



Managing Splunk Clusters

This chapter contains the following sections:

- [Creating an Instant Splunk Cluster, on page 1](#)
- [Creating a Splunk Cluster Using Workflow, on page 5](#)
- [Customizing Splunk Cluster Creation, on page 6](#)
- [Adding Bare Metal Nodes to the Splunk Cluster, on page 10](#)
- [Deleting an Unreachable Cluster Node from Splunk Distribution, on page 13](#)
- [Deploying Splunk Cluster with Archival Node and NFS Support, on page 13](#)
- [Managing a Splunk Cluster, on page 14](#)

Creating an Instant Splunk Cluster

Use this procedure to create an instant Splunk cluster with the predefined values for the UCS Service Profile template. The system creates the QUICK_UCS_SPLUNK template, a new UCS SP Template of container type splunk while creating the instant splunk cluster. You can create a multi-site Splunk cluster or migrate an existing Splunk cluster to a multi-site Splunk cluster. Use the **UCS CPA Migrate Splunk Cluster to Multi-Site** workflow to migrate an existing Splunk cluster to a multi-site Splunk cluster. Until migrations are performed, you cannot completely manage an account in Cisco UCS Director Express for Big Data. **Splunk Cluster Multisite Configuration Generator** task should be modified for the account and site information before executing the workflow.

Step 1 Choose **Solutions > Big Data > Containers**.

Step 2 Click **Cluster Deploy Template**.

Step 3 Click **Instant Splunk Cluster Creation**.

Step 4 On the **Instant Splunk Cluster Creation** screen, complete the following fields:

Name	Description
Big Data Account Name field	The name of the Big Data account.
UCSM Policy Name Prefix field	The UCSM Policy Name prefix.
Monitoring Console Protocol drop-down list	Choose HTTP or HTTPS protocol.

Name	Description
Monitoring Console Port Number field	Enter the port number. Enter an integer between 1024 and 65535. Usage of reserved ports by Linux OS should be avoided so that the web server path is reachable.
SSH (root) Password field	The SSH root password. Special characters such as \$, %, and & are not supported. Note The SSH username pertains to the root user.
Confirm SSH Password field	Enter the SSH root password. Special characters such as \$, %, and & are not supported.
Splunk Manager Password field	The management console password. Special characters such as \$, %, and & are not supported.
Confirm Splunk Manager Password field	Enter the management console password. Special characters such as \$, %, and & are not supported.
OS Version drop-down list	Choose the operating system to be installed on the servers in this cluster.
Splunk Distribution Version drop-down list	Choose the Splunk Enterprise version to be used for this cluster.
Multi-UCSM check box	Check the Multi-UCSM check box if you use multiple UCSM accounts. Note If you use the multiple UCSM accounts option, you can configure the Splunk Server Roles as described in Step 7. You can add UCSM Specific Inputs in the Add Entry to UCSM Specific Inputs table. The following workflows are created during an Instant Splunk Cluster creation and Customized Splunk Cluster creation: <ul style="list-style-type: none">• UCS CPA Multi-UCSM Splunk Cluster WF• Single UCSM Server Configuration WF (This WF is triggered per UCSM Account. For example, UCSM 120, UCSM121)• UCS CPA Node Bare Metal (This WF is triggered per Node)
UCS Manager Account drop-down list	Choose the Cisco UCS Manager account for this cluster.
Organization drop-down list	Choose the organization in which the servers for this cluster are located.

Name	Description
UCS SP Template table	Choose an existing UCS Service Profile Template for cluster creation.
PXE VLAN ID field	Enter the PXE VLAN ID. Enter an integer between 1 and 3967 or between 4048 and 4093.

Step 5

In the **Splunk Server Roles** table, if you want to edit a Splunk Server Role, select the row for that role, and click **Edit**.

Step 6

On the **Edit Splunk Server Roles Entry** screen, complete the following fields and click **Submit**. The fields displayed in the **Edit Splunk Server Roles Entry** screen is based on the server role selection.

Note Admin roles such as deploying roles on a bare metal agent and choosing license master, cluster master, and bare metal of the deployment server are only supported during fresh cluster creation. Also, existing IP addresses for the admin roles are only supported through fresh cluster creation.

