, using the **'Add to Environment'** button at the bottom of the screen.

Alternatively, if this is a new installation, or if for some reason you require a new **Splunk Environment**, create a new Splunk Environment using the **'+ Build Environment'** button (top-right of the screen). This will reveal a four-step wizard that enables the capture of the required detail in order to build the desired Splunk configuration.



A typical Splunk configuration process includes the following stages;

- **Step 1:** Creating or specifying a Splunk Environment name, cluster arrangement, and Splunk binary version to be installed
- **Step 2:** Creating a Splunk cluster for the environment.
- **Step 3:** Organize nodes into a cluster.
- **Step 4:** Specify the site name for the cluster.

Repeat the above steps to create additional Splunk Environments or Clusters.

Note

If required, refer to Splunk documentation architecture Best Practices to understand more about how clustering works in Splunk.

Step 1: Creating the Splunk Environment

In this step there are several attributes to be specified:

- **Deployment Type**

Select **'Deploy Multi-Use Environment'** for a traditional Splunk Indexer and Search Head Cluster configuration.

Select **'Deploy Independent Stream Forwarders'** to create a Splunk Stream forwarder. (Special use-case only).

- **Environment Name**

Create a suitable name for your Splunk Environment. This is simply a label and can therefore include spaces etc.

- **Available Sites**

Traditionally, Splunk clustering is of the single-site or multi-site variety, however, with **Gemini Central** we always begin by creating a multi-site cluster arrangement. This has many advantages for future scaling or data migration and will still enable you to work in a single-site arrangement if required.

For this reason, use the value of **'site1'** here to represent a **'single-site'** cluster arrangement. For **'multi-site'** cluster arrangements, additional site names can be added here using a comma as a delimiter, ie. site1,site2, etc

For further information regarding Splunk clustering, please refer to Splunk documentation.

- **Splunk Software**

Select the version of Splunk required in this environment. Ideally, there should be only one Splunk version used across the entire Environment.

Use the **'Choose the file'** upload link provided to upload a new Splunk binary.

Deployment Type
Choose to deploy Independent Stream Forwarders only or Multi-use Environment

Deploy Multi-Use Environment
This environment will be comprised of some combinations of Search Head Clusters, Indexer Clusters, and/or standalone instances.

Deploy Independent Stream Forwarders
This environment will only be comprised of Independent Stream Forwarders.

Environment Name
Specify a name to identify your Environment.

Gemini Cluster - Appliances

Available Sites
Add a comma-separated list of physical or logical locations. Assigning nodes to sites will be done at a later step.

site1 x add a tag

acceptable sites: site1, site2, ..., site63

Splunk Software
Select or upload the desired Splunk software to be used in this process.

Splunk Enterprise v7.3.1

Drop package here or click to [choose the file](#) from your computer

Select the **'Organize Cluster'** button at the bottom of the screen to reveal the following:

Create Environment 2 Organize Clusters 3 Organize Nodes

This Environment

Gemini Cluster - Appliances
Splunk Enterprise v7.3.1 site1

Organize Clusters

+ New Cluster

Add as many clusters as needed for this environment. Assigning nodes to clusters will be done in a later step.

Step 2: Create Splunk Clusters

Select the **'+ New Cluster'** icon to create a new Splunk Cluster. Note that it is required that you create an **Indexer Cluster** before moving on to add a **Search Head Cluster**.

Creating an Indexer Cluster

Use a suitable **'Name'** for the new **Indexer Cluster**

Ensure that the **'Type'** is set to **'Indexer'**

Create a new **'Splunk secret key'** that will be used to authenticate the cluster members.

Select the **'Organize Nodes'** button to progress in creating the Indexer Cluster

This Environment

Appliance Training Environment
Splunk Enterprise v7.3.3  site1

Organize Clusters
Add as many clusters as needed for this environment. Assigning nodes to clusters will be done in a later step.

Name

Type  

Splunk Secret

 New Cluster

Creating a Search Head Cluster

A Search Head Cluster can only be created once an Indexer Cluster has been provisioned. The following will also be required in order to complete this operation;

- The Cluster Master IP address
- The Indexer Cluster secret key

Use a suitable **'Name'** for the new **Search Head Cluster**.

Ensure that the **'Type'** is set to **'Search Head'**

Enter a **secret key** for the Search head Cluster in the **Splunk Secret** box used for authenticating Search Heads to their cluster. Best practice dictates that this should be different from the secret key used for the Indexer Cluster.

Enter the **Cluster Master** IP address in the **Indexer Master URI** box.

Enter the **Indexer Cluster** secret key in the **Indexer Secret:** box. This is visible in its encrypted form within the **Splunk Secret** entry for the **Indexer Cluster**. Do not be tempted to copy and paste this into the Indexer Secret box. Always use the original secret key assigned to the Indexer Cluster prior to its encryption.

Note

Creation of a **Search Head Cluster** will require identification of a valid **Cluster Master**. If you are building a brand new Splunk environment, an **Indexer Cluster** will need to be provisioned prior to creating a **Search Head Cluster**.

This Environment

Appliance Training Cluster
site1

Organize Clusters
Add as many clusters as needed for this environment. Assigning nodes to clusters will be done in a later step.

Name	gemi-cluster-idx-01	✖
Type	Indexer	▼
Splunk Secret	RENTWlpcWFBZSREOCgULBQI=	
Name	gemi-cluster-sh-01	✖
Type	Search Head	▼
Splunk Secret	shclustersecret	
Indexer Master URI	10.1.5.41	
Indexer Secret	indexerclustersecret	

+ New Cluster

Note

Secret keys for both the **Indexer** and the **Search Head** Clusters should be recorded and held in a secure place. These are fundamental to the successful completion of any future cluster related function.

Select the '**Organize Nodes**' button to progress in creating a Search Head Cluster.

Step 3: Organize the Nodes

From the '**Available Nodes**' section, select nodes to create your Splunk cluster.

For an Indexer Cluster ensure that you include a Cluster Master(Master Node).

For a Search Head Cluster ensure that you include a Deployer.

If this is to be the first Indexer Cluster created in the Management Center, refer to the '**Initial Splunk Environment**' section below.

Use the graphical interface to select the required number of nodes. There are two ways of achieving this as shown below.

- Select multiple nodes and assign them to the appropriate cluster using the '**+ Add To Cluster**' icon

+ Add To Cluster + Assign Standalone Node

<input type="checkbox"/> Name	IP Address	
<input checked="" type="checkbox"/> gemini-002	10.1.5.42	⋮
<input type="checkbox"/> gemini-005	10.1.5.45	⋮
<input type="checkbox"/> gemini-007	10.1.5.47	⋮
<input checked="" type="checkbox"/> gemini-001	10.1.5.41	⋮
<input type="checkbox"/> gemini-004	10.1.5.44	⋮
<input checked="" type="checkbox"/> gemini-003	10.1.5.43	⋮
<input type="checkbox"/> gemini-006	10.1.5.46	⋮

+ Add to Indexer Clusters

gemini-app-idx-cluster1

gemini-app-shc-cluster1

+ Add to Search Head Clusters

- Or select an individual node and assign it to a relevant cluster using the vertical ellipsis icon to the right.

+ Add To Cluster + Assign Standalone Node

<input type="checkbox"/> Name	IP Address	Model	Type
<input type="checkbox"/> gemini-002	10.1.5.42		⋮
<input type="checkbox"/> gemini-005	10.1.5.45		⋮
<input type="checkbox"/> gemini-007	10.1.5.47		⋮
<input type="checkbox"/> gemini-001	10.1.5.41		⋮
<input checked="" type="checkbox"/> gemini-004	10.1.5.44		⋮
<input type="checkbox"/> gemini-003	10.1.5.43		⋮
<input type="checkbox"/> gemini-006	10.1.5.46		⋮

+ Add to Indexer Clusters

+ Add to Search Head Clusters

Splunk requirements will be enforced and if the requested configuration is unsuitable a warning will be shown in the status field.

Refer to the **'Your Clusters'** panel to confirm your selections and specify which nodes should become the **Cluster Master** or **Deployer** using the vertical ellipsis icon.

Your Clusters
Review the configuration and change the type if needed

Name	Address	Model	Type	Status
gemini-app-idx-cluster1				Indexer
gemini-001	10.1.5.41	Master Node		✔ ⋮
gemini-002	10.1.5.42	Indexer		✔ ⋮
gemini-003	10.1.5.43	Indexer		✔ ⋮
gemini-app-shc-cluster1				Indexer
gemini-004	10.1.5.44	Master Node		✔ ⋮
gemini-005	10.1.5.45	Indexer		✔ ⋮
gemini-007	10.1.5.47	Indexer		✔ ⋮
gemini-006	10.1.5.46	Indexer		✔ ⋮

Set as Master Node

Remove from Cluster

When all additions have been made, select the **'Locate Nodes'** button to progress

Step 4: Locate Nodes

In this step, we specify the node location. This is especially important for a multi-site Cluster environment.

For a single-site operation, all nodes should be assigned 'site1'. This is the first site name used as a default by Splunk.

Alternatively, if you are creating a multi-site environment, assign nodes to site1 or site2 accordingly.

Highlight one or more nodes and use the '+ Set Site' button to assign them to a site from the drop-down list. Typically this will just have the 'site1' value.

Locate Nodes

All Nodes have been automatically located to the first site. Adjust the local assignment as needed.

+ Set Site

Name	IP Address	Model	Type	Site	
gemini-app-idx-cluster1			Indexer		
- <input checked="" type="checkbox"/> gemini-001	10.1.5.41		Master Node	📍 site1	⋮
- <input checked="" type="checkbox"/> gemini-002	10.1.5.42		Indexer	📍 site1	⋮
- <input checked="" type="checkbox"/> gemini-003	10.1.5.43		Indexer	📍 site1	⋮
gemini-app-shc-cluster1			Search Head		
- <input checked="" type="checkbox"/> gemini-004	10.1.5.44		Deployer	📍 site1	⋮
- <input checked="" type="checkbox"/> gemini-005	10.1.5.45		Search Head	📍 site1	⋮
- <input checked="" type="checkbox"/> gemini-007	10.1.5.47		Search Head	📍 site1	⋮
- <input checked="" type="checkbox"/> gemini-006	10.1.5.46		Search Head	📍 site1	⋮

Select the '**Deploy**' button to continue with the Splunk environment deployment.

At this point in time, the Splunk Cluster will be built. This will usually take a few minutes to complete but will depend on the complexity and number of Nodes involved.

Initial Splunk Environment

If this is your first installation within the Management Center and you have successfully built an **Indexer Cluster**, you will now have the required **Cluster Master** references that can be used to build a subsequent **Search Head Cluster**.

Return to the **Splunk Environments** Dashboard using the **Splunk Icon** on the Vertical Menu Bar of the Home page.

The **Cluster Master** IP address will be required to set up the **Search Head Cluster**. In order to discover the assigned IP/DNS name of the Cluster Master, select the **Node / Name** option from the vertical menu.

Return to the **Splunk Environments** dashboard and choose the appropriate number of instances from the '**Unassigned Nodes**' panel to form a **Search Head Cluster** (minimum of 4), then select the '**Create New Cluster**' button to reveal the following:

Environment Name

Specify a name to identify your Environment.

Appliance Training Cluster

Available Sites

Add a comma-separated list of physical or logical locations. Assigning nodes to sites will be done at a later step.

site1 x

acceptable sites: site1, site2, ..., site63

At **Step 1** of the Wizard, select the appropriate Splunk **Environment Name** for which you require a Search Head Cluster, and add the **site** detail accordingly. This will usually consist of **'site1'**, unless you are building a multi-site cluster environment, in which case you will need to add **'site2'**, etc.

Select the **'Organize Cluster'** button when done.

At **Step 2**, select the **'+ New Cluster'** button to create a new Cluster separate to the existing Indexer Cluster

Enter an appropriate name for your Search Head Cluster.

Change the **'Type'** to **Search Head**.

Enter a **secret key** used for authenticating the Search Heads to their cluster

Enter the **IP address** of the **Cluster Master** node.

Enter the **Indexer Secret** from the **Indexer Cluster**. Ensure that you use the original secret key here and not the encrypted value visible on this screen.

Select the **'Organize Nodes'** button on completion.

This Environment **Appliance Training Cluster**

site1

Organize Clusters
Add as many clusters as needed for this environment. Assigning nodes to clusters will be done in a later step.

Name	gemini-cluster-idx-01	✖
Type	Indexer	
Splunk Secret	RENTWlpcWFBZSREOCgULBQI=	
Name	gemini-cluster-sh-01	✖
Type	Search Head	
Splunk Secret	shclustersecret	
Indexer Master URI	10.1.5.41	
Indexer Secret	indexerclustersecret	

+ New Cluster

At **Step 3**, you will be presented with the following screen. Select the nodes required from the **'Available Nodes'** presented, and use the **'+ Add To Cluster'** button to assign them to the newly created **Search Head Cluster** listed.

Available Nodes

Select unassigned nodes and assign them with the designated roles

Name	IP Address	
gemini-816fdc	10.1.5.46	⋮
gemini-beba10	10.1.5.47	⋮
gemini-acbab2	10.1.5.44	⋮
gemini-22f35f	10.1.5.45	⋮

+ Add To Cluster

Add to Indexer Clusters

gemini-cluster-idx-01

Add to Search Head Clusters

gemini-cluster-shc-01

The **'Your Clusters'** panel below allows you to choose which node you want to assign as the **Deployer** instance and confirms the Clusters now present.

Your Clusters

Review the configuration and change the type if needed

Name	Address	Model	Type	Status	
gemini-cluster-idx-01				Indexer	✓
gemini-002	10.1.5.42		Cluster Master	✓	
gemini-003	10.1.5.43		Cluster Peer	✓	
gemini-001	10.1.5.41		Cluster Peer	✓	
gemini-cluster-shc-01				Search Head	
gemini-acbab2	10.1.5.44		Deployer	✓ ⋮	
gemini-816fdc	10.1.5.46		Search Head	✓ ⋮	
gemini-beba10	10.1.5.47		Search Head	⋮	
gemini-22f35f	10.1.5.45		Search Head	⋮	

Set as Deployer

Remove from Cluster

At **Step 4**, select the **'Locate Nodes'** button to assign this cluster to a **'site'**. Highlight instances in the cluster, as shown below, and select the **'+ Set Site'** button to select the site number. This will generally be **'site1'** in a single site cluster arrangement. This may change to **'site2'** or **'site3'**, etc, if you are using a multi-site cluster arrangement.

Finally, select the **'Deploy'** button to create this Search Head Cluster using the information provided.

Locate Nodes

All Nodes have been automatically located to the first site. Adjust the local assignment as needed.

Name	IP Address	Model	Type	Site
gemini-cluster-idx-01				Indexer
- gemini-002	10.1.5.42		Cluster Master	📍 site1
- gemini-003	10.1.5.43		Cluster Peer	📍 site1
- gemini-001	10.1.5.41		Cluster Peer	📍 site1
gemini-cluster-shc-01				Search Head
- ✓ gemini-acbab2	10.1.5.44		Deployer	📍 site1 ⋮
- ✓ gemini-816fdc	10.1.5.46		Search Head	📍 site1 ⋮
- ✓ gemini-beba10	10.1.5.47		Search Head	📍 site1 ⋮
- ✓ gemini-22f35f	10.1.5.45		Search Head	📍 site1 ⋮

+ Set Site

site1

Verification of the Splunk Environment

Select **Splunk / Environments** from the vertical menu-bar at any time to obtain an overview.

SPLUNK · Splunk Environments

Environments Clusters Nodes Search for name, IP

Name	Type	Site	Version	Contains
Appliance Splunk Cluster				2 Clusters
gemini-app-idx-cluster1				3 Nodes
gemini-001	Cluster Master	site1	Splunk Enterprise 7.3.1	
gemini-003	Cluster Peer	site1	Splunk Enterprise 7.3.1	
gemini-002	Cluster Peer	site1	Splunk Enterprise 7.3.1	
gemini-app-shc-cluster1				4 Nodes
gemini-004	SHC Deployer	site1	Splunk Enterprise 7.3.1	
gemini-005	SHC Member	site1	Splunk Enterprise 7.3.1	
gemini-007	SHC Member	site1	Splunk Enterprise 7.3.1	
gemini-006	SHC Member	site1	Splunk Enterprise 7.3.1	

Should you see anything other than the expected output here, you may need to destroy the cluster and re-attempt addition. Verify that you have entered the correct site references, which should all be set to **'site1'** if there is to be only one site present. Also ensure that you have entered the correct IP address for the Cluster Master, and secret key for the Indexer Cluster when creating a Search Head Cluster.

Post Splunk Cluster Deployment Features and Tasks

Splunk Activation

If the **Bulk Provisioning** feature of Gemini Central has been used to deploy Splunk Clusters, it is still required to **'Activate'** Splunk on each Gemini Instance in order to access Splunk control and management features via the Gemini web interface. This action does not affect the local Splunk installation in any way, it simply grants access via the Gemini web interface.

It is required to **'Activate'** Splunk on each Gemini instance using the button located on the **Home** dashboard. The Splunk Menu icon will then appear in the Vertical menu bar.

Splunk Web Port (8000)

Any instance other than a Search Head or Cluster Master will have its **splunk-web** port disabled by default, meaning that you will not be able to access the Splunk web interface (8000).

If you wish to enable the web port select **Splunk / Web Interface** from the menu bar and ensure that **Splunk Web Control** is enabled. Use the **Open Splunk Web** heading to open Splunk in a new tab.

Note

This is not recommended for indexers! The instance will need to be restarted - do not restart an individual Indexer that is participating in an Indexer Cluster

Splunk Receiver Port (9997)

Important: Any instance that has been designated a **Splunk Indexer** will not have a default **Receiver Port** setting of 9997. This means that the Indexers will not be able to receive data. Customers must choose a receiver Port value and apply this to all Splunk Indexers. This can be achieved from an **inputs.conf** file setting applied via the **Cluster Master** or **Splunk Deployment Server**.

Splunk Boot-Start

Splunk is installed with the **'boot-start'** feature enabled, using **systemd** control. This preferred service control method will be used unless Splunk has already been installed (remote agent option) and is using the older **initd** method of control, in which case this will be left in place.

Splunk Workload Management features

Splunk's **Workload Management** feature can be used if desired with Splunk on Gemini Central instances, providing the **'systemd'** control is in place (see above).

Splunk admin password

Important: When Splunk is deployed using **Bulk Provisioning**, the default admin password is automatically changed from **'changeme'** to **'gemini123'**. This important action allows for better feature control via the management port (8089). Change this password to conform to your own password if required, this can be done using the **Splunk / Daemon / Advanced Configurations** dashboard.

SSL Passwords

Important: OS User accounts such as **sbox** and **splunk** have a default expiry of 60 days on their accounts. If you have changed the SSH passwords from their defaults, and you wish to freeze them for the future, navigate to the **Settings / Password Policy** dashboard and remove the checkmark from the Password Expiration box. Alternatively, adjust the settings at this dashboard to meet your own company Password Policy.

Heartbeat monitoring feature

Gemini Central contains the ability to monitor and advise on Splunk instances. This feature can be found on the **Splunk Environments** dashboard. The heartbeat period is 60 secs, and if Splunk is detected as unavailable on an instance, the **heartbeat icon** will turn from green to red.

Resource activity

Understanding resource utilization such as disk space and CPU usage can be monitored in real-time from the **Cluster / Manage Nodes** dashboard in Gemini Central.

Using the Splunk Web interface

Verify the formation of the Splunk Environment using Splunk's own **Indexer Clustering** dashboard located on the **Cluster Master** instance.

In order to check the status of your **Splunk** environment, you will first need to access the **Cluster Master** Gemini instance.

Use the following procedure to access the **Cluster Master** web interface, and indeed any other Splunk Node that you wish to gain web access.

1. Note the **Node** name that is associated with the **'Cluster Master'** from the **Splunk Environments** screen.
2. Select **Node / Name** from the vertical menu-bar and locate its IP address

3. Use a browser to access the Splunk UI directly, using `http://<CM_address>:8000`

Alternatively, login to the **Gemini instance** at `https://<CM_address>` using the 'admin' password, and select **Splunk / Web Interface** from the menu-bar to **'Open Splunk Web'**.

1. Login to Splunk at its web interface, and locate the **Indexer Clustering** dashboard from the **Settings** menu.

This should show the status of both the Indexer and Search Head Cluster. Below is an example of the result you can expect. If anything other than green checkmarks, you may need to investigate further.

Please feel free to contact support@gemini.com for any assistance with your Splunk Environment.

The screenshot shows the 'Indexer Clustering: Master Node' dashboard. At the top, there are three green checkmarks indicating that 'All Data is Searchable', 'Search Factor is Met', and 'Replication Factor is Met'. Below these, it displays '2 searchable' and '0 not searchable' Peers, and '3 searchable' and '0 not searchable' Indexes. A table below lists the peers: gemini-001 and gemini-003, both at site1, with a status of 'Up' and 6504 buckets.

Peer Name	Site	Fully Searchable	Status	Buckets
gemini-001	site1	✓ Yes	Up	6504
gemini-003	site1	✓ Yes	Up	6504

Adding more Nodes

If you need to add an **Unassigned Node** to either an existing Indexer or Search Head Cluster, refer to the [Add to cluster](#) section

If you need to add another instance to the Management Center as an Unassigned Node resource, refer to the [Adding a Node](#) section.

Adding a Splunk License Manager

Splunk requires an Enterprise License to be installed and shared among the Indexer and Search head peers. Splunk uses a Licence Manage mechanism to achieve this. For details on how to set this up in a Gemini Central environment, please refer to the [Gemini Central - Create a Splunk License Server](#) document

Adding a Splunk Monitoring Console

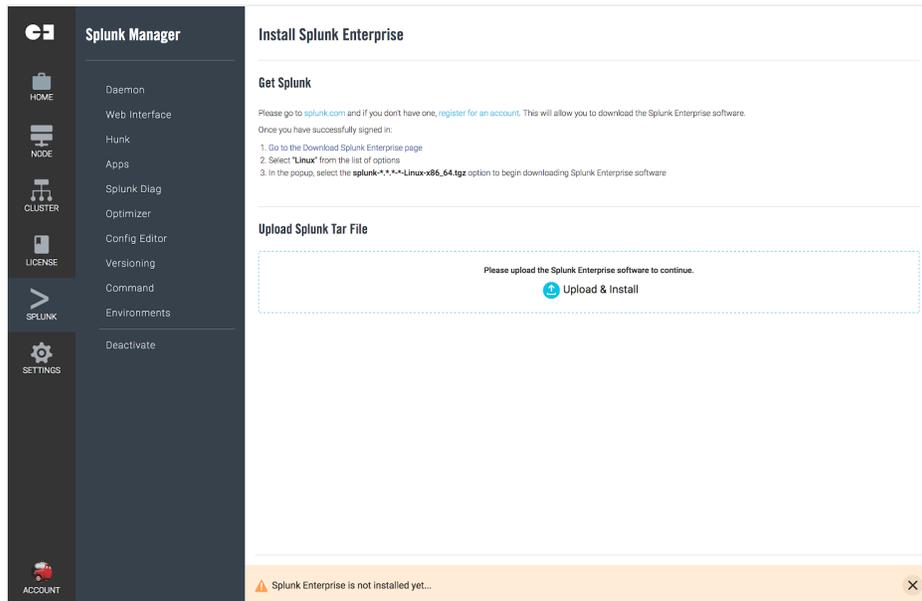
The **Splunk Monitoring Console** is an invaluable tool for maintaining your Splunk Environment and assisting in troubleshooting, however, it is not enabled by default and we strongly recommend that this app is enabled on one of your Gemini instances. We would recommend the Gemini **Management Center** node, or alternatively, if you have less than 20 Instances, the **Cluster Master** instance would be an ideal choice. For larger sites, we would recommend either dedicating a specific stand-alone search head or to use another low usage instance that is perhaps running as a Splunk License Master or Deployment Server.

For details on how to set this up in a Gemini Central environment, please refer to the [Gemini Central - Enable a Splunk Monitoring Console](#) document.

Standalone Splunk Installation

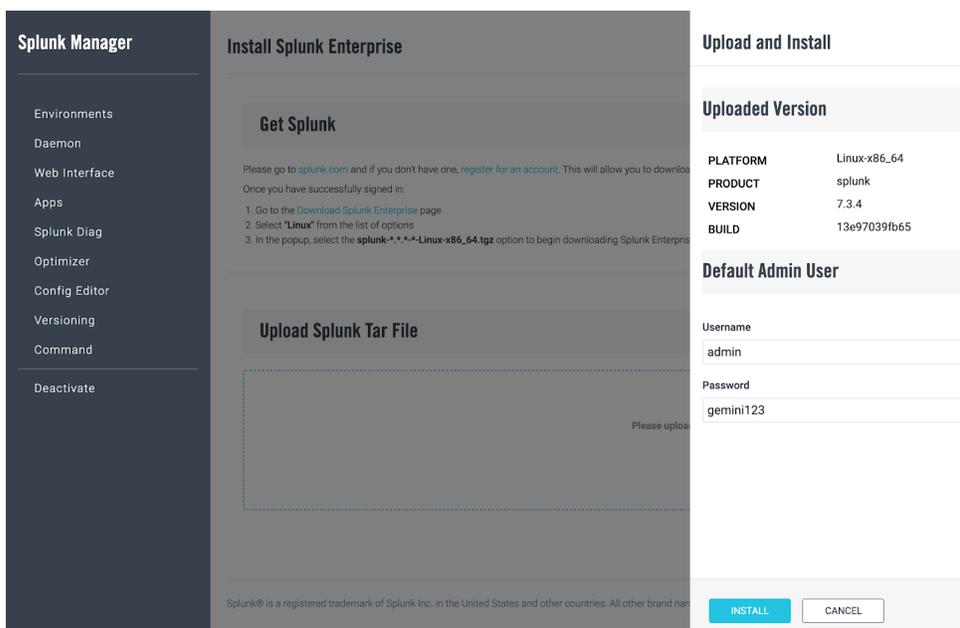
If you want to use the instance as a standalone Splunk instance, you will first need to **'Activate'** the Splunk panel from the **Home** dashboard, and then navigate to the **Splunk / Daemon** dashboard.

From the **'Upload Splunk Tar File'** panel, select the **'Upload & Install'** button to begin the upload process.



When the file has finished uploading, its status is displayed in the subpanel. Change the default admin credentials supplied if desired, then select the **'Install'** button and accept the Splunk Software License Agreement.

The Splunk Enterprise software will now be installed on your Gemini instance.



On completion of the installation, the **Splunk / Daemon** dashboard is displayed;

SPLUNK > Daemon

Splunk Service Control

SPLUNK HOME /opt/splunk
Version Splunk 7.3.4 (build 13e97039fb65)

Boot-Start

Enable Splunk Boot-Start
 Enable BOOT-START in order to start Splunk daemon automatically at boot time.

Run Splunk as a systemd service.

It is highly recommended that the '**Boot-Start**' option is enabled and run as a **systemd** service, as opposed to the older **initd** method of service control.

Ensure that the '**Run Splunk as a systemd service**' checkbox is 'made' prior to switching the **Boot-start** slider to the right to enable (Unless you prefer to use the **initd** service control).

Further options from the **Splunk** menu could be considered at this point including;

Splunk Manager

- Environments
- Daemon
- Web Interface**
- Apps
- Splunk Diag
- Optimizer
- Config Editor
- Versioning
- Command

Web Interface - Turning off the Web port to prevent access.

Optimizer - Select a Splunk best practice template from the options given.

Versioning - Create an initial Splunk configuration status that could be used for Rollback purposes.

Config Editor - The ability to upload, move, modify, copy and extract Splunk configuration files.

Command - The ability to run Splunk CLI commands through the web interface.

Gemini Explore

Gemini Explore is an intuitive visual graph-based data exploration tool that works directly on Splunk, CSV or JDBC data sources.

Using this dynamic multi-layer visualization tool, the user is able to drill-down and interact with their data in a whole new way.

Intuitive to use, as it mimics the way our brains ‘think’. On discovering something interesting, ‘click’ to instinctively locate more detail and how it may relate to other datasets.

Prerequisites:

- Minimum of 200GB available disk space
- Minimum of 4 CPU cores (8 is recommended)
- Minimum of 16GB available RAM (32GB is recommended)
- Port access required: 80:tcp, 9000:tcp
- Public web access is required to acquire the Explore binary file for installation
- A valid Gemini Central **Trial** or **Enterprise** license

Activating Gemini Central

If you are interested in using this product, select the ‘**Activate**’ button from the Featured Platforms panel of the **HOME** dashboard associated with Gemini Explore.

The screenshot shows a 'Featured Platforms' section with three cards:

- Gemini Explore:** Product: Gemini Explore, Status: Not Install. Description: Gemini Explore enables you to diagnose problems and uncover insights visually. You may define the understanding of your data and the relationships between them, and then uncover the complex scenarios or discover the hidden stories in a visual way. You can explore the data across data silos and visualize the relationships.
- Splunk Enterprise:** Product: Splunk Enterprise, Status: Splunk 7.3.4 (build 13e97039fb65). Description: Splunk Enterprise makes it simple to collect, analyze and act upon the untapped value of the big data generated by your technology infrastructure, security systems and business applications—giving you the insights to drive operational performance and business results.
- Tableau Server:** Product: Tableau Server, Status: Not Install. Description: Governed self-service analytics at scale with Tableau Server. Share your data and dashboards to multiply your impact. Whether you keep your Server deployment on-prem or deploy to the public cloud you can keep the management of your server in your hands.

The **Explore** feature is currently available to include in Gemini Central as a ‘**Beta Trial**’ option, available on request from **Gemini Data**(support@gemini.com), in the form of an **explore-x-x.pack** install file.

Once this **Explore** pack has been received, use the ‘**Upload**’ facility to add **Explore** to the Gemini Central **HOME** dashboard.

Install Gemini Explore

The screenshot shows a section titled 'Upload Gemini Explore install pack'. Inside a dashed box, there is a prompt: 'Click here to choose the tarball' with an 'Upload' button below it.

At the install prompt, use the ‘**Install**’ button to proceed.

Install Gemini Explore

Install Components

Component	Installed	Running
frontend		
router		
frontend-ui		
mongodb		
parser		
neo4j		
graphapi		

INSTALL

The installation process collects the required components from a secure internet location and will take several minutes to complete. On completion, the **Explore Components** resource dashboard will be fully populated. This dashboard can be opened at any time, using the **Explore** icon on the vertical menu bar, to view resources currently in use.

Gemini Explore

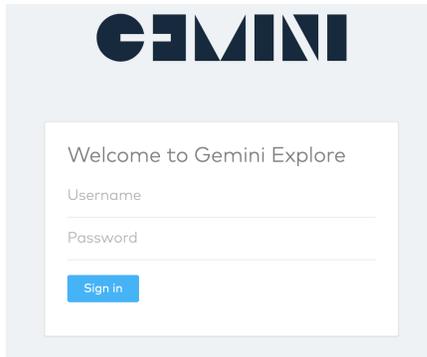
● Stop
 ⊞ Open Gemini Explore
 ✖ Uninstall

Explore Components

Service	Status	CPU	Memory			Network		I/O	
			Used (MB)	Limit (MB)	Percent	Send (KB)	Receive (KB)	Read (KB)	Write (KB)
frontend	✓	0%	198.4	30156.8	0%	35020.8	26316.8	0	0
frontend-ui	✓	0%	19.32	30156.8	0%	38.9	705	0	0
mongodb	✓	0%	32.78	30156.8	0%	26112	34816	0	0
parser	✓	0%	56.86	30156.8	0%	20275.2	8683.52	0	0
redis	✓	0%	1.762	30156.8	0%	379	390	0	0
neo4j	✓	0%	750.1	30156.8	0%	1730.56	1720.32	0	0
router	✓	0%	1.562	30156.8	0%	17408	18432	0	0
graphapi	✓	0%	385	30156.8	0%	21.8	20.9	0	0
websocket-manager	✓	0%	20.64	30156.8	0%	204	191	0	0

Use the **'Open Gemini Explore'** button to open the **Explore** web interface in a new browser tab.

Note: If you receive a **'400 Bad Request'** message, ensure you use the **https://** prefix and continue on to accept the inevitable certificate warning.

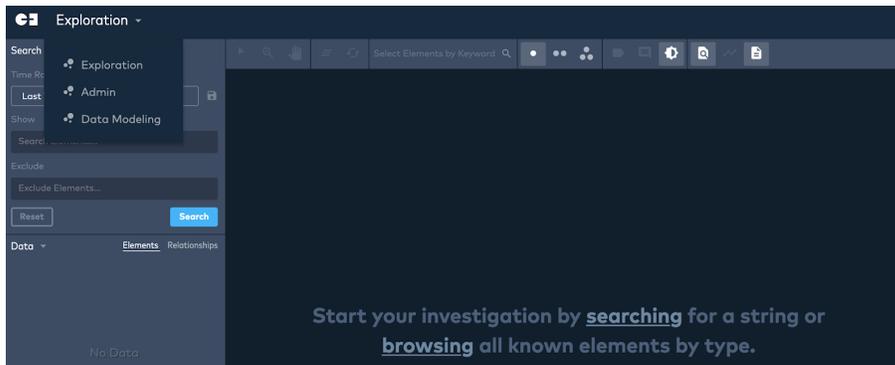


Login using the default credentials:

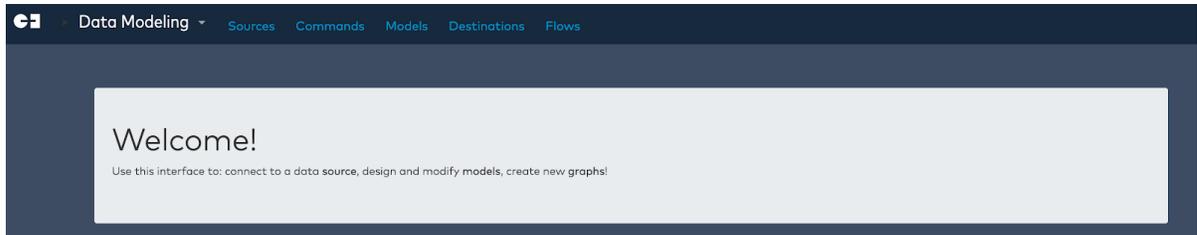
Username: **init@geminidata.com**

Password: **changeme**

This will open the **Explore** canvas, and give access to the **Exploration** menu.



Select the **Data Modeling** menu option to begin your journey with exploring data using this exciting technology.



Licensing Gemini Central (Beta Trial)

In order to use **Gemini Explore** as a fully functioning 'Beta trial' within Gemini Central 2.8, it needs to connect to a valid Manage license. This can be provided from the local Manage instance running Gemini Explore, or from an optional external Manage instance that is running an Enterprise license.

From the **Gemini Central** web interface, navigate to **License / License Status** and verify that either the **Trial** license is in place, or a License Server connected and is still within its expiry date.

Navigate to **License / License Server**, and select the 'Yes' tab to '**Allow Remote Access**'

Add an asterisk (*) in the box marked '**White List**', and select the '**Update**' button (see below)

License Server Settings

Allow Remote Access

Allows another appliance to use this appliance as a license server.

 No

 Yes

Token String

The security token another appliance needs to enter to access this license server.

White List

You could use wildcard * as any character for remote IP address, one line for one rule.

Notes

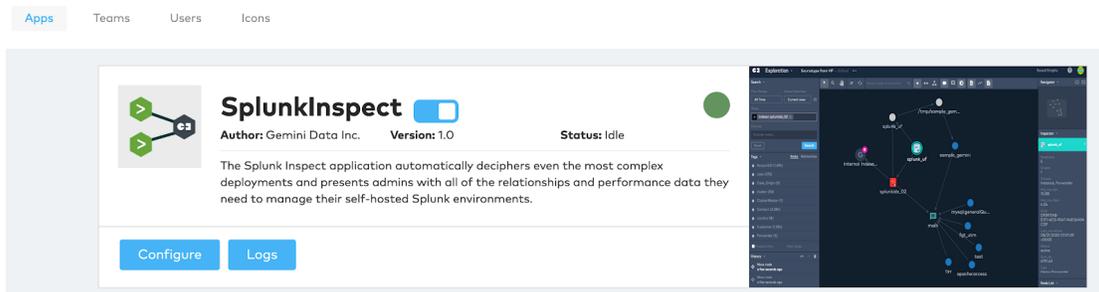
1. **Gemini Explore** must detect a valid **Trial** or **Enterprise License**. The Trial license will expire after 30 days.
2. For help & guidance on this issue, contact support@gemini.com

Integrating Inspect for Splunk Admin

As an introduction to the **Explore** experience and to what this product can offer, we have embedded our **Inspect for Splunk Admin** app for your convenience. This will interrogate any Splunk server running a configured Splunk **Monitoring Console**, and allow you to view details of the Splunk infrastructure, including relationships between Splunk instances and their components.

We would always recommend that a **Monitoring Console** instance is created as part of a Splunk deployment. Within a Gemini **Splunk Environment**, we would recommend use of the **Cluster Master** node as an appropriate choice for the Monitoring Console location.

Once a Monitoring Console source has been established, select the **'Admin'** dashboard from the **Data Modeling** menu and from here select the **'Apps'** sub-menu.



Select the **'Configure'** button to reveal the following Monitoring Console setup screen;

Leave the **Scheme** setting at the default **'Https'**

The **Host** field represents the Splunk instance operating as the Monitoring Console. Enter an IP address or FQDN of such an instance.

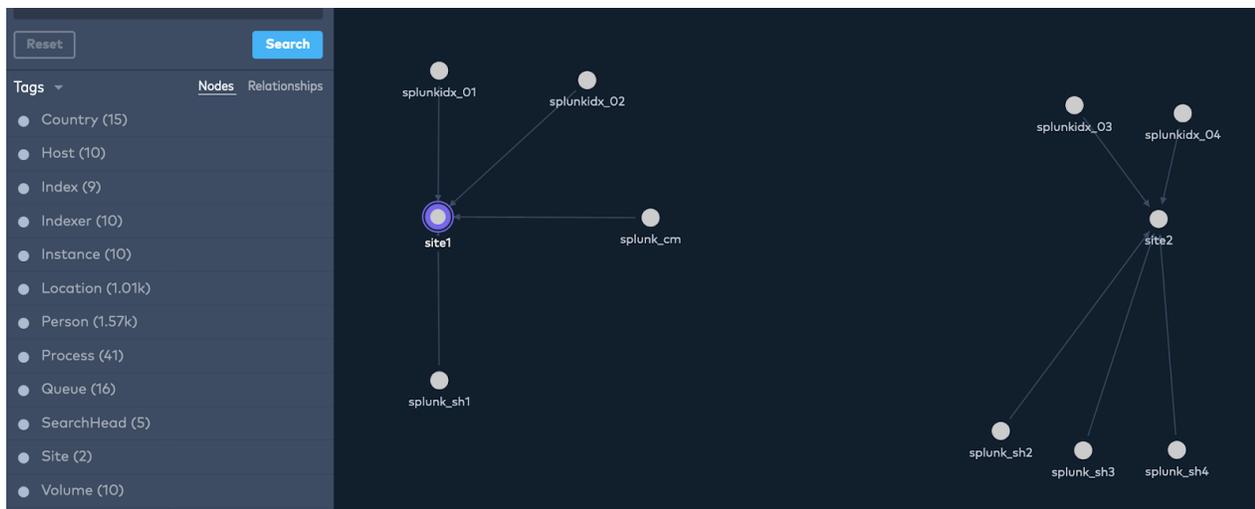
The **Port** field represents the Splunk management port operating at the Monitoring Console instance. This will usually be the default port of 8089.

Use an admin user's credentials for the **Username** and **Password**.

Select the **'Submit'** button on completion.

Return to the **Admin / Apps** dashboard at any to disable **Inspect for Splunk Admin** or check on its data Flow status. This app will take several minutes before anything becomes visible at the Exploration canvas, or the 'completed' Data Flows on the above App indicator starts to increase from zero. This is perfectly normal. When all of the Data Flows will have completed(37), the Exploration canvas can be used.

Open the canvas from the **Exploration** menu, the **Inspect for Splunk Admin** node categories will be visible to the left of the canvas revealing your nominated Splunk environment.



Use the **'+'** icon or the **'Show'** entry box to add elements to the canvas.

Right-click elements to **'Insert Neighbors'** or **'Insert Relationships'**, or **double-click** to explore all the connections automatically.

For more information on getting started with **Gemini Explore**, please refer to the [Gemini Explore - User Guide](#)

For more information on **Inspect for Splunk Admin**, please refer to the [Inspect for Splunk Admin App- Quick Start Guide](#)

Note also, that from the **Exploration** menu, you can select the **Data Modelling** dashboard to reveal that **Inspect for Splunk Admin** has added a new **Source** to Explore that replicates Splunk connection detail supplied in the setup screen.

Sources			Add new
Name	Type	Created At	
SplunkInspect	splunk	2020-08-05T18:05:16.000Z	Delete

Tableau Server

We have offered the ability to run **Tableau Server** as a featured product within Gemini Central. The latest Tableau binary file can be uploaded and integrated within our web interface. If you already have a Tableau license, this may be added at the License prompt.

Prerequisites:

- Minimum of 15GB available disk space (50GB would be a typical working amount)
- Minimum of 4 CPU cores (8 is recommended)
- Minimum of 16GB available RAM (32GB is recommended)
- Root permission to complete the installation and perform administrative Tableau tasks
- Public web access required for dependencies, registration and initialization.

Installation of Tableau Server

If you are interested in running **Tableau** inside Gemini Central for business intelligence or analysis purposes, select the '**Activate**' button from the **Featured Platforms** panel.

Featured Platforms



Product: Gemini Explore Status: Not Install

Gemini Explore enables you to diagnose problems and uncover insights visually. You may define the understanding of your data and the relationships between them, and then uncover the complex scenarios or discover the hidden stories in a visual way. You can explore the data across data silos and visualize the relationships.

[ACTIVATE](#)



Product: Splunk Enterprise Status: Splunk 7.3.4 (build 13e97039fb65)

Splunk Enterprise makes it simple to collect, analyze and act upon the untapped value of the big data generated by your technology infrastructure, security systems and business applications—giving you the insights to drive operational performance and business results.

[ACTIVATE](#)



Product: Tableau Server Status: Not Install

Governed self-service analytics at scale with Tableau Server. Share your data and dashboards to multiply your impact. Whether you keep your Server deployment on-prem or deploy to the public cloud you can keep the management of your server in your hands.

[ACTIVATE](#)

This will open up a Tableau Server Installation page;

If you do not already have the Tableau Server **rpm** file, use the link in **Step 1** to obtain the latest version, direct from the Tableau website.

Upload the appropriate Tableau **rpm** file using the ‘**Upload and Install**’ link at **Step 2**.