Name	Description
Node Type field	Displays the Splunk node role.
Node Count field	The number of nodes in the splunk cluster for the selected node type.
Host Name Prefix drop-down list	Choose the host name prefix for this splunk cluster.
SSD Boot Drives Available for OS check box	<p>Check this check box if you do not want to validate the server disk availability for RAID level OS disks. Ensure that the servers contain SSD.</p> <p>If the check box is not selected, the disk availability for both the OS disk and data disk are validated based on their RAID level.</p> <p>Note This check box is not displayed when the UCSM version is greater than or equal to 3.</p>
Search Head to be part of cluster	By default, this option is checked and disabled. The search head role is added to all Search Head cluster.
Validate Page check box	Check Validate Page to recalculate admin hostnames per given hostname prefix and node count.
Deploy roles on Bare Metal check box	Check Deploy roles on Bare Metal to deploy roles on a bare metal agent. By default, this option is checked. Uncheck this option to deploy admin roles on Live Nodes.
Use Existing License Master check box	Check Use Existing License Master to use the existing license master.
License Master BM drop-down list	Choose the license master bare metal.
Monitoring Console BM drop-down list	Choose the monitoring console bare metal.
Cluster Master BM drop-down list	Choose the cluster master bare metal.

Name	Description
Deployer BMs table	Choose the bare metal of the deployer server.
Deployment Server BMs table	Choose the bare metal of the deployment server.
Current License Master Live IPs	Enter the IP addresses of the current license master. This field is displayed when Use Existing License Master is checked.
New License Master Live IP	Enter the IP address of the new license master. This field is displayed when Deploy roles on Bare Metal is unchecked.
Monitoring Console Live IP	Enter the IP address of the monitoring console. This field is displayed when Deploy roles on Bare Metal is unchecked.
Cluster Master Live IP	Enter the IP address of the new license cluster master. This field is displayed when Deploy roles on Bare Metal is unchecked.
Deployer Live IPs	Enter the IP addresses of the deployer server. This field is displayed when Deploy roles on Bare Metal is unchecked.
Deployment Server Live IPs	Enter the IP addresses of the deployment server. This field is displayed when Deploy roles on Bare Metal is unchecked.
Server Pool table	Enter the server pool that you want to use for the cluster for the selected node type. The Cisco UCS Manager account and the organization that you choose determine which server pools are displayed in this area.

Note All Live IPs provided for admin roles of a Splunk cluster except for Existing Licensing server and running OS should be same as the Splunk Indexer or Search Head cluster.

Note Hostnames separated by comma or IP addresses can be provided and the hostname resolution should happen from the Cisco UCS Director appliance.

Step 7 In the **vNIC Template** table, review and, if desired, edit the vNIC templates available for the cluster.

Step 8 If you want to edit a vNIC template, select the row for that template and click **Edit**.

Step 9 On the **Edit vNIC Template Entry** screen, complete the following fields and click **Submit**.

Name	Description
vNIC Name drop-down list	The vNIC name in the selected template. This field is for your information only.
IP Pool field	Choose the Big Data IP pool that you want to use for IP addresses assigned to this vNIC.

Name	Description
MAC Address Pool drop-down list	Choose the MAC address pool that you want to use for this cluster. (This drop-down list is disabled if an existing UCS SP Template is selected.)
VLAN ID field	The VLAN ID for this cluster. (This field is disabled if an existing UCS SP Template is selected.)

Step 10 In the **Site Preferences** table, click **Add (+)** to add one or more sites.

Step 11 On the **Add Entry to Site Preferences** screen, complete the following fields and click **Submit**.

Name	Description
Site Name drop-down list	Choose the site in which the servers for this cluster are located.
Indexers field	Click Select to choose the indexers for the site and click Select .
Search Heads field	Click Select to choose the search heads for the site and click Select .
Replication Factor drop-down list	Choose replication factor for the site.
Search Factor drop-down list	Choose search factor for the site. The search factor must be less than or equal to the replication factor.

Step 12 Click **Submit**.

Step 13 Specify the origin and total site replication factors.

Step 14 Specify the origin and total site search factors.

Step 15 Choose a mater site from **Master Site Name**.