On completion of the upload, the **Tableau Server Installation** screen will be displayed giving details of the version.

A default Admin user for Tableau is suggested;

Username: tableau

Password: tableau

Edit these details accordingly if preferred, but record them securely.

Note that this process also creates a Gemini **OS User** with the same credentials, that can be used for SSH access.

Select the ‘**Start**’ button at the bottom of the panel to begin the Server Installation.

The Tableau Server will take a few minutes to install, and on completion will revert back to the **Installation Wizard** dashboard;

Step 2: Upload and Install Tableau Server Software

 Upload & Install

Step 3: Run Initialize Wizard on Tableau Server

After the Tableau Server is installed, it will lead you to the Initialize Wizard: [Tableau Server Manager](#)

Run through the Tableau initialisation process by selecting the [Tableau Server Manager](#) link at **Step 3**: This will take you to the **Tableau Services Manager** screen which has been made available on the following port:

https://<gemini_instance>:8850

Sign in using the credentials entered at **Step 2**:



Sign In to Tableau Services Manager

Enter administrator credentials. [Learn more...](#)

Sign In

Enter your Tableau License key or optionally register for the 14-day trial option;



Enter your license product key to get started with Tableau Server.

Product Key

The key has 20 characters

0000-0000-0000-0000-000

[I can't find my product key.](#)

Activate License

Try it free for 14 days

Start Tableau Server Trial

When initializing **Tableau**, care should be taken to select the correct Identity Store and Gateway options, as these can not be changed following installation, and the default **Gateway Port** is already in use within Gemini Central.

- Select the **'Local' Identity store**, unless you want to use Active Directory
- Change the default **Gateway Port** from '80' to **'8888'** (port 80 is already in use)



The settings below are all you need to get started.

Identity Store

You cannot change the identity store after initializing.

Local

Active Directory

Gateway Port

Port Number: 8888

Select the **'Initialize'** button to instigate the install, and monitor as it moves through the installation process.



Initializing...

Step 11 of 34

Waiting for services to reconfigure.

```

11:22:35 AM succeeded: Updating Configuration.
11:22:35 AM succeeded: Validating that there are no pending changes.
11:22:36 AM succeeded: Generating passwords.
11:22:36 AM succeeded: Generating Unique Cluster Identifier.
11:22:38 AM succeeded: Generating search server ssl certificate.
11:22:39 AM succeeded: Generating Elastic Server SSL certificate.
11:22:41 AM succeeded: Generating ActiveMQ Server SSL Certificate.
11:22:42 AM succeeded: Generating key store.
11:22:43 AM succeeded: Promoting configuration.
    
```

This process could take a while to finish.
 Click [Learn more](#) about configuring your server deployment with Tableau Services Manager. The server will be running after the initialization is complete.
You will need to create a Tableau Server Administrator account when this process finishes.

Installation can take several minutes, and note that a **Tableau Server Administrator** account will need to be created on completion of this process.



Initialization Complete

You will need to create a Tableau Server Administrator account. [Learn more](#)

[Continue](#)

When the **Initialization Complete** message is displayed, return to the **Gemini Installation Wizard** dashboard.

Step 4: should now be visible (a browser refresh may be required at this point);

Step 4: Create Default User Account

Now the Tableau Server is installed and online. Create the first Tableau Server user account in below to complete the installation.

Username

Password

CREATE

Use this to create the required **Tableau Server Administrator** account to enable the login to Tableau. Edit the details with your preferred Username and Password credentials, and select the **'Create'** button.

Login to **Tableau** using the Gateway port set during initialization ie.

https://<gemini_instance>:8888

Following integration with **Gemini Central**, some useful Server controls are available from the **Tableau / Service Control** menu, including an optional Boot-Start feature.

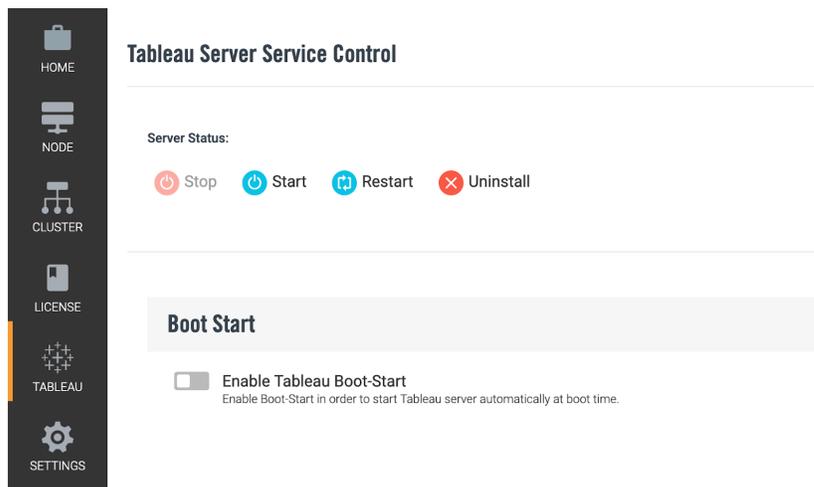


Tableau Server Service Control

Server Status:

Boot Start

Enable Tableau Boot-Start
 Enable Boot-Start in order to start Tableau server automatically at boot time.

MinIO Object Storage

Gemini have teamed up with the S3 object storage company MinIO to offer an alternative Splunk SmartStore option for your storage solution. For details on this solution please contact Gemini on contact@gemini.com

LOG Menu

The **LOG** tab is the starting point for building a **Syslog-NG** environment, otherwise known as Gemini **Log Receiver**, using Gemini instances.

The benefits of a centralized log receiver have been well documented over the years, and **syslog** has gained near-universal support across most platforms.

The **Log Receiver** feature in **Gemini Central** includes 'log splitting' features based on **syslog-ng**, allowing a granular approach to the logging of network-related products and equipment. The inclusion of an HEC receiver as a 'destination' was new in Gemini Enterprise V2.9.

In order to deal with all your syslog requirements, we recommend adding **Log Receiver** as an integral part of the **Splunk Environments** dashboard in **Gemini Central**. This could incorporate a single standalone instance, or a small cluster of instances to offer 'high availability'.

By using our **Manage Group** and **Failover** features, **Load Balancing** and **High Availability** can be achieved. Each Gemini instance will feature both **Splunk** and the **Log Receiver** working together in a **Manage Group** to collate the events and forward them on to your Splunk Indexers or HEC collector.

This section is designed to help you create the necessary rules to receive, filter and store your incoming network-based logs and forward this to your Splunk environment.

The features of Gemini Log Receiver

The Gemini **Log Receiver** dashboard has been designed to offer a simple visual experience that makes it easy to create, view and troubleshoot your Syslog rules.

Other key features are listed below;

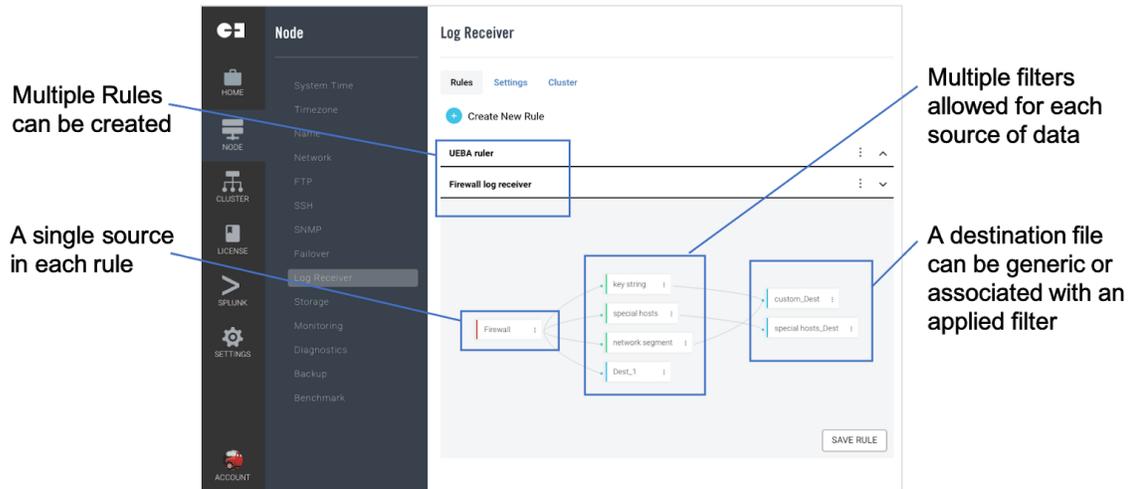
- Multiple rules allowed for various data sources.
- Powerful filters to split Syslog into different log file destinations or HEC collectors.
- Easily integrate with Splunk.
- Integral log rotation for better housekeeping.
- Rules can be replicated and distributed to other nodes if required.

Rule Manager

The **Rule Manager** dashboard can be viewed from the **LOG** menu within the **Gemini** web interface.

A required log receiver '**Rule**' can be divided into three potential sections; **Source**, **Filter**(optional) and **Destination**.

Each **Rule** has flexibility, for instance, it does not necessarily need to contain a **Filter**, and its **Destination** can either be a generic receiver file or be separated into a more granular file, as desired.



Log Receiver - Rule Manager Dashboard

The first step in building a Log Receiver environment in Gemini is to create the Syslog-NG rules required.

To create a new 'Rule' select the '+ Create New Rule' button on the **Log Receiver** dashboard, and assign it a logical name. Select the 'Save' button to move through the process of adding a Source and Destination definition.

Note

It is recommended that you create a naming convention for your log receiver components, ie. Rule, Source, Filter and Destination naming.

If at any time you need to edit or delete a Rule, use the vertical ellipsis menu located adjacent to each Rule.

NODE > Log Receiver

Rules Settings Cluster

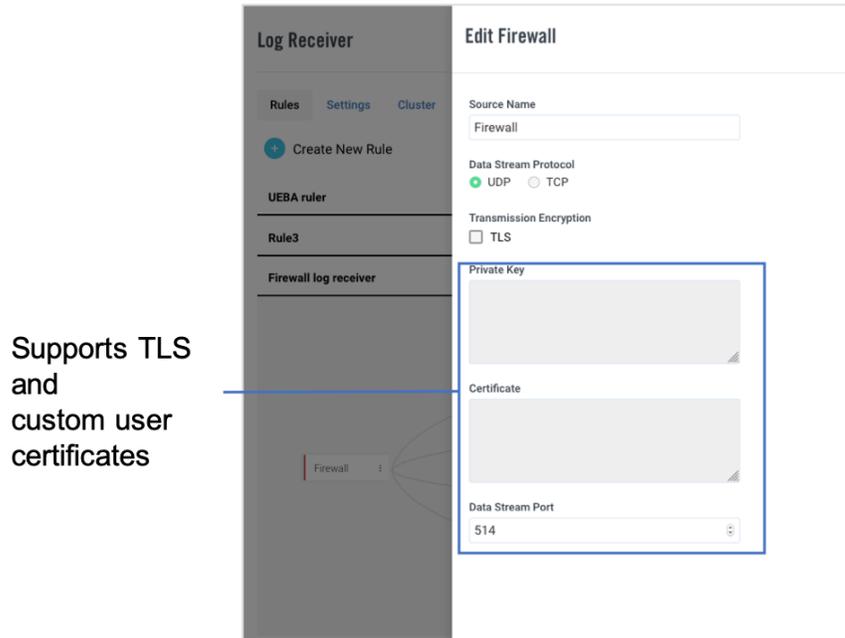
+ Create New Rule

firewall

⋮
Edit
Delete

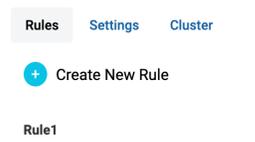
Log Receiver - Source

Enter a logical **Source Name** of your choice to define the source host, and select from either the UDP or TCP protocol and select the port required for the source host.

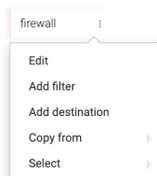


Add encryption in the form of TLS or a custom provided certificate. Note that encryption is only available for the TCP protocol.

Select the **'Save'** button to create the first part of the **Rule**. This will add the **Rule** to the dashboard for the addition of a **Destination** or optional **Filter**.



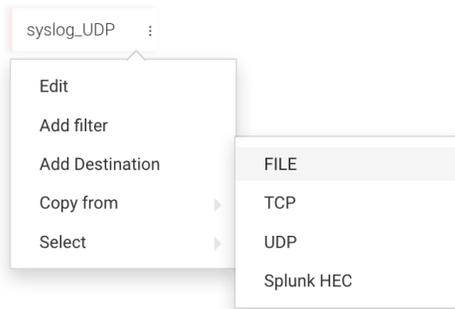
To access the options available for completing or modifying a **Rule**, use the small **vertical ellipsis** button associated with each one, shown in the picture opposite.



- **Edit** - Rules can be edited at any time using this option
- **Add filter** - Enable and configure an optional log receiver filter
- **Add destination** - Each Rule must have a destination applied before it can be saved.
- **Copy from** - Enables a quick way to replicate similar Rule components.

Log Receiver - Destination

A **Destination** is required in order for the **Rule** to be saved. To add an appropriate destination file to the log receiver, select the '**Add destination**' option from the **vertical ellipsis** menu associated with the chosen **Rule** to reveal a submenu.



- When choosing **Add Destination**, a submenu appears offering the choice between a log file, TCP/UDP port, or a Splunk HEC destination.

Note that if a **Filter** option is required, this can be added later (see next section).

Add a logical destination name - preferably using your own naming convention document - for the **Destination Name** entry.

Destination Log File Splitting - Overview

This feature enables a **Rule** for a single **data source** to create multiple destination files based on various criteria, such as host value, Facility or protocol.

Without any log file splitting in place, the destination location and filename of receiver files will be created in the `/opt/sbox/data/<rule_name>/` directory, with a filename dictated by the '**Destination File**' entry box. The exact directory location for the **Destination File** will depend on your choice of **Log File Splitting** applied.

For instance, if the 'Facility' option is chosen separate sub-directories containing events from different syslog daemon facility values will be created (Note: there are case sensitive)

If a syslog message does not fall into another facility value, it will default to the 'user' facility. Otherwise, authentication events will find themselves in the 'auth' directory, kernel events in the 'kern' directory, etc. Examples of resultant directories are given below;

```
/opt/sbox/data/syslog_UDP/user/syslog_UDP.log
```

```
/opt/sbox/data/syslog_UDP/auth/syslog_UDP.log
```

```
/opt/sbox/data/syslog_UDP/kern/syslog_UDP.log
```

Split logs files by specific condition

Store log files with custom path and file name

Log rotation support (24h/size)

Integrate with Splunk by specifying sourcetype & index

Decide on how you want log rotation to operate on each data source.

Selecting a specific Splunk sourcetype and index will help with Splunk’s data input procedure. It is essential that the **Index** or **Indexes** described here have been created on the Splunk Indexers before the Log Receiver is active.

Indexes should be created using a Cluster Master app where Indexer Clustering is used.

Note The **‘Monitor in Splunk’** feature shown here will only be visible if **Splunk** has been **‘Activated’** on this instance.

This option will allow the setting of a Sourcetype used to create a 'monitored input' in the `/etc/system/local/inputs.conf` file

IMPORTANT: It is crucial that any **Index** specified here has been created at **all** your Production Indexers before forwarding is enabled. It will also need to exist locally, if you are testing the rules locally before enabling 'Splunk Forwarding'.

Note

The '**Final**' Log path flags checkbox should not normally be used. This is a special case scenario sometimes required if rules created are in conflict with one another.

Check with support@gemini.com if you are considering using this feature.

Destination Log File Splitting - By host

This option allows us to split the network feed by the originating host.

The host can be identified by either IP address or DNS hostname. Confirm this selection using the '**Settings**' panel located at the top of the dashboard.

Edit syslog_UDP_dest

Destination Name

Log path flags

 Final

Log File Splitting

 None (User Defined)

 Host

 Facility

 Level (Severity)

 Program

User Custom Path

Destination File

Full Path

`/opt/sbox/data/syslog_UDP/$HOST/syslog_UDP.log`

Log Receiver requirement: To receive events over **UDP:514** from various devices on the network and split by '**Host IP address**'.

Using the **Log Receiver** dashboard, we have created;

- a **Rule** called **Syslog Server**
- a **Source** called **syslog_UDP**
- a **Destination** called **syslog_UDP_dest**
- a **Destination Filename** called **syslog_UDP.log**

We have enabled the '**Host**' option from the **Log File Splitting** selector to create separate sub-directories containing events from different devices.

If a syslog message came from the host 10.1.1.12, it would create the following file in the following location;

`/opt/sbox/data/syslog_UDP/10.1.1.12/syslog_UDP.log`

Note: We have chosen to split by IP address(default), not DNS name

Destination Log File Splitting - By Facility

This option allows us to split the network feed by the ‘**selector**’ field of the **syslogd daemon**.

Select the ‘**Log File Splitting**’ value of ‘**Facility**’ to filter on the part of the system **generating the message**, enabling you to split by one of the following keywords;

- | | | |
|---|---|--|
| <ul style="list-style-type: none"> • auth • authpriv • cron • daemon • kern • lpr • mail | <ul style="list-style-type: none"> • mark • news • syslog • user • uucp • local0 through local7 | <p>All these keywords (with the exception of mark) correspond to the similar “LOG_” values specified to the <code>openlog()</code> and <code>syslog()</code> routines</p> |
|---|---|--|

Edit syslog_UDP_dest

Destination Name

Log path flags

 Final

Log File Splitting

 None (User Defined)

 Host

 Facility

 Level (Severity)

 Program

User Custom Path

Destination File

Full Path

/opt/sbox/data/syslog_UDP/\$FACILITY/syslog_UDP.log

Log Receiver requirement: To receive events over **UDP:514** from various devices on the network and split by ‘**Facility**’

Using the **Log Receiver** dashboard, we have created;

- a **Rule** called **Syslog Server**
- a **Source** called **syslog_UDP**
- a **Destination** called **syslog_UDP_dest**
- a **Destination Filename** called **syslog_UDP.log**

We have enabled the ‘**Facility**’ option from the **Log File Splitting** selector to create separate sub-directories containing events from different syslog daemon facility values (Note: there are case sensitive)

If the syslog message does not fall into another facility value, it will default to the ‘**user**’ facility. Otherwise, **authentication** events will find themselves in the ‘**auth**’ directory, **kernel** events in the ‘**kern**’ directory, etc. Example directories below;

/opt/sbox/data/syslog_UDP/user/syslog_UDP.log

/opt/sbox/data/syslog_UDP/auth/syslog_UDP.log

/opt/sbox/data/syslog_UDP/kern/syslog_UDP.log

Destination Log File Splitting - By Level(Severity)

This option allows us to split the network feed by the 'action' field of the **syslogd daemon**, commonly known as '**Severity**'.

Select the '**Log File Splitting**' value of '**Level(Severity)**' to filter on *severity of the message*, enabling you to split by one of the following - listed in order of most critical to least critical;

- emerg
- alert
- crit
- err
- warning
- notice
- info
- debug

These keywords also correspond to the similar "LOG_" values specified to the *syslog()* routine

Edit syslog_UDP_dest

Destination Name

Log path flags

Final

Log File Splitting

None (User Defined)

Host

Facility

Level (Severity)

Program

User Custom Path

Destination File

Full Path

/opt/sbox/data/syslog_UDP/\$LEVEL/syslog_UDP.log

Log Receiver requirement: To receive events over **UDP:514** from various devices on the network and split by '**Severity**'.

Using the **Log Receiver** dashboard, we have created;

- a **Rule** called **Syslog Server**
- a **Source** called **syslog_UDP**
- a **Destination** called **syslog_UDP_dest**
- a **Destination Filename** called **syslog_UDP.log**

We have enabled the '**Level(Severity)**' option from the **Log File Splitting** selector to create separate sub-directories containing events with different **Severity** values, for example;

/opt/sbox/data/syslog_UDP/alert/syslog_UDP.log

/opt/sbox/data/syslog_UDP/crit/syslog_UDP.log

/opt/sbox/data/syslog_UDP/warning/syslog_UDP.log

Destination Log File Splitting - By Program

This option allows us to split the network feed by the **Program** or process involved as defined in the message. This can be useful to segregate by **sshd**, **ftp**, **docker**, etc. if that is something that is required.

Edit syslog_UDP_dest

Destination Name

Log path flags

 Final

Log File Splitting

- None (User Defined)
 Host
 Facility
 Level (Severity)
 Program

User Custom Path

Destination File

Full Path

/opt/sbox/data/syslog_UDP/\$PROGRAM/syslog_UDP.log

Log Receiver requirement: To receive events over **UDP:514** from various devices on the network and split by 'program/process'.

Using the **Log Receiver** dashboard, we have created;

- a **Rule** called **Syslog Server**
- a **Source** called **syslog_UDP**
- a **Destination** called **syslog_UDP_dest**
- a **Destination Filename** called **syslog_UDP.log**

We have enabled the 'Program' option from the **Log File Splitting** selector to create separate sub-directories containing events with different **program/process** values, for example;

/opt/sbox/data/syslog_UDP/sshd/syslog_UDP.log

/opt/sbox/data/syslog_UDP/ftp/syslog_UDP.log

/opt/sbox/data/syslog_UDP/dockerd/syslog_UDP.log

Destination Log File Splitting - User Custom Path

For any of the above options, or at any time during the creation or modification of rules, a separate 'Customer defined' sub-directory can be formed. This can be used to further segregate events perhaps.

The result of adding a '**User Custom Path**' would create one or more subdirectory levels as required that follow the '**Rule**' name.

For example, if an entry of '**mycustomerdir**' was added to the **User Custom Path** input box, the result would become;

```
/opt/sbox/data/<rule_name>/mycustomerdir/<log_file_split>/<Destination_File_name>
```

Note

It is recommended that you create a naming convention for your log receiver components, ie. Rule, Source, Filter and Destination naming.

Log Receiver Settings - Filter

Despite all the options so far discussed, it is often required to enable another layer of filtering to the collection of log files, and this can be achieved by the addition of a '**Filter**'.

Together with splitting the network feed using the '**Log File Splitting**' methods described, further filtering can be achieved by the following three methods;

- Network Segment
- Hostname
- Regular expression

This would, for example, enable us to filter by both **Host** and **Severity** if we required, and as shown in the example below. Notice that color has been added to visually distinguish between the **Source**(red), **Filter**(green) and **Destination**(blue).



Filter by Host

In order to create a **Filter**, select the '**Add Filter**' option from the vertical ellipsis menu at the **Source** of the rule in question. The following example shows the addition of a Filter called 'firewall1'.

Edit firewall1

Filter Name

Type

- Match
Use a regular expression to filter messages based on a specified header or content field.
- Host
Match a regular expression to the text of the log message.
- Netmask
Select only messages sent by a host whose IP address belongs to the specified IPv4 subnet.
For example: 192.168.5.0/255.255.255.0 or 192.168.5.0/24.

Filter

Edit firewall1_Dest

Destination Name

Log path flags

 Final

Log File Splitting

- None (User Defined)
 Host
 Facility
 Level (Severity)
 Program

User Custom Path

Destination File

Full Path

/opt/sbox/data/syslog_UDP/\$LEVEL/firewall1_Dest.log

The filter has been created to specifically locate the **Host 'firewall1'** using a regular expression against the source network feed.

As can be seen here, the **Destination** for this filter is further split by the **'Severity'** value.

The destination file can be seen in the value of the **'Full Path'**.

Filter by Netmask

Edit Prod_network

Filter Name

Type

- Match**
Use a regular expression to filter messages based on a specified header or content field.
- Host**
Match a regular expression to the text of the log message.
- Netmask**
Select only messages sent by a host whose IP address belongs to the specified IPv4 subnet. For example: 192.168.5.0/255.255.255.0 or 192.168.5.0/24.

Filter

By choosing the **Netmask** option, filters can be used to segregate between different Networks by adding notation in the form of **network_address/network_mask** or by **CIDR** notation.

For example, by selecting the **'Netmask'** filter type, and adding **'10.1.5.0/24'** to the Filter entry box, we can segregate the events from this network from others.

Filter by Match

Edit PIX-firewall

Filter Name

Type

- Match**
Use a regular expression to filter messages based on a specified header or content field.
- Host**
Match a regular expression to the text of the log message.
- Netmask**
Select only messages sent by a host whose IP address belongs to the specified IPv4 subnet. For example: 192.168.5.0/255.255.255.0 or 192.168.5.0/24.

Filter

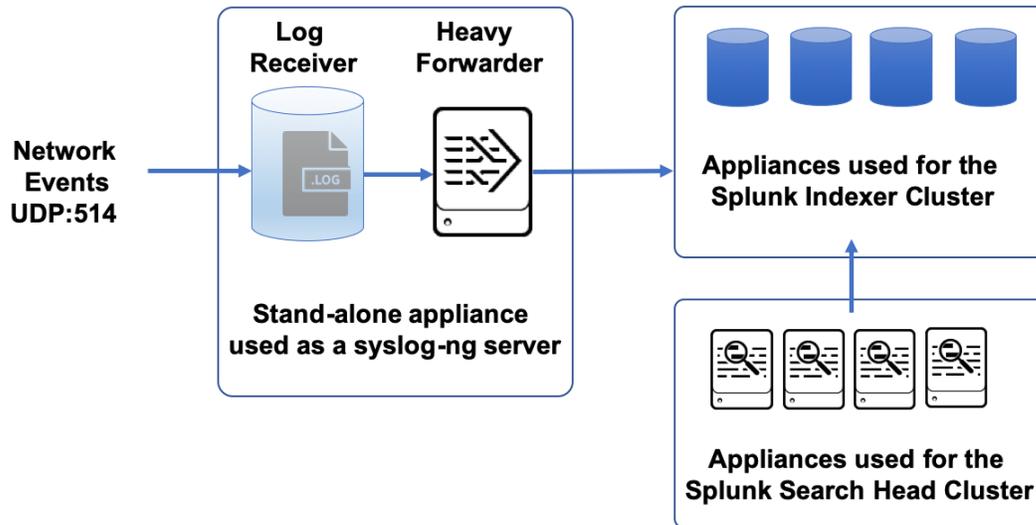
By choosing the **Match** option, filters can be setup for any number of categories by using a series of regular expressions.

For example, by selecting the **'Match'** filter type, and adding **'%PIX'** to the Filter entry box, we can filter specifically for Cisco-PIX firewall messages.

Forwarding Log Receiver events into Splunk

Forwarding the granular logs created by the **Log Receiver** into **Splunk** for analysis and reporting purposes is made easier in **Gemini Central** by the use of its integral **Splunk** instance which can be repurposed as a **Heavy Forwarder** as shown in the diagram below.

Once Splunk has been activated, **Log Receiver** rules will automatically be created as Splunk **Monitored Inputs** in an inputs.conf file. (Note: Splunk will need to be restarted to input any changes).



Enable Splunk Indexers to receive events from the Log Receiver

The following tasks are all to be completed on Production Indexers in receipt of Log Receiver data. This is usually achieved using a 'Base app' from the Cluster Master instance as detailed below.

In order for the **Splunk Indexer Cluster** to receive logs from the **Log Receiver**, all Production Indexers must;

- Have their **Receiver Port** open (default 9997)
- Contain the **Indexes** used in **Log Receiver Rule** settings

Verify the Receiver Port at the Indexers

If you already have data in Splunk, the **Receiver Port**, usually set to 9997, is probably already open to receive events from **Universal Forwarders**. If in doubt, or if this is a new installation, verify this using Gemini's **Splunk / Command** dashboard at any Indexer instance. The following command will confirm the status with a message, 'Receiving is enabled' or 'Receiving is disabled'.

```
display listen -auth admin:<password>
```

If the result 'Receiving is disabled' is displayed, use the following process at the Splunk **Cluster Master** instance;

- Login to the Gemini web interface of the **Cluster Master** instance and navigate to the **Splunk / Config Editor** dashboard.
- Using the config editor screen, click through to locate the **/opt/splunk/etc/master-apps/_cluster/local** directory
- Select the 'Create New File' button, and enter the name **inputs.conf** at the prompt (take care that this is spelled correctly!). Select the 'Add' button to confirm.
- Select the newly created **inputs.conf** file to reveal a simple editor, and copy and paste the following into the box. Select the 'Save' button to confirm

```
[splunktcp://9997]
```

Creating an Index at the Indexers

It is crucial that any **Index** specified in the creation of **Log Receiver Rules** has been created at **all** the **Splunk** indexers before forwarding is enabled. This is usually achieved at the **Cluster Master** by a **Base App** setting in an **indexes.conf** file. Please verify that this has been achieved and that the required indexes exist before proceeding.

Note

This process can be achieved using **Deployment Server** or a similar log management tool. Please refer to your Splunk Admin if in any doubt.

If, as an example, you had created an index destination called 'syslog' when creating your Log Receiver rules, the following **inputs.conf** file would need to be created at the Cluster Master.

- Login to the Gemini web interface of the **Cluster Master** instance and navigate to the **Splunk / Config Editor** dashboard.
- Using the config editor screen, click through to locate the **/opt/splunk/etc/master-apps/_cluster/local** directory
- Select the 'Create New File' button and enter the name **indexes.conf** at the prompt (take care that this is spelled correctly!). Select the 'Add' button to confirm.
- Select the newly created **indexes.conf** file to reveal a simple editor, and copy and paste the following into the box. Select the 'Save' button to confirm

```
[syslog]
homePath = $SPLUNK_DB/syslog/db
coldPath = $SPLUNK_DB/syslog/colddb
thawedPath = $SPLUNK_DB/syslog/thaweddb
repFactor = auto
```

Any changes or additions such as these made at the **Cluster Master** should be followed by a **'cluster-bundle push'** in order to distribute to the Indexers that form the Indexer Cluster. Complete the following procedure after such changes.

- Login to the **Splunk** web interface of the **Cluster Master** node (ie. `http://<cluster_master_IP>:8000`)
- Navigate to the **Settings / Indexer Clustering** dashboard.
- Select '**Configuration Bundle Actions**' from the **Edit** menu (top right of the dashboard)
- Conduct a **'Push'** of the configuration bundle.

A Rolling Restart will probably not be required on this occasion, but leave Splunk to advise.

Testing Log Receiver Rules before enabling forwarding (optional)

This step is purely optional, and you may wish to omit this step if you are familiar with both syslog and the Log Receiver feature.

Note

It is imperative to create a *local* Splunk **index** referred to in any of the Log Receiver rules, for this test to be successful. This test is only recommended if this is a new Log Receiver environment and you want to test the facility and/or rules.

As **Log Receiver Rules** are saved at the **Rule Manager** dashboard, the necessary Splunk input stanzas are automatically added to the local Splunk instance via an **inputs.conf** file stored in the `/opt/splunk/etc/system/local` directory, although please note that a *restart* of Splunk will be required to activate any changes to monitor input(s).

Login to the **Splunk** web interface at your **Log Receiver** instance, and run a search at the index (ie. `index=syslog`) to verify that the **Rules** are working correctly (see below for an example).

Note

If the search does not show results, check that the required Indexes have been created and complete a restart of this instance.

index=syslog

✓ 13 events (before 3/24/20 6:13:48.000 AM) No Event Sampling

Events (13) Patterns Statistics Visualization

Format Timeline - Zoom Out + Zoom to Selection X Deselect

source

5 Values, 100% of events Selected Yes No

Reports

Top values Top values by time Rare values

Events with this field

Values	Count	%
/opt/sbox/data/syslog_UDP/crit/syslog_UDP.log	9	69.231%
/opt/sbox/data/syslog_UDP/10.195.50.250/PIX-firewall.log	1	7.692%
/opt/sbox/data/syslog_UDP/crit/firewall1_dest.log	1	7.692%
/opt/sbox/data/syslog_UDP/ftp/syslog_UDP.log	1	7.692%
/opt/sbox/data/syslog_UDP/user/syslog_UDP.log	1	7.692%

SELECTED FIELDS

- a host 1
- a source 5
- a sourcetype 1

INTERESTING FIELDS

- # date_hour 2
- # date_mday 2
- # date_minute 6
- a date_month 1
- # date_second 9
- a date_wday 2
- # date_year 1
- a date_zone 1
- a index 1
- # linecount 1
- a punct 3
- a splunk_server 1
- # timeendpos 1
- # timestartpos 1

+ Extract New Fields

Create a Heavy Forwarder to forward Log Receiver data

The following tasks are all to be completed from the Splunk web interface acting as a Heavy Forwarder at the Log Receiver instance.

In order to create a **Heavy Forwarder** from the local Splunk instance, we need to complete three tasks;

- Setup **forwarding** of all logs to the Indexer Cluster.
- **Delete** any local **Index** used in the testing of syslog rules (if used for testing)
- Change the **license mode** of this instance to that of a '**Forwarding Licence**'

To set up forwarding of the logs to the Clustered Indexers. Open the **Forwarding and Receiving** dashboard located in the **Settings** menu of Splunk (see below), and select the '**+ Add New**' button.

Forwarding and receiving

Forward data
Set up forwarding between two or more Splunk instances.

Forwarding defaults

Configure forwarding + Add new

Receive data
Configure this instance to receive data forwarded from other instances.

Configure receiving + Add new

Add each **Indexer** and its receiving port to the '**Host**' input box one-by-one, until all Clustered Indexers have been added. If there are many Indexers and you have been granted access to the CLI, it may be easier to edit the `/etc/system/local/outputs.conf` file directly.

Add new
Forwarding and receiving > Forward data > Add new

Enter host:port to forward data to. Data will be auto load balanced to each host:port.

Host *

Set as host:port or IP:port.
You must also enable receiving on this host.

If you have conducted testing of the **Log Receiver** Rules on this instance, delete the Index(s) used throughout the testing process. This action will reset the 'fishbucket' index, allowing the events received during testing to be resent to the Production Indexers.

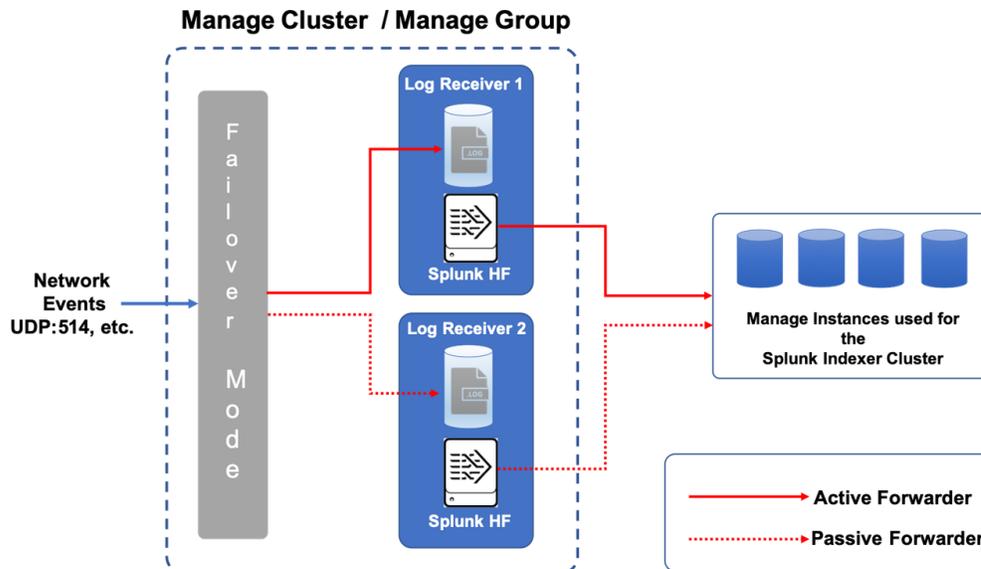
In order to complete the process of making this instance a **Heavy Forwarder**, open the **Settings / Licensing** menu, select the '**Change Licensing Group**' button and choose the '**Forwarder License**' option.

Restart the Splunk instance to commit all of these changes made at the Heavy Forwarder.

Creating a High-Availability Syslog environment

If only one **Log Receiver** instance exists within a network, a single point of failure exists for the collection of Syslog and other network-related events. It is therefore highly recommended to have at least two Log Receiver instances operating in a '**failover mode**'.

As the following diagram suggests, we would recommend creating at least two **Log Receiver** instances working together in a **Gemini Central Group**, sharing the same Log Receiver rules, and using Gemini's built-in **Failover** function to maintain a consistent working presence.



To provide this environment for **Log Receiver**, a **Manage Cluster** and subsequent **Manage Group** will need to be created involving two or more Log Receiver instances. Log Receiver rules must first be manually recreated within any additional members of the **Node Group**.

In order to complete a true **High Availability(HA)** environment for syslog, complete the following;

Step 1: Set up a Log Receiver

Create an initial Log Receiver instance in **Manage** that utilizes **Splunk** as a Heavy Forwarder. Refer to the [Log Receiver](#) section for details.

Step 2: Create a Manage Cluster and Manage Group

Combine two or more Log Receivers together to form a group suitable for High Availability. Refer to [Creating a Manage Cluster](#) section for details.

Step 4: Create Failover Groups between members

Use the Failover feature to create two virtual IP Failover Groups between Log Receivers to provide a proper HA environment. Refer to the [Failover](#) section for details.

Login to the web interface of any additional Log Receiver instances required, navigate to the **LOG / Rule Manager** dashboard and verify that the exact same **Rules** exist on all the instances.

Once verification of the Log Receiver rules has been established, it is important to also verify that the local **Splunk** platform of other Log Receivers has been set up correctly and act in the exact same way as the original Log Receiver. This includes switching to the **Splunk Forwarder Licence** and the setting up of **Indexer Forwarding**.

- Login to the **Splunk** web interface on the Child node, and navigate to the **Settings / Data Inputs / Files & Directories** dashboard.
- Scroll to the bottom of the list and observe the Data Inputs pointing to the **/opt/sbox** directories. If there are none present, the Splunk server needs to be restarted following replication of the syslog Rules.
- Navigate to the **Settings / Server Controls** dashboard, **Restart Splunk** and return to the **Data Inputs** dashboard to confirm that the syslog monitor inputs are present and correct.

Note

Do not be tempted to edit the Data Inputs in Splunk. Any changes should be made at each Log Receiver node, prior to a Splunk restart on each node.

- Navigate to the **Splunk Settings / Licensing** dashboard, and change the License type to, '**Forwarder License**'. Restart Splunk when prompted.
- To complete the creation of a **Splunk Heavy Forwarder** on a Child node, navigate to the **Settings / Forwarding and receiving** dashboard and select the '**Configure Forwarding**' option.
- Using the '**New Forwarding Host**' button, enter the **Indexer** values required for your Indexer Cluster, one by one, in the form of **<indexer_address>:9997**

If possible, test each Log Receiver/Heavy Forwarder instance to ensure that they work correctly on their own, prior to enabling the **Failover** feature to achieve full High Availability.

Load Balancing a syslog feed

By adding a [reciprocal Failover Group](#), the ability to **Load Balance** a syslog feed between two or more servers is granted.

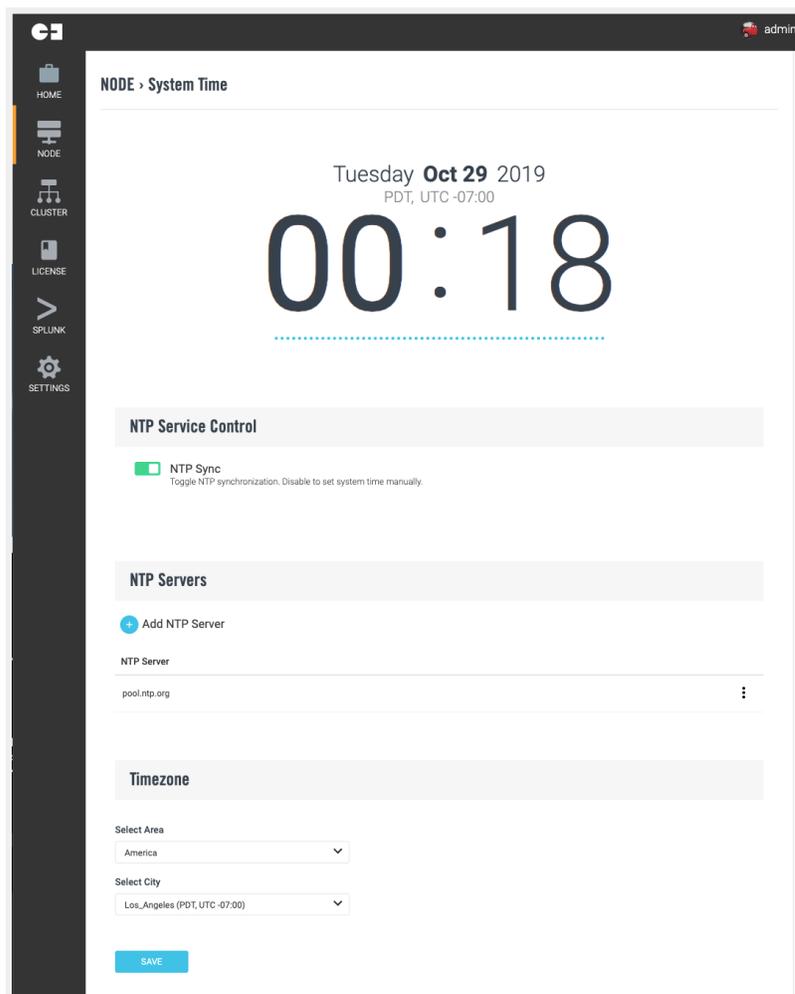
Load Balancing can be facilitated by a third-party load balancer, like **F5**, or manually created using a number of techniques including; the sending of IP addresses with an even-numbered ending octet to one VIP and odd-numbered ending octets to another VIP.

NODE Menu

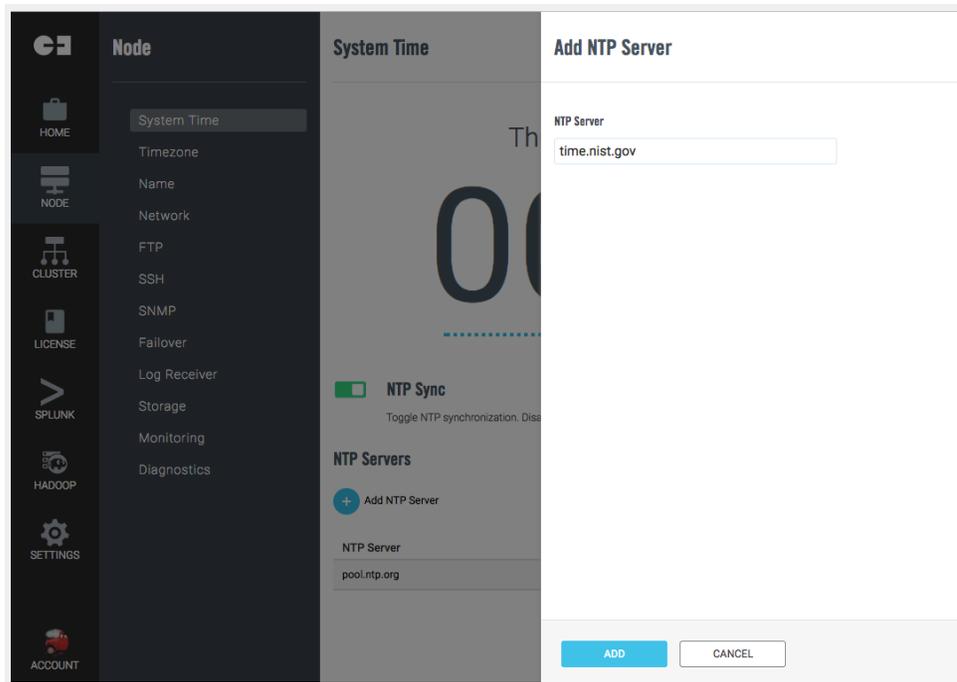
The **NODE** tab is the starting point for the configuration of the host and server functions related to Gemini instances.

System Time

Accurate timekeeping is vital to ensure the correct event order. If distributed Splunk environments become out of sync, then transactional searches may return inaccurate results from inaccurate event timestamping.

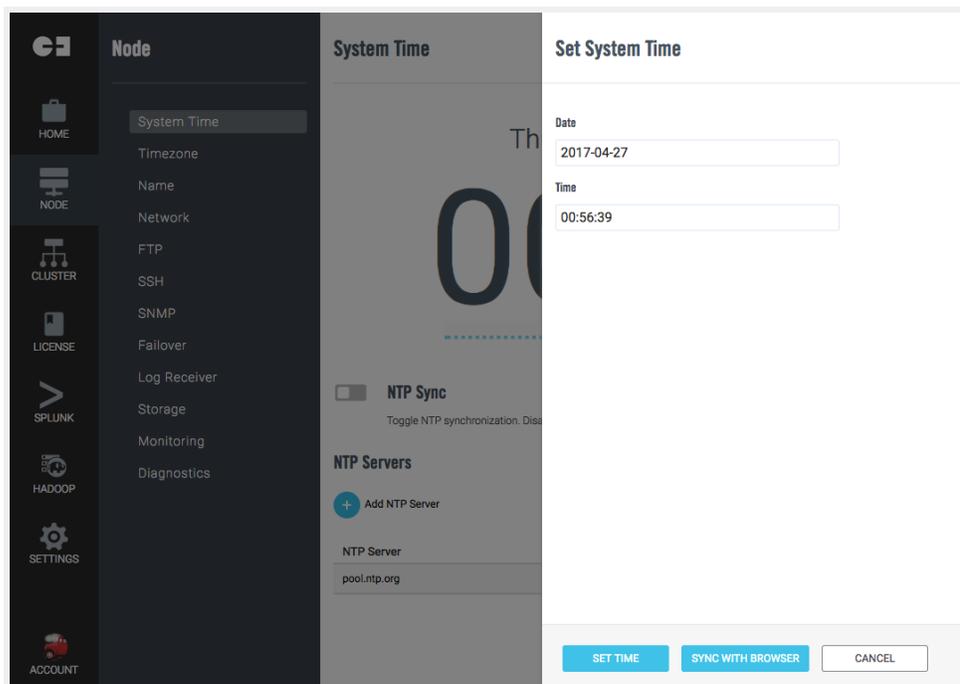


Gemini Central uses `pool.ntp.org` as a default time source. Additional network time sources, either external or internal may be added by selecting 'Add NTP Server'



Setting the **NTP Sync** toggle to the **'OFF'** position will halt further network time updates and allow for manual editing of the system time. This may be required under special circumstances, but is not advisable for general operations.

Select the **'Set Time'** option to correct the DateTime manually, or select **'Sync with Browser'** to update the DateTime settings with the local client PC.

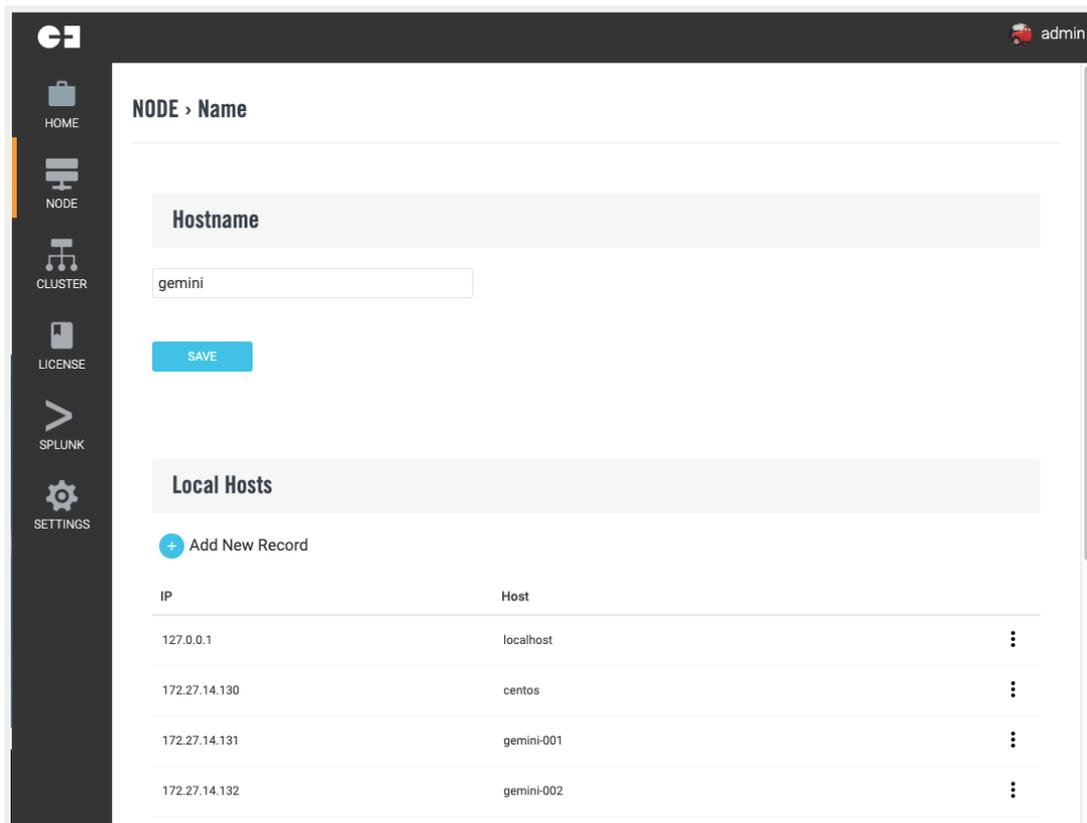


Name

Hostname

To prevent conflicts in distributed Splunk environments as well as declare the source path of received events, Manage requires that each device has a unique hostname.

Splunk will use this unique hostname as a default value to populate both `server.conf` and `inputs.conf` when it is started for the first time.



Local Hosts

While not required in normal operation, manually configuring local hosts can ensure connectivity between hosts in either the absence or failure of a DNS server.

High latency DNS servers or networks may also benefit from this manual configuration.

The manual configuration of hosts is not considered best practice and should **only** be used in exceptional cases as multiple static configurations can be complex to manage.

Note

DNS settings should be configured separately on each network interface using the **Network** tab.

To add a static host, select **'Add New Record'** and specify the new host **IP address** and **name**.

The screenshot shows the 'Add Local Host Record' dialog box in the Gemini Central interface. The dialog has two input fields: 'IP' with the value '10.10.94.87' and 'Host' with the value 'gemini-lab87'. Below the input fields are 'ADD' and 'CANCEL' buttons. In the background, a table lists existing host records:

IP	Host
127.0.0.1	localhost
172.31.32.194	ec2-52-27-118
172.31.42.136	ec2-52-27-204
172.31.42.136	gemini-2

Network

Manage network interface configurations may be reviewed and edited here.

Manage supports multiple network interface cards (NICs) and Gemini appliances each contain four or six NICs depending on the model.

NIC bonding and port redirects may also be configured here.

The screenshot shows the 'Network' configuration page in the Gemini Central interface. The page is divided into several sections:

- Ethernet**: Shows details for network interfaces eth0, eth1, and eth2.

Interface	Configure	DHCP	Link	Interface State	MAC Address	IP	Netmask	Gateway	MTU
eth0	Configure	Connected	Up	08:00:27:0b:58:6c	192.168.56.101	255.255.255.0		1500	
eth1	Configure	Disabled	Connected	Up	08:00:27:bb:99:4d			1500	
eth2	Configure							1500	
- Name Servers**: A section with a message 'No name server configured.' and an 'Add Name Server' button.
- Routes**: A section with a message 'There is no static route existed.' and an 'Add Route' button.

Each NIC may be configured with either a manually assigned IP address or via DHCP.

Advanced configurations like MTU and TX queue length can be configured to improve network performance where appropriate.

Note

Set MTU to a value larger than 1,500 to enable **Jumbo Frame** if the ethernet interface has an **iSCSI** connection. Consult your NAS vendor for more details.

Select the 'Edit Configuration' icon to make any changes, and select the 'Save' button to exit.

The screenshot shows the 'Network' configuration page with a modal for editing 'eth.name'. The modal includes the following fields and options:

- Configure IPv4:** Radio buttons for Disable, DHCP, and Manually.
- IP:** Text input field containing '192.168.56.101'.
- Netmask:** Text input field containing '255.255.255.0'.
- Domain:** Text input field.
- Other Settings:**
 - MTU:** Spin-down menu set to '1500'.
 - txqueuelen:** Spin-down menu set to '1000'.

Buttons for 'SAVE' and 'CANCEL' are located at the bottom of the modal.

Static routes may be added to a specific network interface in order to communicate with networks not directly connected to the Gemini appliance.

The screenshot shows the 'Add Route - eth0' configuration modal. The fields are as follows:

- Network:** Text input field containing '192.168.94.0'.
- Maskbits:** Text input field containing '24'.
- Gateway:** Text input field containing '192.168.94.87'.

Buttons for 'ADD' and 'CANCEL' are located at the bottom of the modal.

NIC Bonding

Gemini Central provides support for 'link aggregation'. It is possible to bind multiple physical NICs into one '**virtual interface**', in order to increase throughput beyond that of a single connection whilst at the same time providing redundancy in the event of a single NIC failure.

Select '**+ Create Virtual Interface**' to create a new NIC arrangement.

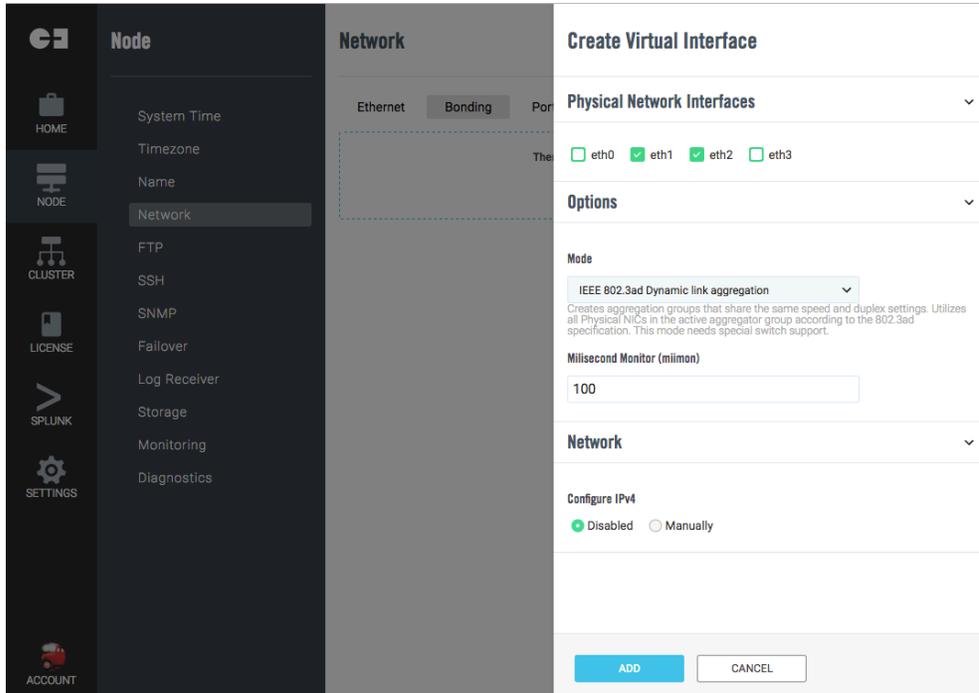
Select from the available **Physical Network Interfaces**, and using the table below as a guide, choose a '**load balancing and fault tolerance**' option from the **Mode** drop-down menu.

Mode	How it works	Fault Tolerance	Load Balancing
Round Robin	Packets are sequentially transmitted/received through each interface one by one.	No	Yes
Active-Backup	One NIC is active while another NIC is asleep. If the active NIC goes down, another NIC becomes active.	Yes	No
XOR	The MAC address of the slave NIC is matched up against the incoming request's MAC and once this connection is established the same NIC is used to transmit/receive with the destination MAC.	Yes	Yes
Broadcast	All transmissions are sent on all slaves.	Yes	No
Dynamic Link Aggregation	Aggregated NICs act as one NIC which results in a higher throughput whilst providing failover in the case of a NIC failure. This requires switch hardware that supports the IEEE 802.3ad protocol	Yes	Yes
Adaptive Transmit Load Balancing	Outgoing traffic is distributed depending on the current load at each NIC. Incoming traffic is received by the current slave. If the receiving slave fails, another slave takes over the MAC address of the failed slave.	Yes	Yes
Adaptive Load Balancing	Unlike Dynamic Link Aggregation, Adaptive Load Balancing does not require any particular switch configuration. The receiving packets are load-balanced through ARP negotiation. Adaptive Load Balancing is only supported in x86 environments.	Yes	Yes

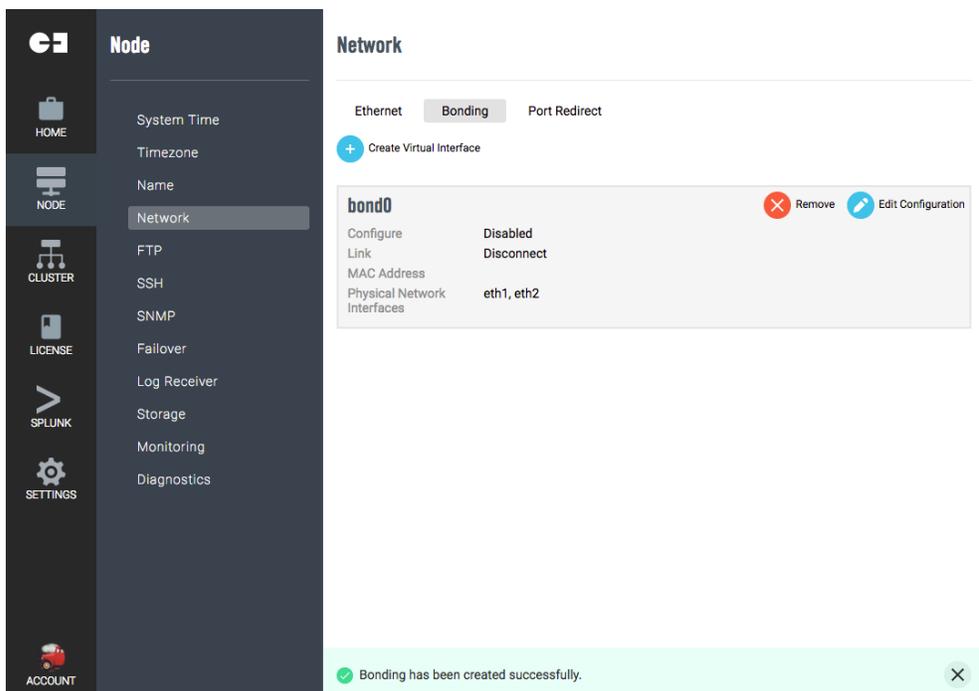
We have a built-in media-independent interface (MII) to confirm and verify the status of the network interface.

Specify the frequency of monitoring by entering a value in the 'Millisecond Monitor' box. The default value is 100ms.

Select the 'Add' button to complete the process.



Once created, the new **Virtual Interface** will be listed in the UI (see the example below). For further configuration or to remove the bonded group, use the appropriate icons.



Port Redirect

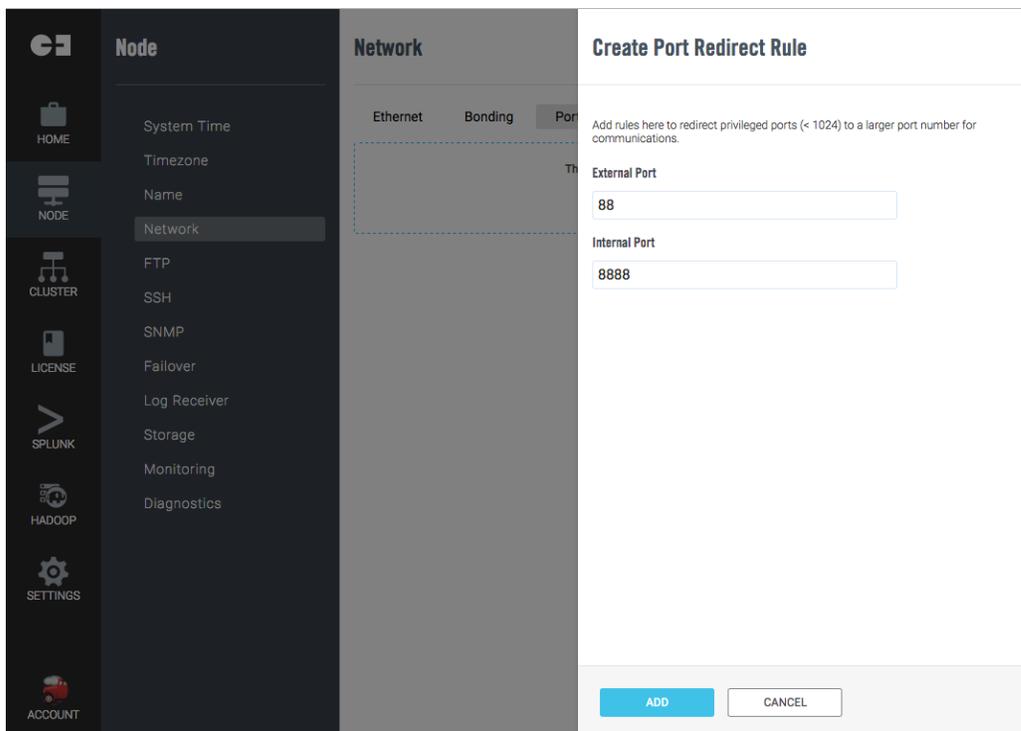
As with any other application running as a **non-root** user on a Linux/Unix platform, Splunk will be unable to bind and listen to any privileged port (< 1024).

Port Redirect allows you to define rules to redirect incoming connections on privileged ports to a port above 1024.

By default Splunk uses port 9997 to receive data from Forwarders to avoid this issue, but if for instance you had a Syslog server that did not have a Splunk Forwarder, this **Port Redirect** feature could help.

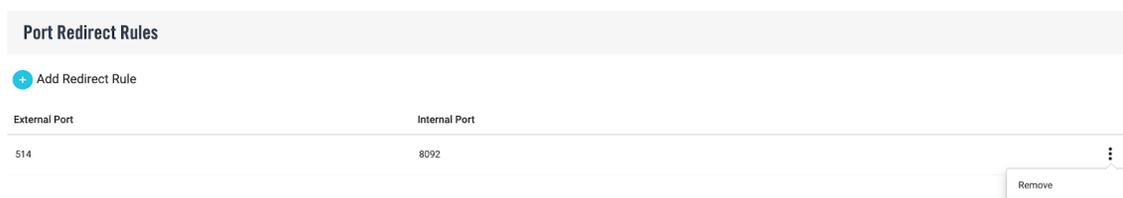
To apply a **Port Redirect**, Select the **'Add redirect Rule'** button, and enter the **source** (External Port) and **Destination** (Internal Port) in the boxes provided.

Select the **'Add'** button to complete the process.



Redirected ports will be listed in the UI (see the example below). Use the **'Add redirect Rule'** button to create other rules.

To remove the rule, locate the vertical ellipsis icon at the end of the row, and choose the **'Remove'** option.



OS Users

To clarify, Gemini Central uses two types of Users; **OS User** accounts and **Manage User** accounts.

- **OS User** accounts are created for secure SSH access to the instance and **do not** give rights to login to the Gemini web interface.
- **Manage Users** are created for access to the Gemini web interface, and are discussed in the **Authentication** section.

Management of **OS User (ssh)** accounts including the addition of SSH public keys and the unlocking of passwords can be achieved from the **OS Users** dashboard.

To unlock a locked **OS User** account, select 'Yes' in the 'Allow Login' section.

The screenshot shows the Gemini Central interface. On the left is a navigation sidebar with icons for HOME, NODE, CLUSTER, EXPLORE, LICENSE, SPLUNK, TABLEAU, and SETTINGS. The main content area is titled 'NODE > OS User' and contains three sections: 'OS User Accounts', 'Replicate OS User Settings', and a 'SUBMIT' button. The 'OS User Accounts' section displays a table with the following data:

OS Users	User Group	Allow Login
sbox	gemini,sbox	Yes
splunk	splunk	Yes
tableau	wheel,tableau,tadmin	Yes

The 'Replicate OS User Settings' section has a 'Select Node Group' dropdown set to 'ALL NODES' and a 'SUBMIT' button. On the right side, the 'Create OS User' form is visible, with the following fields and options:

- User Name:** example
- FullName:** admin
- Password:** [masked]
- Retype Password:** [empty]
- User Group:** users, gemini, sbox, splunk, tableau
- Allow Login:** Yes, No
- SSH public key:** [text area containing a public key]

At the bottom of the form are 'CREATE' and 'CANCEL' buttons.

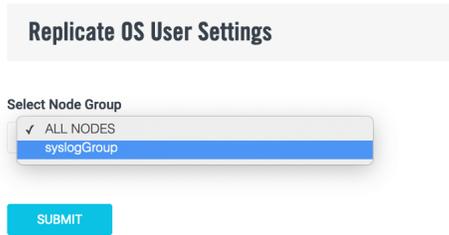
Note For security reasons, 'disallow' any **OS User** accounts that are unused.

Gemini Central has two built-in **OS users** as standard, 'sbox' and 'splunk'.

- Use the **sbox** OS user when dealing with Gemini issues such as instance initialization or recovery.
- Use the **splunk** OS User account for any manual intervention required in the /etc/splunk directory, if this can not be achieved using the [Config Editor](#) feature.
- Activation of **Tableau** on the instance will automatically create another OS user, 'tableau'

In some cases you might need a dedicated **OS User** account to run scripts or applications. Assign this dedicated user to appropriate groups for access permissions to other accounts.

OS Users that have been created at the **Management Center**, can be conveniently exported to other Gemini instances using **Manage Groups** if desired.



In order to use this option, a **Manage Group** would need to exist, then simply select it from the drop-down box and select the '**Submit**' button.

This result of this action can be monitored at the **Cluster / Execute Jobs** dashboard

FTP

Adding data to Splunk is always best achieved with the use of Universal or Heavy Forwarders.

If for some reason, this is not possible, one option could be to enable the **FTP service** allowing data to be written to a file in the **/opt** directory, which can then be monitored in Splunk.

There are two stages required to enable this feature; the first is to configure the FTP service, and the second is to set up a monitored input (inputs.conf) in Splunk.

Note

The FTP protocol is not natively encrypted and should **only** be used when security practices allow.

FTP Service

To enable the FTP service, use the **'FTP Service'** toggle slider and select the desired port on which you want the service to run (defaults to 2121).

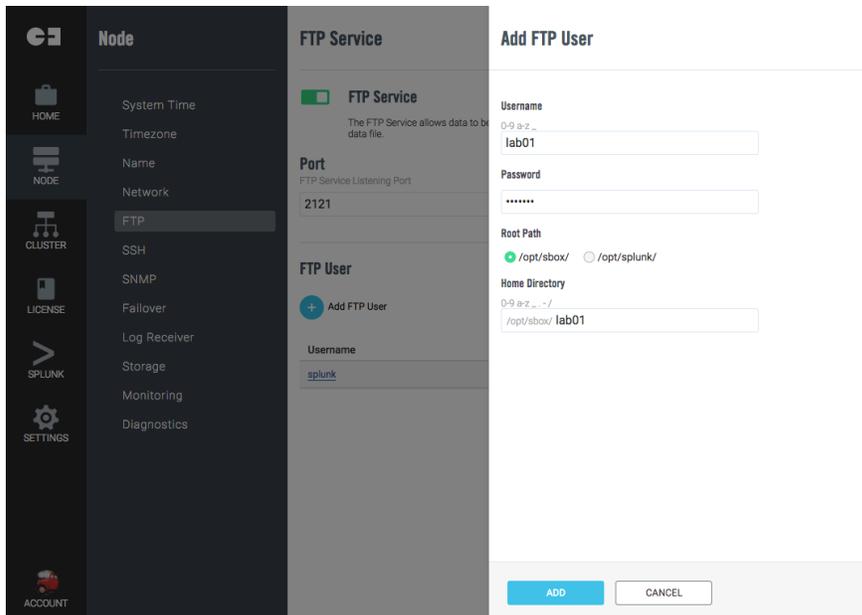
FTP User

The FTP protocol requires both user credentials and a directory to store received files as part of the configuration.

Gemini Central creates a default **FTP User** named **'splunk'** with a home directory of **'/opt/splunk'**.

For additional FTP accounts, select the **'Add FTP User'** button and provide the desired username, password and Home directory folder.

To edit an existing account, including the default Splunk user, simply select the user from the **'Username'** column and modify accordingly.



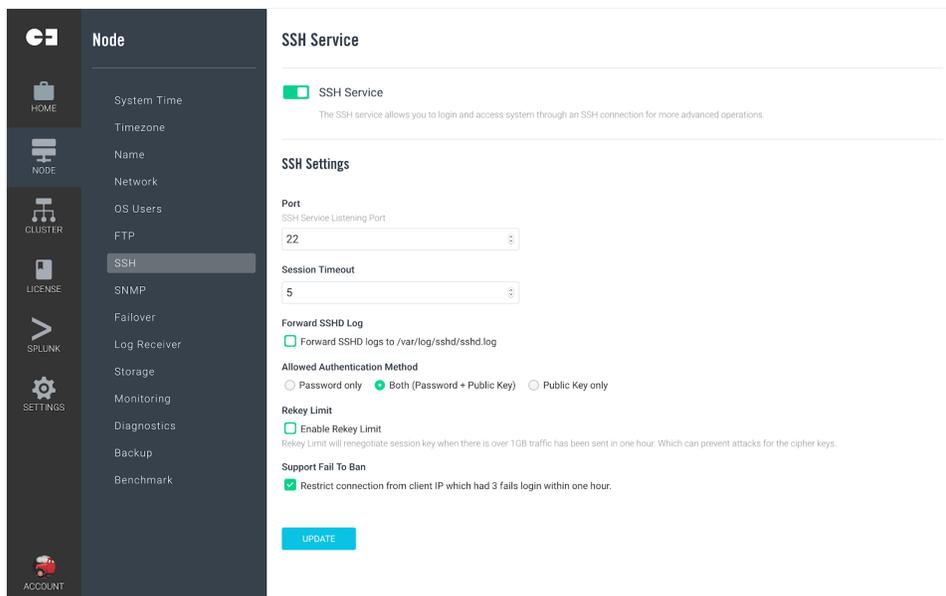
SSH

The **SSH** service (natively encrypted) is enabled by default on each Gemini Central instance.

Refer to the **OS Users** menu for available user accounts and their group access. Reserved user accounts of **'sbox'** and **'splunk'** included by default as detailed below:

- **sbox** : facing jet function drive
- **splunk**: think adventure kitchen chest

Note that both accounts have an enforced password change at initial login.



SSH settings can be modified using the following information;

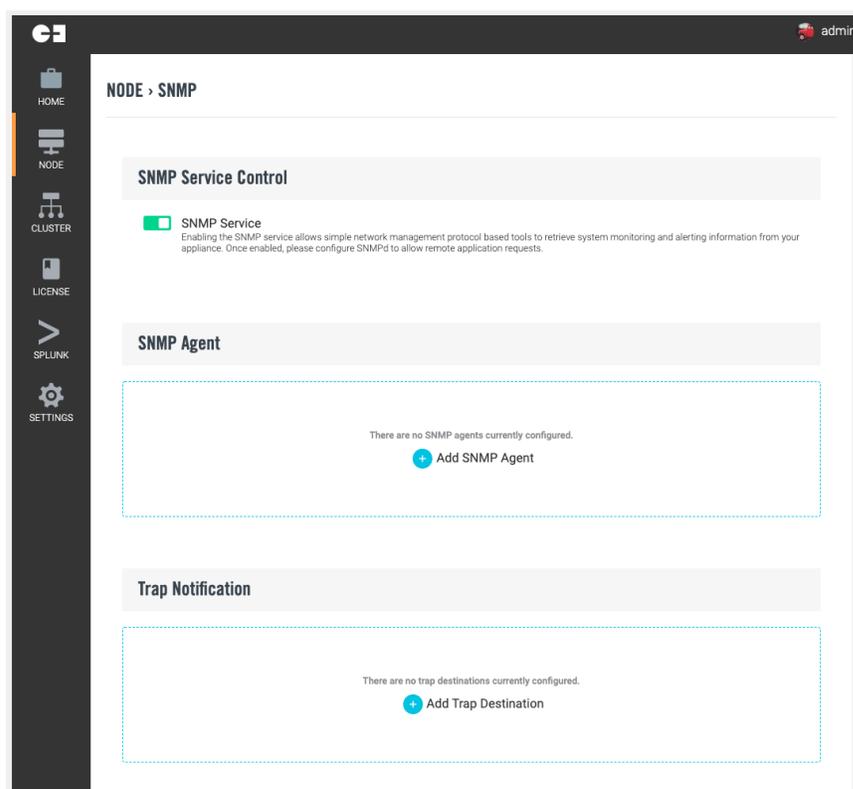
Port:	Listening port of SSH service (default 22)
Session Timeout:	Timeout interval (in minutes)
Forward SSHD Log:	When enabled a copy of SSHD logs will be sent to /var/log/sshhd/sshhd.log for further use.
Allowed Authentication Method:	SSH login with password or authorized private key. Note: SSH keys are only applicable when Manage is running on AWS
Enable Rekey Limit:	If enabled, this will renegotiate a new key after traffic reaches 1GB. This will prevent against the key being cracked and traffic being decrypted by attackers. Note: AWS only
Support the Fail to Ban:	Enable this to restrict a client that has failed to connect 3 times, for a period of one hour. Note: This is enabled by default.

SNMP

SNMP Service

If you require Simple Network Management Protocol (SNMP) data from the **Gemini Central** instance for reasons of monitoring or alerting, an internal **SNMP service** will need to be enabled.

There are two possible modes of operation available to an external **SNMP Management Host**; the **polling** method, or the **trapping** method. To enable this instance to offer either SNMP option, use the toggle slider for the '**SNMP Service**'



SNMP Agent (polling option)

Once configured, this agent will allow polling of the instance by an **SNMP Agent**.

Verify that the **SNMP Service** has been activated. Select the '**Add SNMP Agent**' button to create a new SNMP agent entry. Multiple SNMP Agents can be configured if required.

Select a unique name for each SNMP Agent and choose an appropriate agent version from the options presented; Version 1, Version 2c or Version 3

Note

Only alphanumeric, dot, hyphen, and underscore characters are allowed in the input fields.

SNMP Agent version 1

SNMP Version 1 is not encrypted and authentication will happen in plain text. This version should therefore only be used when other, more secure versions are not possible. SNMP v1 supports a maximum of 32 bits per counter.

SNMP Agent version 2c

SNMP Version 2c is also non-encrypted and authentication occurs in plain text. This version should only be used when other, more secure versions are not possible. SNMP V2c supports a maximum of 64 bits per counter.

For either of these two options complete the **Network** and **Maskbit** (Subnet Mask) entries for the host network, and enter a **Community String** for SNMP authentication.

Note: The default string 'public' should be avoided.

The screenshot displays the 'Create SNMP Agent' configuration window. On the left, a navigation sidebar shows 'Node' selected, with 'SNMP' highlighted under the 'Node' section. The main content area is split into two panes: 'SNMP Service' and 'SNMP Agent'. The 'SNMP Service' pane shows a toggle switch for 'SNMP Service' which is turned on, with a note: 'Enabling the SNMP service allows alerting information from your agent'. The 'SNMP Agent' pane is currently empty. The 'Create SNMP Agent' form on the right contains the following fields:

- Name:** Home Brew
- Agent Version:** Radio buttons for 1, 2c (selected), and 3.
- Network:** 10.10.9.0
- Maskbit:** 255.255.255.0
- Community String:** never_ever_public

At the bottom of the form are two buttons: 'ADD' (highlighted in blue) and 'CANCEL'.

SNMP Agent version 3

SNMP Version 3 supports authentication, encryption and 64 bit counters. This would therefore be the optimum choice if you need SNMP alerting.

Select the most appropriate authorization method under the **Authorization Algorithm** section. Gemini Central supports either **MD5** or **SHA** authentication. Enter the desired authentication password.

Manage supports **DES** or **AES128** encryption methods. Select the desired method and enter the encryption password. Select the ADD button to complete the process.

Note: AES128 is considered to be the more secure of the two options.

SNMP Trap Destinations (trapping option)

Once configured, this option will send SNMP information to an external SNMP Manager host. Provide the address of the SNMP Manager Host and follow up with specific trap thresholds required for this Gemini instance.

Enter the **IP address** of your SNMP Host

Select the protocol you prefer from the following:

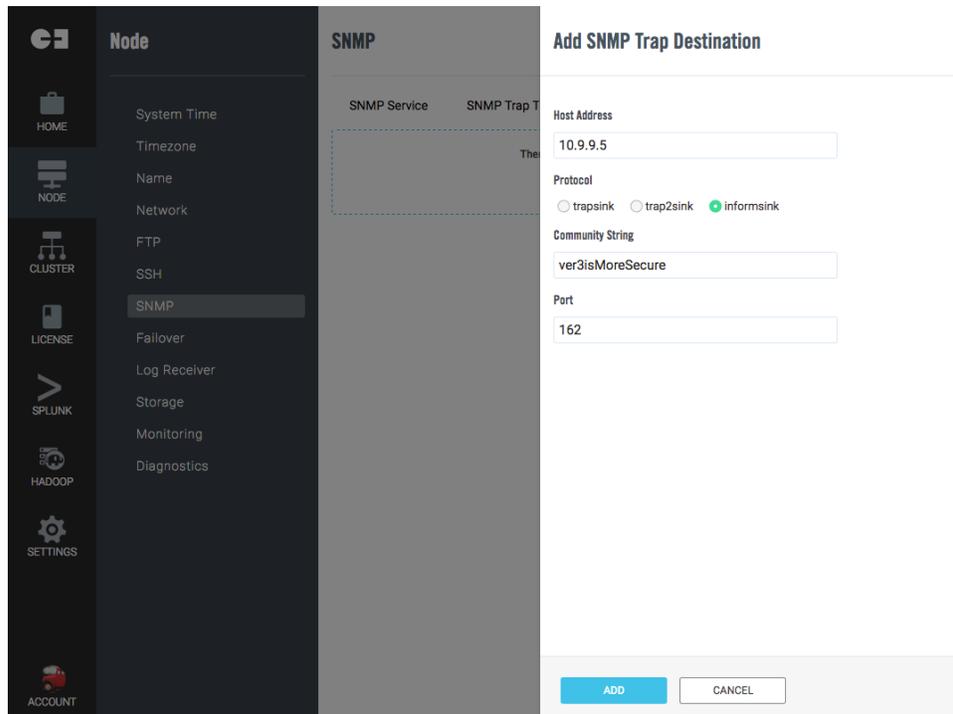
- **trapsink** - send **SNMPv1** traps
- **trap2sink** - send **SNMPv2** traps
- **informsink** - send **'inform'** notifications

Enter the **'Community String'** (see Note below)

Enter the chosen **'Port'** over which to send information.

Note

Only alphanumeric, dot, hyphen, and underscore characters are allowed in the input fields.



SNMP Trap Thresholds

Enable the desired SNMP trap frequency and threshold values required for the instance performance metrics.

SNMP Traps may be enabled for:

- **Processes** - A multi-choice offering including; ftp, splunk, ssh and syslog-ng
- **Disk** usage
- **Network Link**
- **CPU** usage
- **Memory** usage

SNMP Trap Thresholds

PROCESS

Frequency (sec) 600

Process

ftp splunk ssh syslog-ng

DISK

Frequency (sec) 600

Threshold (< %) 5

LINK

Frequency (sec) 600

CPU

Frequency (sec) 600

Threshold (> %) 95

MEMORY

Frequency (sec) 600

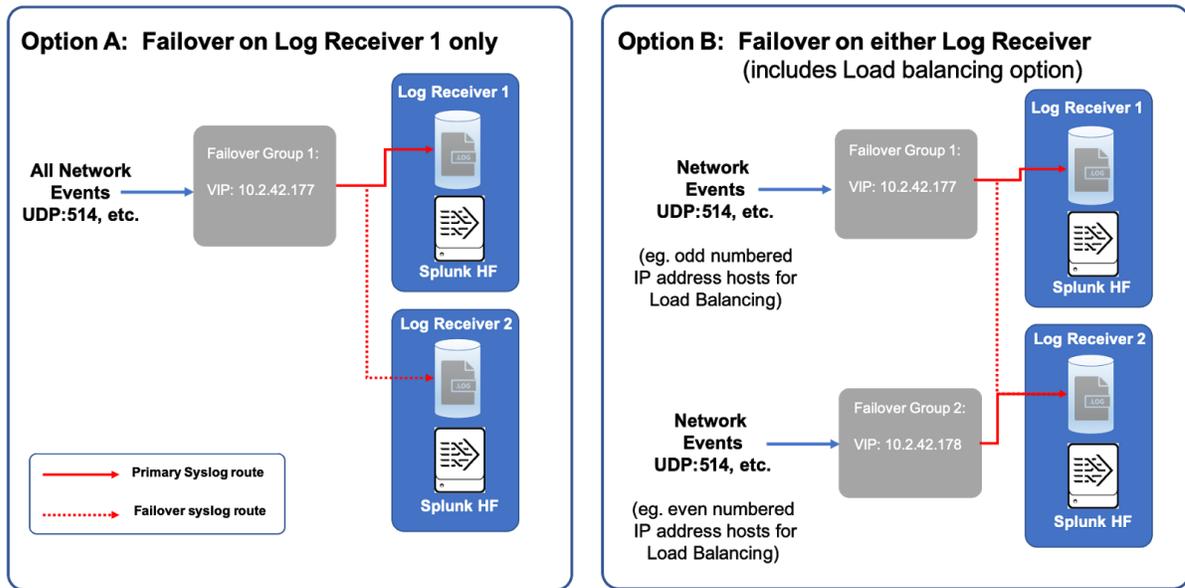
Threshold (< MB) 1

SAVE

Failover

This feature involves one or more **Failover Groups** offering **High Availability**. This is mainly intended for the appliance version of Gemini Central, as Cloud and VM infrastructure usually employ their own technology to cope with network related issues.

Adding failover would be a natural extension to the use of **Log Receiver** in Gemini Central. If you are using the Log Receiver feature, the following diagram offers two options involving the creation of one or two **Failover Groups** as used specifically with two Log Receiver instances. Although not explicit to Log receiver instances, these examples are intended to show how this feature can be used in one of two modes of operation, Option A - only one server involved, or Option B - reciprocal mode



The first example, **Option A**, describes a method that is controlled entirely by one single Log Receiver, in this case allowing the failover of Log Receiver 1 to Log Receiver 2, should Log Receiver 1 fail.

In the second example, **Option B**, two failover groups have been created, each controlled by one of the Log Receivers. In this scenario, either device could fail and the other one will take over allowing a reciprocal failover for each device. This option has the advantage, should it be required, of offering an additional **Load Balancing** feature as both **VIP addresses** are available. For example, syslog hosts ending with an odd IP address could be directed to one VIP address, and hosts ending with an even IP address to the other. This is just an idea, it is up to you whether or not you use this feature..

Each **Failover Group** has one 'active' master node using a **virtual IP address**, and one or more standby slave nodes that are ready to take over for a failed master. Each Gemini appliance can be part of a different Failover Group and each group should be provisioned using a different port number.

Creating a Failover Group

Before creating a new failover group, you will need a VIP address that you can use for each Failover Group. This static IP address is usually provided by your Network Administrator. Please ensure you have the necessary VIP addresses before you begin the following process.

Option A - The process to set up a single Failover Group

Login to the Manage web interface of your primary instance, in our example this is 'Log Receiver 1', and navigate to the **Node / Failover** dashboard (shown below).

NODE > Failover

Failover Groups

There are no failover groups currently configured.

[+ Create New Failover Group](#) [Join Existing Group](#)

Create New Failover Group

Virtual NIC - IP Address

Monitor

Detect Splunk

Add node to this group

Remote Node

 [+](#)

[ADD](#) [CANCEL](#)

Select the **'+ Create New Failover Group'** button to reveal the setup screen opposite;

Generally, the Network Administrator would allocate a static IP address for the VIP address, which should be entered in the **'Virtual NIC - IP Address'** box.

The Monitor drop-down box reveals a **'Detect Splunk'** value that should always be selected. This invokes a **'keepalived'** daemon to monitor Splunk on this device.

The **Remote Node** box requires manual entry of the device(s) that you wish to 'failover' too. The **'+'** button is only required if more than one device is involved

Select the **'Add'** button to save the changes to reveal the following dashboard.

NODE > Failover

Failover Groups

[+ Create New Failover Group](#) [Join Existing Group](#)

Virtual IP	Local	Members
10.2.42.177	10.2.42.155	1

Note

There is currently only **one member** of the Failover Group, this is of course incomplete. In order to complete the process, another device, in our case Log Receiver 2 will need to join the Failover Group.

Joining a Failover Group

To complete the Failover Group, login to the Manage web interface of another instance, in our example, **Log Receiver 2**, and navigate to its **Node / Failover** dashboard.

Select the **'Join Existing Group'** button and enter the **VIP address** used for this Failover Group. Select the **'Join'** button at the bottom of the dashboard to make the connection and complete the 'Group'.

Join Existing Group

Virtual NIC - IP Address

SCAN

Virtual IP

10.2.42.177	+
-------------	---

Alternatively, the **'Scan'** button can be used to invoke a search for the Virtual IP address and populate the entry box as shown below.