Step 16 Click **Submit**.

Creating a Splunk Cluster Using Workflow

In Cisco UCS Director Express for Big Data, administrator can map the advanced catalog option to Splunk cluster creation workflow, with limited user inputs, so that the service end user can trigger cluster creation. See [Cisco UCS Director End User Portal Guide](#).

Before you begin

- Create a service profile template
- Create a server pool in the Cisco UCS Manager account that you plan to use for this cluster. See [Cisco UCS Director Management Guide for Cisco UCS Manager](#).
- Create a MAC address pool
- Create a user with user role as service end user.

-
- Step 1** Log into Cisco UCS Director Express for Big Data using admin credentials.
- Step 2** Choose **Orchestration** and click **Workflows**.
- Step 3** Click **Add Workflow**.
- Step 4** On the **Add Workflow Details** page, enter the workflow name and choose a folder. Click **Next**.
- Step 5** On the **Add User Inputs** page, enter the required details and click **Next**.
- Step 6** On the **Add User Outputs** page, enter the required details and click **Submit**.
- Step 7** Double-click the workflow in the **Workflow Designer**.
- Step 8** Add the Initiate Splunk Cluster task.
- Step 9** Select the attributes that you want to map to the workflow input fields. Check the **Map to User Input** check box to provide user inputs, if required.
- Step 10** Enter required details in the **Splunk Service Role** table, **vNIC Template** table, and **Site Preferences** table, and click **Submit**.
- Step 11** Choose **Policies > Catalogs** and click **Add Catalog**.
- Step 12** On the **Add Catalog** page, choose the catalog type as Advanced and select a workflow. Click **Submit** to map the workflow to the catalog.
- Step 13** Log into Cisco UCS Director Express for Big Data using service end user credentials.
- Step 14** Choose **Catalogs**. The **Catalogs** page displays the list of catalogs available for the service end user.
- Step 15** Select a catalog and click **Create Request**. The **Create Server Request** page displays the mapped user inputs.
- Step 16** Specify the required details.
- Step 17** Click **Next** and enter the cluster details in the **Customize Workflow** page.
- Step 18** Click **Next** and view the cluster details in the **Summary** page.
- Step 19** Click **Submit** to trigger a workflow for creating a Splunk cluster.
-

Customizing Splunk Cluster Creation

You can create a multi-site Splunk cluster or migrate an existing Splunk cluster to a multi-site Splunk cluster. Use the **UCS CPA Migrate Splunk Cluster to Multi-Site** workflow to migrate an existing Splunk cluster to a multi-site Splunk cluster. Until migrations are performed, you cannot completely manage an account in Cisco UCS Director Express for Big Data. **Splunk Cluster Multisite Configuration Generator** task should be modified for the account and site information before executing the workflow.

Before you begin

- Create a UCS Service Profile Template.
- Create a Cluster Deploy Template.

-
- Step 1** Choose **Solutions > Big Data > Containers**.
- Step 2** Click **Cluster Deploy Templates**.
- Step 3** Click **Add** to create a cluster deploy template for the Splunk cluster. See [Adding a Cluster Deployment Template](#).

Step 4 Click **Customized Splunk Cluster Creation**.

Step 5 On the **Customized Splunk Cluster Creation** screen, complete the following fields.

Name	Description
Big Data Account Name field	The name of the Big Data account.
UCSM Policy Name Prefix field	The UCSM Policy Name prefix.
Monitoring Console Protocol drop-down list	Choose HTTP or HTTPS protocol.
Monitoring Console Port Number field	Enter the port number. Enter an integer between 1024 and 65535.
SSH (root) Password field	The SSH root password. Special characters such as \$, %, and & are not supported. Note The SSH username pertains to the root user.
Confirm SSH Password field	Enter the SSH root password. Special characters such as \$, %, and & are not supported.
Splunk Manager Password field	The management console password. Special characters such as \$, %, and & are not supported.
Confirm Splunk Manager Password field	Enter the management console password. Special characters such as \$, %, and & are not supported.
OS Version drop-down list	Choose the operating system to be installed on the servers in this cluster.
Splunk Distribution Version drop-down list	Choose the Splunk distribution version to be used for this cluster.
Multi-UCSM check box	Check the Multi-UCSM check box if you use multiple UCSM accounts. Note If you use the multiple UCSM accounts option, you can configure the Splunk Server Roles as described in the Step 8. You can add UCSM Specific Inputs in the Add Entry to UCSM Specific Inputs table. The following workflows are created during an Instant Splunk Cluster creation and Customized Splunk Cluster creation: <ul style="list-style-type: none"> • UCS CPA Multi-UCSM Splunk Cluster WF • Single UCSM Server Configuration WF (This WF is triggered per UCSM Account. For example, UCSM 120, UCSM121) • UCS CPA Node Bare Metal (This WF is triggered per Node)