Note that if the scan fails to detect the IP address, add the VIP manually as directed above, and select the **'Join'** button

A dashboard similar to that below should follow, suggesting that there are now two members in this Failover Group

Failover Groups		
+ Create New Failover Group Join Existing Group		
Virtual IP	Local	Members
10.2.42.177	☁ 10.2.42.156	2

For extra detail regarding the Failover Group, select the **Virtual IP** address listed to display the following dashboard;

NODE > Failover

Failover Groups

[+ Create New Failover Group](#)
[Join Existing Group](#)

Virtual IP	Local
10.2.42.177	☁ 10.2.42.156

Failover Group - 10.2.42.177

Group Properties

Virtual IP: 10.2.42.177

Monitor: Detect Splunk

Local

Role: Master

IP: 10.2.42.155

Members

Role	IP
☁ Slave	10.2.42.156
📍 Master	10.2.42.155

[LEAVE GROUP](#)
[REMOVE GROUP](#)
[CLOSE](#)

Note

Notice that out of the two members, one is a **Master** - the current device in use, and the other a **Slave**, the failover option for this server.

We have now completed the failover scenario as depicted in **Option A** of the diagram shown at the start of this section.

Option B - Creating a reciprocal failover Group with optional Load Balancing

To recreate **Option B** - adding a second **Failover Group** managed by **Log Receiver 2** with potential **Load Balancing** (should this be required) - the following additional tasks would be required given using our example scenario;

Login to the Manage web interface of **Log Receiver 2**, and navigate to the **Node / Failover** dashboard.

Select the **'+ Create New Failover Group'** button to enter the details of a second **VIP Address** and select the **'Add'** button to save the changes.

Log back into the Manage web interface of the **Log Receiver 1** instance, and navigate to the **Node / Failover** dashboard to **'Join Existing Group'**. Select the **'Scan'** button to bring back the **VIP Address** entry, and select the **'Join'** button at the bottom of the dashboard to make the connection.

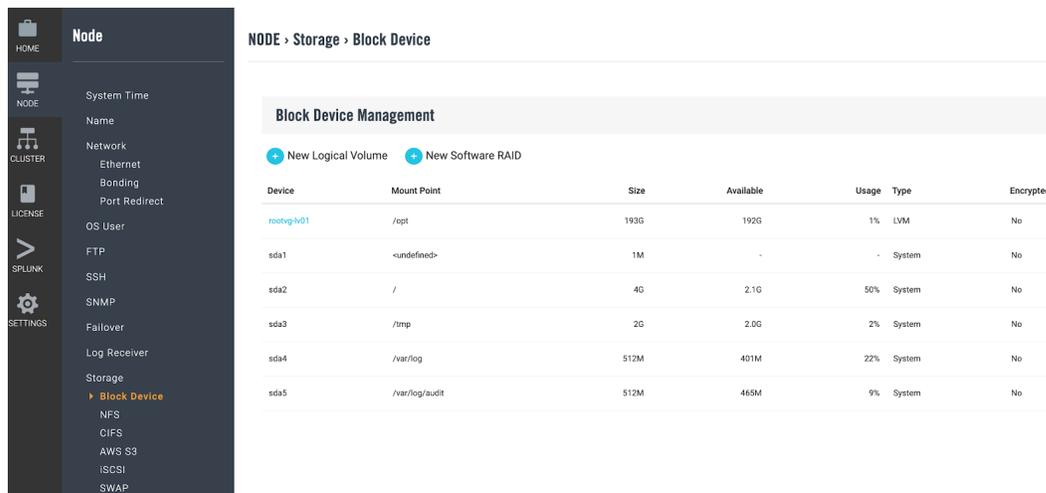
A dashboard similar to that below should follow, suggesting that there are two members in each Failover Group.

Failover Groups		
+ Create New Failover Group Join Existing Group		
Virtual IP	Local	Members
10.2.42.177	 10.2.42.155	2
10.2.42.178	 10.2.42.155	2

Storage

The Storage section allows administrators to manage both local and attached storage. This includes direct and network-attached storage, used to extend the disk capacity for data applications such as Splunk and Tableau.

This will allow the volume of an existing system to be extended and the mount point for Splunk indexes may also be defined. This feature will also allow the capability to read files from network storage.



Block Device Management

[New Logical Volume](#)
[New Software RAID](#)

Device	Mount Point	Size	Available	Usage	Type	Encrypted
rootvg-h01	/opt	193G	192G	1%	LVM	No
sda1	<undefined>	1M	-	-	System	No
sda2	/	4G	2.1G	50%	System	No
sda3	/tmp	2G	2.0G	2%	System	No
sda4	/var/log	512M	401M	22%	System	No
sda5	/var/log/audit	512M	465M	9%	System	No

Storage Devices

All detected attached storage is listed here giving the following opportunities:

- Create a RAID disk from multiple storage devices
- Create a new logical volume for grouping storage devices as one
- Merge storage devices with the existing logical volume to extend disk capacity, or mount it to a designated mount point.

Plan your storage use by considering future data growth and potential expansion. Some actions are not reversible, so good planning is essential before taking action.

Mount disk and mount points

New storage devices can be added as a user custom mount point under `/opt/mnt/`. The owner of this mount point is `sbox` and permission is open to all. You may choose to maintain owners and permissions of files and folders under this mount point yourself.

If this storage device is entirely for use with Splunk, you may choose to mount it to `/opt/splunk` directly, as shown below.

The screenshot shows the Gemini Central interface with a 'Mount to File System - md127' dialog box open. The background shows a table of storage devices with columns for Device, Mount Point, Size, and Availability. The dialog box has three main sections: 'Mount Point', 'Encryption', and 'Key'. In the 'Mount Point' section, two input fields are visible: the first contains '/opt/splunk' and the second contains '/opt/mnt/ abc' with a green checkmark. In the 'Encryption' section, the 'No' radio button is selected. In the 'Key' section, the 'Use Existing Key File' radio button is selected. At the bottom of the dialog are 'MOUNT' and 'CANCEL' buttons.

Device	Mount Point	Size	Avail
NewVG-NewLV	<undefined>	3G	
md127	<undefined>	4G	
rootvg-lv01	/opt	9G	8
sda1	<undefined>	1M	
sda2	/	4G	2
sda3	/tmp	2G	2
sda4	/var/log	512M	47
sda5	/var/log/audit	512M	48
sdb	<undefined>	1G	
sdc	<undefined>	1G	
sdd	<undefined>	1G	
sde	<undefined>	1G	

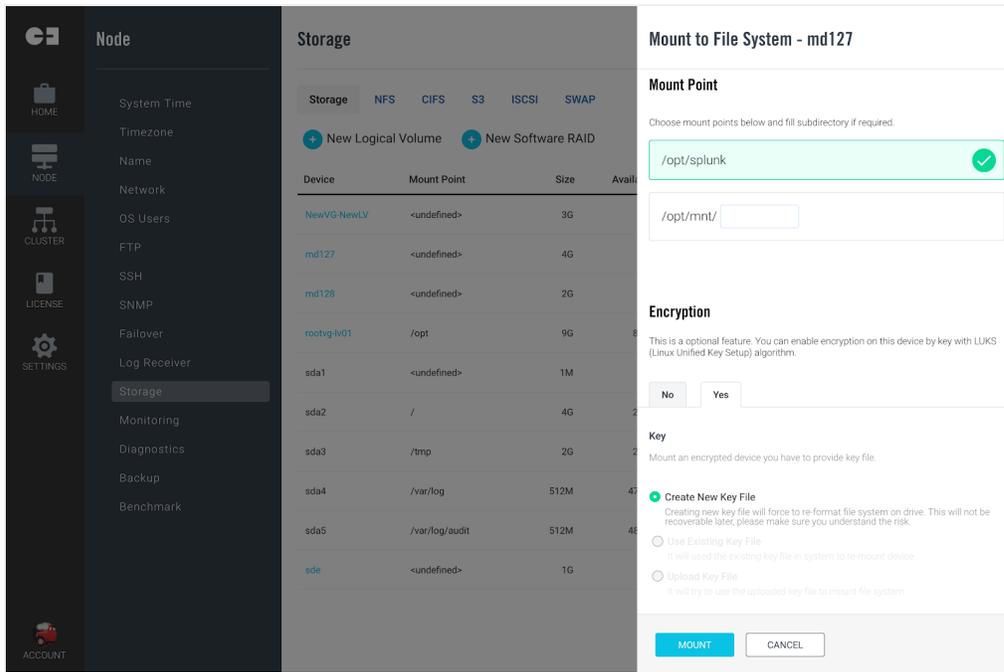
Note

The custom path `/opt/sbox/mount` is deprecated and has been removed from selections. Existing mounts will continue without any impacts until unmounted.

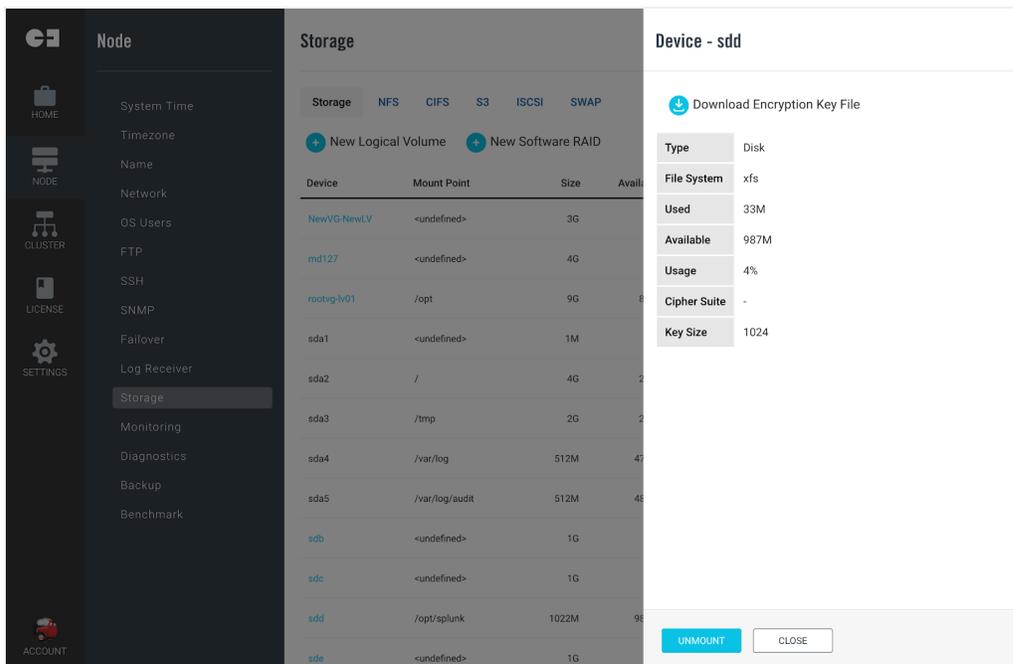
Encryption and Decryption

Gemini Central supports disk encryption which has been simplified and implemented as an option while mounting disks. This is optional and disabled to all disks by default. You may encrypt a disk and mount it with a new key, or mount it with an existing key.

- **Create New Key File:** This will encrypt the disk with a new key file. All the data on this disk will be erased.
- **Use Existing Key File:** If the disk was encrypted from this machine, this allows the disk to be mounted again with the existing key.
- **Upload Key File:** If the disk was encrypted somewhere else, this allows the disk to be mounted again with a provided key file.



Once mounted, you should create a backup of the encryption key file. This is highly recommended.



Notes

- An encrypted disk can not be used for creating a RAID disk or merging into a logical volume. Decrypt before any new allocation.
- Encrypting a logical volume is not supported.
- Encryption with a new key followed by decryption will erase all data.
- A backup of the key file is highly recommended.

Create Software RAID Disk

With RAID, you can group more than one storage device in a disk array to create redundancy or efficiency, depending on the RAID level chosen. Refer to this guide for more understanding about RAID and RAID levels: <https://en.wikipedia.org/wiki/RAID>.

You should select the most appropriate RAID level for your use cases:

- **RAID 0(Striping):** When disk redundancy doesn't matter, and cares about disk performance.
- **RAID 1(Mirroring):** When there are only 2 disks, and data integrity / availability are important.
- **RAID 5:** When there are more than 3 disks, cares about data integrity and availability as well as the performance. This is balanced in performance, capacity, and availability.

The screenshot shows the Gemini Central interface. On the left is a navigation menu with options like HOME, NODE, CLUSTER, LICENSE, SETTINGS, and ACCOUNT. The main area is divided into 'Node' and 'Storage' sections. The 'Storage' section contains a table of storage devices. Overlaid on this is a 'Create New Software RAID' dialog box. The dialog has three radio buttons for RAID Level: RAID 0, RAID 1, and RAID 5. Below this is a table titled 'Choose Physical Disk Device' with columns for 'Device' and 'Size'. The table lists several devices: /dev/sdb (1G), /dev/sdc (1G), /dev/sdd (1G), /dev/sde (1G), and /dev/sdi (3G). Each device has a checkbox next to it. At the bottom of the dialog are 'ADD' and 'CLOSE' buttons.

Device	Mount Point	Size	Availab
md127	/opt/splunk	4G	4.0
rootvg-lv01	/opt	9G	8.5
sda1	<undefined>	1M	
sda2	/	4G	2.2
sda3	/tmp	2G	2.0
sda4	/var/log	512M	477
sda5	/var/log/audit	512M	482
sdb	<undefined>	1G	
sdc	<undefined>	1G	
sdd	<undefined>	1G	
sde	<undefined>	1G	

Notes

- This is specifically to benefit instances that do not have a hardware RAID controller, e.g. VMware, Hyper-V, and AWS. Disk drives on Gemini Appliance are already supported and managed by a RAID controller.
- Merging a RAID disk into a logical volume is not supported.
- The size of each storage device can be different when selecting RAID 5, but this might create wasted disk space.
- Mixing various storage types, e.g. SSD, HDD and iSCSI disks together in one RAID array is not recommended. It will slow down RAID disk performance and increase latency.

Create a Logical Volume

The main advantage of a **logical volume** is the ability to extend disk space when required. More storage devices may be added into an existing logical volume at any time, to extend overall disk capacity.

The screenshot shows the 'Storage' section of the Gemini Central interface. A 'Create New Logical Volume' dialog is open, allowing the user to create a new logical volume. The dialog includes fields for 'Volume Group' (NewVG) and 'Logical Volume' (NewLV). Below these fields, there is a section titled 'Choose Physical Disk Device' with a table of available devices and their sizes. The 'ADD' button is highlighted in blue.

Device	Mount Point	Size	Available
md127	/opt/splunk	4G	4.0G
rootvg-lv01	/opt	9G	8.5G
sda1	<undefined>	1M	
sda2	/	4G	2.2G
sda3	/tmp	2G	2.0G
sda4	/var/log	512M	477M
sda5	/var/log/audit	512M	483M
sdb	<undefined>	1G	
sdc	<undefined>	1G	
sdd	<undefined>	1G	
sde	<undefined>	1G	

Device	Size
<input type="checkbox"/> /dev/sdb	1G
<input type="checkbox"/> /dev/sdc	1G
<input checked="" type="checkbox"/> /dev/sdd	1G
<input checked="" type="checkbox"/> /dev/sde	1G
<input checked="" type="checkbox"/> /dev/sdi	3G

Notes

- A logical volume can be created with one or more storage devices.
- The size of each storage device can be different.
- There is not a way to split storage devices from an existing logical volume, except by the entire removal of the logical volume. Plan storage devices carefully.
- The default logical volume rootvg-lv01 can not be removed.
- Mixing various storage types, e.g. SSD, HDD and iSCSI disks together in one RAID array is not recommended. It will slow down RAID disk performance and increase latency.

Merge Disk

Merge storage devices into a logical volume - You may select a target logical volume if more than one logical volume exists.

The screenshot shows the Gemini Central interface with a sidebar on the left containing navigation options like HOME, NODE, CLUSTER, LICENSE, SETTINGS, and ACCOUNT. The main content area is titled 'Node' and displays a table of storage devices. A dialog box titled 'Merge into logical volume - /dev/sdb' is open, asking for confirmation to merge the device /dev/sdb into LVM. The dialog includes a warning: 'Note that the device cannot be removed once complete otherwise it would crash the system.' Below the warning is a dropdown menu labeled 'Logical Volume' with 'NewVG' selected. At the bottom of the dialog are 'MERGE' and 'CANCEL' buttons.

Device	Mount Point	Size	Avail
NewVG-NewLV	<undefined>	3G	
md127	/opt/splunk	4G	4
rootvg-lv01	/opt	9G	8
sda1	<undefined>	1M	
sda2	/	4G	2
sda3	/tmp	2G	2
sda4	/var/log	512M	47
sda5	/var/log/audit	512M	48
sdb	<undefined>	1G	
sdc	<undefined>	1G	
sdd	<undefined>	1G	
sde	<undefined>	1G	

Notes

- Once a storage device has been merged into the default logical volume “rootvg-lv01”, this action will not be able to be reverted.
- When a device has been merged into a logical volume, it must keep attached unless the partition might be corrupted and data will be lost.
- Merging a RAID disk into a logical volume is not supported.
- Merging an encrypted disk into a logical volume is not supported.

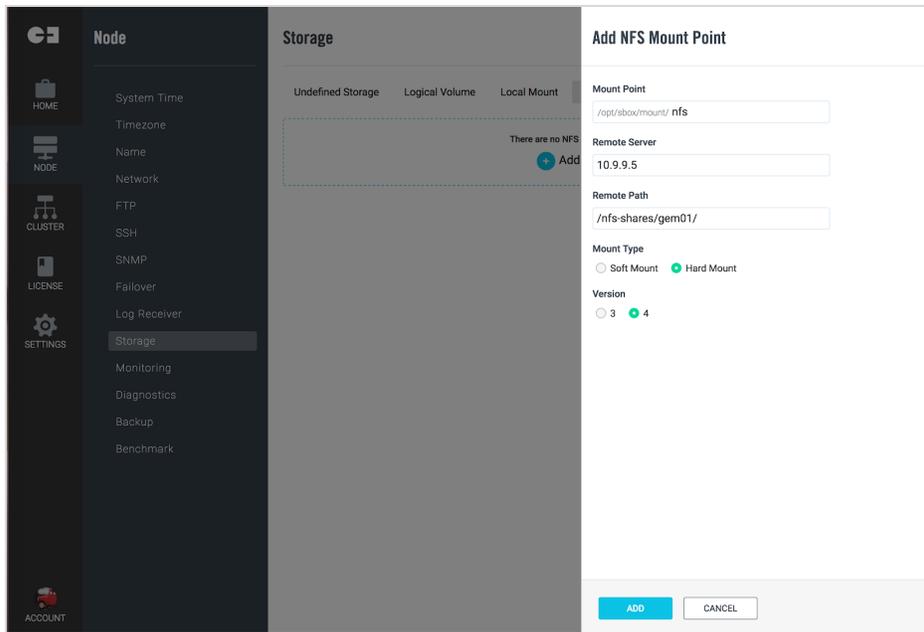
Add an NFS Mount

To define an NFS Mount Point:

- Enter the **local mount point** (located at the `/opt/sbox/data` folder),
- Enter the **IP address** of the remote server
- Enter the **remote folder** (must start with a leading `'/`).
- Select the mount type. A **'Hard mount'** is recommended by Splunk when the mount point is used for cold buckets.
- Select the NFS version. This must match the NFS server version.
- Select the **'Add'** button to add the new NFS mount.

Notes

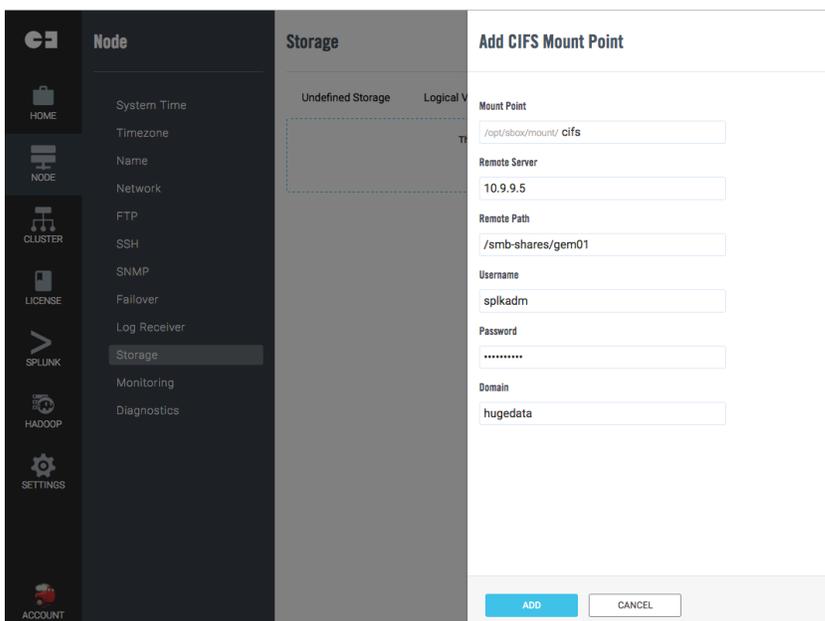
A mount point will not be detected and validated until you enable the configuration by clicking the 'mount' button. Once enabled, Gemini Central will automatically mount the NFS Mount Point upon boot.



Add a CIFS Mount

- To define a CIFS Mount Point:
- Enter the **local mount point** (located at **/opt/sbox/data**)
- Enter the **IP address** of the remote server
- Enter the **remote folder** (must start with a leading ' / ')
- Enter the **Username**
- Enter the **Password**
- Select the **'Add'** button to add the new CIFS mount.

Please note that a mount point will not be detected and validated until you enable the configuration. When enabled, Gemini Central will automatically mount the CIFS Mount Point upon boot.



Add an S3 Mount

To define an **Amazon S3** Mount Point:

- Enter the **S3 bucket name** you want to mount and the local mount point will locate at **/opt/sbox/data/s3/<bucket name>** folder.
- Enter the **IAM Access Key ID**
- Enter the **IAM Secret** Access Key
- If you want all the data stored in the S3 bucket to be encrypted, enable **Server-Side Encryption(SSE)**, selecting a proper key option.

To obtain your S3 Access credentials, log in to your **AWS Console**, open the **Users** section in the **IAM Service** area and select the desired user.

Create an Access key in the Security credentials tab.

Please note that access to S3 storage requires a connection to the public internet from the node.

Note

S3 is designed for data archival and not applicable to Splunk indexing. Specifying hot/warm/cold buckets to an S3 mount mounts will cause Splunk to malfunction.

The screenshot displays the 'Add S3 Mount Point' configuration window in the Splunk interface. The interface is dark-themed with a sidebar on the left containing navigation icons for HOME, NODE, CLUSTER, LICENSE, SPLUNK, HADOOP, SETTINGS, and ACCOUNT. The main panel is titled 'Storage' and shows 'Undefined Storage' and 'Logical V'. The configuration form on the right includes the following fields and options:

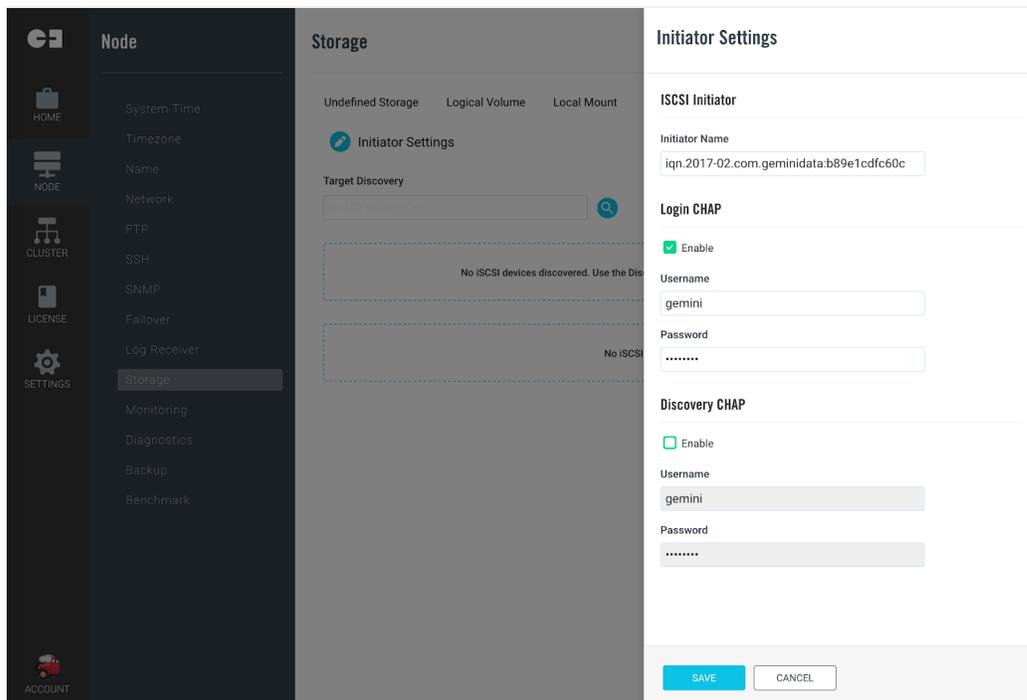
- Bucket:** icebox
- Mount Point:** /opt/sbox/mount/s3/icebox
- IAM Access Key ID:** ENFEFMKFNLEWAWDUJFDW
- IAM Secret Access Key:** iMY3/6WryCJxYUJdXThJaqxBTBECgYEA5x6/I
- Server-Side Encryption:** Enable
- Key Option:**
 - Amazon S3-Managed Keys (SSE-S3)
 - AWS KMS-Managed Keys (SSE-KMS)
 - Customer-Provided Keys (SSE-C)

At the bottom of the form, there are two buttons: 'ADD' (highlighted in blue) and 'CANCEL'.

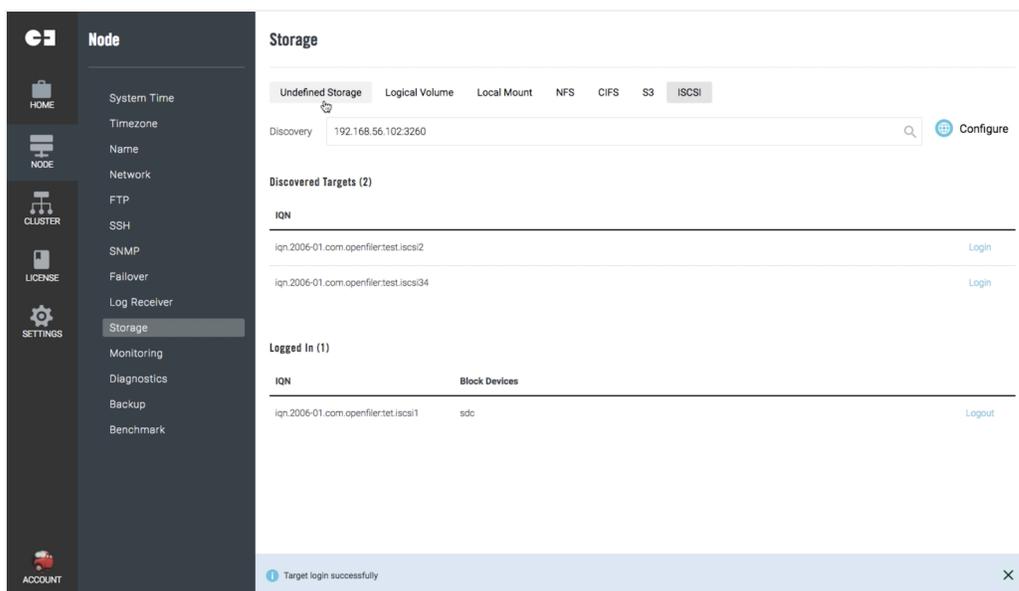
Add an iSCSI Target

To add an iSCSI target:

- Modify the **Initiator Settings** and specify the **Login CHAP** and **Discovery CHAP** details. These must match the settings exactly on the **iSCSI** target.



- In the **Target Discovery** field, input the iSCSI target **IP address** and **port**, eg. **192.168.1.100:3260**. Note the default discovery port is 3260/tcp.
- Once the iSCSI targets have been found, they will be listed.
- Select **“Login”** to connect to the iSCSI target.
- Once connected, there is a new block device detected and listed within the **‘Undefined Storage’** area (see below)



- Select the **‘Undefined Storage’** tab and mount it from here.

Ask your **NAS Administrator** to obtain the **iSCSI target** information and CHAP credentials.

Please note that connected iSCSI target only means there are new block devices available. Do not forget to **mount** them in ‘Undefined Storage’.

Note

Set MTU to a value larger than 1,500 to enable Jumbo Frames across the ethernet interface used for iSCSI connection. This will improve iSCSI performance. Consult your NAS vendor for more details.

Manage Swap space

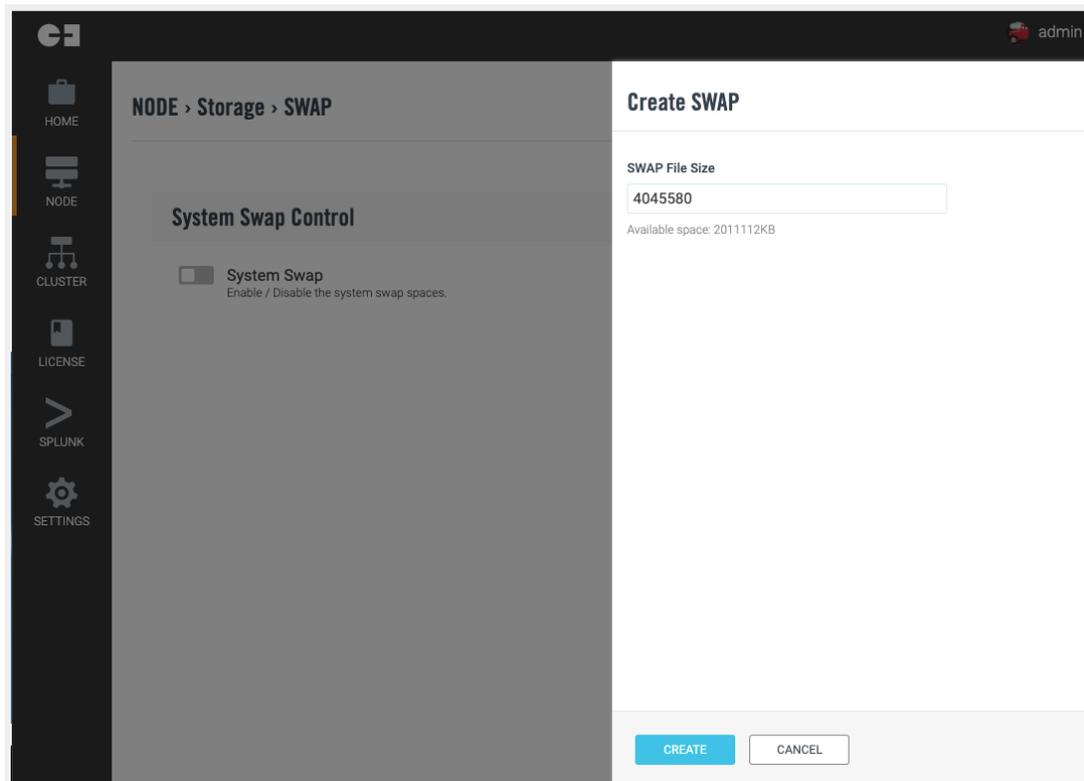
Swap space is disabled by default on Gemini Central as this is optimal for performance but we stress *only* when there is enough physical RAM on the instance.

This setting generally has the most effect on an Appliance based instance.

Reasons for enabling the swap space could be;

- Useful for appliances with a particularly heavy load
- To prevent 'Out of Memory' (OOM) errors from occurring
- For Appliances that have a limited amount of RAM.
- Where the emphasis is more on reliability than performance.

Enable using the slider button, to allow for a specific SWAP file size to be entered. A value that is equal to or greater than the amount of physical RAM installed, is recommended.



Log Forwarding

The local Linux **system** logs together with the **Gemini Central logs** are both stored in a safe place with restricted access. In order to access these more readily, add a **Log Forwarding** rule using the Log Forwarding dashboard. The destination for these logs could be local storage (this instance) or perhaps a Syslog server.

Rules can be set for either the Gemini Central logs or the local System logs. Use the appropriate **'+ Add Forwarding Rule'** button to enable the process.

To send the log file to a physical file on this instance;

- Enter a filename (eg. `system_file_<instance>.log`) to create a local readable log file which can be read by any OS user.
- Enable Log Rotation to prevent the file from growing too large and consuming all the disk space.

The screenshot displays the Gemini Central interface for configuring log forwarding. On the left, a sidebar contains navigation icons for HOME, NODE, CLUSTER, LICENSE, SPLUNK, and SETTINGS. The main content area is titled 'NODE > Log Forwarding' and is divided into two sections: 'Forward Manage Logs' and 'Forward System Logs'. Both sections currently show 'There are no admin log forwarding rules' and 'There are no system log forwarding rules' respectively, with a '+ Add Forwarding Rule' button in each. A modal window titled 'Add System Log Forwarding Rule' is open on the right, containing the following configuration options:

- Name:** SystemLogs
- Destination Protocol:** File (selected), UDP, TCP
- Destination File & Path:** /opt/sbox/data/system.log
- Log Rotate:** Rotate the log file
- Frequency:** Daily, Weekly, Monthly, Yearly
- Number of Copies:** 4
- Size:** 10
- Unit:** KBytes, MBytes, GBytes
- Compress:** Compress the rotated files

At the bottom of the modal are 'ADD' and 'CANCEL' buttons.

To send the logs to a syslog server

- Select the Destination Protocol, UDP or TCP, and enter the IP address of your syslog server. The service defaults to port 514, which can be customized.

The screenshot displays the Gemini Central administration interface. On the left is a dark sidebar with navigation icons for HOME, NODE, CLUSTER, LICENSE, SPLUNK, and SETTINGS. The main content area is titled 'NODE > Log Forwarding' and is divided into two sections: 'Forward Manage Logs' and 'Forward System Logs'. Both sections show a message 'There are no admin log forwardings' and 'There are no system log forwardings' respectively, with an 'Add Forwarding Rule' button. A modal dialog box titled 'Add Admin Log Forwarding Rule' is open on the right. It contains the following fields and options:

- Name:** A text input field containing 'GEM_Logs'.
- Destination Protocol:** Three radio buttons labeled 'File', 'UDP', and 'TCP'. 'UDP' is selected.
- Destination Host:** A text input field containing '10.9.9.5'.
- Port:** A text input field containing '514'.
- Transmission Encryption:** A checked checkbox labeled 'TLS'.
- TLS Key:** A text input field.
- TLS Certification:** A text input field.

At the bottom of the dialog are two buttons: 'ADD' (highlighted in blue) and 'CANCEL'.

Diagnostics

The **Diagnostics Panel** provides useful access to **network tools** without the need to access the command-line (CLI) interface.

The following commands can be executed with the resulting outputs shown below;

- **PING**

The screenshot shows the Gemini Central interface with the **Diagnostics** panel selected in the left sidebar. The **Ping** tool is active, displaying the following output:

```

Start to send ICMP ping packets...
PING 8.8.8.8 (8.8.8.8) 56(84) bytes of data:
64 bytes from 8.8.8.8: icmp_seq=1 ttl=48 time=7.59 ms
64 bytes from 8.8.8.8: icmp_seq=2 ttl=48 time=7.62 ms
64 bytes from 8.8.8.8: icmp_seq=3 ttl=48 time=7.57 ms
64 bytes from 8.8.8.8: icmp_seq=4 ttl=48 time=7.58 ms
64 bytes from 8.8.8.8: icmp_seq=5 ttl=48 time=7.65 ms
64 bytes from 8.8.8.8: icmp_seq=6 ttl=48 time=7.66 ms
64 bytes from 8.8.8.8: icmp_seq=7 ttl=48 time=7.63 ms
64 bytes from 8.8.8.8: icmp_seq=8 ttl=48 time=7.61 ms
64 bytes from 8.8.8.8: icmp_seq=9 ttl=48 time=7.62 ms
64 bytes from 8.8.8.8: icmp_seq=10 ttl=48 time=7.69 ms
64 bytes from 8.8.8.8: icmp_seq=11 ttl=48 time=7.62 ms
64 bytes from 8.8.8.8: icmp_seq=12 ttl=48 time=7.61 ms
64 bytes from 8.8.8.8: icmp_seq=13 ttl=48 time=7.61 ms
64 bytes from 8.8.8.8: icmp_seq=14 ttl=48 time=7.63 ms
64 bytes from 8.8.8.8: icmp_seq=15 ttl=48 time=7.62 ms
64 bytes from 8.8.8.8: icmp_seq=16 ttl=48 time=7.81 ms
64 bytes from 8.8.8.8: icmp_seq=17 ttl=48 time=7.39 ms
64 bytes from 8.8.8.8: icmp_seq=18 ttl=48 time=7.66 ms
64 bytes from 8.8.8.8: icmp_seq=19 ttl=48 time=7.65 ms
64 bytes from 8.8.8.8: icmp_seq=20 ttl=48 time=7.63 ms
64 bytes from 8.8.8.8: icmp_seq=21 ttl=48 time=7.66 ms
64 bytes from 8.8.8.8: icmp_seq=22 ttl=48 time=7.61 ms
64 bytes from 8.8.8.8: icmp_seq=23 ttl=48 time=7.63 ms
64 bytes from 8.8.8.8: icmp_seq=24 ttl=48 time=7.57 ms
64 bytes from 8.8.8.8: icmp_seq=25 ttl=48 time=7.63 ms
64 bytes from 8.8.8.8: icmp_seq=26 ttl=48 time=7.61 ms
64 bytes from 8.8.8.8: icmp_seq=27 ttl=48 time=7.60 ms
  
```

A **STOP (ESC)** button is visible at the bottom of the output area.

- **TCP Connect**

The screenshot shows the Gemini Central interface with the **Diagnostics** panel selected in the left sidebar. The **TCP Connect** tool is active, displaying the following interface:

Diagnostics

Ping **TCP Connect** NSLOOKUP TRACEROUTE TCPDUMP

Host:

Port:

Connected successfully.

TEST

• NSLOOKUP

The screenshot shows the 'Diagnostics' section of the Gemini Central interface. On the left is a dark sidebar with navigation icons for HOME, NODE, CLUSTER, LICENSE, SPLUNK, HADOOP, SETTINGS, and ACCOUNT. The main area is titled 'Node' and lists various system services like System Time, Timezone, Name, Network, FTP, SSH, SNMP, Failover, Log Receiver, Storage, Monitoring, and Diagnostics (which is highlighted). The 'Diagnostics' panel has tabs for Ping, TCP Connect, NSLOOKUP (selected), TRACEROUTE, and TCPDUMP. Under 'Domain Name', there is a text input field containing 'geminidata.com'. Below this, the results of the NSLOOKUP are displayed in a dark box: Server: 172.31.0.2, Address: 172.31.0.2#53, and a 'Non-authoritative answer' with Name: geminidata.com, Address: 54.68.61.203, Name: geminidata.com, and Address: 52.88.238.73. A blue 'TEST' button is located below the results.

• Traceroute

The screenshot shows the 'Diagnostics' section of the Gemini Central interface, similar to the NSLOOKUP screenshot. The 'Diagnostics' panel has tabs for Ping, TCP Connect, NSLOOKUP, TRACEROUTE (selected), and TCPDUMP. Under 'Host', there is a text input field containing 'www.geminidata.com'. Below this, the results of the traceroute are displayed in a dark box: 'traceroute to www.geminidata.com (54.68.61.203), 30 hops max, 60 byte packets'. The results show a sequence of hops from 1 to 30, with hop 1 showing a delay of 4.826 ms and all other hops marked with an asterisk (*). A blue 'TEST' button is located below the results.

• TCP Dump

Node

HOME

NODE

CLUSTER

LICENSE

SPLUNK

HADOOP

SETTINGS

ACCOUNT

System Time

Timezone

Name

Network

FTP

SSH

SNMP

Failover

Log Receiver

Storage

Monitoring

Diagnostics

Diagnostics

- Ping
 TCP Connect
 NSLOOKUP
 TRACEROUTE
 TCPDUMP

Network Interface

- Any
 eth0

Protocol

- Both
 TCP
 UDP

Host

Optional. The source host IP address, domain or hostname.

Port

Optional. The destination port of localhost.

TEST

Node

HOME

NODE

CLUSTER

LICENSE

SPLUNK

HADOOP

SETTINGS

ACCOUNT

System Time

Timezone

Name

Network

FTP

SSH

SNMP

Failover

Log Receiver

Storage

Monitoring

Diagnostics

Diagnostics

- Ping
 TCP Connect
 NSLOOKUP
 TRACEROUTE
 TCPDUMP

```

21:04:45.949821 IP 61.216.95.25.gprs-data > 172.31.42.136.https: tcp 0
21:04:46.154500 IP 61.216.95.25.gprs-data > 172.31.42.136.https: tcp 597
21:04:46.171491 IP 172.31.42.136.https > 61.216.95.25.gprs-data: tcp 344
21:04:46.171547 IP 172.31.42.136.https > 61.216.95.25.gprs-data: tcp 363
21:04:46.336090 IP 61.216.95.25.gprs-data > 172.31.42.136.https: tcp 0
21:04:46.336537 IP 61.216.95.25.gprs-data > 172.31.42.136.https: tcp 0
21:04:46.541397 IP 61.216.95.25.gprs-data > 172.31.42.136.https: tcp 597
21:04:46.558603 IP 172.31.42.136.https > 61.216.95.25.gprs-data: tcp 344
21:04:46.558664 IP 172.31.42.136.https > 61.216.95.25.gprs-data: tcp 363
21:04:46.722578 IP 61.216.95.25.gprs-data > 172.31.42.136.https: tcp 0
21:04:46.723049 IP 61.216.95.25.gprs-data > 172.31.42.136.https: tcp 0
21:04:46.928907 IP 61.216.95.25.gprs-data > 172.31.42.136.https: tcp 597
21:04:46.945808 IP 172.31.42.136.https > 61.216.95.25.gprs-data: tcp 344
21:04:46.945871 IP 172.31.42.136.https > 61.216.95.25.gprs-data: tcp 363
21:04:47.150829 IP 61.216.95.25.gprs-data > 172.31.42.136.https: tcp 0
21:04:47.151042 IP 61.216.95.25.gprs-data > 172.31.42.136.https: tcp 0
21:04:47.361119 IP 61.216.95.25.gprs-data > 172.31.42.136.https: tcp 597
21:04:47.377890 IP 172.31.42.136.https > 61.216.95.25.gprs-data: tcp 344
21:04:47.377984 IP 172.31.42.136.https > 61.216.95.25.gprs-data: tcp 365
21:04:47.542672 IP 61.216.95.25.gprs-data > 172.31.42.136.https: tcp 0
21:04:47.542898 IP 61.216.95.25.gprs-data > 172.31.42.136.https: tcp 0
21:04:47.747986 IP 61.216.95.25.gprs-data > 172.31.42.136.https: tcp 597
21:04:47.765152 IP 172.31.42.136.https > 61.216.95.25.gprs-data: tcp 344
21:04:47.765231 IP 172.31.42.136.https > 61.216.95.25.gprs-data: tcp 365
21:04:47.932092 IP 61.216.95.25.gprs-data > 172.31.42.136.https: tcp 0
21:04:47.932304 IP 61.216.95.25.gprs-data > 172.31.42.136.https: tcp 0
21:04:48.137011 IP 61.216.95.25.gprs-data > 172.31.42.136.https: tcp 597
21:04:48.154593 IP 172.31.42.136.https > 61.216.95.25.gprs-data: tcp 344
21:04:48.154668 IP 172.31.42.136.https > 61.216.95.25.gprs-data: tcp 363
21:04:48.319976 IP 61.216.95.25.gprs-data > 172.31.42.136.https: tcp 0
21:04:48.320170 IP 61.216.95.25.gprs-data > 172.31.42.136.https: tcp 0
21:04:48.525380 IP 61.216.95.25.gprs-data > 172.31.42.136.https: tcp 597
21:04:48.542693 IP 172.31.42.136.https > 61.216.95.25.gprs-data: tcp 344
21:04:48.542748 IP 172.31.42.136.https > 61.216.95.25.gprs-data: tcp 363
21:04:48.709444 IP 61.216.95.25.gprs-data > 172.31.42.136.https: tcp 0
        
```

STOP (ESC)

• IOSTAT

Node

- HOME
- NODE
- CLUSTER
- LICENSE
- SETTINGS
- ACCOUNT

Diagnostics

[Ping](#)
[TCP Connect](#)
[NSLOOKUP](#)
[TRACEROUTE](#)
[TCPDUMP](#)
[IOSTAT](#)

Device
 LVM /dev/sda

Interval

Count

Detail

[TEST](#)

Node

- HOME
- NODE
- CLUSTER
- LICENSE
- SETTINGS
- ACCOUNT

Diagnostics

[Ping](#)
[TCP Connect](#)
[NSLOOKUP](#)
[TRACEROUTE](#)
[TCPDUMP](#)
[IOSTAT](#)

```

Linux 3.10.0-693.17.1.el7.x86_64 (gemini-270b386c) 02/21/2018 _x86_64_ (3 CPU)
02/21/2018 09:24:33 PM
Device: rrgm/s wrqm/s r/s w/s rMB/s wMB/s avgrq-sz avqqu-sz await_r_await_w_await svctm %util
sda 0.00 0.03 617.72 155.73 2.43 0.75 8.41 0.79 1.01 1.13 0.56 0.06 4.54
02/21/2018 09:24:34 PM
Device: rrgm/s wrqm/s r/s w/s rMB/s wMB/s avgrq-sz avqqu-sz await_r_await_w_await svctm %util
sda 0.00 0.00 0.00 4.00 0.00 0.02 8.00 0.01 1.50 0.00 1.50 1.50 0.60
02/21/2018 09:24:35 PM
Device: rrgm/s wrqm/s r/s w/s rMB/s wMB/s avgrq-sz avqqu-sz await_r_await_w_await svctm %util
sda 0.00 0.00 0.00 4.00 0.00 0.02 8.00 0.00 0.25 0.00 0.25 0.25 0.10
02/21/2018 09:24:36 PM
Device: rrgm/s wrqm/s r/s w/s rMB/s wMB/s avgrq-sz avqqu-sz await_r_await_w_await svctm %util
sda 0.00 0.00 0.00 4.00 0.00 0.02 8.00 0.00 1.00 0.00 1.00 1.00 0.40
02/21/2018 09:24:37 PM
Device: rrgm/s wrqm/s r/s w/s rMB/s wMB/s avgrq-sz avqqu-sz await_r_await_w_await svctm %util
sda 0.00 0.00 0.00 4.95 0.00 0.02 7.60 0.00 0.80 0.00 0.80 0.80 0.40
02/21/2018 09:24:38 PM
Device: rrgm/s wrqm/s r/s w/s rMB/s wMB/s avgrq-sz avqqu-sz await_r_await_w_await svctm %util
sda 0.00 0.00 0.00 4.95 0.00 0.02 8.00 0.00 0.80 0.00 0.80 0.80 0.40
                    
```

[RESET \(ESC\)](#)

Rsync Backup

There are many modern network-attached storage devices that now support backup using **rsync**. With this feature you can backup Splunk configurations and data in the `/opt/sbox` folder, to the remote storage regularly.

To enable rsync backup, you need to do the following:

- Complete the SSH key exchange process between Gemini Central and the remote server, and allow this remote server use SSH login using a public key.
 - Select '**Download SSH Public Key**' to download the SSH public key from Manage. The default file name should be `id_rsa.pub`.
 - Add this **public key** into the authorized list of the remote server, usually located at `~/.ssh/authorized_keys`. For your convenience, use the following command to add it into the authorized list on the remote server:

```
cat id_rsa.pub >> ~/.ssh/authorized_keys
```

- Configure remote server information. There are four field values required:

- Remote Hostname/IP.**
- Remote Port** - The SSH listening port on the remote server. Default 22/tcp.
- Destination Path** - The folder name the backup data will be sent to.
- User Name** - Should match the one used for creation of the SSH public key.

Backup Rsync

[Download SSH Public Key](#)

Backup Configuration Through Rsync

Download SSH public key and add it to remote server's `~/.ssh/authorized_keys` file for connecting without password.
Complete configuring below settings for connecting to remote server and then then click on toggle button to enable backing up.

Setting

Remote Hostname/IP
172.27.11.11

Remote Port
22

Destination Path
nas_backup/

User Name
backup

Backup Scope

Splunk Configuration
 /opt/sbox

Backup Plan

- Determine the backup scope. There are two options available here: **Splunk Configuration** and folders in `/opt/sbox`. Within the `/opt/sbox/` option, you can specify which folders would you like to backup.

4. Configure the backup plan. In this section you need to determine the backup strategy, including the policy and schedule required:
 - If you select '**Always create a new full copy**', disk space may be quickly consumed. Monitor the free disk space of the remote server regularly.
 - If you select '**Keep a single copy up-to-date**', then there will only be one copy that exists, which should be the latest. However, you will not be able to restore data from older copies.
5. Select the '**Save**' button, to commit the configuration.

6. Select the '**Backup Configuration Through Rsync**' toggle button to enable rsync backup. This will verify for a successful exchange of the public key and also add it to the authorized key list of the remote server.

Benchmark

Use this feature to evaluate if the hardware specifications are suitable in taking on the role of running intensive disk I/O tasks, eg. Splunk Indexer.

Here you can run a Disk Benchmark on specific devices, monitor the disk IOPS(Input and Output Operations Per Second) in real time, and download the results. The following detailed benchmark methodology link can be found in the Gemini Support Portal.

<https://support.geminidata.com/learn/article/benchmarking-methodology/>

How to complete a disk benchmark:

- Select **'+ Run Benchmark'** and choose the target device to benchmark from the **'Select Device'** drop-down panel.
- Read the **'Notes'** provided with care before you proceed with the **'Run Benchmark'** button.

Note: When initiated, there is no way to cancel or stop the benchmark test.

The screenshot shows the Gemini Central interface for running a disk benchmark. On the left is a dark sidebar with navigation icons and labels: HOME, NODE, CLUSTER, LICENSE, SETTINGS, and ACCOUNT. The main panel is titled 'Disk Benchmark' and contains a '+ Run Benchmark' button. Below this button is a 'Select Device' dropdown menu with '/opt (dm-0)' selected. A 'Note' section lists several important warnings: disk benchmark drains system resources, it should not be run on production environments, it requires significant disk space, each job runs 5 times for 10 minutes, and running applications should be closed for accurate results. At the bottom of the panel are 'RUN BENCHMARK' and 'CANCEL' buttons. The current IOPS is shown as 'N/A'.

- During the benchmark process, it will monitor the operating system and display the IOPS in real-time. It will also record the max IOPS on screen.

Node

- System Time
- Timezone
- Name
- Network
- FTP
- SSH
- SNMP
- Failover
- Log Receiver
- Storage
- Monitoring
- Diagnostics
- Backup
- Benchmark**

Disk Benchmark

In progress (10:27)

You can run disk benchmark on specific devices, monitor the disk IOPS in real time, and download the result. The detailed benchmark methodology is available in the Gemini Support Portal.

Current: **12769** IOPS Maximum: **24299** IOPS

No recorded benchmark report.

- When the benchmark test has been completed, you will be given the result as an average value. This can be downloaded for deeper analysis.

Node

- System Time
- Timezone
- Name
- Network
- FTP
- SSH
- SNMP
- Failover
- Log Receiver
- Storage
- Monitoring
- Diagnostics
- Backup
- Benchmark**

Disk Benchmark

+ Run Benchmark

You can run disk benchmark on specific devices, monitor the disk IOPS in real time, and download the result. The detailed benchmark methodology is available in the Gemini Support Portal.

Current: **N/A** IOPS Maximum: **N/A** IOPS

Timestamp	Device Name	Avg IOPS
2018-02-22 13:10:19 +0800	dm-0	17705

Download
Delete

Disk benchmark testing finished.

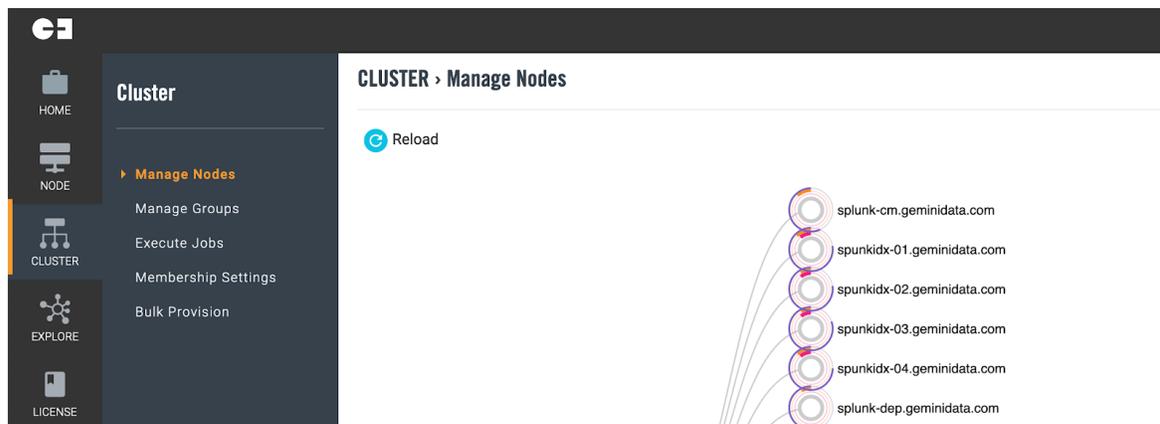
Cluster

The Gemini **Cluster** tab is the starting point for managing individual groups of Gemini Central nodes. This might include small groups for **High Availability**, **Failover** or instances that require collective **jobs** to be executed.

Cluster group registration is completed from the central Management Center instance operating as a 'Parent' node which registers nodes directly as 'Child nodes'.

Note

This process was initiated from the 'Child' node in Gemini Central versions prior to 2.8.



Manage Nodes

The **Manage Nodes** dashboard of the Parent node, ie. the **Management Center** for example, is the central location for the enhanced monitoring functionality of Gemini Central.

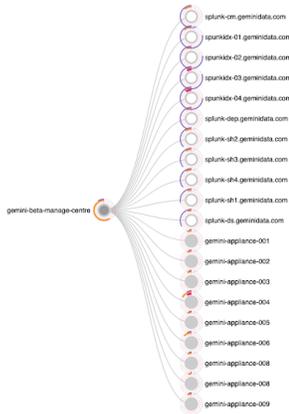
As suggested, this is particularly of use at the **Management Center** instance, as represented in the example below;

Note

This dashboard was previously known as the 'Topology' dashboard in Gemini Central versions prior to 2.8.

CLUSTER - Manage Nodes

Reload



Add Node

Manage Nodes				
Hostname	IP Address	Node Type	Joined Node Group	Assigned Jobs
gemini-beta-manage-centre	-	Management Center	1	0
splunk-om.gemini-beta.com	10.1.5.193	Gemini Agent	-	-
splunk-01.gemini-beta.com	10.1.5.190	Gemini Agent	-	-

With this dashboard it is now possible for the **Management Center** to show the status of *all* your Gemini instances whether they be appliances, cloud-based instances, virtual nodes, or even remote Splunk clusters (using Gemini Agents).

This dashboard will show the status of each ‘Child’ node registered to the Management Center and will include important metrics like CPU, RAM and disk space all measured at one-minute intervals.

If the **Bulk Provision** method was used to create multiple instances, this feature will be created automatically with the Management Center as Parent of the cluster instances.

The example below represents one such instance from the dashboard. Hover over the instance icon with your mouse to see the details panel presented to the right.



gemini-2785f16b

IP	127.0.0.1
CPU Cores	2
CPU Percent	0%
Memory Usage	555MB
Memory Total	1998MB
Memory Percent	35.4%
Disk Usage	2690MB
Disk Total	16343MB

The circular icon consists of a central roundel that differs to represent the instance type, surrounded by three coloured circles. Each circle gives an immediate representation of a key metric.



- **Blue** outside circle = Disk Space (% used)
- **Orange** circle = RAM usage (% used)
- **Red** inner circle = CPU usage (% used)

Identification of the instance type can be determined by the inner roundel;



Adding instances to an existing Manage Group

The addition of other Gemini nodes into an existing **Manage Group**, must be achieved from the **'Parent'** node. A parent node is the main control node for each unique Gemini Group. The most common parent node is the **Management Center**.

To add another instance to the parent cluster, whether it be an appliance, instance or remote agent, use the **'Add Node'** button on the **Cluster / Manage Nodes** dashboard.

Enter the **IP address** of **Hostname** of the instance, and select the **'ADD NODE'** button to confirm.

Manage Groups

Manage Groups enables the creation of smaller sub-groups of instances. This could be for reasons that include:

- Two standalone Gemini **Log Receiver** instances that are required to work together in a **Failover Group** or perform **Load Balancing** for Syslog.
- A group of instances that require a common '**Job**' to be scheduled or completed

The creation of a **Manage Group** can only be completed from the '**Parent**' node. The Management Center is the ultimate parent node, and can be used to create a **Manage Group** from any of the instances within its cluster.

Use the '+ **Create Node Group**' button to open the **Create Node Group** panel (see below), and enter the appropriate details.

The **Node Group** in the example below is called **SyslogGroup**, and consists of gemini-syslog1 & gemini-syslog2.

CLUSTER > Manage Groups

Node Groups

There are no groups currently configured.

[+ Create Node Group](#)

Create Node Group

Jobs can be dispatched to node groups only. Create a group and add a node to the group before assign jobs to it.
Note that the child node can be added to multiple groups.

Group Name
SyslogGroup

Description
Two Syslog Servers that will be operating in failover mode to provide syslog services to our network

Members
gemini-syslog1 x gemini-syslog2 x

ADD **CANCEL**

Select the '**Add**' button at the bottom of the dashboard to create the **Node Group**. The confirmation screen should resemble that below.

CLUSTER > Manage Groups

Node Groups

[+ Create Node Group](#)

Group Name	Description	Members	Jobs
SyslogGroup	Two Syslog Servers that will be operating in failover mode to provide syslog services to our network	2	0

Execute Jobs

The **Execute Jobs** dashboard allows you to execute predefined 'Jobs' such as starting/stopping or enabling services for multiple Manage instances.

It also allows the viewing of the current status of a Job, and any associated report detail.

Jobs can be assigned to all nodes or to a previously defined **Node Group**, and executed at a specific time.

Select the **'+ Create Job'** button to add or create a new Job.

Subject	Node Group	Task	Status
test	ALL NODES	Get component's version info	Reported
test2	ALL NODES	Get system hardware information	Reported

Select the Job itself from the tabular list to learn more details regarding the Job. Each Job has various states in which it can reside; Dispatched, Received, Fail, or Finished (successful).

← [Execute Jobs - Replicate Syslog Settings](#)

Remove

Subject	Replicate Syslog Settings	Status	Finished
Description	Replication Log Receiver rule to cluster nodes	Created At	2020-06-11 13:30:41
Task	(2201) Replicate Syslog Settings.	Received At	-
Node Group	syslogServers	Finished At	2020-06-11 13:31:18

Dispatches

Node	Status	Last Updated
gemini-syslog1	Finished	2020-06-11 13:30:44
gemini-syslog2	Finished	2020-06-11 13:31:18

Backup Center

This feature is designed to leverage the Gemini **Management Center** as a central repository for the storage of essential Gemini and Splunk configuration files. It will only work from a 'Parent' node, such as the Management Center.

Backup Center

Backup Job Manager

This will help you to backup the current Splunk and Manage configurations from a target node group. You may specify the scope, interval, copies, and execution time for each backup job.

Also, you may restore any one of the backups to a target node in the Gemini Cluster when needed.

[+ Add Backup Job](#)

Name	Description	Node Group	Interval	Status
Backup_All_Instances		allnodes	once	Success

- Edit
- Schedule
- Remove

Prior to creating a Backup Job, consider creating smaller groups of nodes for backup purposes, such as Indexers, Heavy Forwarders and Log Receivers. To create a Node Group, select the **+Create Node Group** button from the **Cluster / Manage Groups** dashboard and add the desired instances to the group.

CLUSTER > Manage Groups

Node Groups

[+ Create Node Group](#)

Group Name	Description	Members	Jobs
SyslogNG_Server_Group	two syslog-ng servers that will operate in a failover group	2	0
allnodes	Group created for a Global Backup or update job	9	0
SplunkIndexers	All Splunk Indexers including the CM	3	0

To create either a one-off or scheduled backup, use the **+ Add Backup Job** button on the **Cluster / Backup Center** dashboard to reveal the following options;

Add Backup Job

Job Name

Description

Node Group

Scope

 Both Splunk Manage

Interval

 Daily Weekly Monthly Run Immediately

Schedule Time

Copies

- Enter a logical name for the Job, ie. Splunk_Indexers
- Enter a description of the Job
- Use the **Node Group** dropdown menu to select the relevant group. If you do not find a suitable group here refer to the last paragraph on how to create a Manage Group.
- Use the **Scope** to specifically include either **Gemini** or **Splunk** configuration files, or have them backed up together within different directories of the same zip file.
- Use the **Interval** option to ensure that a regular backup is taken, based on the 'Schedule Time' and the maximum number of 'Copies' that are required to be kept.
- There is also an option here to complete a one-off backup in the form of the 'Run Immediately' option.

Select the 'ADD' button to commit this Job to run.

Once the Job is saved, locate the vertical ellipsis menu on the Backup Center dashboard to Edit or Remove any of the Jobs listed.

Backup Job Detail

To find more detail regarding a backup job, select the backup job itself from the **Backup Center** dashboard, to reveal the **Execution History** panel. Alongside each Job the vertical ellipsis menu will offer the ability to view 'Detail' of the backup job, and the ability to 'Download' the actual backup zip file.

Restoring Backups

To 'restore' a backup that has been completed, select the appropriate backup Job from the **Backup Center** dashboard to reveal an **Execution History** panel.

Backup Job - Backup_All_Instances

[← Back to The Job Index](#)

Execution History

Start (UTC)	End (UTC)	Status	Message
2020-11-25 10:19:48	2020-11-25 10:21:32	Success	Complete running job through ansible, playbook=backup

- Detail
- Restore
- Download

Restore Jobs

Start (UTC)	End (UTC)	Target	Status	Message
2020-11-25 10:22:37	2020-11-25 10:23:22	172.27.14.171	Success	

From the vertical ellipsis menu to the right of the relevant Job locate the **'Restore'** option. This will reveal the option to select the Gemini instance that is required to be restored.

Restore from Backup File

Nodes

172.27.14.134

Choose the Node to be restored from the drop-down list. If the **'both'** option was used to backup both the Gemini and Splunk configuration files, then both will be restored after confirming with the **'OK'** button.

For the ability to select either a Splunk or a Gemini restoration, we recommend the use of the Manage Group feature to create and split more backup jobs appropriately.

Membership Settings

The **Membership Settings** feature allows you to enable and configure Parent/Child relationships between Manage nodes.

If the **Bulk Provision** method was used to create multiple instances, relationships will automatically have been created with the Management Center as the **Parent** of all other **'Child'** node instances.

CLUSTER > Membership Settings

Reset Node Settings

Node Identifier

Hostname

Join Cluster
 Accept another Manage Node add this node to its cluster.
 No Yes

White List IP
 Only hosts listed below will be allowed to add local node into cluster.
 Leave empty to deny all, or use a wildcard (*) to allow all hosts.
 Multiple entries must be comma separated.

SAVE

Bulk Provision

This option has been detailed earlier in the Admin Guide. This will invoke the step-by-step **Bulk Provisioning** wizard to guide you through the provisioning of multiple nodes.

The node that initiates this process will become the Parent node and Licence Server for any instances provisioned.

Refer to the [Bulk Provisioning](#) section for detailed configuration steps.

License

Understanding Gemini Central Licensing

The **License** tab allows you to configure your Gemini Central Licenses, nominate License Servers and attach License Agents.

The Manage software license can be in one of three states;

1. A **Trial License** (30 Days)
2. A valid **Enterprise License**
3. A **Free License** (restricted features)

Changes to licensing can be completed at the Manage web interface at any time. A valid Enterprise License should be added within 30 days of Trial License activation.

Note

The Trial License will convert to the Free (restricted) license if not converted to an Enterprise License during the trial license period.

Initiating a Gemini license for your Manage environment can be achieved;

1. During the **Bulk Provision** process of Manage.

Select the '**Enterprise Edition (Purchased License)**' option, when the licensing prompt appears and follow the instructions titled '**Generate a License Request**'.

On receipt of the **License file**, it can be installed at any time within a 30 day period using the web interface.

2. Using the Manage web interface

Login to the Gemini Central web interface at any time to request an Enterprise License.

Navigate to the **License / License Status** menu, and follow the on-screen instructions starting with '**Step 1 - Generate a License Request**'.

License Status

License Status presents the current active license, including; the type, volume (number of nodes), and expiration date. Select the **Product** listed in the **Active License** panel to display more detail on the license.

If you need to request a License File, follow the steps on the License Status page to request and attach your License File in a three-step process. This ideally should be achieved within 30 days of installation.

If the **Trial** license is converted to an Enterprise License *after* the 30 day period, all slave nodes will remain at **Free** license status. All slave instances will function correctly in this state, but if you wish to correct the visible status on the dashboard, or if you wish to install another Featured Platform, you must restart Enterprise Services on all Slave nodes by running the following CLI command as the SBOX User:

```
sbox service --restart
```

Remote Licenses

The **Remote Licenses** dashboard enables you to link to a known Manage License Server.

Use the **'Add License Server'** button to enter the credentials of a Gemini Central instance that contains a valid **Trial** or **Enterprise** license.

Add Remote License Server

Host

Token String

Weight

Must between 0 and 99

Enter the IP address or FQDN of the chosen License Server in the Host entry box.

Enter the **Token String** value located on the **License / License Server** dashboard of the chosen License Server.

Enter a **'Weight'** value if you have two License Server destinations present, to determine which has the higher priority.

When connected successfully, a checkmark should be visible in both the **'Connected'** and **'Authenticated'** columns. Use the **'Refresh & Verify'** button if you have just set up this connection.

Manage License

License Status

Remote Licenses

Inventory

License Server

LICENSE > Remote License

Remote License Servers

+ Add License Server
 ↻ Refresh & Verify

	Host	Weight	Connected	Authenticated
☰	10.2.70.57	0	✓	✓
☰	10.2.70.56	4	✓	✓

Inventory

The **Inventory** dashboard will list all the licenses present as **Active** or **Inactive**. It also allows you to Revoke the Trial license when an Enterprise License has been installed.

LICENSE > Inventory

Active Licenses

↑ Upload License File
 ↓ Download License Request
↓ Revoke Trial License

Product	Upload Time	Expiration	Volume
Gemini Manage Trial	2020-07-28 07:35:49	2020-08-27 07:35:49	1

If for any reason the **License Inventory** is incorrect, perhaps following the application of a new Enterprise License for example, it may be necessary to restart Enterprise Services on the Slave nodes by running the following CLI command as the SBOX User:

```
sbox service --restart
```

License Server

The **License Server** dashboard allows you to enable and configure this Gemini Central instance as a **Manage License Server**.

A Manage License Server can manage both **Trial** and **Enterprise** licenses, granting permissions to other nodes connected to this instance.

LICENSE › License Server**License Server Settings****Allow Remote Access**

Allows another appliance to use this appliance as a license server.

 No Yes**Token String**

The security token another appliance needs to enter to access this license server.

White List

You could use wildcard * as any character for remote IP address, one line for one rule.

To use this instance as a Manage License Server, ensure that you '**Allow Remote Access**' by selecting the '**Yes**' tab.

Configure a suitable **Token String** that will be used by all remote nodes for registration.

Simply use an asterix (*) to allow all Manage instances to connect.

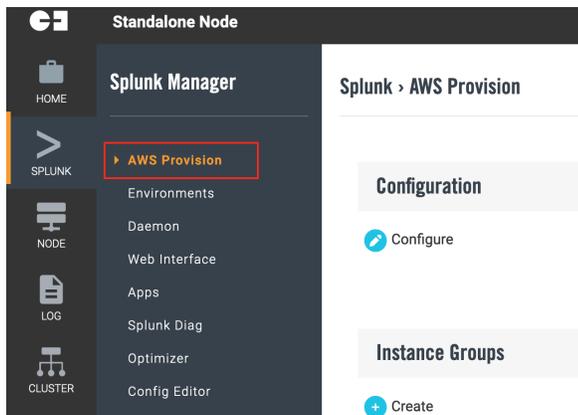
Or, if required, restrict nodes allowed to register by creating a '**White List**' controlled by either IP address or Hostname (multiple entries must be separated by a comma).

Splunk

The **Splunk** tab contains various areas of management related to the local Splunk platform installation. It will allow you to perform common tasks that would otherwise require access to the Splunk web interface or terminal screen.

Note that some of the functions here may not be suitable if the instance forms part of a Splunk cluster ie. **Upgrade Splunk(Daemon** dashboard) is achieved centrally using the **Rolling Upgrade** feature, and the editing of conf files (**Config Editor**) may be better suited to the function of Deployment Server.

New Feature: New to **Central 3.0** is an **AWS Provisioning** tool for **Splunk** clusters.



AWS Provisioning, enables central provisioning of complete Splunk Indexer and Search Head clusters based on **Splunk AMI's** and **AWS EC2 instances**. This is a departure from the use of our hardened OS, but will enable customers to get up and running quickly with complex Splunk environments on **AWS** in just a few minutes. This comes with the added benefit of Gemini Central offering central observation and a fast Splunk upgrade facility.

AWS Provisioning

Prerequisites for AWS Provisioning

In order to successfully complete the AWS Provisioning process, you will need access to valid **AWS** credentials for your environment such as the **AWS_IAM_ACCESS_ID** and **AWS_IAM_SECRET_KEY**. It may also be helpful to have access to the **AWS EC2** web console to review instances created.

AWS Provisioning Procedure

AWS Provisioning enables central provisioning of complete Splunk Indexer and Search Head clusters based on Splunk AMI's and AWS EC2 instances. A Gemini Agent will be automatically added to each instance during the process to give the added benefit of central observation and a speedy Splunk upgrade facility.

To instigate AWS provisioning to provide Splunk cluster resources, follow the procedure below;

Step 1.0 From Gemini Central's **Home** dashboard , select the '**Activate**' button from the Splunk Featured Platform, if not already activated, to reveal the Splunk Icon at the vertical menu bar.

Step 2.0 Select the **Splunk** icon to reveal a new addition to the menu. Select the '**AWS Provision**' option to open and begin the AWS provisioning process.

Step 3.0 In order to be able to deploy AWS instances, select the **'Configure'** option and enter the following information;

- AWS ACCESS KEY ID:
- AWS SECRET ACCESS KEY:
- AWS Region:
- Allow List: (allows public access to the AWS environment from your Gemini Central provisioning tool)

Step 4.0 When complete, select the **'Configure'** button to reveal the screen shown below. At this time a connection and initial configuration to AWS is attempted, using the above credentials.

Note: This process may take several minutes, so please be patient.

Configure

AWS ACCESS KEY ID

AKIAQB[REDACTED]5Z22A4

AWS SECRET ACCESS KEY

X8aQ7cmsEX[REDACTED]jqQ2yTunBYmMiAp

Region

us-west-2

Allow list

61.216.95.30/32

Ip of the host required to access the instances created

CONFIGURE

CANCEL

The screenshot shows the Splunk interface for 'Standalone Node' with the user 'admin'. The main content area is titled 'Splunk > AWS Provision' and displays a progress bar. The progress bar has two segments: a blue segment on the left labeled 'Arranging Aws Resources' and a grey segment on the right labeled 'Complete'. A red box highlights the 'Arranging Aws Resources' segment. Below the progress bar, there is a message: 'Get your SSH private key when key pair is ready.' with a 'DOWNLOAD' button. At the bottom of the screen, there is an 'OK' button.

Step 5.0 Assuming the AWS connection has been successful using the details provided, a confirmation screen will appear.

Use the **'Download'** button to obtain your **SSH private key**, required in order to **SSH** into any of the instances.

Select the **'OK'** button to return to the AWS Provisioning dashboard.

Step 6.0 From the **AWS Provisioning** dashboard, select the **+ Create** button to begin the creation of Splunk clusters using AWS instances.

This action will open a configuration screen for the creation of a **Splunk Indexer cluster** followed by an optional **Search Head cluster**.

Create Instance Group

Name
Gemini_AWS_test

Sizing Plan
 Instance type: c5.4xlarge (16 vCPU, 32G RAM), EBS: 8T, Total Instances: 3
 Instance type: c5.18xlarge (72 vCPU, 144G RAM), EBS: 16T, Total Instances: 3

Splunk

Splunk Cluster Name
gemini_aws_idx

Splunk Cluster Type
Indexer Cluster

Splunk Version (AMI)
splunk_AMI_8.1.3_2021-03-23

Daily Volume (GB/Day)
10

Retention (Day)
365

Splunk Admin Password
password@1

Splunk Secret
idxcluster

CREATE **CANCEL**

Name: Use a unique name that references the location of Gemini Central and contains the use case for the cluster.

Sizing Plan: There are only two instances that may become a Splunk Indexer (as recommended by Splunk), listed here in the drop-down box.

- C5.4 xLarge
- C5.18 xLarge

The 'Total Instance' count is purely dependent on the Daily Volume and Retention settings provided. In this way, the number and type of instances can be varied accordingly.

Splunk Cluster Name: Use a name that reflects the use case for the Indexer cluster being created.

Splunk Cluster Type: Choose between an Indexer or Search Head cluster. An Indexer cluster must first be provisioned before an SHC can be created.

Splunk Version (AMI): AMI's listed here are maintained by Splunk. Select one that meets your requirements.

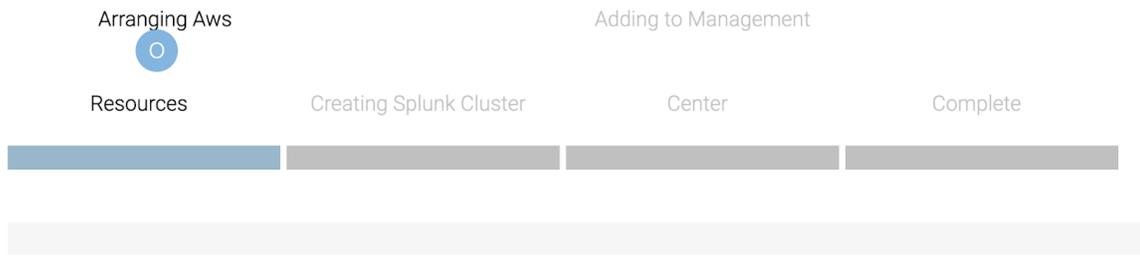
Daily Volume and Retention: Use to determine the number and type of AWS instances required. This links dynamically to the 'Sizing Plan' entry.

Splunk Admin Password: Create a suitable password to be used for the admin account of all instances. Ensure to keep a record of this.

Splunk Secret: Create a suitable secret key for communication within your Splunk Indexer cluster. Ensure to keep a record of this.

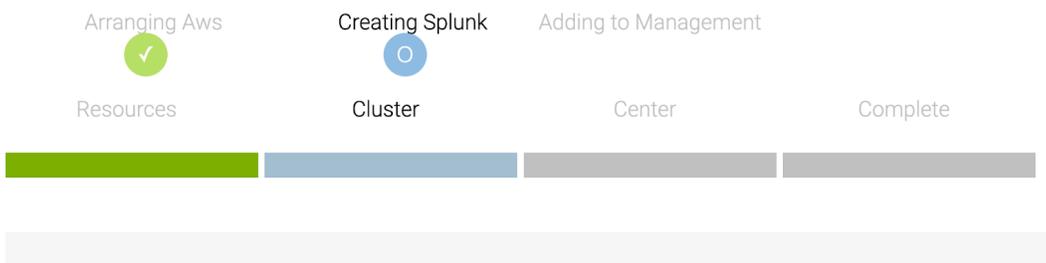
Step 7.0 When all the details have been correctly added, select the **'Create'** button to reveal the AWS Provisioning progress dashboard which begins looking like the image below;

Splunk > AWS Provision



The process begins at the ‘**Resources**’ stage as it provisions the required **AWS instances**. This process could take some time to complete especially if many resources are required. On completion however, the ‘**Creating Splunk Cluster**’ stage will begin to form an **Indexer Cluster**.

Splunk > AWS Provision

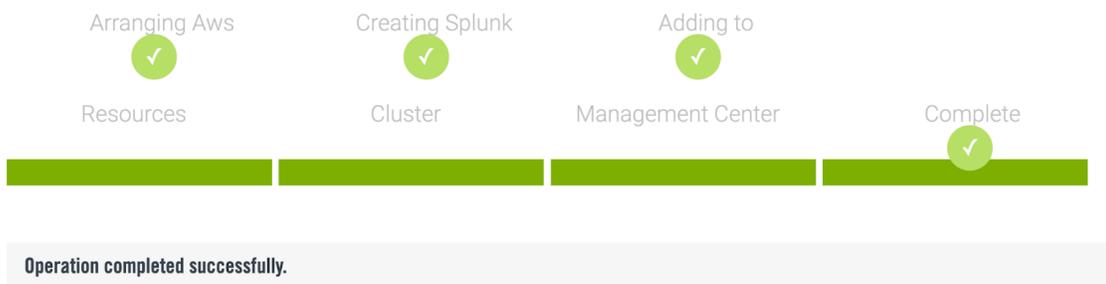


A **Splunk Indexer Cluster** is formed together with a **Cluster Master** using the detail offered during setup.

Using **Splunk approved AMI’s** keeps the process as authentic as possible, and with the addition of a **Gemini Agent** on each instance, this enables our **Gemini Central** platform to monitor and potentially upgrade the Splunk environment from one central dashboard on the **Management Center** node.

The final part of the process creates the connection between the **Gemini Agent** and **Gemini Central**. If all goes well the following confirmation screen should follow;

Splunk > AWS Provision



Selecting the ‘**OK**’ button will close the provisioning wizard, and reveal a summary screen of the Instance Groups present.

Selecting the **Indexer Cluster** label at any time will open the **Splunk Environments** dashboard. The **vertical ellipsis** menu offers options to **Destroy** or **Upgrade** the AWS Splunk environment.

Splunk > AWS Provision

Configuration

[Configure](#)

Instance Groups

[+ Create](#)

gemini-AWS-demo5 - Indexer Cluster - gemini-demo5-idx

Instance Name	Splunk role	Instance ID	Instance Type	Public Ip	Public DNS	
gemini-demo5-idx_1	CLUSTER_MASTER	i-0d705db0f6b8a8464	c5.4xlarge	54.214.136.141	ec2-54-214-136-141.us-west-2.compute.amazonaws.com	
gemini-demo5-idx_2	CLUSTER_PEER	i-0cf7732ec3a1435b6	c5.4xlarge	34.218.241.143	ec2-34-218-241-143.us-west-2.compute.amazonaws.com	
gemini-demo5-idx_3	CLUSTER_PEER	i-025eb9e70460e31bc	c5.4xlarge	18.237.110.165	ec2-18-237-110-165.us-west-2.compute.amazonaws.com	

Step 8.0 Verification of the provisioning process can be completed in three ways;

- The **Splunk Environments** dashboard in **Gemini Central**
- The **AWS EC2** dashboard
- The **Splunk web** interface of the **Cluster Master** instance (Indexer Clustering dashboard)

Navigate to the Splunk Environments dashboard of **Gemini Central** using either the cluster label on the **AWS Provisioning** dashboard or using Gemini Central’s **Splunk** icon.

From this dashboard, we can determine the Cluster Master node, IP addresses assigned, and also note the green circular icons which confirm that the status of splunkd on each instance is ‘running’.

SPLUNK > Splunk Environments

Environments
 Clusters
 Nodes

[+ Build Environment](#)
[+ Add Node](#)

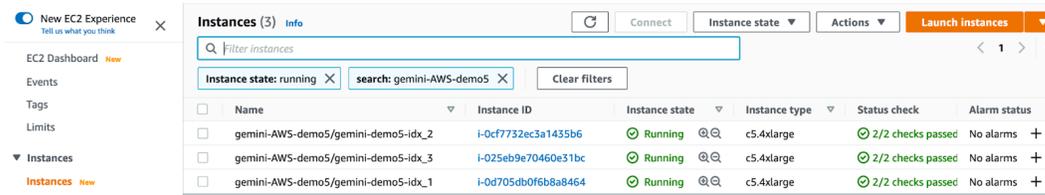
No Environments

Unassigned Nodes (3)

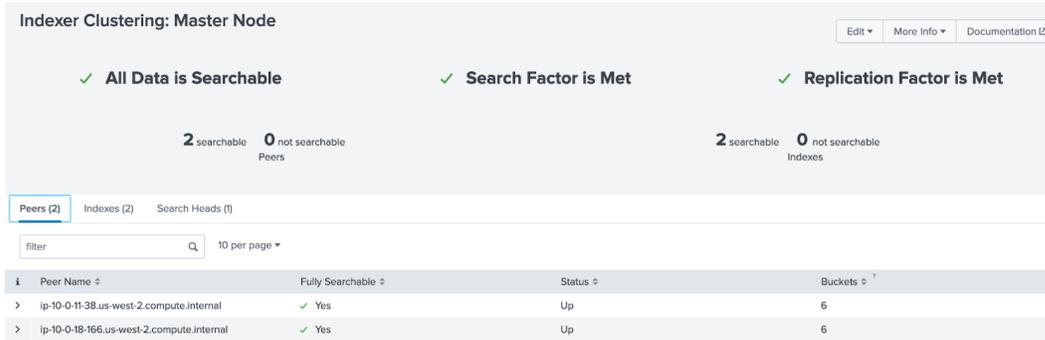
<input type="checkbox"/> Name	IP	Type	Splunk Software	Site	Deployment Type	Last Job
<input checked="" type="checkbox"/> gemini-demo5-idx	-	-	-	-	-	<input checked="" type="checkbox"/>
<input checked="" type="checkbox"/> ec2-54-214-136-141.us-west-2.compute.amazonaws.com	54.214.136.141	Gemini Agent	<input checked="" type="checkbox"/> Splunk Enterprise 8.1.3	default	Cluster Master	<input checked="" type="checkbox"/>
<input checked="" type="checkbox"/> ec2-34-218-241-143.us-west-2.compute.amazonaws.com	34.218.241.143	Gemini Agent	<input checked="" type="checkbox"/> Splunk Enterprise 8.1.3	default	Cluster Peer	<input checked="" type="checkbox"/>
<input checked="" type="checkbox"/> ec2-18-237-110-165.us-west-2.compute.amazonaws.com	18.237.110.165	Gemini Agent	<input checked="" type="checkbox"/> Splunk Enterprise 8.1.3	default	Cluster Peer	<input checked="" type="checkbox"/>

[ASSIGN TO ENVIRONMENT](#)
[CREATE NEW CLUSTER](#)
[ADD TO CLUSTER](#)

If you have login credentials to your AWS console, open the **EC2 Dashboard** and select ‘**running instances**’ to view details such as internal IP addresses assigned, VPC ID’s, and Subnet information, etc.



Alternatively, using the IP addresses given, login to the Splunk web interface of the **Cluster Master** instance to check the status of the **Indexer Clustering** dashboard.



Once verification has been completed, it is possible that you will also want to create a **Search Head Cluster** that links to the **Indexer Cluster**. If this is indeed the case, follow **Steps 9.0 - 12.0** for details on the process.

Step 9.0 Creating a Splunk **Search Head Cluster** on **AWS** uses a very similar process to that used for the Indexer Cluster.

Step 10.0 Select the **+ Create** button from the **AWS Provision** dashboard of **Gemini Central**.

Step 11.0 Carefully enter details of the required Search Head Cluster using the following as a guide;

Create Instance Group

gemini-AWS-demo5-shc

Sizing Plan
Instance type: c5.4xlarge (16 vCPU, 32G RAM), EBS: 1T, T

Splunk

Splunk Cluster Name
gemini-demo5-shc

Splunk Cluster Type
Search Head Cluster

Splunk Version (AMI)
splunk_AML_8.1.3_2021-03-23

Concurrent Users
16

Concurrent Searches
32

Master URI
10.0.17.139

Indexer Secret
idxcluster

Splunk Admin Password
gemini@123

Splunk Secret
shcluster

Name: Use a unique name that references the location of Gemini Central and contains the use case for the cluster.

Sizing Plan: There are only two instances that may become a Splunk SHC (as recommended by Splunk), listed here in the drop-down box.

- C5.4 xLarge
- C5.9 xLarge

The 'Total Instance' count is purely dependent on the Concurrent Users and Concurrent Searches settings provided. In this way, the number and type of instances can be varied accordingly.

Splunk Cluster Name: Use a unique name that reflects the use case for the Indexer cluster being created.

Splunk Cluster Type: Select the Search Head Cluster option. An Indexer cluster must first be provisioned before an SHC can be created.

Splunk Version (AMI): AMI's listed here are maintained by Splunk. Select one that meets your requirements.

Concurrent Users: and **Concurrent Searches:** These metrics are used by Splunk to gauge the number of CPU cores required for efficient searching. These entries link dynamically to the 'Sizing Plan' entry.

Master URI: Notice how this value is automatically populated with the internal IP address used on the AWS subnet. DO NOT modify this setting.