Name	Description
UCS Manager Account drop-down list	Choose the Cisco UCS Manager account for this cluster.
Organization drop-down list	Choose the organization in which the servers for this cluster are located.
UCS SP Template table	Choose an existing UCS Service Profile Template for cluster creation.
PXE VLAN ID field	Enter the PXE VLAN ID.

Step 6

In the **Splunk Server Roles** table, if you want to edit a Splunk Server Role, select the row for that role, and click **Edit**.

Step 7

On the **Edit Splunk Server Roles Entry** screen, complete the following fields and click **Submit**. The fields displayed in the **Edit Splunk Server Roles Entry** screen is based on the server role

Note Admin roles such as deploying roles on a bare metal agent and choosing license master, cluster master, and bare metal of the deployment server are only supported during fresh cluster creation. Also, existing IP addresses for the admin roles are only supported through fresh cluster creation.

Name	Description
Node Type field	Displays the Splunk node role.
Node Count field	The number of nodes in the splunk cluster for the selected node type.
Host Name Prefix drop-down list	Choose the host name prefix for this splunk cluster.
SSD Boot Drives Available for OS check box	<p>Check this check box if you do not want to validate the server disk availability for RAID level OS disks. Ensure that the servers contain SSD.</p> <p>If the check box is not selected, the disk availability for both the OS disk and data disk are validated based on their RAID level.</p> <p>Note This check box is not displayed when the UCSM version is greater than or equal to 3.</p>
Search Head to be part of cluster	By default, this option is checked and disabled. The search head role is added to all Search Head cluster.
Validate Page check box	Check Validate Page to recalculate admin hostnames per given hostname prefix and node count.
Deploy roles on Bare Metal check box	Check Deploy roles on Bare Metal to deploy roles on a bare metal agent. By default, this option is checked. Uncheck this option to deploy admin roles on Live Nodes.
Use Existing License Master check box	Check Use Existing License Master to use the existing license master.
License Master BM drop-down list	Choose the license master bare metal.

Name	Description
Monitoring Console BM drop-down list	Choose the monitoring console bare metal.
Cluster Master BM drop-down list	Choose the cluster master bare metal.
Deployer BMs table	Choose the bare metal of the deployer server.
Deployment Server BMs table	Choose the bare metal of the deployment server.
Current License Master Live IPs	Enter the IP addresses of the current license master. This field is displayed when Use Existing License Master is checked.
New License Master Live IP	Enter the IP address of the new license master. This field is displayed when Deploy roles on Bare Metal is unchecked.
Monitoring Console Live IP	Enter the IP address of the monitoring console. This field is displayed when Deploy roles on Bare Metal is unchecked.
Cluster Master Live IP	Enter the IP address of the new license cluster master. This field is displayed when Deploy roles on Bare Metal is unchecked.
Deployer Live IPs	Enter the IP addresses of the deployer server. This field is displayed when Deploy roles on Bare Metal is unchecked.
Deployment Server Live IPs	Enter the IP addresses of the deployment server. This field is displayed when Deploy roles on Bare Metal is unchecked.
Server Pool table	Enter the server pool that you want to use for the cluster for the selected node type. The Cisco UCS Manager account and the organization that you choose determine which server pools are displayed in this area.

Step 8 In the **vNIC Template** table, review and, if desired, edit the vNIC templates available for the cluster.

Step 9 If you want to edit a vNIC template, select the row for that template and click **Edit**.

Step 10 On the **Edit vNIC Template Entry** screen, complete the following fields and click **Submit**.

Name	Description
vNIC Name drop-down list	The vNIC name in the selected template. This field is for your information only.
IP Pool field	Choose the big data IP pool that you want to use for IP addresses assigned to this vNIC.

Name	Description
MAC Address Pool drop-down list	Choose the MAC address pool that you want to use for this cluster. (This drop-down list is disabled if an existing UCS SP Template is selected.)
VLAN ID field	The VLAN ID for this cluster. (This field is disabled if an existing UCS SP Template is selected.)

Note When you use vNIC bonding, ensure that you assign IP Pool, MAC Address Pool, and VLAN ID to the first vNIC in the **vNIC Template** table.

Step 11 In the **Site Preferences** table, click **Add (+)** to add one or more sites.

Step 12 On the **Add Entry to Site Preferences** screen, complete the following fields and click **Submit**.

Name	Description
Site Name drop-down list	Choose the site in which the servers for this cluster are located.
Indexers field	Click Select to choose the indexers for the site and click Select .
Search Heads field	Click Select to choose the search heads for the site and click Select .
Replication Factor drop-down list	Choose replication factor for the site.
Search Factor drop-down list	Choose search factor for the site. The search factor must be less than or equal to the replication factor.

Step 13 Click **Submit**.

Step 14 Specify the origin and total site replication factors.

Step 15 Specify the origin and total site search factors.

Step 16 Choose a mater site from **Master Site Name**.

Step 17 Click **Submit**.

Adding Bare Metal Nodes to the Splunk Cluster

To add a Bare Metal node to a single-site Splunk cluster, cluster should be migrated to multi-site Splunk cluster using the **UCS CPA Migrate Splunk Cluster to Multi-Site** workflow.



Note To add bare metal nodes to the Splunk clusters using RHEL 7.4 or CentOS7.4 (created prior to Release 6.6.0.1), create a service profile template in Cisco UCS Manager with UEFI boot option.

Step 1 Choose **Solutions > Big Data > Accounts**.

Step 2 Click **Splunk Accounts**.

Step 3 Double-click the Splunk account.

You can see only the **Hosts** tab.

Step 4 Click **Add Bare Metal Nodes**.

Step 5 Create a service profile template in Cisco UCS Manager with UEFI boot option, if you want to add bare metal nodes to the Splunk clusters using RHEL 7.4 or CentOS7.4 (created prior to Release 6.6.0.1).

Step 6 On the **Add Bare Metal Nodes** screen, complete the following fields:

Name	Description
Big Data Account Name field	The name of the Big Data account.
UCSM Policy Name Prefix field	The UCSM Policy Name prefix.
Monitoring Console Port Number field	Enter the port number. Enter an integer between 1024 and 65535. Usage of reserved ports by Linux OS should be avoided so that the web server path is reachable.
Monitoring Console Protocol drop-down list	Choose HTTP or HTTPS protocol.
OS Version drop-down list	Choose the operating system to be installed on the servers in this cluster.
Splunk Version drop-down list	Choose the Splunk version.
UCS Manager Account drop-down list	Choose the Cisco UCS Manager account for this cluster.
Organization drop-down list	Choose the organization in which the servers for this cluster are located.
UCS SP Template	Choose an existing UCS Service Profile Template for the cluster creation.
PXE VLAN ID field	Enter the PXE VLAN ID.
UCSTemplate Name table	Choose the UCS Service Profile Template for Splunk.

Step 7 In the **Splunk Server Roles** table, if you want to edit a Splunk Server Role, select the row for that role, and click **Edit**.

Step 8 On the **Edit Splunk Server Roles Entry** screen, complete the following fields and click **Submit**.

Name	Description
Node Type field	Displays the Splunk node role.
Node Count field	The number of nodes in the splunk cluster for the selected node type.
Host Name Prefix drop-down list	Choose the host name prefix for this splunk cluster.