Indexer Secret: This should match the secret key used for the formation of the Indexer Cluster to which you wish to attach.

Splunk Admin Password: Create a suitable password to be used for the admin account of all instances. Ensure to keep a record of this.

Splunk Secret: Create a suitable secret key for communication within your Search Head Cluster. Ensure to keep a record of this.

Select the 'Create' button to begin the AWS provisioning process.

On successful completion of AWS provisioning, select the 'OK' button to reveal the additional cluster.

Instance Groups						
+ Create						
gemini-AWS-demo5 - Indexer Cluster - gemini-demo5-idx						
Instance Name	Splunk role	Instance ID	Instance Type	Public Ip	Public DNS	
gemini-demo5-idx_1	CLUSTER_MASTER	i-0d705db0f6b8a9464	c5.4xlarge	54.214.136.141	ec2-54-214-136-141.us-west-2.compute.amazonaws.com	
gemini-demo5-idx_2	CLUSTER_PEER	i-0cf7732ec3a1435b6	c5.4xlarge	34.218.241.143	ec2-34-218-241-143.us-west-2.compute.amazonaws.com	
gemini-demo5-idx_3	CLUSTER_PEER	i-025eb9e70460e31bc	c5.4xlarge	18.237.110.165	ec2-18-237-110-165.us-west-2.compute.amazonaws.com	
gemini-AWS-demo5-shc - Search Head Cluster - gemini-demo5-shc						
Instance Name	Splunk role	Instance ID	Instance Type	Public Ip	Public DNS	
gemini-demo5-shc_1	SHC_MEMBER	i-0aecd6ad5df0b2f0b	c5.4xlarge	54.191.75.131	ec2-54-191-75-131.us-west-2.compute.amazonaws.com	
gemini-demo5-shc_2	SHC_MEMBER	i-03dcd06a70279cb3b	c5.4xlarge	54.188.65.244	ec2-54-188-65-244.us-west-2.compute.amazonaws.com	
gemini-demo5-shc_3	SHC_DEPLOYER	i-057877062d1d3d29a	c5.4xlarge	54.245.141.187	ec2-54-245-141-187.us-west-2.compute.amazonaws.com	
gemini-demo5-shc_4	SHC_MEMBER	i-04b84be315e8204ff	c5.4xlarge	34.211.228.199	ec2-34-211-228-199.us-west-2.compute.amazonaws.com	

Step 12 Verification and observation of the cluster can be carried out by one of the same three methods as used for the Indexer Cluster;

- The **Splunk Environments** dashboard in **Gemini Central**
- The **AWS EC2** dashboard
- The Splunk web interface of the **Cluster Master** instance (Indexer Clustering dashboard)

Navigate to the **Splunk Environments** dashboard of **Gemini Central** using either the cluster label on the AWS Provisioning dashboard or using Gemini Central's Splunk icon. The screen will show the addition of a Search Head Cluster but due to the use of Gemini Agents within the provisioning wizard, Nodes from both clusters are present only as '**Unassigned Nodes**'. It is possible to leave them in this position if desired, but it is more common to add these clusters into a Splunk Environment space.

The process behind this is to create a '**shell**' environment into which the remote clusters can be imported. For details on this process, refer to the section entitled, [Creating a 'shell environment'](#).

Daemon

Allows you to review and modify settings related to Splunk Enterprise's **splunkd** process without requiring a command-line interface. Examples include;

- Stop or restart Splunk
- Upgrade the Splunk version (standalone Node only)
- Reset the Splunk Admin password
- Enable/Disable automatic boot-start and choose between 'initd' or 'systemd'

When Splunk is activated during a **Bulk Provisioning** process it is configured to enable an automatic boot-start using the **systemd** method of service control: ie. '**-systemd-managed 1**'

Existing installations of Splunk on Manage instances used the older 'initd' system of management control, that if discovered will be left in this state. Use the 'systemd service' checkbox to migrate to the newer system.

The Splunk **Workload Management** feature requires **systemd** to be enabled.

For a standalone version of Manage, running Splunk, Boot-Start is disabled by default.

- Advanced configurations such as changing the Splunk **Server Name** and the default directory used for Splunk indexes.

Note

If changing the **Server Name** here, the **default-hostname** value is *not* changed. Please change the default-hostname param manually to match the server name

The screenshot displays the 'SPLUNK > Daemon' configuration page. On the left is a dark sidebar with navigation icons for HOME, NODE, CLUSTER, LICENSE, SPLUNK, and SETTINGS. The main content area is titled 'SPLUNK > Daemon' and contains two sections:

- Splunk Service Control:** A horizontal bar with four buttons: 'Stop Splunk' (power icon), 'Restart Splunk' (refresh icon), 'Upgrade Splunk' (upward arrow icon), and 'Destroy Splunk Instance' (cross icon).
- Splunk Home Information:** A table showing 'SPLUNK HOME' as '/opt/splunk' and 'Version' as 'Splunk 7.3.4 (build 13e97039fb65)'.
- Boot-Start:** A section with two checkboxes:
 - Enable Splunk Boot-Start**: Enable BOOT-START in order to start Splunk daemon automatically at boot time.
 - Run Splunk as a systemd service.**

Web Interface

Allows you to review and modify settings related to Splunk Enterprise's Web Interface, Splunk Web. Here you may

- Disable or enable Splunk Web

Following a Bulk Provision process, only Search Heads and the Cluster Master have the Splunk web port enabled. This is standard practise for a secure Indexer Cluster.

- Launch Splunk web in a separate browser tab
- Review and modify advanced configurations such as enabling encryption and the default web port.

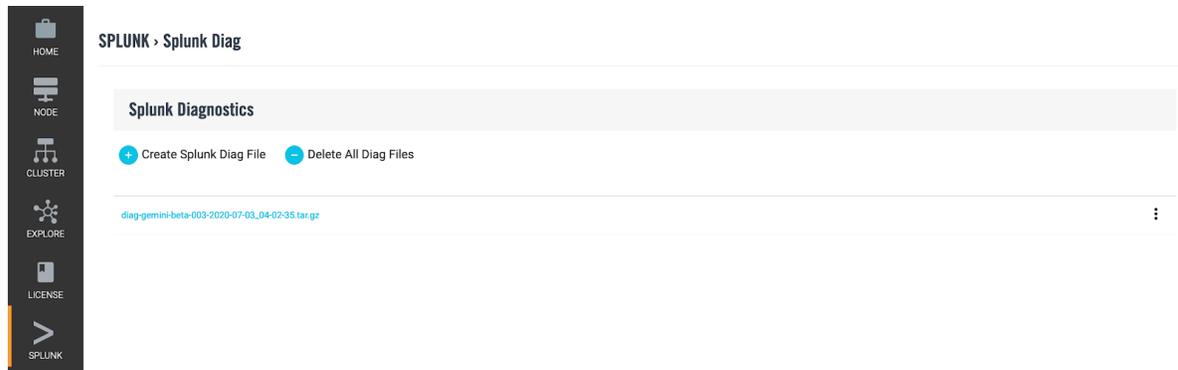
Apps

The **Apps** dashboard provides a list of all currently installed Splunk apps on the Gemini instance. Each app may be downloaded to your desktop as a tarball file using the vertical ellipsis button. Alternatively, select an app to access its directory structure using the Gemini **Config Editor** interface.

Splunk Diag

This dashboard allows you to quickly create a Splunk diag file [./splunk diag] from this, or any other Gemini instance that belongs to the Manage Cluster.

If this is conducted from the **Management Center**, the Splunk diag files from all Splunk instances controlled by this 'Parent' node, including those on Gemini Agents can be gathered to create a central repository for Diag files.



Use the '+ Create Splunk Diag File' button to open a selection panel;

Generating Splunk Diag

Remote Nodes

Please choose a node

The drop-down box will offer a selection from any Splunk instance that the **Management Center** (or parent node) is aware of, including remote Gemini Agents!

Optimizer

The **Optimizer** will set various configuration options for Splunk that suit various Splunk instance types. Settings are based on Splunk best practice recommendations, and will be applied from the `opt/splunk/etc/system/local` directory.

- Splunk Default (settings that result from a standard install of Splunk)
- Indexer
- Heavy Forwarder
- Search Head
- All In One

Splunk Optimizer

Please selected the desired instance type

Search Head

The following settings will be applied to your Splunk configurations.
Each item below will be written directly in the associated config file in \$SPLUNK_HOME/etc/system/local, overwriting existing keys if they exist. Enable the Versioning tool first in case settings rollback is required.

Stanza	Key	Recommended Value	Description
default	srchDiskQuota	1000	
default	srchJobsQuota	10	
default	rtSrchJobsQuota	12	
default	srchMaxTime	0	
monitor:///sbox/admin/var/log/	disabled	true	Add SBOX logs m
monitor:///sbox/admin/var/log/	index	default	Ingest SBOX logs
monitor:///sbox/admin/var/log/	sourcetype	sbox	Add SBOX logs m
settings	startwebservice	true	Set whether or nc
settings	enableSplunkWebSSL	true	
settings	sslVersions	tlse,tlse1.0	Allow TLS only fo
settings	privKeyPath	etc/auth/splunkweb/privkey.pem	
settings	max_upload_size	1024	Set maximum siz
search	status_cache_size	20000	The number of se
join	subsearch_maxout	70000	Maximum result r
lookup	max_memtable_bytes	20000000	Maximum 200ME

It is highly recommended that if using this feature, that you first apply our ['Versioning'](#) feature. This will ensure that if unintended effects are noticed, this application of settings can be reversed using the **'Rollback'** option.

Note

Warning: Ensure that you review the settings before selecting the **'Apply'** button, as there is no easy way to undo settings applied using this feature.

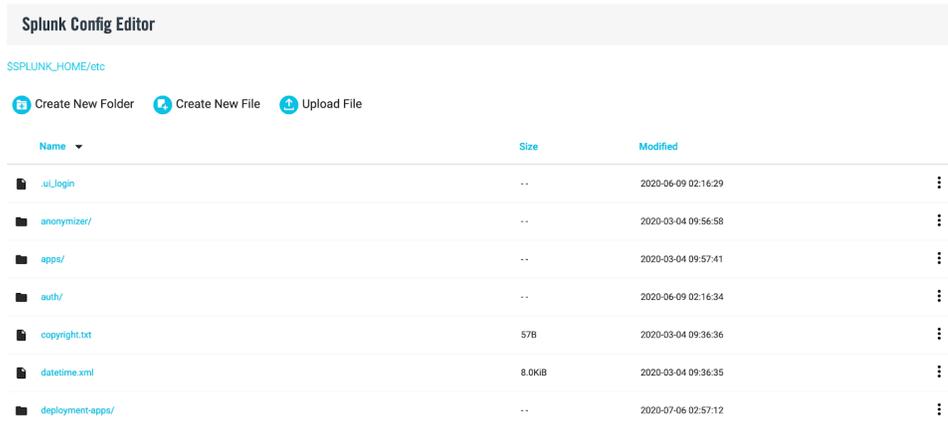
1. These settings will take precedence over **all** other settings due to their location (`/etc/system/local`)
2. Re-applying the **'Splunk Default'** template will **NOT** recover the settings.
3. Original settings can only be recovered using the **Versioning / Rollback** facility or by a manual edit of files in the `/splunk/etc/system/local/` directory

Config Editor

The **Config Editor** feature allows you to edit, create, upload and unzip files within the `$SPLUNK_HOME/etc/` directory path using the convenience of the web interface.

Use the **Icon buttons** to create new directories for config files or apps, create a new file in an on-screen editor, and upload files from your workstation.

SPLUNK > Config Editor



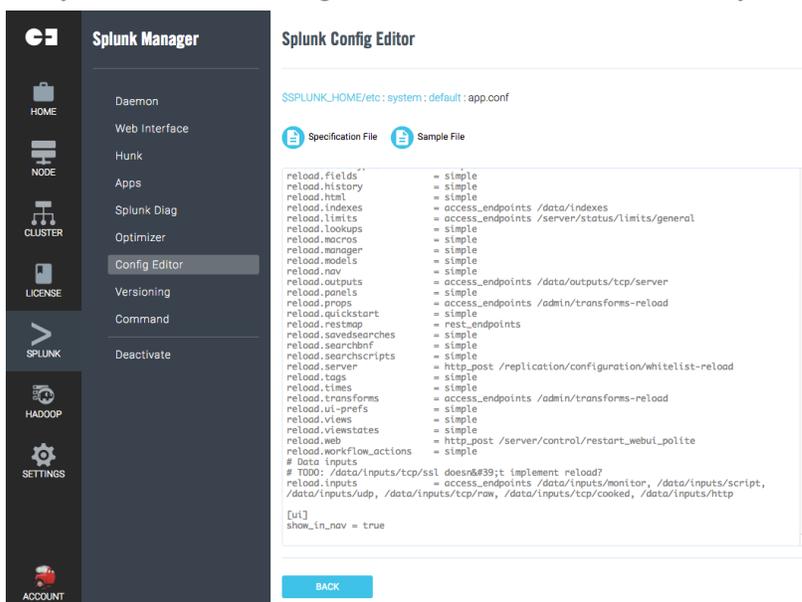
Using the **vertical ellipsis** menu to the right of the file or directory reveals; **copy**, **move**, **rename**, **remove** or **extract** functions.

To create a new folder/directory, use the **'Move'** open to **'Add New folder'**.

To extract a tarball or zip file, use the **Extract** option (see below)



Take care when editing Splunk `.conf` files. Remember that clustered instances control configuration files centrally. Never make changes to files in a **'Default'** directory, *a/ways* use the **'Local'** directory.



To **rollback** any changes, use this in conjunction with the **Versioning** feature.

Versioning

This feature allows use of our **Splunk Configuration Repository**, which in turn allows you to maintain version control of your Splunk configuration changes. Provided you are running the Enterprise version of Gemini Central, an unlimited number of configuration versions can be retained.

This feature actually uses the 'git' mechanism behind the scenes to provide a simple capture, roll-back and recover facility, following any changes that may have had unintended consequences.

This feature is **'off'** by default. Enable the feature using the slider button on the **Splunk / Versioning** dashboard. Once enabled, an **'Initial commit'** version of the current configuration will automatically be made that will become version (1).

Use the **'Create New Revision'** button to record any changes made to Splunk configuration files on this instance. When initiating a new revision, it is mandatory - and extremely useful - that you create a **'Description'** regarding the changes made to the Splunk configuration for reference purposes.

Each revision will capture incremental changes for the addition, amendment or deletion of Splunk files contained in the **splunk/etc** directory, as compared to the previous numbered **'commit'** version.

Splunk Manager

- HOME
- NODE
- CLUSTER
- LICENSE
- SPLUNK
- HADOOP
- SETTINGS
- ACCOUNT

Splunk Versioning

Splunk Configuration Repository
Allows you to save the current state of configuration files.

Revisions

Create new revision to save status, or click on a revision to roll back.

Revision	Description
2	Optimized for se
1	Initial commit

New Revision

Description: Optimized for heavy forwarder

File	Change
system/local/inputs.conf	Modified
system/local/limits.conf	Modified
system/local/outputs.conf	Modified
system/local/server.conf	Modified
system/local/web.conf	Modified
system/local/default-mode.conf	Added

SAVE REVERT CANCEL

Rollback Option

Use the **'Rollback'** option located at the vertical ellipsis menu at any time you need to temporarily rollback your changes to any one of the previous revision states.

Revisions				
Create New Revision Create new revision to save status, or click on revision number to check detail information and changes.				
Head	Revision	Description	User	Created
	2	Addition of Deployment Server feature	admin	2020-07-06 03:01:37
	1	Initial commit	admin	2020-07-06 02:35:24

On initiation of the **Rollback**, a message is displayed advising that an earlier revision has been selected and is now in use by Splunk.

Rollback State

Splunk Configuration is temporarily in revision 4. You can go back to the latest saved revision by clicking the 'Cancel Rollback' icon. Creating a 'new revision' in this state will overwrite revisions after revision 4.

Information

Behind the scenes this command initiates a **git-checkout master <@HEAD>** command. (use the **Cancel** option to perform git revert)

For logging information, see the **master** file in; /opt/splunk/etc/.git/logs/refs/heads/

Recover option

The **Recover** option offers a more permanent change to the committed history of a version. This command can be found by selecting the revision number listed under the Revision column. The process will completely eliminate the history regarding this revision, although any additional files or changes are not restored automatically. If additional files or changes were created in this version that are no longer required, manually amend those changes in the /etc directory of Splunk.

On initiation of **Recover**, the following message will be displayed confirming complete removal of the revision history.

Are you sure?

Recover will change Splunk Configuration to revision 5. Revision after revision 5 will be permanently deleted.

Information

Behind the scenes this command initiates a **git-reset[1]** command to remove the commit history (Note: use the **Cancel** option to perform git revert)

For logging information, see the **master** file in; /opt/splunk/etc/.git/logs/refs/heads/

Command

The **Splunk Command** feature allows you to issue Splunk commands directly from your browser.

Some commands like **status** or **version** do not require authentication and can be run as shown in the example below.

Other commands must have an authentication parameter added in order to run;

ie. **list licenses -auth admin:gemini123**

Notice that the 'splunk' command is already implied when using this interface



Installing a Splunk App

An example use-case for the **Command** feature, could be to add a new Splunk App.

1. Use either the **Splunk Config Editor** or SCP to upload the Splunk App to your Manage instance. We will assume the use of SCP to the /tmp directory.
2. Use the following command in the command entry box and select the **'Execute'** button

install app /tmp/NewAppName.spl -auth admin:<password>

SPLUNK > Command

Command Helper

Enter a Splunk command to run

E.g. "version" or "status" for Splunk information.

Note:

1. For security reason, some characters and "cmd" command are restricted.
2. User interactive commands like "apply shcluster-bundle" are not supported.

splunk > install app /tmp/NewAppName.spl -auth admin:password

EXECUTE

3. A message should confirm successful installation of the App.

Splunk Environments

Within Gemini Central, a **'Splunk Environment'** generally contains an **Indexer** and **Search Head Cluster** created in one or more locations, perhaps with one or more Splunk standalone instances. Several such 'environments' could exist making the management of multiple Splunk environments easier to maintain centrally at the **Environments** dashboard of the **Management Center**.

If you require one or more **Splunk Environments** to be created using existing Gemini instances, or if you want to add more instances to grow your existing Splunk environment, use the **Splunk Environments** dashboard of Gemini Central.

Using **Gemini Agents**, Splunk Environments support the adoption of existing 'remote' Splunk instances, including complete Splunk clusters. Refer to the [Gemini Agents](#) section for more detail.

When a requirement exists to upgrade Splunk to a new version, you are able to upgrade entire environments, including remote Splunk clusters with a few simple clicks of the mouse.

The **Splunk Environments** dashboard has a built-in heartbeat monitoring feature that works across **all** Gemini instances including those remote Splunk instances that contain Gemini Agents. This feature gives a real-time status of your entire Splunk deployment and is represented by the green icon. This will immediately change to red should Splunk fail.

To access the **Environments** dashboard from the vertical menu bar, select **Splunk / Environments** (see below for an example).

Appliance Cluster 01			2 Clusters	✓	⋮	
gemini-appliance-009			Splunk Enterprise 7.3.4	✓	⋮	
gemini-appliance-008			Splunk Enterprise 7.3.4	✓	⋮	
gemini-cluster-idxcluster		Indexer		3 Nodes	✓	⋮
gemini-appliance-001	Cluster Master	site1	Splunk Enterprise 7.3.4	✓	⋮	
gemini-appliance-002	Cluster Peer	site1	Splunk Enterprise 7.3.4	✓	⋮	
gemini-appliance-003	Cluster Peer	site1	Splunk Enterprise 7.3.4	✓	⋮	
gemini-cluster-shcluster		Search Head		4 Nodes	✓	⋮
gemini-appliance-004	SHC Deployer	site1	Splunk Enterprise 7.3.4	✓	⋮	
gemini-appliance-007	SHC Member	site1	Splunk Enterprise 7.3.4	✓	⋮	
gemini-appliance-005	SHC Member	site1	Splunk Enterprise 7.3.4	✓	⋮	
gemini-appliance-006	SHC Member	site1	Splunk Enterprise 7.3.4	✓	⋮	
Remote Splunk Cluster				2 Clusters	✓	⋮
splunk-sh1.geminiidata.com			Splunk Enterprise 7.3.3	✓	⋮	
splunk-ds.geminiidata.com			Splunk Enterprise 7.3.3	✓	⋮	
Remote_prod_SHC_01	Search Head			4 Nodes	✓	⋮

Prerequisites:

If you want to create an **Indexer Cluster** from instances, the following conditions must be met:

- At least three nodes are required for an indexer cluster, one being assigned as the Cluster Master and other two as peer nodes.

- If multi-site clustering is required, there must be at least 2 indexers on each site. Check the **Nodes / Name** dashboard for naming and IP detail.

If you want to create a **Search Head Cluster** from instances, the following conditions must be met:

- A Splunk Indexer Cluster must already exist, the Cluster Master IP address and Indexer secret key must both be known. These are all requirements for creating a Search Head Cluster.
- At least 4 nodes are required for a Search Head Cluster, one designated as a Deployer and the others as Search Peer nodes. The Cluster Master and Deployer must be separate instances.

If you want to adopt an **existing Splunk cluster** or environment, including a **remote Splunk cluster**, the following conditions must be met:

- A complete cluster must be added in one operation, observing the correct number and roles(as above). Incomplete cluster members or incorrect cluster information will cause a failure during assignment into an environment. Cluster Master and Deployer can not exist on the same node.
- If the existing Splunk cluster is not running on Gemini instances, ie, a remote Splunk environment, ensure that Gemini Agents have been correctly installed on all the target hosts.
- A new Splunk Environment must exist in Gemini Central prior to the adoption of another cluster. To facilitate this, a **'shell'** environment can be created as the container for the adoption of an existing Splunk instance or cluster within the Splunk Environments dashboard.

Refer to the [Creating a 'shell environment'](#) section for details.

The screenshot shows the 'SPLUNK > Splunk Environments' dashboard. At the top, there are filters for 'Environments', 'Clusters', and 'Nodes', along with a search bar for name and IP. Below the filters, there are buttons for 'Build Environment' and 'Add Node'. The main content area shows 'No Environments' in a dashed box. Below that, there is a section for 'Unassigned Nodes (10)' with a table listing nodes.

Name	IP	Type	Splunk Software	Site	Deployment Type	Status
gemini-004	172.27.14.134	Software Appliance	Not Installed	-		✓
gemini-001	172.27.14.131	Software Appliance	Not Installed	-		✓
centos	172.27.14.130	Manage Agent	Not Installed	-		✓
gemini-003	172.27.14.133	Software Appliance	Not Installed	-		✓
gemini-005	172.27.14.135	Software Appliance	Not Installed	-		✓

Adding a Node (Unassigned Nodes)

An additional Node includes any **Manage Instance(s)** or **Gemini Agent(s)** that are required to be added into the **Management Center** node of Gemini Central.

In order to add standalone Splunk Instances or complete Splunk clusters into the **Management Center**, they will first need to be made available as **'Unassigned Nodes'**. Only then can they be re-assigned to an existing Splunk Environment, or used for the creation of a new Splunk Environment.

There are several ways of adding to the list of Unassigned Nodes depending on the number and type involved;

- [Add a single Node](#) (Splunk need not be present)
- [Add a pre-existing Splunk Indexer Cluster](#) (instances or Gemini Agents)
- [Add a pre-existing Splunk Search Head Cluster](#) (instances or Gemini Agents)
- [Add a group of standalone Splunk Nodes](#) (instances or Gemini Agents)

If you already have unassigned node(s), and you wish to assign them to a Cluster or Environment, refer to the [Assigning Unassigned Nodes and Clusters to a Splunk Environment](#) section.

Add a single Node

To begin the process select the '+ Add Node' button from the Splunk Environments dashboard of the Management Center, to reveal the following:

Add Node

Please input FQDN to discover

FQDN

IP

Enter a single **FQDN name** and **IP address** into the appropriate boxes. Note that both of these are required to be completed. If you do not know the name of the host or indeed if you want to rename it, add the chosen name to the **FQDN** box and select the '**Add**' button. This will add the Node with the chosen name.

Notes

Appliance addition: It is important that the version of Manage used is the same as that at the Management Center

Remote node addition: The Gemini Agent must first be installed on the remote Splunk instance.

Add a pre-existing Splunk Indexer Cluster

This method is ideally suited to the addition of external Splunk Environments running our **Gemini Agents**.

For details on the installation of Gemini Agents, refer to the [Gemini Agents](#) section.

Select the '+ Add Node' button from the **Splunk Environments** dashboard of the **Management Center** and locate the **JSON** and **CSV** manifest templates.

Or upload a list of nodes or configuration of a cluster in the conf file to add nodes

Drop package here or click to [choose the file](#) from your computer

[Sample Indexer Cluster Manifest\(.json\)](#)
[Sample Search Head Cluster Manifest\(.json\)](#)
[Sample Standalone Nodes Manifest\(.json\)](#)
[Sample Indexer Cluster Manifest\(.csv\)](#)
[Sample Search Head Cluster Manifest\(.csv\)](#)
[Sample Standalone Nodes Manifest\(.csv\)](#)

If you want to add a **Splunk Indexer Cluster** that exists externally, first ensure that each instance of that cluster has the **Gemini Agent** installed.

Select the chosen format(CSV or JSON) for the '**Sample Indexer Cluster Manifest**' to download the appropriate manifest file. This should be opened in a suitable text editor.

JSON Manifest Option

After reading the instructions given at the beginning of the manifest file, complete the manifest using the correct JSON format where;

'**cluster_name**' is the name of the cluster. Give it a unique name.

'**type**' Use "**SPLUNK_INDEXER**" for an Indexer Cluster.

'**nodes**' is the list of cluster members. There are 7 attributes in each node:

'**hostname**' is used to identify the node. Please keep it the same with the node.

'**ip**' is the IPv4 address of this node. Use the IP address the management node can connect to.

'**role**' is the role it has in the cluster; In an indexer cluster, this could be either 'CLUSTER_MASTER' or 'CLUSTER_PEER'.

'**splunk_home**' is the home directory of the Splunk service (/opt/splunk by default)

'**splunk_user**' is an OS user created by Gemini to run Splunk.

'**admin_username**' is the account used for administrator privilege.

'**admin_password**' is the password of the admin account. Not stored in Manage.

'**secret**' is the secret cluster key, used to communicate between cluster members.

An example of this JSON manifest can be seen below.

```
{
  "cluster_name": "Remote_prod_IDXCluster",
  "type": "SPLUNK_INDEXER",
  "nodes": [
    {
      "hostname": "splunk_cm.example.com",
      "ip": "192.168.1.1",
      "role": "CLUSTER_MASTER",
      "splunk_home": "/opt/splunk",
```

```

        "splunk_user": "splunk",
        "admin_username": "admin",
        "admin_password": "password"
    },
    {
        "hostname": "splunkidx_01.example.com",
        "ip": "192.168.1.2",
        "role": "CLUSTER_PEER",
        "splunk_home": "/opt/splunk",
        "splunk_user": "splunk",
        "admin_username": "admin",
        "admin_password": "password"
    },
    {
        "hostname": "splunkidx_02.example.com",
        "ip": "192.168.1.3",
        "role": "CLUSTER_PEER",
        "splunk_home": "/opt/splunk",
        "splunk_user": "splunk",
        "admin_username": "admin",
        "admin_password": "password"
    },
    {
        "hostname": "splunkidx_03.example.com",
        "ip": "192.168.1.4",
        "role": "CLUSTER_PEER",
        "splunk_home": "/opt/splunk",
        "splunk_user": "splunk",
        "admin_username": "admin",
        "admin_password": "password"
    }
},
"secret": "idxcluster_key"
}

```

Verify the following issues carefully before saving the manifest.

- You have used the correct **cluster_name** as found in **server.conf** on the Cluster Master
- You have included all members of the Indexer Cluster
- All entries have been added using valid JSON formatting.
- The **'secret'** references the **Indexer Clustering secret** key

CSV Manifest Option

This template is simpler and should be completed using the same criteria as described above in the JSON option.

cluster_name	type	role	secret	ip	hostname	splunk_home	splunk_user	admin_username	admin_password
Remote_prod_IDXCluster	SPLUNK_INDEXER	CLUSTER_MASTER	idxcluster	192.168.1.1	splunk_cm.example.com	/opt/splunk	splunk	admin	password
Remote_prod_IDXCluster	SPLUNK_INDEXER	CLUSTER_PEER	idxcluster	192.168.1.2	splunkidx_01.example.com	/opt/splunk	splunk	admin	password
Remote_prod_IDXCluster	SPLUNK_INDEXER	CLUSTER_PEER	idxcluster	192.168.1.3	splunkidx_02.example.com	/opt/splunk	splunk	admin	password
Remote_prod_IDXCluster	SPLUNK_INDEXER	CLUSTER_PEER	idxcluster	192.168.1.4	splunkidx_03.example.com	/opt/splunk	splunk	admin	password

Verify the following issues carefully before saving the manifest.

- You have used the correct **cluster_name** as found in **server.conf** on the Cluster Master
- You have included all members of the Indexer Cluster
- The file is correctly formatted as a CSV file. Download and use the **dos2unix** utility to ensure the correct format if required.
- The **'secret'** references the **Indexer Clustering secret** key

Once the detail has been added correctly, the saved file can simply be dropped into the box marked, **'Drop package here or click to choose the file from your computer'**.

Alternatively, use the highlighted **'click to choose'** link to locate and upload the file.

Select the **'Add'** button to complete the process. **Refresh** the browser. Confirm the new list of **'Unassigned Nodes'**, an example is given below.

Unassigned Nodes (5)

<input type="checkbox"/> Name	IP	Type	Splunk Software	Site	Deployment Type	Status
<input type="checkbox"/> Remote_prod_IDXCCluster		Software Appliance		-		
<input type="checkbox"/> splunk_cm	10.1.5.193	Software Appliance	Splunk Enterprise 7.3.3	-	Cluster Master	
<input type="checkbox"/> splunkidx_01	10.1.5.190	Software Appliance	Splunk Enterprise 7.3.3	site1	Cluster Peer	
<input type="checkbox"/> splunkidx_02	10.1.5.191	Software Appliance	Splunk Enterprise 7.3.3	site1	Cluster Peer	
<input type="checkbox"/> splunkidx_03	10.1.5.20	Software Appliance	Splunk Enterprise 7.3.3	site2	Cluster Peer	
<input type="checkbox"/> splunkidx_04	10.1.5.21	Software Appliance	Splunk Enterprise 7.3.3	site2	Cluster Peer	

Add a pre-existing Splunk Search Head Cluster

This method is ideally suited to the addition of external Splunk Environments running our **Gemini Agents**.

For details on the installation of Gemini Agents, refer to the [Gemini Agents](#) section

Select the **'+ Add Node'** button from the **Splunk Environments** dashboard of the **Management Center** and locate the **JSON** and **CSV** manifest templates.

If you want to add a **Splunk Search Head Cluster** that exists externally, first ensure that each instance of that cluster has the **Gemini Agent** installed.

Select the **'Sample Search Head Cluster Manifest'** to download the appropriate manifest file. This should be opened in a suitable text editor.

JSON Manifest Option

After reading the instructions given at the beginning of the file complete the manifest using the correct JSON format where;

'cluster_name' is the name of the cluster. Give it a unique name.

'type' Use **'SPLUNK_SHC'** for a Search Head Cluster.

'nodes' is the list of cluster members. There are 7 attributes in each node:

'hostname' is used to identify the node. Please keep it the same with the node.

'ip' is the IPv4 address of this node. Use the IP address the management node can connect to.

'role' is the role it has in the cluster; this could be either **'SHC_DEPLOYER'** or **'SHC_MEMBER'**.

'splunk_home' is the home directory of the Splunk service (**/opt/splunk** by default)

'splunk_user' is an OS user created by Gemini to run Splunk.

'**admin_username**' is the account used for administrator privilege.

'**admin_password**' is the password of the admin account. Not stored in Manage.

'**secret**' is the secret cluster key, used to communicate between SHC members.

'**indexer_secret**' is the secret indexer cluster key

An example of this JSON manifest can be seen below;

```
{
  "cluster_name": "Remote_prod_SHC",
  "type": "SPLUNK_SHC",
  "nodes": [
    {
      "hostname": "splunk_dep.example.com",
      "ip": "192.168.1.1",
      "role": "SHC_DEPLOYER",
      "splunk_home": "/opt/splunk",
      "splunk_user": "splunk",
      "admin_username": "admin",
      "admin_password": "password"
    },
    {
      "hostname": "splunk_sh2.example.com",
      "ip": "192.168.1.2",
      "role": "SHC_MEMBER",
      "splunk_home": "/opt/splunk",
      "splunk_user": "splunk",
      "admin_username": "admin",
      "admin_password": "password"
    },
    {
      "hostname": "splunk_sh3.example.com",
      "ip": "192.168.1.3",
      "role": "SHC_MEMBER",
      "splunk_home": "/opt/splunk",
      "splunk_user": "splunk",
      "admin_username": "admin",
      "admin_password": "password"
    },
    {
      "hostname": "splunk_sh4.example.com",
      "ip": "192.168.1.4",
      "role": "SHC_MEMBER",
      "splunk_home": "/opt/splunk",
      "splunk_user": "splunk",
      "admin_username": "admin",
      "admin_password": "password"
    }
  ],
  "secret": "shcluster_key",
  "indexer_secret": "idxcluster_key"
}
```

Tip

To ensure that your JSON is correct, use a third party validator such as <https://jsonlint.com>

Verify the following issues before saving the manifest.

- You have used the correct Search Head Cluster Name as found in the server.conf of any search head in the cluster.
- You have used the correct secret keys. These usually differ between the SHC and Indexer cluster
- All entries have been made using valid JSON formatting.

CSV Manifest Option

This template is simpler and should be completed using the same criteria as described above in the JSON option.

cluster_name	type	role	secret	indexer_secret	ip	hostname	splunk_home	splunk_user	admin_user	admin_password
Remote_prod_SHC	SPLUNK_SHC	SHC_DEPLOYER	shcluster	idxcluster	192.168.1.1	splunk_dep.examp	/opt/splunk	splunk	admin	changeme
Remote_prod_SHC	SPLUNK_SHC	SHC_MEMBER	shcluster	idxcluster	192.168.1.2	splunk_sh2.exempl	/opt/splunk	splunk	admin	changeme
Remote_prod_SHC	SPLUNK_SHC	SHC_MEMBER	shcluster	idxcluster	192.168.1.3	splunk_sh3.exempl	/opt/splunk	splunk	admin	changeme
Remote_prod_SHC	SPLUNK_SHC	SHC_MEMBER	shcluster	idxcluster	192.168.1.4	splunk_sh4.exempl	/opt/splunk	splunk	admin	changeme

Verify the following issues carefully before saving the manifest.

- You have used the correct Search Head Cluster Name as found in the server.conf of any search head in the cluster.
- You have used the correct secret keys. These usually differ between the SHC and Indexer cluster
- The file is correctly formatted as a CSV file. Download and use the **dos2unix** utility to ensure the correct format if required.

When complete, the saved file can simply be dropped into the box marked,

‘Drop package here or click to choose the file from your computer’.

Alternatively, use the highlighted **‘click to choose’** link to locate and upload the file.

Select the **‘Add’** button to complete the process. **Refresh** the browser. Confirm the new list of **‘Unassigned Nodes’**.

Unassigned Nodes (9)

Name	IP	Type	Splunk Software	Site	Deployment Type	Status
Remote_prod_IDXCluster		Software Appliance		-		✔
splunk_cm	10.1.5.193	Software Appliance	Splunk Enterprise 7.3.3	-	Cluster Master	✔
splunkidx_01	10.1.5.190	Software Appliance	Splunk Enterprise 7.3.3	site1	Cluster Peer	✔
splunkidx_02	10.1.5.191	Software Appliance	Splunk Enterprise 7.3.3	site1	Cluster Peer	✔
splunkidx_03	10.1.5.20	Software Appliance	Splunk Enterprise 7.3.3	site2	Cluster Peer	✔
splunkidx_04	10.1.5.21	Software Appliance	Splunk Enterprise 7.3.3	site2	Cluster Peer	✔
Remote_prod_SHC		Software Appliance		-		✔
splunk_dep	10.1.5.25	Software Appliance	Splunk Enterprise 7.3.3	-	SHC Deployer	✔
splunk_sh2	10.1.5.22	Software Appliance	Splunk Enterprise 7.3.3	-	SHC Member	✔
splunk_sh3	10.1.5.23	Software Appliance	Splunk Enterprise 7.3.3	-	SHC Member	✔
splunk_sh4	10.1.5.24	Software Appliance	Splunk Enterprise 7.3.3	-	SHC Member	✔

Add a Group of standalone Splunk Nodes

This could be used for bringing other instances such as; standalone Search Head, Heavy Forwarder, License Manager or Deployment Server into a Splunk Environment.

Select the **'+ Add Node'** button from the Splunk Environments dashboard of the Management Center and locate the **JSON** and **CSV manifest** templates.

Select the **'Sample Standalone Nodes Manifest'** to download the appropriate manifest file. This should be opened in a suitable text editor.

JSON Manifest Option

After reading the instructions given at the beginning of the file complete the manifest using the correct JSON format where;

'hostname' is used to identify the node. Please keep it the same with the node.

'ip' is the IPv4 address of this node. Use the IP address the management node can connect to.

'splunk_home' is the home directory of the Splunk service (/opt/splunk by default)

'splunk_user' is an OS user created by Gemini to run Splunk.

'admin_username' is the account used for administrator privilege.

'admin_password' is the password of the admin account. Not stored in Manage.

An example of this manifest is shown below for two potential Splunk instances;

```
{
  ],
  "type": "STANDALONE_NODE",
  "nodes": [
    {
      "hostname": "gemini-ds.example.com",
      "ip": "192.168.56.114",
      "splunk_home": "/opt/splunk",
      "splunk_user": "splunk",
      "admin_username": "admin",
      "admin_password": "password"
    },
    {
      "hostname": "gemini-mc.example.com",
      "ip": "192.168.56.103",
      "splunk_home": "/opt/splunk",
      "splunk_user": "splunk",
      "admin_username": "admin",
      "admin_password": "password"
    }
  ]
}
```

Tip

To ensure that your JSON is correct, use a third party validator such as <https://jsonlint.com>

CSV Manifest Option

This template is simpler and should be completed using the same criteria as described above in the JSON option.

type	ip	hostname	splunk_home	splunk_user	admin_username	admin_password
STANDALONE_NODE	192.168.56.114	gemini-ds	/opt/splunk	splunk	admin	password
STANDALONE_NODE	192.168.56.115	gemin-mc	/opt/splunk	splunk	admin	password

When complete, the saved file can simply be dropped into the box marked,

‘**Drop package here or [click to choose](#) the file from your computer**’.

Alternatively, use the highlighted ‘[click to choose](#)’ link to locate and upload the file.

Select the ‘**Add**’ button to complete the process. **Refresh** the browser before confirming the new list of ‘**Unassigned Nodes**’.

Assigning Unassigned Nodes and Clusters to a Splunk Environment

When **Unassigned Nodes** have been made available we have three choices highlighted by the buttons at the bottom of the screen to work with an existing Splunk Environment.

- [Assign to Environment](#) - Assigns Nodes or Clusters to a specific Environment
- [Create New Cluster](#) - Used to create a new Splunk Cluster
- [Add to Cluster](#) - Used to increase the number of Indexers or Search Heads

Before using these options you may need to create a new ‘**Splunk Environment**’. For instance, it is always recommended to assign a remote **Gemini Agent cluster** to a different **Splunk Environment**, in order to separate it from existing instance clusters and other potential Splunk Clusters in Manage.

If required, refer to the [Creating a ‘shell environment](#)’ section for help on creating a new **Splunk Environment** before assigning nodes or clusters.

For assistance on adding remote Splunk instances and clusters, refer to the [Gemini Agents](#) section.

Assign to Environment

If a **Splunk Environment** already exists, this option can be used to create standalone Splunk instances such as a **Deployment Server** or **Management Console**, or perhaps a Gemini **Log Receiver** instance, within that environment.

Alternatively, in the case of a **remote Splunk cluster**, once a suitable ‘**shell environment**’ exists, this option can be used to locate a Splunk cluster into that **Splunk Environment** shell, as in the example image below.

If you require simply to add an additional **Indexer** or **Search Head** to your existing **Splunk Environment**, then please refer to the [Add to Cluster](#) section.

Unassigned Nodes (3)

Name	IP	Type	Splunk Software	Site	Deployment Type	Status
<input checked="" type="checkbox"/> splunk-remote-cluster		Software Appliance		-		<input checked="" type="checkbox"/>
<input checked="" type="checkbox"/> splunkidx_01	10.1.5.190	Software Appliance	Splunk Enterprise 7.3.1	-	Cluster Peer	<input checked="" type="checkbox"/>
<input checked="" type="checkbox"/> splunkidx_02	10.1.5.191	Software Appliance	Splunk Enterprise 7.3.1	-	Cluster Peer	<input checked="" type="checkbox"/>
<input checked="" type="checkbox"/> splunk_com	10.1.5.193	Software Appliance	Splunk Enterprise 7.3.1	-	Cluster Master	<input checked="" type="checkbox"/>

Select the required node from the **Unassigned Nodes** panel of the **Splunk Environments** dashboard.

Select the **'Assign to Environment'** button at the bottom of the screen.

Form the panel that opens on the right, choose an appropriate **'Splunk Environment'** to which this node should be added, and select the **'Assign'** button.

Create New Cluster

This option is aimed at standalone Unassigned Nodes. Use the built-in Wizard to create a new Splunk **Indexer** or **Search Head Cluster**.

Select the appropriate number of instances from the **'Unassigned Nodes'** panel to form either an **Indexer** (minimum of 3) or **Search Head Cluster** (minimum of 4), and choose the **'Create New Cluster'** button to reveal the start of the Wizard:

The screenshot shows two configuration fields. The first is labeled 'Environment Name' with the instruction 'Specify a name to identify your Environment.' The dropdown menu is set to 'Appliance Training Cluster'. The second field is labeled 'Available Sites' with the instruction 'Add a comma-separated list of physical or logical locations. Assigning nodes to sites will be done at a later step.' The input field contains 'site1' with a small 'x' icon to its right. Below the input field, it says 'acceptable sites: site1, site2,, site63'.

At **Step 1** of the wizard, select the appropriate Splunk **Environment Name** into which you require to create this new Splunk cluster, and add the **site** detail accordingly. This will usually consist of **'site1'**, unless you are building a multi-site cluster environment, in which case you will need to add **'site2'**, etc.

Notes

1. If the **Splunk Environment** does not yet exist for this cluster, you will need to first create one, please refer to the [Creating a 'shell environment'](#) section for details.
2. If you are wanting to add both an Indexer Cluster and a Search Head Cluster, create the Indexer Cluster first, as a Cluster Master reference and secret key is required.

Select the **'Organize Cluster'** button when done.

At **Step 2**, select the **'+ New Cluster'** button to add a new cluster.

Add an Indexer Cluster prior to that of a Search Head Cluster. Therefore, perform this as two separate tasks in two separate clusters as suggested below;

Creating an Indexer Cluster

Use a suitable **'Name'** for the new **Indexer Cluster**

Ensure that the **'Type'** is set to **'Indexer'**

Create a new **'Splunk secret key'** that will be used to authenticate the cluster members.
Select the **'Organize Nodes'** button to progress in creating the Indexer Cluster

This Environment

Appliance Training Environment
Splunk Enterprise v7.3.3 site1

Organize Clusters
Add as many clusters as needed for this environment.
Assigning nodes to clusters will be done in a later step.

Name: gemini-cluster-idx01

Type: Indexer

Splunk Secret: newSecretKey

+ New Cluster

Creating a Search Head Cluster

A Search Head Cluster can only be created once an Indexer Cluster has been provisioned. The following will also be required in order to complete this operation;

- The Cluster Master IP address
- The Indexer Cluster secret key

Use a suitable **'Name'** for the new **Search Head Cluster**

Ensure that the **'Type'** is set to **'Search Head'**

Enter a **secret key** for the Search head Cluster in the **Splunk Secret** box used for authenticating Search Heads to their cluster. Best practice dictates that this should be different from the secret key used for the Indexer Cluster.

Enter the **Cluster Master** IP address in the **Indexer Master URI** box.

Enter the **Indexer Cluster** secret key in the **Indexer Secret:** box. This is visible in its encrypted form within the **Splunk Secret** entry for the **Indexer Cluster**. Do not be tempted to copy and paste this into the Indexer Secret box. Always use the original secret key assigned to the Indexer Cluster prior to its encryption.

Note

Creation of a **Search Head Cluster** will require identification of a valid **Cluster Master**. If you are building a brand new Splunk environment, an **Indexer Cluster** will need to be provisioned prior to the **Search Head Cluster**.

This Environment

Appliance Training Cluster

 site1
Organize Clusters

Add as many clusters as needed for this environment.
Assigning nodes to clusters will be done in a later step.

Name	<input type="text" value="gemini-cluster-idx-01"/>	
Type	<input type="text" value="Indexer"/>	
Splunk Secret	<input "="" type="text" value="RENTWlpcWFBZSREOCgULBQI="/>	
Name	<input type="text" value="gemini-cluster-sh-01"/>	
Type	<input type="text" value="Search Head"/>	
Splunk Secret	<input type="text" value="shclustersecret"/>	
Indexer Master URI	<input type="text" value="10.1.5.41"/>	
Indexer Secret	<input type="text" value="indexerclustersecret"/>	

 New Cluster
Note

The **secret keys** for both the **Indexer** and the **Search Head** Clusters should be recorded and held in a secure place. These are fundamental to the successful completion of any future cluster related function.

Select the **'Organize Nodes'** button on completion.

At **Step 3**, we are presented with the following screen, allowing you to designate which nodes will make up the Indexer or Search Head Cluster, and which one will be made the Cluster Master or Deployer instance.

Available Nodes

Select unassigned nodes and assign them with the designated roles

Name	IP Address	
<input checked="" type="checkbox"/> gemini-816fdc	10.1.5.46	⋮
<input checked="" type="checkbox"/> gemini-beba10	10.1.5.47	⋮
<input checked="" type="checkbox"/> gemini-acbab2	10.1.5.44	⋮
<input checked="" type="checkbox"/> gemini-22f35f	10.1.5.45	⋮

+ Add To Cluster

Add to Indexer Clusters

gemini-cluster-idx-01

Add to Search Head Clusters

gemini-cluster-shc-01

Your Clusters

Review the configuration and change the type if needed

Name	Address	Model	Type	Status	
gemini-cluster-idx-01				Indexer	<input checked="" type="checkbox"/>
gemini-002	10.1.5.42		Cluster Master	<input checked="" type="checkbox"/>	
gemini-003	10.1.5.43		Cluster Peer	<input checked="" type="checkbox"/>	
gemini-001	10.1.5.41		Cluster Peer	<input checked="" type="checkbox"/>	
gemini-cluster-shc-01				Search Head	

Select the nodes required from the ‘**Available Nodes**’ presented, and use the ‘**+ Add To Cluster**’ button to assign them to the newly created **Search Head Cluster** listed.

Refer to the ‘**Your Clusters**’ panel to choose which node you want to assign as a **Cluster Master** or **Deployer** instance using the vertical ellipsis menu.

Your Clusters

Review the configuration and change the type if needed

Name	Address	Model	Type	Status	
gemini-cluster-idx-01				Indexer	<input checked="" type="checkbox"/>
gemini-002	10.1.5.42		Cluster Master	<input checked="" type="checkbox"/>	
gemini-003	10.1.5.43		Cluster Peer	<input checked="" type="checkbox"/>	
gemini-001	10.1.5.41		Cluster Peer	<input checked="" type="checkbox"/>	
gemini-cluster-shc-01				Search Head	
gemini-acbab2	10.1.5.44		Deployer	<input checked="" type="checkbox"/> ⋮	
gemini-816fdc	10.1.5.46		Search Head	<input checked="" type="checkbox"/> ⋮	
gemini-beba10	10.1.5.47		Search Head	<input checked="" type="checkbox"/> ⋮	
gemini-22f35f	10.1.5.45		Search Head	<input checked="" type="checkbox"/> ⋮	

Set as Deployer

Remove from Cluster

At **Step 4**, select the ‘**Locate Nodes**’ button to assign this cluster to a ‘**site**’. Highlight instances in the cluster, as shown below, and select the ‘**+ Set Site**’ button to select the site number. This will generally be ‘**site1**’ in a single site cluster arrangement. This may change to ‘**site2**’ or ‘**site3**’, etc, if you are using a multi-site cluster arrangement.

Finally, select the **'Deploy'** button to create this Cluster using the information provided.

Locate Nodes

All Nodes have been automatically located to the first site. Adjust the local assignment as needed.

+ Set Site
site1

Name	IP Address	Model	Type	Site
gemini-cluster-idx-01			Indexer	
- gemini-002	10.1.5.42		Cluster Master	📍 site1
- gemini-003	10.1.5.43		Cluster Peer	📍 site1
- gemini-001	10.1.5.41		Cluster Peer	📍 site1
gemini-cluster-shc-01			Search Head	
- <input checked="" type="checkbox"/> gemini-acbab2	10.1.5.44		Deployer	📍 site1 ⋮
- <input checked="" type="checkbox"/> gemini-916fdc	10.1.5.46		Search Head	📍 site1 ⋮
- <input checked="" type="checkbox"/> gemini-beba10	10.1.5.47		Search Head	📍 site1 ⋮
- <input checked="" type="checkbox"/> gemini-22f35f	10.1.5.45		Search Head	📍 site1 ⋮

Verification of the Splunk Environment

Select **Splunk / Environments** from the vertical menu-bar at any time to obtain an overview.

SPLUNK > Splunk Environments

Environments
 Clusters
 Nodes

+ Build Environment
+ Add Node

Name	Type	Site	Version	Contains	Last Job
Appliance Cluster 01				2 Clusters	<input checked="" type="checkbox"/> ⋮
gemini-appliance-009		📍	<input checked="" type="checkbox"/> Splunk Enterprise 7.3.4		<input checked="" type="checkbox"/> ⋮
gemini-appliance-008		📍	<input checked="" type="checkbox"/> Splunk Enterprise 7.3.4		<input checked="" type="checkbox"/> ⋮
gemini-cluster-idxcluster		Indexer		3 Nodes	<input checked="" type="checkbox"/> ⋮
gemini-appliance-001	Cluster Master	📍 site1	<input checked="" type="checkbox"/> Splunk Enterprise 7.3.4		<input checked="" type="checkbox"/> ⋮
gemini-appliance-002	Cluster Peer	📍 site1	<input checked="" type="checkbox"/> Splunk Enterprise 7.3.4		<input checked="" type="checkbox"/> ⋮
gemini-appliance-003	Cluster Peer	📍 site1	<input checked="" type="checkbox"/> Splunk Enterprise 7.3.4		<input checked="" type="checkbox"/> ⋮
gemini-cluster-shcluster		Search Head		4 Nodes	<input checked="" type="checkbox"/> ⋮
gemini-appliance-004	SHC Deployer	📍 site1	<input checked="" type="checkbox"/> Splunk Enterprise 7.3.4		<input checked="" type="checkbox"/> ⋮
gemini-appliance-007	SHC Member	📍 site1	<input checked="" type="checkbox"/> Splunk Enterprise 7.3.4		<input checked="" type="checkbox"/> ⋮

Should you see anything other than the expected output here, you may need to destroy the cluster and re-attempt addition. Verify that you have entered the correct site references, which should all be set to **'site1'** if there is to be only one site present. Also ensure that you have entered the correct IP address for the Cluster Master when setting up a Search Head Cluster.

Add to Cluster

This option is used to expand the capacity of Splunk **Indexer** or **Search Head** Clusters. This process can be used for Gemini instances or remote Splunk Environments, although in both cases, Splunk should **NOT** be installed. The installation of Splunk is taken care of during integration to a Splunk Environment.

Add an additional Indexer or Search Head to a Splunk Cluster

Highlight the required node(s) from the **Unassigned Nodes** panel of the **Splunk Environments** dashboard.

Select the **'Add to Cluster'** button to invoke a three-step Wizard to bring Unassigned Nodes into a cluster with the correct settings.

Unassigned Nodes (1)							
<input type="checkbox"/> Name	IP	Type	Splunk Software	Site	Deployment Type	Status	
<input checked="" type="checkbox"/> gemini-008	10.1.5.48	Software Appliance	Not Installed	-		✔	⋮

Note

Splunk must not be installed on the **Unassigned Node** prior to being added to an existing Splunk cluster. Gemini Central will install the correct version of Splunk to the Node when it is added to the cluster.

At **Step 1**, we need to add the Node(s) to the relevant **Splunk Environment**. If there is only one environment, this will already be highlighted, otherwise use the drop-down list for alternatives.

The site **name** will also need to be provided here. The default site name for a single-site cluster, or multi-site cluster acting as a single site, is **'site1'**.

In a multi-site environment, ensure you add all the sites, ie. **site1, site2**

Do not forget to select **'Enter'** after adding each name to register it correctly.

Add Nodes to Cluster

1 Select Environment

2 Organize Nodes

3 Locate Nodes

Environment Name
Specify a name to identify your Environment.

Appliance Splunk Cluster
▼

Available Sites
Add a comma-separated list of physical or logical locations. Assigning nodes to sites will be done at a later step.

site1 x
⌵

acceptable sites: site1,site2, ..., site63

Select the **'Organize Nodes'** button at the bottom, to reveal the following screen.

This Environment

Appliance Splunk Cluster

Available Nodes

Select unassigned nodes and assign them with the designated roles

+ Add To Cluster

There are no available Nodes.

Your Clusters

Review the configuration and change the type if needed

Name	Address	Model	Type	Status
gemini-app-idx-cluster1			Indexer	✓
gemini-008	10.1.5.48		Indexer	✓
gemini-app-shc-cluster1			Search Head	✓

At **Step 2**, select each node from the ‘**Available Nodes**’ box and use the ‘**+ Add To Cluster**’ button to assign them into the appropriate **Indexer** or **Search head Cluster**.

Verify the selection in the ‘**Your Clusters**’ panel.

Note that it only displays the nodes being added, it does not show nodes already inside the Cluster.

Select the ‘**Locate Nodes**’ button to display the following screen.

+ Set Site

✓ Name	IP Address	Model	Type	Site
gemini-cluster-idx-01			Indexer	
- gemini-002	10.1.5.42		Master Node	📍 site1
- gemini-003	10.1.5.43		Indexer	📍 site1
- gemini-001	10.1.5.41		Indexer	📍 site1
- ✓ gemini-008	10.1.5.48		Indexer	📍 site1

At **Step 3**, select the correct site name. Select the Node(s) with a checkmark, and use the ‘**+ Set Site**’ button or the vertical **ellipsis icon** to set the correct site name.

This is usually ‘**site1**’ for a single-site Cluster, or could be ‘**site2**’ if you are adding to a multi-site cluster.

Finally, select the ‘**Deploy**’ button to carry out the addition of your Node(s) as instructed.

The following screen should verify the successful addition of your Node(s) into the appropriate Splunk Cluster. The current version of Splunk being used by Gemini Central is automatically installed to the new node(s) and all relevant Splunk cluster settings applied.

gemini-app-idx-cluster1		Indexer	4 Nodes		✓	⋮
gemini-001	Cluster Master	site1	Splunk Enterprise 7.3.1		✓	⋮
gemini-003	Cluster Peer	site1	Splunk Enterprise 7.3.1		✓	⋮
gemini-002	Cluster Peer	site1	Splunk Enterprise 7.3.1		✓	⋮
gemini-008	Cluster Peer	site1	Splunk Enterprise 7.3.1		✓	⋮

Verify that Splunk has successfully integrated the new instance by observing the **Indexer Clustering** dashboard at the **Cluster Master** node, as shown below.

Indexer Clustering: Master Node

[Edit](#) [More Info](#) [Document](#)

✓ **All Data is Searchable**

3 searchable 0 not searchable
Peers

✓ **Search Factor is Met**

✓ **Replication Factor is Met**

3 searchable 0 not searchable
Indexes

Peers (3) Indexes (3) Search Heads (4)

filter 10 per page

Peer Name	Site	Fully Searchable	Status	Buckets
gemini-001	site1	✓ Yes	Up	5454
gemini-003	site1	✓ Yes	Up	5452
gemini-008	site1	✓ Yes	Up	2120

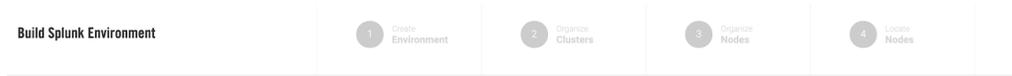
If a new **Splunk Indexer** has been added to the cluster in order to expand the capacity of Splunk, it is recommended that you perform a **Data Rebalance** function located at the **Edit** button of this dashboard. This will distribute primary data buckets from other Indexer peers across to the new one, making the entire system more efficient. Note that this process will take some time to complete and have a slight impact on overall performance whilst being completed. For this reason, it is advised to complete Data Rebalancing during hours of low usage.

Also, if you have a Splunk **Monitoring Console** in use, any additional Node(s) will initially appear as a 'New' item and will need configuring. Navigate to the **Settings / General Setup** menu of the Monitoring Console and use the **Edit** button to verify any 'New' instances listed with its correct Server Roles. Complete the process by selecting the **Apply Settings** button. An example of this dashboard is shown below.

Edit Selected Instances 25 Per Page											
i	Instance (host)	Instance (serverName)	Machine	Server roles	Custom groups	Indexer Cluster(s)	Search Head Cluster(s)	Monitoring	State	Problems	Actions
>	gemini-002	gemini-002	gemini-002	Indexer License Master		gemini-cluster-idx-01		✓ Enabled	● Configured		Edit
>	gemini-003	gemini-003	gemini-003	Indexer License Master		gemini-cluster-idx-01		✓ Enabled	● Configured		Edit
>	gemini-85693e	gemini-85693e	gemini-85693e	Indexer License Master		gemini-cluster-idx-01		✓ Enabled	● New		Edit

New Instance install - Clusters and Standalone Instances

If this is a new instance installation, a **Splunk Environment** will first need to be created using the ‘+ **Build Environment**’ button (top right of the screen). This will reveal a four-step wizard that enables the required detail to build the desired Splunk configuration.



A typical configuration includes the following stages:

- Step 1: Specify a Splunk Environment name, cluster arrangement, and Splunk binary version to be installed
- Step 2: Create a Splunk cluster for the environment.
- Step 3: Organize nodes into a cluster.
- Step 4: Specify the site name for the cluster.

Repeat the above to create additional Splunk Clusters.

For further information on this process, please refer to the [Building the Splunk Environment](#) section of this Administration Guide.

Adopting remote Splunk Instances and Clusters

Nodes that have been added and listed as ‘**Unassigned Nodes**’ may require the creation of a new **Splunk Environment**, especially if they have been added via **Gemini Agents**.

If a suitable environment does not already exist, create a ‘**shell environment**’ to act as a new repository for external Splunk clusters. Once this environment has been created simply migrate the external Nodes and Clusters into this environment.

Creating a ‘shell environment’

This process is similar to building a new environment, with the exception that we just create a ‘shell’ environment, minimum detail is therefore required.

Select the ‘+ **Build Environment**’ button from the Splunk Environments screen of the Management Center.

- **Deployment Type**

Select ‘**Deploy Multi-Use Environment**’

- **Environment Name**

Create a suitable name for the external Splunk cluster, maybe use words which locate the cluster, ie. Building 1, NY_East, Remote_SiteA. This is simply a label and can therefore include spaces etc.

- **Available Sites**

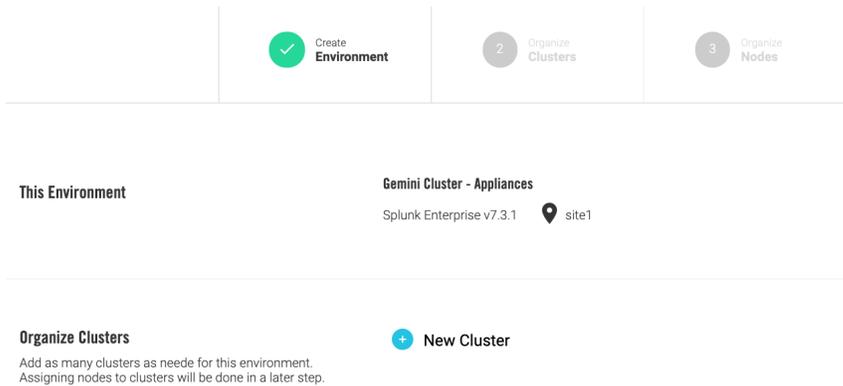
Enter the appropriate site names for this remote cluster. Typically, 'site1' is used for a single-site cluster, multi-site clusters will have additional site names. Select 'enter' to register each site here.

- **Splunk Software**

Select the version of Splunk required in this environment. Ideally, there should be only one Splunk version used in the entire environment. Note that the selected Splunk version will not take effect until the next upgrade operation.

Use the upload link provided to bring in a new Splunk binary.

Select the 'Organize Cluster' button at the bottom of the dashboard to reveal the following screen, but make no changes here.



Select the 'Organize Nodes' button to progress to the 'Available Nodes' dashboard where once again we make no changes.

Select the 'Locate Nodes' button, leave all settings unchanged, and select the 'Deploy' button to create our new Splunk Environment shell.

Return to the **Splunk Environments** dashboard to confirm the existence of a new **Splunk Environment**.

Assigning an external Splunk Cluster to an Environment

We now have a new **Splunk Environment** into which we can add our unassigned instances and Splunk Clusters.

Open the **Splunk Environments** dashboard from the **Management Center** and highlight the required **Unassigned Splunk Cluster(s)**.

Select the 'Add to Environment' button to reveal a list of available Splunk Environments in the right-hand pane. Choose the appropriate environment, and select the 'Assign' button to complete the process.



The **Splunk Environments** dashboard should now show the complete remote clusters inside our new Splunk Environment.

Remote Splunk Cluster				2 Clusters		⋮	
splunk-sh1.geminidata.com			Splunk Enterprise 7.3.3			⋮	
splunk-ds.geminidata.com			Splunk Enterprise 7.3.3			⋮	
Remote_prod_SHC_01				Search Head	4 Nodes		⋮
splunk-dep.geminidata.com	SHC Deployer	site2	Splunk Enterprise 7.3.3			⋮	
splunk-sh2.geminidata.com	SHC Member	site2	Splunk Enterprise 7.3.3			⋮	
splunk-sh3.geminidata.com	SHC Member	site2	Splunk Enterprise 7.3.3			⋮	
splunk-sh4.geminidata.com	SHC Member	site2	Splunk Enterprise 7.3.3			⋮	
Remote_prod_IDXCluster				Indexer	5 Nodes		⋮
splunk-cm.geminidata.com	Cluster Master	site1	Splunk Enterprise 7.3.3			⋮	
splunkidx-01.geminidata.com	Cluster Peer	site1	Splunk Enterprise 7.3.3			⋮	
splunkidx-02.geminidata.com	Cluster Peer	site1	Splunk Enterprise 7.3.3			⋮	
splunkidx-03.geminidata.com	Cluster Peer	site2	Splunk Enterprise 7.3.3			⋮	
splunkidx-04.geminidata.com	Cluster Peer	site2	Splunk Enterprise 7.3.3			⋮	

Assigning Splunk Standalone Nodes

Select the required target **'Unassigned Nodes'** and then select the **'Assign to Environment'** button to assign them to a relevant environment.

Assigning a new remote Indexer or Search Head to a Cluster

Highlight the required target **'Unassigned Node(s)'** and select the **'Add to Cluster'** button to invoke the 3 stage wizard assigning them to a relevant cluster.

At this stage it is important the Splunk has **NOT** been installed to the node. Splunk will be installed as part of the Cluster integration process. See below for an example of the Unassigned Node before it can be assigned to a Cluster.

Unassigned Nodes (1)							
<input type="checkbox"/> Name	IP	Type	Splunk Software	Site	Deployment Type	Status	
<input type="checkbox"/> splunk_spare	10.1.5.27	Gemini Agent	Not Installed	-			⋮

At **Step 1**, we need to add the Node(s) to the relevant **Splunk Environment**. From the drop-down list, select the appropriate 'Remote Splunk Environment'.

The site **name(s)** will also need to be provided here. The default site name for a single-site cluster, or multi-site cluster acting as a single site, is **'site1'**.

In a multi-site environment, ensure you add all the sites, ie. **site1, site2**

Do not forget to select **'Enter'** after adding each name to register it correctly.

Environment Name
Specify a name to identify your Environment.

Splunk Remote Cluster

Available Sites
Add a comma-separated list of physical or logical locations. Assigning nodes to sites will be done at a later step.

site1 x site2 x

acceptable sites: site1, site2,, site63

Select the **'Organize Nodes'** button at the bottom, to reveal the following screen.

This Environment **Splunk Remote Cluster**

Available Nodes
Select unassigned nodes and assign them with the designated roles

+ Add To Cluster

<input type="checkbox"/> Name	IP Address	Model	Type
<input checked="" type="checkbox"/> splunk_spare	10.1.5.27		⋮

At **Step 2**, select each node from the **'Available Nodes'** box and use the **'+ Add To Cluster'** button to assign them into the appropriate remote **Indexer** or **Search head Cluster**.

Verify the selection in the **'Your Clusters'** panel.

Select the **'Locate Nodes'** button to display the following screen.

Locate Nodes

All Nodes have been automatically located to the first site. Adjust the local assignment as needed.

 + Set Site

<input type="checkbox"/> Name	IP Address	Model	Type	Site
Remote_prod_SHC			Search Head	
- splunk_dep	10.1.5.25		Deployer	📍
- splunk_sh2	10.1.5.22		Search Head	📍
- splunk_sh3	10.1.5.23		Search Head	📍
- splunk_sh4	10.1.5.24		Search Head	📍
<input type="checkbox"/> splunk_spare	10.1.5.27		Search Head	📍
Remote_prod_IDXCluster			Indexer	
- splunk_cm	10.1.5.193		Cluster Master	📍
- splunkidx_01	10.1.5.190		Cluster Peer	📍 site1
- splunkidx_02	10.1.5.191		Cluster Peer	📍 site1
- splunkidx_03	10.1.5.20		Cluster Peer	📍 site2
- splunkidx_04	10.1.5.21		Cluster Peer	📍 site2

At **Step 3**, assign the correct site name. Select the Node(s) with a checkmark, and use the ‘+ Set Site’ button or the vertical **ellipsis icon** to set the correct site name.

This is usually ‘**site1**’ for a single-site Cluster, or could be ‘site2’ if working in a multi-site environment.

Finally, select the ‘**Deploy**’ button to carry out the addition of your Node(s) as instructed.

The following screen should verify the successful addition of your Node(s) into the appropriate Splunk Cluster. The current version of Splunk being used by Gemini Central is automatically installed to the new node(s) and all relevant Splunk cluster settings applied.

gemini-app-idx-cluster1		Indexer	4 Nodes		📍	⋮
gemini-001	Cluster Master	📍 site1	Splunk Enterprise 7.3.1		✅	⋮
gemini-003	Cluster Peer	📍 site1	Splunk Enterprise 7.3.1		✅	⋮
gemini-002	Cluster Peer	📍 site1	Splunk Enterprise 7.3.1		✅	⋮
gemini-008	Cluster Peer	📍 site1	Splunk Enterprise 7.3.1		✅	⋮

Verify that Splunk has successfully integrated the new instance by observing the **Indexer Clustering** dashboard at the **Cluster Master** node.

Deploy Independent Stream Forwarder

This feature has been deprecated.

Operations and Administration of Splunk

Splunk Environment Level Options

Several Splunk administrative tasks can be achieved centrally once resident in Gemini's **Management Center**. This includes both Gemini Central Splunk Environments and remote Splunk Environments controlled via Gemini Agents.

Use the vertical ellipsis menu located adjacent to the **Splunk Environment** name.

- **Rolling Upgrade** - Use this feature, to perform an upgrade of the Splunk software for all cluster members one-by-one. This will ensure that Splunk has minimal downtime. Standalone nodes in the environment will also be upgraded in parallel.
- **Delete** - Delete the entire environment. All Splunk instances in this environment will be removed. This is a destructive option.

The screenshot shows the 'SPLUNK > Splunk Environments' page. It features a sidebar with navigation icons for HOME, NODE, CLUSTER, LICENSE, and SPLUNK. The main content area has a search bar and buttons for 'Build Environment' and 'Add Node'. Below is a table with columns: Name, Type, Site, Version, Contains, and Status. The table lists several environments, and a context menu is open over the 'env-azQ8t' row, showing 'Rolling Upgrade' and 'Delete' options.

Name	Type	Site	Version	Contains	Status
env-azQ8t				2 Clusters	✓
gnmi-008		📍			✓
gnmi-009		📍			✓
gemini-cluster-cl-gG2Go	Search Head			4 Nodes	✓

Notes

The **'Delete'** option involves the removal of Splunk instances in this environment. If you want to keep the Splunk instances, first **'remove'** them before the Environment is deleted.

Upgrading Splunk using the Rolling Upgrade Feature

Important notes to consider:

- Versions of Splunk within clusters should be consistent. Splunk should not be upgraded on individual instances, hence this feature.
- The recommended Splunk upgrade procedure is followed, where only one node will go down for upgrade at any one time. This may take a long time when the environment is large.
- The rolling upgrade is only available to clusters. Standalone nodes in the same environment will be upgraded in parallel when you begin a rolling upgrade.

- Care should be taken for specific upgrades especially before a major version upgrade, ie. 6.x to 7.x has specific requirements. Please check the Splunk documentation to understand the requirements. References for Version 8.0:
 - For a search head cluster: <https://docs.splunk.com/Documentation/Splunk/8.0.0/DistSearch/UpgradeaSHC>
 - For an indexer cluster: <https://docs.splunk.com/Documentation/Splunk/8.0.0/Indexer/Upgradecluster>

Rolling Upgrade Procedure

Return to the **Splunk / Environments** dashboard and locate the Splunk Environment that you require to upgrade. This will include 'remote' Splunk Environments using Gemini Agents.

From the **vertical ellipsis** menu adjacent to the environment name, select the **'Rolling Upgrade'** option.

Enter the version of software you wish to use for the upgrade from the drop-down list, or use the **'Choose the file'** option to upload a new version of Splunk to the Management Center before selecting the **'Upgrade'** button

Upgrade Environment

Splunk Enterprise v8.0.0

Drop package here or click to [choose the file from your computer](#)
Extension must be tar.gz or tgz. Package must be Linux x64

UPGRADE CANCEL

The Splunk Environment will be placed into **'Maintenance Mode'** during the process and upgraded on a one-by-one basis in the following order:

- Cluster Master
- Indexer Peers
- Deployer
- Search Heads

Monitor the **Indexer Clustering** Dashboard from the **Cluster Master** for more detail. Once all the instances have restarted, there may be a delay until replication settles down (see below). This is perfectly normal following an upgrade.

Indexer Clustering: Master Node

✓ All Data is Searchable
⚠ Search Factor is Not Met
⚠ Replication Factor is Not Met

2 searchable 0 not searchable Peers
 3 searchable 0 not searchable Indexes

Peers (2) **Indexes (3)** Search Heads (4)

Bucket Status filter Q 10 per page

Index Name	Fully Searchable	Searchable Data Copies	Replicated Data Copies	Buckets	Cumulative Raw Data Size
.audit	✓ Yes	2	2	5	< 0.01 GB
.internal	✓ Yes	2	2	9	0.10 GB

The Gemini Central Splunk Environments dashboard, should show the new Splunk version throughout.

Manual Upgrade of a Splunk Instance

This would not normally be required, as each Splunk instance should be part of an existing Splunk Environment benefiting from the 'Rolling Upgrade' feature. However there may be some circumstances where an upgrade is required on an individual instance.

Select the **'Upgrade Splunk'** icon from the **Splunk / Daemon** dashboard of the Management Center (see below). This will reveal the **Upload and install** panel enabling the option to upload a new version of Splunk.

SPLUNK / Daemon

Splunk Service Control

⏻ Stop Splunk
🔄 Restart Splunk
🔄 Upgrade Splunk
✖ Destroy Splunk Instance

SPLUNK HOME /opt/splunk
Version Splunk 7.3.3 (build 7af3758d0d5e)

Boot-Start

Enable Splunk Boot-Start
Enable BOOTSTART in order to start Splunk daemon automatically at boot time.

Upload and Install

Upload Splunk File

Click here to choose the tarball

Choose Uploaded File

Choose...

When the upload has been completed, an **Upload and Install** panel will be presented (see below)

If this node is the **Management Center**, you may not have Splunk currently installed, so you can ignore this **'Install'** panel and just select the **'Cancel'** button.

Complete the default **admin** user credentials as required and select the **'Install'** button to continue.

Splunk Cluster Level Options

Once the Splunk Environment has been built and deployed, the following administration tasks can be achieved using the vertical ellipsis menu located adjacent to the **Splunk Cluster** name.

- **Remove from Environment** - This option will remove this cluster entirely from the current Splunk Environment. Individual Splunk instances will retain their configuration and data. The nodes can then be found in the Unassigned Nodes panel.
- **Destroy Cluster** - The entire cluster will be deleted. All the installed Splunk instances in this cluster will be removed.
- **Remove from Cluster** - This option is available from the vertical ellipsis menu for each specific node, which when used will remove it from the cluster. The Splunk instance on this node will also be removed.
- **Generate Diag** - This option is available from the vertical ellipsis menu for each node, which when instigated will; request, run, and import a **splunk diag** from the instance concerned. The diags can be stored centrally at the Management Center and made available from the **Splunk / Splunk Diag** dashboard.

Remove from Environment option

The screenshot shows the Splunk Management Center interface. The main content area is titled "SPLUNK > Splunk Environments". It features a navigation sidebar on the left with icons for HOME, NODE, CLUSTER, LICENSE, SPLUNK, and SETTINGS. The main area has a search bar and buttons for "Build Environment" and "Add Node". Below this is a table with columns: Name, Type, Site, Version, Contains, and Status. The table lists several environments and nodes. A context menu is open over the "gemini-cluster-cl-gG2Go" cluster, showing options "Remove from Environment" and "Destroy Cluster".

Name	Type	Site	Version	Contains	Status
env-azQ8t				2 Clusters	✓
gnmi-008		📍			✓
gnmi-009		📍			✓
gemini-cluster-cl-gG2Go	Search Head			4 Nodes	✓
gnmi-007	SHC Deployer	📍 site1			
gnmi-006	SHC Member	📍 site1			
gnmi-004	SHC Member	📍 site1			✓

Remove from Cluster option

The screenshot shows the Splunk Environments administration console. The table below represents the data visible in the interface:

Name	Type	Site	Version	Contains	Status
env-azQ8t				2 Clusters	✓
gnmi-009		📍	Splunk Enterprise 7.1.1		✓
gnmi-008		📍	Splunk Enterprise 7.1.1		✓
gemini-cluster-cl-vhrh7	Indexer			3 Nodes	✓
gnmi-003	Cluster Peer	📍 site1	Splunk Enterprise 7.1.1		✓
gnmi-001	Cluster Master	📍 site1	Splunk Enterprise 7.1.1		✓
gnmi-002	Cluster Peer	📍 site1	Splunk Enterprise 7.1.1		✓
gemini-cluster-cl-gG2Go	Search Head			4 Nodes	✓
gnmi-007	SHC Deployer	📍 site1	Splunk Enterprise 7.1.1		✓

Notes

Removing a node from an indexer cluster can take considerable time as it completes an elegant **'offline'** process. This process dictates that all data in the node will first be offloaded to other peer nodes, prior to it shutting down.

Removing a key role from a cluster is forbidden, e.g. Deployer and Cluster Master.

It is also forbidden to remove members should this reduce their number below the minimum requirements for the cluster.

Splunk Standalone Options

Once the Splunk Environment has been built and deployed, the following administration tasks can be achieved using the vertical ellipsis menu located on each **Splunk Node**.

- **Leave Environment** - This will remove the standalone node from the Splunk Environment. It will then be re-assigned to the Unassigned Nodes panel. Splunk will not be removed from this node using this option.

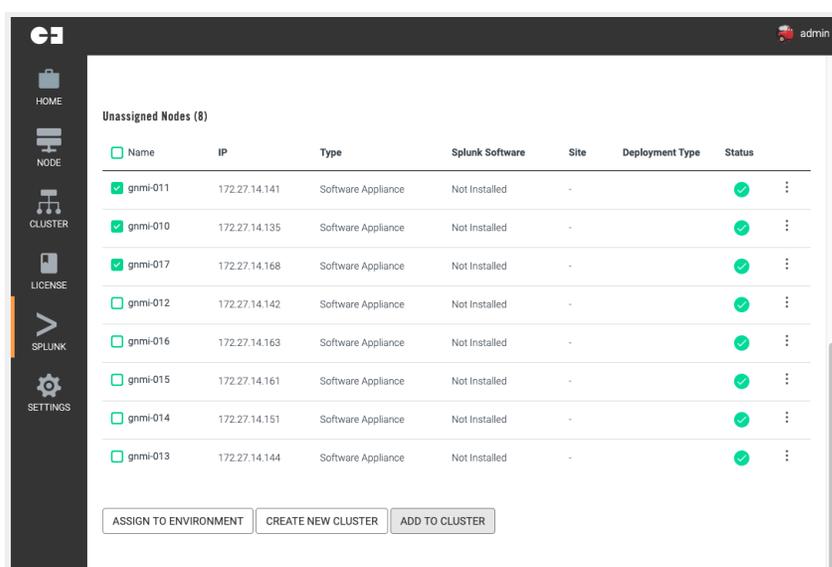
The screenshot shows the Splunk Environments administration console with a context menu open over the 'gemini-cluster-cl-vhrh7' cluster. The menu options are:

- Remove from Cluster
- Leave Environment

Unassigned Nodes Panel Options

- **Assign to environment** - This option will assign selected nodes into a specific Splunk Environment. For standalone nodes that do not have Splunk installed, Splunk installation will be performed when nodes are assigned into the environment. Useful to create ancillary Splunk instances, ie. Deployment Server
- **Create New Cluster** - This option will create a new cluster from selected nodes in the Unassigned Nodes panel. This is achieved using a four-step wizard and can be used to create either an Indexer or Search Head Cluster for addition into an existing Splunk Environment. If a new cluster is required in a new Splunk Environment, this will first need to be created.
- **Add to Cluster** - Add selected nodes into an existing cluster. Useful to expand the Indexer or Search Head cluster with additional nodes.

For full details of all the above operations, please refer to the [Assigning Unassigned Nodes and Clusters to a Splunk Environment](#) section of this Admin Guide.



Gemini Agents

A **Gemini Agent** is stand-alone software installed to remote Splunk instances in order to use the management features and convenience of Gemini Central.

Your nominated **Management Center** can be used to administer and view all Gemini Agents alongside Manage instances.

Installing a Gemini Agent to an existing Splunk environment will enable the following benefits of **Gemini Central** software.

- Ability to create new Splunk indexer and Search Head clusters
- Ability to add/remove Splunk cluster members

- Ability to install, maintain and upgrade Splunk software from a central location.
- Ability to adopt existing Splunk instances and clusters

Known Issues and Restrictions - Read before progressing

Customer installations of Splunk can vary and there can be certain configurations that will prevent the Gemini Agent from operating correctly. It is important that you read and understand these issues before you proceed with the Agent installation.

- If the **Cluster Master** and **Deployer** are sharing the same instance, adoption of the remote cluster will fail.

For further details on this issue, or if you are unsure about the suitability of your remote installation of Splunk, please contact support@gemini.com

Splunk Interface Detail

A **Gemini Agent** is designed to be installed on a target host over tcp port 4444. It can be installed by a standard user, ie. user=splunk, providing this account has elevated privileges enabling installation of the agent. The presence and location of Splunk, will be discovered automatically during installation.

Prerequisites for Installation

Supported Operating Systems

Currently, only the Linux OS is supported by the **Gemini Agent**. The supported Linux OS must meet the following criteria:

- x86 64bit architecture.
- Linux 3.x and 4.x kernel version.
- **systemd** used as system and service management

The following Linux distributions have been tested and are therefore recommended:

- RHEL 7
- CentOS 7
- Ubuntu 18.04 LTS
- Ubuntu 19.10
- Debian 10

Note

The Gemini Agent is not currently available for Windows OS.

System Requirements on the Target Host

The following resources are required by **Gemini Agent**:

- Memory: 20 MB RAM
- Disk Space: 200 MB

Minimum resource recommended for the system environment:

- Memory: 4 GB RAM
- Disk Space: 1GB (required for installation and temp space during upgrades)

- Access to port 4444 from the Management Center instance.

An account with privileged access is required for **Gemini Agent** installation and service control, this will be determined during installation.

Note

It is highly recommended that the host is running time synchronization, ie. the NTP service

Verify that your Splunk host has the above dependencies before continuing with the installation of a **Gemini Agent** on your target machine.

Installation of Gemini Agent

The Gemini Agent installation can only take place if the Management Center is operating at Version 2.7 or higher, and the **agent distribution** mechanism has been enabled.

- **Step 1** - The **agent distribution** system is enabled at the Gemini Management Center
- **Step 2** - Includes prerequisites for a successful installation of the agent on the host machine.
- **Step 3** - Describes the **agent installation** process itself.

Step 1: Management Center - Enable Gemini Agent distribution

The ability to distribute **Gemini Agents** is disabled at the Management Center by default.

To enable the **Gemini Agent distribution** feature and receive confirmation of the download URI;

Login at the terminal interface of your **Management Center** instance using the '**sbox**' account, and type the following command;

```
sbox agent --download-link
```

This will return the **download URI** link which can be used at any remote instance. Be sure to copy this for easy reference. Naming convention of the agent includes the date of release and the version number as follows; gemini-agent-<YY>.<MM>-<Ver>

```
[sbox@gemini-1c8d22 ~]$ sbox agent --download-link
https://10.2.x.x:4444/download/agent/gemini-agent-20.06-15.tar.gz
```

Step 2: Gemini Agent - Prerequisites (Splunk Host)

There are a number of checks to make to ensure that this is a suitable environment for Splunk and the Gemini Agent to work effectively.

- Ensure the **Management Center** has the **Gemini Agent** feature enabled (Step 1)
- Ensure that port 4444 is an allowed port
- Ensure that this instance conforms to an **NTP** standard or similar (Splunk best practice)

Login at your chosen Splunk host(Linux), as the **Splunk** owner account for this purpose, and **sudo** if elevated privileges are required.

Verify that the firewall on this host device is either non active or allows TCP in/out on **Port 4444** using a combination of the following commands

```
systemctl status firewalld
```

Use either of these commands to verify open ports or to verify the route through to the Management Center

```
netstat: netstat -na (check for 0.0.0.0:4444 LISTEN)
nmap utility: nc -vz <ip_address_of_management_center> 4444
```

Note: you may need to install net-tools or nmap to use these commands.

Verify that this instance is using a suitable form of **time synchronization**, vital for Splunk. Investigate whether the NTP service is running and the contents of the ntp.conf file in use and if necessary start the **ntpd** service. Use the following commands to assist with this;

```
systemctl status ntpd.service
cat /etc/ntp.conf
systemctl start ntpd.service
```

For assistance with any of the above prerequisites, please refer to your SysAdmin, or contact support@gemini.com

Step 3: Installing the Gemini Agent - Splunk Host

From the terminal of the Splunk host machine, navigate to the **/opt** directory and using **curl**, transfer the **gemini-agent** tarball from the Management Center using the download-link URI (refer to Step 1)

```
cd /opt

curl -k -O <Download-Link URI>
```

```
ie. curl -k -O
https://10.2.70.65:4444/download/agent/gemini-agent-20.06-15.tar.gz
```

Alternatively, download the **gemini-agent** tarball from the Management Center instance and use `scp` to copy it to the `/opt/` directory of your Splunk host

Unpack the **gemini-agent** tarball into a suitable destination folder, we recommend using the `/opt` directory as shown below which will unpack the Gemini Agent into the `/opt/gemini` folder.

```
tar -zxvf gemini-agent-<YY>.<MM>-<Ver>.tar.gz -C /opt
```

Run the following command with **root** privilege to complete the installation:

```
sudo /opt/gemini/agent/bin/agent start
```

Two possible output screens will follow depending on whether Splunk has already been installed on the instance.

If Splunk is either not installed or not running;

The Gemini Agent automatically detects that Splunk is non-operational. It proceeds to create a gemini service on the instance, and then prompts for further information regarding Splunk;

```
+ Generate UUID...
+ Creating SystemD service...
+ Configuring...
Has Splunk installed? [yes]:
```

- **Has Splunk installed? [yes]:**

Select 'no' as a response here.

- **Installed \$SPLUNK_HOME? [/opt/splunk]:**

Select 'enter' to confirm that this will be the Splunk install directory, or provide an alternative.

If Splunk is already installed and running;

The Gemini Agent detects the presence and location of Splunk automatically and prompts only for confirmation of the required admin account and password.

```
+ Configuring...
+ Splunk is running by splunk
+ Splunk home is /opt/splunk
Splunk admin user [admin]:
```

- **Splunk admin user [admin]:**
Select 'enter' to confirm the user 'admin' or provide another account with admin rights.
- **Splunk admin password:**
Type the password for the above account.

On receipt of a valid password, the screen should resemble the example below.

```
Splunk admin password:  
+ Configuration set successfully  
+ Gemini Agent is running.
```

Gemini Agent - CLI options at the Splunk host

agent status

If at any time you want to verify whether the **Gemini Agent** service is active, use the following command;

```
sudo /opt/gemini/agent/bin/agent status
```

A typical response from this command would be the following message

```
+ Gemini Agent is running.
```

agent --version

If at any time you wish to know which **Gemini Agent** version is active on this instance, run the following command;

```
sudo /opt/gemini/agent/bin/agent --version
```

The output will return the date and version of the Agent in the format: **<YY>.<MM>-<Version>**

agent restart

If you wish to restart the existing Gemini Agent service, run the following command;

```
sudo /opt/gemini/agent/bin/agent restart
```

agent configure

If you wish to go through the initial Gemini Agent configuration script again, for instance if the local Splunk admin password has been changed, run the following command;

```
sudo /opt/gemini/agent/bin/agent configure
```

agent uninstall

If for any reason you need to uninstall the **Gemini Agent** use the following command. Note that in order to upgrade the Gemini Agent it is first required that the existing Gemini Agent is first uninstalled.

```
sudo /opt/gemini/agent/bin/agent uninstall
```

agent stop/start

If you need to stop or start the agent manually, use the following commands;

```
sudo /opt/gemini/agent/bin/agent stop  
sudo /opt/gemini/agent/bin/agent start
```

The above process is also applicable to host machines without Splunk already installed. Splunk will be installed with the attributes specified during the Gemini Agent configuration.

When **Gemini Agents** have been installed in your host Splunk environment successfully, it is possible to register these as '**Unassigned Nodes**' in the **Splunk Environments** section of the **Management Center**.

This can be achieved using the '**+ Add Node button**' and adding them as a; standalone node, Indexer Cluster, or Search Head Cluster, using one of the supplied **JSON** or **CSV Manifest** templates.

For full details on this process, refer to the [Adopting external Splunk Instances and Clusters](#) section.

For details on other features and functions available for Splunk remote instances and clusters, such as upgrading Splunk, refer to the [Operations and Administration of Splunk](#) section.

Gemini Agent - Troubleshooting

For any local issues regarding the Gemini Agent on the Splunk host, the CLI commands above should help.

If required, more assistance can be obtained from log files in the `/opt/gemini/agent/admin/var/log` directory.

For issues regarding the ingestion of Splunk Clusters using JSON Manifest files, locate the **FATAL.log** file found on the **Management Center** instance in the `/var/log/gemini/admin/` directory.

Upgrading the Gemini Agent

Gemini Agents may need to be upgraded when new features are required. Gemini Data will issue new Gemini Agent binaries from time-to-time announced in our [Support Portal](#) and can be obtained on request.

On receipt of a new agent binary, navigate to the **Cluster / Manage Nodes** menu at the **Gemini Management Center**.

Scroll down to the last section titled, **Gemini Agent Binary**, and select the **'+ Upload Agent Binary'** button to locate and add the binary file to the Management Center for distribution.

Once the new binary is present on the Management Center, the following process should be carried out at each remote Splunk node.

Step 1: Login to the console of your remote Splunk agent host using SSH.

Stop and uninstall the existing Gemini Agent using the following command:

```
sudo /opt/gemini/agent/bin/agent uninstall
```

A confirmation screen should follow to confirm that the agent has been removed;

```
Would you like to uninstall Gemini Agent? [y/N]:
y
+ Stopping Gemini Agent...
+ Disabling service...
+ Removing Gemini Agent...
+ Gemini Agent has been uninstalled.
```

Step 2: At the `/opt` directory, use `'curl'` to obtain the latest binary from the **Management Center**.

Substitute the latest `<agent_binary>` with your filename in the following command;

```
cd /opt

curl -k -O
https://<Management_center>:4444/download/agent/<agent_binary>

ie. curl -k -O
https://10.2.70.65:4444/download/agent/gemini-agent-20.06-15.tar.gz
```

Step 3: Unpack the **gemini-agent** tarball into a suitable destination folder, we recommend the `/opt` directory as shown below which will unpack the Gemini Agent into the `/opt/gemini` folder.

```
tar -zxvf gemini-agent-<agent-version>.tar.gz -C /opt
```

Step 4: Run the following command with **root** privilege to complete the installation:

```
sudo /opt/gemini/agent/bin/agent start
```

The Gemini Agent detects the presence and location of Splunk automatically, and prompts only for confirmation of the required admin account and password.

```
+ Configuring...
+ Splunk is running by splunk
+ Splunk home is /opt/splunk
Splunk admin user [admin]:
```

- **Splunk admin user [admin]:**
Select 'enter' to confirm the user 'admin' or provide another account with admin rights.
- **Splunk admin password:**
Type the password for the above account.

On receipt of a valid password, the screen should resemble the example below.

```
Splunk admin password:
+ Configuration set successfully
+ Gemini Agent is running.
```

Refer to the [Installation of Gemini Agent](#) section for more details if required.

The latest details and information on Gemini Agents can be obtained from;

<http://support.geminidata.com/docs>

support@geminidata.com

Settings

The **Settings** menu allows you to view important information and control how Gemini Central behaves.

It also controls Gemini software upgrades, handles authentication options, including password policy enforcement, and enables reboot or shutdown operations.

- **System Admin** - How to add an SSL cert, complete backups, and generate diags
- **System Upgrade** - Instructions to install Gemini Central upgrade packs
- **Information** - Version information on Gemini software and hardware
- **Authentication** - Control User accounts, User Roles, LDAP and SSO settings.
- **Password Policy** - Enables the setting of specific password criteria.
- **HTTP Proxy** - Allows the use of a proxy server for specific tasks.
- **Login Banner** - Allows the sending of a broadcast message to all connected users
- **Reboot & Shutdown**

System Admin

- Admin Web
- Backup and restore
- Diagnose

Admin Web

Allows the installation of a 'custom SSL cert' which can be used to comply with your existing enterprise security policy.

SETTINGS › System Admin › Admin Web

Admin Web Settings

 Install Custom SSL Certificate

Listening Port

Language

SAVE

To use a SSL certificate from an external PKI, select the '**Install Custom SSL Certificate**' button.

Install Custom SSL Certificate

SSL Private Key

SSL Certificate Chain

Add the complete Certificate Chain including the Root-, Intermediate- and Server Certificates in PEM format using the following order:

1. Root Certificate (optional)
2. Intermediate Certificate(s) (optional)
3. Server Certificate

Paste the **Private Key** as a Base64 encoded DER - PEM certificate - to the '**SSL Private Key**' field, and the certificate in the field below, again PEM formatted.

Certificate Chain is supported, including **Root** and **Intermediate** Certificates of related Certificate Authorities. To use this option, paste the entire chain into this area in the correct order as shown below:

1. Root Certificate (if present)
2. Intermediate Certificate (if present)
3. Server Certificate

Note

Ensure that the passphrase is removed from the private key.

Select the '**Apply**' button to install the certificates.

The Manage web interface will restart, and the new certificate will be presented. In some cases, it will be necessary to refresh the browser window (F5).

For security reason, the following certificate guidelines are recommended;

- Certificate length; at least 2048, preferably 4096.
- Key pairs are generated using AES256.
- Signed with SHA-2(SHA-256 or SHA-384), no SHA-1, no MD5

Backup & Restore

It is important that the configuration detail and license information is secured with a backup. This is especially important following initial installation or following an upgrade.

Use the '**Download Backup File**' button to complete the backup. We would recommend the (**Select All Packages**) option.

SETTINGS › System Admin › Backup & Restore

Download Backup File

- (Select All Packages)
- System Settings
- License
- Splunk Manager
- Other Settings

[DOWNLOAD BACKUP FILE](#)

Upload & Restore

Click here to choose file.

 Upload

Diagnose

Use this option when directed by Gemini Support. This will give the support team a full view of the Gemini Central environment and status aiding a speedy resolution of a support issue.

If Manage is installed on a Dell Appliance, a separate Dell diagnostic file will also be available to send to support.

SETTINGS › System Admin › Diagnose

System Diagnose Files

 Generate System Diagnostic File

Hardware Diagnostic File

 Create Dell System E-Support Tool Report

System Update

Use this option to upgrade your Manage software with Update Packs or full Manage Upgrades, as directed by Gemini Support.

Please contact Gemini Support, support@geminidata.com for confirmation of suitability of any Update pack or Manage Upgrade or any other questions or issues related to updating Manage.

Updating older versions of Manage to the latest version should definitely be referred to Gemini Support as there are certain upgrade paths that need to be followed.

A full history of Upgrade Packs installed to this instance will be listed on this dashboard.

Information

Information displayed here details software and hardware information from your Gemini Central instance.

- Software
- Hardware
- Listen Port
- Audit Report

Software

Here you may review the currently installed software versions of

- Gemini Central version
- Linux kernel version
- Java version

SETTINGS > Information > Software

Software Version

Appliance	2.7-266
Linux Kernel	3.10.0-1062.4.3.el7.x86_64
JAVA	java version "1.8.0_192" Java(TM) SE Runtime Environment

Environment Variable

JAVA_HOME	/usr/lib/jvm/jre/
JAVA Directory	/usr/lib/jdk1.8.0_192

Hardware

Detailed hardware information on your Gemini instance hardware includes:

- CPU
- Memory
- NIC
- Chassis

SETTINGS > Information > Hardware

CPU

Model Name	Intel(R) Xeon(R) Gold 6130 CPU @ 2.10GHz
Architecture	x86_64
CPU(s)	12
Core(s) per socket	1

Memory

Virtual Machine

Network Interface Controllers

eth0	VMware VMXNET3 Ethernet Controller (rev 01) Capacity: 10000baseT
------	--

Chassis

Service Tag	VMware-42 14 bc 66 88 13 51 7d-b7 f9 93 54 f1 1c ec e3
Model Name	vmware
UEFI	False

Listen Port

This shows the entire list of ports that are currently open on the Manage instance for protocols specified,

SETTINGS > Information > Listen Port

Listening Ports

Protocol	Local Address	Port
TCP	*	443
TCP	*	4444
TCP	*	8797
TCP	*	35273
TCP	*	9321
TCP	127.0.0.1	8686
TCP	*	111
TCP	*	22
TCP	*	8888
TCP	::	8889
TCP	::	41455
TCP	::	111
TCP	::	22
TCP	::ffff:127.0.0.1	59000

This information may be requested by Gemini Support to assist with your support case.

Audit Report

The Audit Report tab allows you to create downloadable audit reports which will include a list of all the libraries that Gemini Central use together with their version.

The current listening ports will also be included within the report.

Authentication

Gemini Central offers administration access either via the web console or by running CLI commands following a successful SSH login.

- Manage User
- User Permissions
- LDAP
- Single sign-on

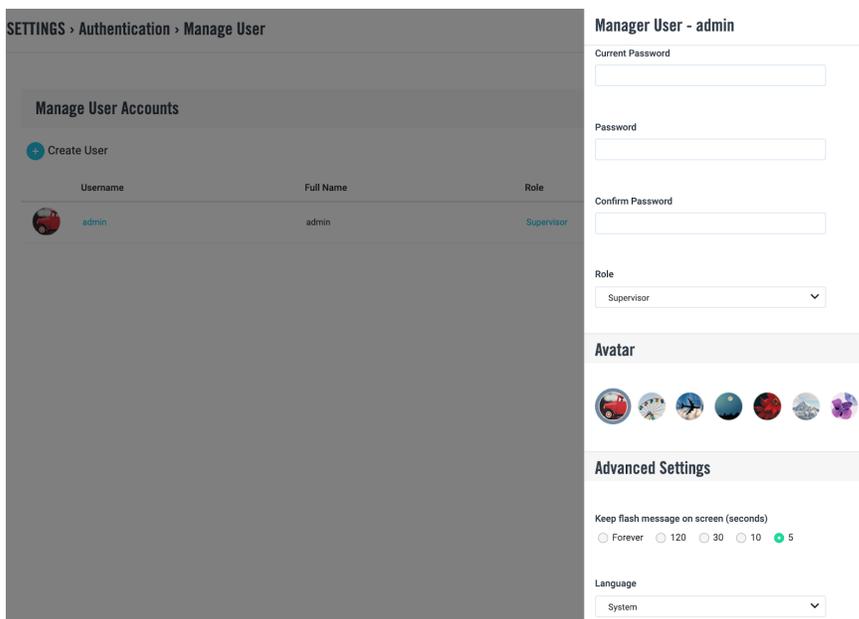
Manage User

Controls access to the Manage web interface. Configure as many local 'Admin' (Supervisor) accounts as required.

Other Roles such as **Splunk Users**, **Splunk Admins** and **Manage Users** can also be created here.

Passwords will need to comply with the **Password Policy** (see section below)

Language choices for users consist of English, Chinese, German and Japanese.



SETTINGS > Authentication > Manage User

Manage User Accounts

[+ Create User](#)

Username	Full Name	Role
admin	admin	Supervisor

Manager User - admin

Current Password

Password

Confirm Password

Role

Avatar



Advanced Settings

Keep flash message on screen (seconds)

Forever 120 30 10 5

Language

User Permissions

This enables you to add to the existing **User Role** templates available in Gemini Central. Each of these built-in Roles has unique user permissions set to control behavior and access across Manage. Additional Roles can be created to inherit existing permissions using these templates, or with customized permissions.

SETTINGS > Authentication > User Permissions

User Role		
+ Create Role		
User Role	Default	Applied Users
Cloudera Admin	Yes	0
Group Admin	Yes	0
Manager Admin	Yes	0
Gemini Manager Auditor	Yes	0
Manager User	Yes	0
Splunk Admin	Yes	0
Splunk User	Yes	0
Supervisor	Yes	1

By selecting an existing **User Role**, a panel will be opened allowing you to edit permissions for various aspects of Manage.

Permissions

LICENSE

License Status	<input checked="" type="radio"/> None	<input type="radio"/> Read	<input type="radio"/> Write
Remote Licenses	<input checked="" type="radio"/> None	<input type="radio"/> Read	<input type="radio"/> Write
Inventory	<input checked="" type="radio"/> None	<input type="radio"/> Read	<input type="radio"/> Write
License Server	<input checked="" type="radio"/> None	<input type="radio"/> Read	<input type="radio"/> Write

SYSTEM

System Time	<input checked="" type="radio"/> None	<input type="radio"/> Read	<input type="radio"/> Write
Timezone	<input checked="" type="radio"/> None	<input type="radio"/> Read	<input type="radio"/> Write
Name	<input checked="" type="radio"/> None	<input type="radio"/> Read	<input type="radio"/> Write
Network	<input checked="" type="radio"/> None	<input type="radio"/> Read	<input type="radio"/> Write
FTP	<input checked="" type="radio"/> None	<input type="radio"/> Read	<input type="radio"/> Write

SAVE
CANCEL

Notes regarding the setting of permissions:

- If Gemini Central is upgraded from versions prior to Manage 2.4, existing users will be granted 'Supervisor' rights following the upgrade.
- One role needs to be applied to each Manage user.
- If using LDAP authentication, specify a default role for each LDAP resource.
- The Roles of 'Supervisor' and 'Manage User' are default roles that cannot be removed.
- A Role may be deleted only if there are no users assigned to that role.
- Permission of 'Read' allows the user read-only access to the status and settings, no other actions are permitted.
- A permission of 'Write' also implies 'Read'.

LDAP

If desired a link to an existing LDAP resource may be established here to support LDAP authentication for the Manage interface. When LDAP resources are configured successfully, users will be able to login to Gemini Central using their LDAP account.

Note that Gemini Central connects using simple BIND and Search/Bind requests and LDAP server access is only used for authentication not for accessing roles or permissions.

Use the toggle slider to enable LDAP authentication if required.

SETTINGS › Authentication › LDAP

LDAP Authentication Control

Enable LDAP Authentication
Authenticate user by LDAP server.

LDAP Resource

There is no LDAP resource currently configured.

[+ Add LDAP Resource](#)

Configure an LDAP resource using **Simple BIND**, as follows;

1. Select the '**Add LDAP Resource**' button to create a new LDAP resource.
2. Configure the **Host** and **Port**, use FQDN (not the IP address). Enable SSL if needed.
3. Select **Simple BIND**.
4. Carefully configure the **User BaseDN**. The BaseDN should be able to locate users who can access Manage. You may want to create a special LDAP group for this.
5. Configure the **Login Attribute**. This should be a real attribute that currently exists within the LDAP directory and can be used for a Manage account username, ie. uid, CN or name.

Configure **Role**. Select a default User Role for new Manage users to be assigned during LDAP authentication.

Note

Please ensure that your BaseDN includes only those users who should have access to the instance administration screens.

The following example shows an LDAP connection profile for a group of Gemini Central users;
New Feature: The addition of the ‘**Validate Connection**’ button will be a welcome addition to this process

Create LDAP Resource

LDAP_Manage_users

Connection

Host
ad_server1.acme.com

Port
389

Create LDAP Resource

SSL

LDAP Authentication ▾

Simple Bind Search/Bind

User Name
admin

Password
.....

User DN Template
uid=\$(username),ou=it_dept, dc=acme, dc=com

User Base DN
ou=it_dept, dc=acme, dc=com

Login Attribute
uid

Role
Splunk User ▾

VALIDATE CONNECTION

Configure an LDAP resource using the alternative **Simple BIND**, as follows;

1. Select the ‘**Add LDAP Resource**’ button to create a new LDAP resource.

2. Configure the **Host** and **Port**, use FQDN (not IP address). Enable SSL if needed.
3. Select **Search/BIND**.
4. Configure the **Lookup DN** and **Lookup Password**. This is used to login to the LDAP server and fetch the LDAP trees. The whole LDAP tree will be cached on the system for further use.
5. Carefully configure the **User BaseDN**. The BaseDN should be for users who can access Manage. You may want to create a special LDAP group for this.
6. Configure the **User Search Filter**.
7. Configure the **Login Attribute**. The login attribute should be a real attribute that exists within the LDAP directory and can be used for a Manage account username, ie. uid, CN or name.
8. Configure **Role**. Select a default User Role for new Manage users to be assigned during LDAP authentication.
9. Select the '**Validate Connection**' button for a real-time connection check to verify the settings.

Single Sign-on (SSO)

Gemini Central Single Sign-on (SSO) provides the ability to use an HTTP Reverse Proxy Server to handle Manage authentication. Once a user is successfully logged into the proxy, they can access the Manage interface without having to login directly.

Gemini Central expects a specific HTTP Request Header from the Reverse Proxy. The name of the HTTP Header field can be configured in the Single Sign-On configuration screen.

Select the '**Automatically Create User**' option when the username from an authenticated request through the Reverse Proxy does not exist as a local Manage admin user. If this option is not selected and the username from the request does not exist in Manage, the request will fail and the Manage login prompt will be shown.

For added security, authentication requests can be restricted to a specific set of IP addresses, and only requests having the Username Field in the HTTP Header.

SETTINGS › Authentication › Single Sign-On

Enable Single Sign-On

SSO Parameters

Username Field in HTTP Header

It only accept alphabet, numeric and hyphen.

Automatically Create User

Yes No

Trusted Remote Address

Leave blank to trust all

UPDATE

Once SSO is authenticated, it will take precedence over other authentication methods such as LDAP.

Password Policy

Setting a **Password Policy** allows you to enforce password requirements to meet your security needs including complexity and duration before a due reset.

SETTINGS › Password Policy

Password complexity

- Minimum password length:
- Prevent password reuse - Number of characters must be changed.
 - Require at least one number.
 - Require at least one upper-case character.
 - Require at least one lower-case character.
 - Require at least one non-alphanumeric character.
 - Require less than 3 repeating consecutive characters.
 - Require less than 3 repeating characters in the same character class.

Password Duration

- Password expiration period (in days)
- Number of used passwords to remember

APPLY

Note

Password complexity applies to both Web admins and OS users.
Password duration only applies to OS users.

HTTP Proxy

This allows you to set a Proxy Server for specific services, ie. to connect to the Tableau website.

SETTINGS > HTTP Proxy

HTTP Proxy Control

Enable HTTP Proxy
This will enable/disable the usage of a Proxy for outgoing connections to download packages for Integration Center and Cloudera.

Additional Settings

Protocol

http

Proxy Server

Hostname or IP Address of Proxy Server

example.com

Port

Port Number of Proxy Server

80

Username

Username for authenticated Proxy Requests

Password

Password for authenticated Proxy Requests

SAVE

VERIFY

Login Banner

Use this to enable and edit the banner message presented to users when accessing either the instance web interface or access via SSH.

This could be used to send out a broadcast message about upcoming maintenance for example.

Simply type your message in the box provided. Do not forget to remove the message when it is no longer appropriate.

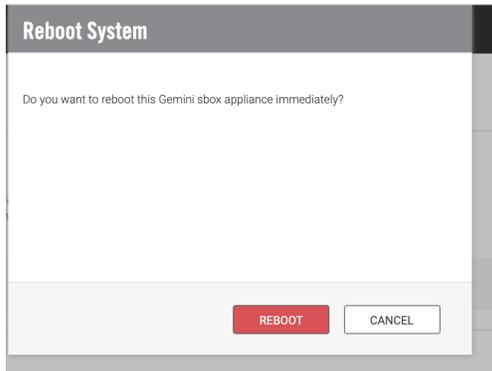
Prompt Login Banner Message
Display message when OS user log in system through console and terminal session.

Banner Content

SAVE

Reboot

Allows you to reboot your Gemini instance immediately. Selecting this option will bring up the following splash screen, enabling you to **'cancel'** if this was done in error.

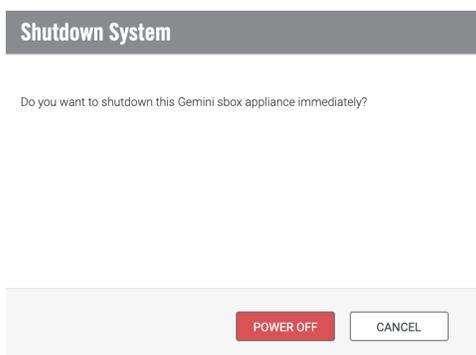


Shutdown

Allows you to shutdown and power off your Gemini instance immediately.

Note that Splunk services will be stopped prior to shutdown in order to prevent unexpected errors.

Selecting this option will bring up the following splash screen, enabling you to **'cancel'** if this was done in error.



CLI Commands

Gemini Central supports a series of shell commands that can be executed locally from a terminal.

Authentication using one of two built-in User accounts (see [Default Passwords for CLI Operations](#) for details) is required in order to use these commands over SSH.

We recommend that you use the **'sbox'** OS user account for the majority of the commands listed here.

For **Splunk** specific issues, you should use the **'splunk'** OS user account.

Once authenticated, simply type **'sbox'** at the command prompt for an overview of the **top-level** commands available in Gemini Central as shown below.

```
sbox
```

Usage: sbox [OPTIONS] [COMMAND]

[OPTIONS]

--version	Show Gemini Central version.
--service-tag	Display service tag of server.
--model	Display model name of server.

[COMMAND]

help	Display help for each of command.
admin	Configure Gemini Central.
agent	Manage Gemini Agent.
config	Configure appliance.
diag	Generate diag information.
isf	Control Independent Stream Forwarder service.
license	Manage Gemini Central license.
network	Setup for specific network interface.
server	Manage appliance server to reboot or shutdown.
service	Control services.
splunk	Manage Splunk deployed on current node.
system	Display system information and patch.

The help Command

To display more detail regarding the top-level commands use the **'help'** command;

```
sbox help
```

Usage: sbox admin [OPTIONS]

[OPTIONS]

<code>--reset-password</code>	Reset password of web UI admin user.
<code>-set-password</code>	Set password for web UI admin user.
<code>--gen-ssl</code>	Regenerate SSL keys for web server.
<code>--installed-packages</code>	Display installed packages in web server.
<code>--skip-wizard</code>	Skip the setup wizard on access.
<code>--restart</code>	Restart sbox-admin service.

Usage: sbox agent [OPTIONS]

[OPTIONS]

<code>--download-link</code>	Enable Agent distribution and display the download URI.
------------------------------	---

Usage: sbox config [OPTIONS]

[OPTIONS]

<code>--accept-eula</code>	To accept End User License Agreement.
<code>-hostname</code>	Set hostname.
<code>-timezone</code>	Set system timezone.
<code>--trial-license</code>	Start trial license for web admin.
<code>-license-file</code>	Import license file. Argument: <path_to_license_file>
<code>-license-server</code>	Indicate remote license server. Argument: <server_host>:<token_string>

Usage: sbox diag [OPTIONS]

[OPTIONS]

<code>--generate</code>	Generate diagnostics zip file.
-------------------------	--------------------------------

Usage: sbox license [OPTIONS]

[OPTIONS]

<code>--trial-license</code>	Grant trial license for Gemini Central.
<code>--revoke-trial</code>	Revoke trial license.
<code>-license-file</code>	Import license file. Argument: <path_to_license_file>
<code>-license-server</code>	Indicate remote license server. Argument: <server_host>:<token_string>

Usage: sbox network [OPTIONS]

[OPTIONS]

<code>--reset</code>	Reset all network interfaces to their default value.
<code>-nic</code>	Setup specific network interface, required for the options below
<code>--disable</code>	Disable specific NIC, when given it ignores <code>--dhcp</code> , <code>-ip</code> , <code>-netmask</code> and <code>-gateway</code> options.
<code>--dhcp</code>	Config the specific NIC as DHCP.
<code>-ip</code>	Set IP address for specific NIC.
<code>-netmask</code>	Set subnet mask for specific NIC, required when set IP address.
<code>-gateway</code>	Set gateway on specific NIC. (optional param).

Usage: sbox service [OPTIONS]**[OPTIONS]**

<code>--reboot</code>	Reboot server.
<code>--shutdown</code>	Shutdown server.

Usage: sbox service [OPTIONS]**[OPTIONS]**

<code>--status</code>	Display status of services.
<code>--listen-port</code>	Display the listening ports of services.
<code>--restart</code>	Restart all services.

Usage: sbox splunk [OPTIONS]**[OPTIONS]**

<code>--kill</code>	Remove Splunk instance.
<code>--reset_environments</code>	Clean up Splunk environment settings.
<code>--backup_setting <file></code>	Backup Splunk settings.
<code>--restore_setting <file></code>	Restore Splunk settings.

Usage: sbox system [OPTIONS]**[OPTIONS]**

<code>--info</code>	Display system information.
<code>-patch <file_path></code>	Apply patch file with <code><file_path></code> .

Commands for initial setup

The network Command

The '**sbox network**' command allows you to complete the basic network settings, including both DHCP and static network settings.

We recommend that you create a permanent static IP address for Gemini Central.

If necessary, identify the name of the device network interface using the following command at the terminal:

```
ip a
```

Output from this command is shown below to reveal in this case; an **interface name** of '**nic0**', and current **ip address** of **192.168.1.100**.

```
[sbox@sboxnode1 ~]$ ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: nic0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP qlen 1000
    link/ether 08:00:27:9e:96:c0 brd ff:ff:ff:ff:ff:ff
    inet 192.168.1.100/24 brd 192.168.1.255 scope global dynamic nic0
        valid_lft 68091sec preferred_lft 68091sec
```

If you wish to configure or change the **static network** settings, use the following command;

```
sbox network -nic <network_interface_name> -ip <ip_address>
-netmask <netmask> -gateway <gateway_ip>
```

An example of setting the Network Interface '**nic0**' to create a static ip address is shown below;

```
sbox network -nic nic0 -ip 192.168.1.100 -netmask 255.255.255.0
-gateway 192.168.1.1
```

Alternatively, if you wish to configure the network interface to use **DHCP** use the following command;

```
sbox network -nic <network_interface_name> --dhcp
```

For security reasons, consider omitting the **-gateway** param for a host-only IP address that has no access to the Public environment.

A summary of the **'network'** command and options is given below, and can be shown at the terminal at any time by typing **'sbox network'**.

Usage: sbox network [OPTIONS]

[OPTIONS]

--reset	Reset all network interface to default value.
-nic	Setup specific network interface, required for below options.
--disable	Disable specific NIC, when given it ignores --dhcp, -ip, -netmask and -gateway options.
--dhcp	Config the specific NIC as DHCP.
-ip	Set IP address for specific NIC.
-netmask	Set subnet mask for specific NIC, required when set IP address.
-gateway	Set gateway on specific NIC. (optional param)

The config Command

In addition to the **'network'** command, the **'config'** command provides additional configuration at deployment, such as setting the instance hostname and timezone. This would normally be set up in the **Manage UI**, but this offers an alternative.

It can also be used to apply a Licence file, and set a Licence Server.

A summary of the **'config'** command and options is given below, and can be shown at the terminal at any time by typing **'sbox config'**.

Usage: sbox config [OPTIONS]

[OPTIONS]

--accept-eula	To accept End User License Agreement.
-hostname	Set hostname.
-timezone	Set system timezone.
--trial-license	Start trial license for web admin.
-license-file	Import license file. Argument: <path_to_license_file>
-license-server	Indicate remote license server. Argument: <server_host>:<token_string>

Note

Note that the license file has to be uploaded to the instance before it can be applied.

The admin Command (setup options)

This command can be used to control the **Manage web interface** administration. This includes resetting a custom installed SSL certificate or disabling the web-based setup wizard when all settings have been applied using the CLI.

A summary of the 'admin' command and options is given below, and can be shown at the terminal at any time by typing 'sbox admin'.

Usage: sbox admin [OPTIONS]

[OPTIONS]

--reset-password	Reset password of web admin user.
--set-password	Set password for web admin user.
--gen-ssl	Regenerate SSL keys for web server.
--installed-packages	Display installed packages in web server.
--skip-wizard	Skip the setup wizard on access.
--restart	Restart sbox-admin service.

Note

Use these commands with caution as some may restart the web UI.

The agent Command - Management Center Node

Gemini Agents are a relatively new feature of Gemini Central available from Version 2.7 and above. For more details, refer to the '[Gemini Agents](#)' section of this Admin Guide, or the separate '**Gemini Central - Gemini Agents Quick Start Guide**'.

The ability to distribute **Gemini Agents** is disabled at the Management Center by default.

To enable the **Gemini Agent distribution** feature and receive confirmation of the download URI;

Login at the terminal interface of your **Management Center** instance using the 'sbox' account, and type the following command;

```
sbox agent --download-link
```

This will return the **download URI** link which can be used at any remote instance. Be sure to copy this for easy reference. Naming convention of the agent includes the date of release and the version number as follows; gemini-agent-<YY>.<MM>-<Ver>

```
[sbox@gemini-1c8d22 ~]$ sbox agent --download-link
https://10.2.x.x:4444/download/agent/gemini-agent-20.06-15.tar.gz
```

Alternatively, please obtain the latest **Gemini Agent** binary from support@gemini.com

*The following ‘agent’ commands are relevant to the **remote** Splunk host running the **Gemini Agent** binary and these should be run from a Terminal session local to the Splunk installation.*

agent status

If at any time you want to verify whether the **Gemini Agent** service is active, use the following command;

```
sudo /opt/gemini/agent/bin/agent status
```

A typical response from this command would be the following message

```
+ Gemini Agent is running.
```

agent --version

If at any time you wish to know which **Gemini Agent** version is active on this instance, run the following command;

```
sudo /opt/gemini/agent/bin/agent --version
```

The output will return the date and version of the Agent in the format: **<YY>.<MM>-<Version>**

agent restart

If you wish to restart the existing Gemini Agent service, run the following command;

```
sudo /opt/gemini/agent/bin/agent restart
```

agent configure

If you wish to go through the initial Gemini Agent configuration script again, for instance if the local Splunk admin password has been changed, run the following command;

```
sudo /opt/gemini/agent/bin/agent configure
```

agent uninstall

If for any reason you need to uninstall the **Gemini Agent** use the following command. Note that in order to upgrade the Gemini Agent it is first required that the existing Gemini Agent is first uninstalled.

```
sudo /opt/gemini/agent/bin/agent uninstall
```

agent stop/start

If you need to stop or start the agent manually, use the following commands;

```
sudo /opt/gemini/agent/bin/agent stop  
sudo /opt/gemini/agent/bin/agent start
```

The above process is also applicable to host machines without Splunk already installed. Splunk will be installed with the attributes specified during the Gemini Agent configuration.

Commands for Information Gathering

The version Operator

Used with the `sbox` command, this operator will display the currently installed version of Gemini Central

```
sbox --version
```

The model Operator

Used with the `sbox` command, this operator will acquire the model of this instance

```
sbox --model
```

On virtualized environments or on public clouds, the returned string represents the Hypervisor type (ie. vmware) or on Amazon EC2, “HVM domU” will be returned.

The service-tag Operator

Used with the `sbox` command, this operator returns the unique service tag of the appliance. This could be useful during a Support issue, and Customer Support may ask for this value when contacted.

```
sbox --service-tag
```

Notes

Please ensure you include all these details, when opening a Customer Support Request.
This information is automatically included within the **Diagnostic Report** created using the Manage web UI.

The admin Command (installed-packages operator)

The ‘**admin**’ command can be used with the ‘**--installed-packages**’ operator to produce a list of installed packages and their versions. This could be useful for audit purposes.

Use the command with the ‘**--installed-packages**’ operator as shown below;

```
sbox admin --installed-packages
```

The screen below shows a typical output from this command;

```

[sbox@gemini-08800277c2740 ~]$ sbox admin --installed-packages
=====
Installed package
=====
sbox-core-2.2-126.x86_64
sbox-uno-2.2-62.x86_64
sbox-hadoop-2.2-228.x86_64
sbox-driver-2.2-21.x86_64
sbox-pepsi-2.2-27.x86_64
sbox-admin-2.2-116.x86_64
sbox-license-2.2-50.x86_64
sbox-console-2.2-63.x86_64
sbox-splmgr-2.2-64.x86_64
sbox-os-2.2-68.x86_64
sbox-failover-2.2-29.x86_64
sbox-solution-2.2-124.x86_64
sbox-theme-2.2-29.x86_64

```

The service Command (status operator)

The **'service'** command can be used with the **'status'** operator to obtain the status of Gemini components.

Use the command with the **'status'** operator as shown below;

```
sbox service --status
```

The screen below shows a typical output from this command;

```

[sbox@gemini ~]$ sbox service --status
=====
Service                Status    Boot-start
=====
sbox-pepsi              True     True
gemini-intcenter-client True     True
gemini-deployment-manager True    True
gemini-deployment-client True     True
sbox-hadoop             True     True
sbox-core               True     True
sbox-admin              True     True
=====

```

The service Command (listen-port operator)

This administrative command operator could be used to find exposed network ports. Exposed ports would normally include the Web UI TCP port exposed to the connected network. To obtain a full list of the open ports and their exposure, run the **'--listen-port'** operator as shown below;

```
sbox service --listen-port
```

The screen below shows a typical output from this command;

```
[sbox@gemini ~]$ sbox service --listen-port
=====
Process                Port      Host
-----
Admin Web              443      *
Admin API              4444     *
Hadoop Manager API    8797     *
Core API               8686     127.0.0.1
Deployment Client      8888     *
Deployment Manager    8889     *
PEPSI                 9321     *
Integration Center Client 59000    127.0.0.1
=====
```

Notes

The wildcard '*' character means that the related network port is open on **all** active network interfaces.

If '127.0.0.1' is shown in the Host column, it means that the port is not exposed to any externally connected network and will allow only 'Host' based communication.

The system Command (info operator)

The 'system' command can be used with the '--info' operator to display hardware and software information. This is useful for collecting system information.

Use the command with the '--info' operator as shown below;

```
sbox system --info
```

The screen below shows a typical output from this command;

```
[sbox@gemini-0f2fd9 ~]$ sbox system --info
java_home      /usr/lib/jvm/jre/
memory_usage
  total        3079756
  cache        734144
  used         1031596
  free         1308720
  buffers      5296

virt_what      vmware
java_version   java version "1.8.0_192"
Java(TM) SE Runtime Environment (build 1.8.0_192-b12)
Java HotSpot(TM) 64-Bit Server VM (build 25.192-b12, mixed mode)

boot_id        a130a300-86d1-4e64-9c9a-ff250e9f4c82
java_directory /usr/lib/jdk1.8.0_192
open_ports     22
               111
               323
               443
               760
               768
               2534
               4444
```

Commands for Troubleshooting

The admin Command (Troubleshooting)

A summary of the 'admin' command and options is given below, and can be shown at the terminal at any time by typing 'sbox admin'.

Usage: sbox admin [OPTIONS]

[OPTIONS]

--reset-password	Reset password of web UI admin user.
-set-password	Set password for web UI admin user.
--gen-ssl	Regenerate SSL keys for web server.
--installed-packages	Display installed packages in web server.
--skip-wizard	Skip the setup wizard on access.
--restart	Restart sbox-admin service.

The admin Command (reset-password operator)

The 'admin' command can be used with the '--reset-password' operator to reset the Manage Web UI 'admin' password.

This can be invaluable if the Customer has forgotten their web UI password. It will unlock the account and set a randomly generated password.

Use the command with the '--reset-operator' operator as shown below;

```
sbox admin --reset-password
```

The screen below shows a typical output from this command;

```
[sbox@gemini-0800277c2740 ~]$ sbox admin --reset-password
+ Set admin password...
Now you can login Gemini Enterprise Manager by 'admin' user and password '92b8bf67'.
```

The admin Command (set-password operator)

The 'admin' command can be used with the '-set-password' operator to set a new password instead of using a randomly generated string as with the '--reset-password' operator. **Note:** that only one hyphen (-) is used with this operator as opposed to two (--).

Use the command with the '-set-password' operator as shown below;

```
sbox admin -set-password <new_password>
```

where '<new_password>' has to be replaced with the desired password.

It is generally recommended to change the admin password using the Web UI, and then to change it using the CLI as shown here.

The admin Command (gen-ssl operator)

The **'admin'** command can be used with the **'--gen-ssl'** operator to reset the SSL certificate used by the web UI.

This can be useful if the web UI is unavailable due to certificate issues (e.g. expired or invalid certificate).

Use the command with the **'--gen-ssl'** operator as shown below to reset the SSL certificate;

```
sbox admin --gen-ssl
```

```
[sbox@gemini-0800277c2740 ~]$ sbox admin --gen-ssl
+ Generating admin SSL key file...
Generating RSA private key, 4096 bit long modulus
.....++
.....++
e is 65537 (0x10001)
writing RSA key
+ Restart admin web service...
Done!
```

Important: This will overwrite any custom private key and certificate installed using the Manage Web GUI. It is recommended to backup private keys and certificates using the Manage backup feature from the Web UI before performing this operation.

The network Command (reset operator)

The **'network'** command can be used with the **'--reset'** operator to reset the network interface settings and remove any IP bondings.

Following the running of this command, basic network settings will be set to default and will have to be configured again (see [The network Command](#) section above for details);

```
sbox network --reset
```

The service Command (restart operator)

The **'service'** command can be used with the **'--restart'** operator to restart the administrative services of Manage. This may be requested by the Gemini Customer Support department during a technical issue.

Additionally, this action could be performed if the Web UI is unresponsive for some reason.

Use the command with the **'--restart'** operator as shown below to restart Gemini services with immediate effect;

```
sbox service --restart
```

The screen below shows a typical output from this command;

```
[sbox@gemini-0800277c2740 ~]$ sbox service --restart
+ Restarting services.....
+ All services have been restarted.
```

The splunk Command

A summary of the **'splunk'** command and options is given below, and can be shown at the terminal at any time by typing **'sbox splunk'**.

Usage: sbox splunk [OPTIONS]

[OPTIONS]

--kill	Remove Splunk instance.
--reset_environment_db	Clean up Splunk environment settings.
--backup_setting <file>	Backup Splunk settings.
--restore_setting <file>	Restore Splunk settings.

The **'splunk'** command can be used with the **'--kill'** operator to remove the installed Splunk instance in its entirety; including the binary file, configurations, and all ingested data.

```
sbox splunk --kill
```

Notes

Warning!: All Splunk configs and data will be deleted.
This is not a recoverable action. Use with caution.

Reset Splunk Environments

The **'splunk'** command can be used with the **'--reset_environment_db'** operator to reset the Splunk Environments database.

This is designed to remove the Splunk Environments information from Manage. During the process, you will be asked if you also want to remove Splunk. This is an important question to which you would normally respond **'NO'**.

This might be necessary if unintended actions have in some way corrupted the **Splunk Environments** dashboard.

Use the command with the **'--reset_environment_db'** operator as shown below to reset the Splunk Environments database;

```
sbox splunk --reset_environment_db
```

Note

Use with care, and watch for the prompts at the console. If answered incorrectly, all Splunk installations in the environment will also be removed.

Read any prompt messages with care and act accordingly.

This command replaces the previous commands (Ver <2.6)

```
sbox splunk --kill
```

```
sbox splunk --undo-manager
```

The system Command (patch operator)

The **'system'** command can be used with the **'--patch'** operator to apply patches to Manage. This could be in the case that the Manage Web UI is inaccessible.

Upload the patch to Gemini Central before running this command.

Use the command with the **'--patch'** operator as shown below to apply a patch without relying on the web interface;

```
sbox system -patch <patch_file>
```

The screen below shows a typical output from this command;

```
[sbox@gemini-0f2fd9 ~]$ sbox system
Usage: sbox system [OPTIONS]

[OPTIONS]
  --info          Display system information.
  -patch <file_path> Apply patch file with <file_path>.

[sbox@gemini-0f2fd9 ~]$ sbox system -patch all_match.patch
.....Patching completed.
[sbox@gemini-0f2fd9 ~]$
```

Commands for System Operations

Additional commands are available specific to interactive shells from the console.

Note

Note: The commands listed below related to System Operations are restricted from being used through SSH sessions and have to be typed at the terminal!

System Reboot

To reboot the Gemini instance, just type;

```
reboot
```

System Power Off

The shut down the Gemini instance, just type;

```
poweroff
```

Default Passwords for CLI Operations

Only two accounts are provisioned for command-line access to the instance by default. If **Tableau** has been installed the default suggested password is suggested below:

OS account	default password	Description
sbox	facing jet function drive	Used for Manage administration
splunk	think adventure kitchen chest	Used for Splunk administration
tableau	tableau	(when Tableau has been installed)

All users will be required to change the default password upon initial login;

```
sbox login: sbox
Password:
You are required to change your password immediately (root enforced)
Changing password for sbox.
(current) UNIX password:
New password:
Retype new password:
[sbox@sbox ~]$_
```

All OS User accounts have a default expiry of 60 days on their accounts. If you have changed the SSH passwords from their defaults, and you wish to freeze them for the future, navigate to the **Settings / Password Policy** dashboard and remove the checkmark from the relevant box.