Name	Description
SSD Boot Drives Available for OS check box	<p>Check this check box if you do not want to validate the server disk availability for RAID level OS disks. Ensure that the servers contain SSD.</p> <p>If the check box is not selected, the disk availability for both the OS disk and data disk are validated based on their RAID level.</p> <p>Note This check box is not displayed when the UCSM version is greater than or equal to 3.</p>
Search Head to be part of cluster	By default, this option is checked and disabled. The search head role is added to all Search Head cluster.
Server Pool table	<p>Enter the server pool that you want to use for the cluster for the selected node type.</p> <p>The Cisco UCS Manager account and the organization that you choose determine which server pools are displayed in this area.</p>

Step 9 In the **vNIC Template** table, review and, if desired, edit the vNIC templates available for the cluster.

Step 10 If you want to edit a vNIC template, select the row for that template and click **Edit**.

Step 11 On the **Edit vNIC Template Entry** screen, complete the following fields and click **Submit**.

Table 1:

Name	Description
vNIC Name drop-down list	This field is for your information only.
IP Pool drop-down list	Choose the Big Data IP pool that you want to use for IP addresses assigned to this vNIC.
MAC Address Pool drop-down list	Choose the MAC address pool that you want to use for this cluster. (This drop-down list is disabled if an existing UCS SP Template is selected.)
First MAC Address field	Choose the MAC address pool that you want to use for this cluster.
Size field	Enter the size.
VLAN ID field	The VLAN ID for this cluster.

Step 12 Click **Submit**.

Note By default, the hardware default is used as UUID pool for the servers in the cluster.

Step 13 In the **Site Preferences** table, click **Add (+)** to add one or more sites.

Note Click **Edit** to add a node in the existing site.

Step 14 On the **Add Entry to Site Preferences** screen, complete the following fields and click **Submit**.

Name	Description
Site Name drop-down list	Choose the site in which the servers for this cluster are located.
Indexers field	Click Select to choose the indexers for the site and click Select .
Search Heads field	Click Select to choose the search heads for the site and click Select .
Replication Factor drop-down list	Choose replication factor for the site.
Search Factor drop-down list	Choose search factor for the site. The search factor must be less than or equal to the replication factor.

Step 15 Click **Submit**.

Step 16 Specify the origin and total site replication factors.

Step 17 Specify the origin and total site search factors

Step 18 Click **Submit**.

Deleting an Unreachable Cluster Node from Splunk Distribution

In a Splunk distribution, when a node is unreachable and the node status is displayed as **Unknown**, you can delete a node by clicking **Delete Node to Bare Metal**. The node gets deleted from the Splunk user interface and the status is not updated in the Cisco UCS Director Express for Big Data user interface (refer CSCvg90939 bug). You should click **Delete Node** to delete the node from the Cisco UCS Director Express for Big Data user interface.

Deploying Splunk Cluster with Archival Node and NFS Support

The following are the scenarios to deploy a Splunk cluster along with Archival node:

- Configure Archival Node along with a Splunk Cluster—Archival node is configured along with the cluster automatically.
- Configure Archival Node on a Bare Metal—You can use the add Bare Metal option along with archival node settings like hostname prefix, number of archival nodes, and server pool. When the node comes up, the UCS CPA Splunk Add Live Archival Node workflow is used to configure NFS related setting on the node and configure mount point on indexers.
- Configuring Archival Node on a Live Node—You can use this to configure NFS related setting on the node and add it to the cluster.

For more information on how archival node disks are allocated to indexers, see the latest *Cisco UCS Integrated Infrastructure for Big Data with Splunk Enterprise*.

Managing a Splunk Cluster

You can manage the Splunk cluster from the **Hosts** tab.

Step 1 Choose **Solutions > Big Data > Accounts**.

Step 2 Click **Splunk Accounts**.

Step 3 Click **Summary** to view the statistics data report for the selected Splunk Account and the high-level report on the cluster and node account.

Step 4 Click **Hosts** to perform the following actions:

Name	Description
Refresh	Refreshes the page.
Favorite	Adds the page to Favorites.
Add Bare Metal Nodes	Add bare metal nodes to the Splunk cluster. You can add Indexer, Search Head, or Administrative node through Add Bare Metal workflow. You need to provide the Replication Factor based on the Indexer count.

Note You can also start, stop, or restart the Splunk cluster.

Step 5 Select a host that allows you to perform the following actions:

Name	Description
View Details	Displays the summary of the CPU usage, the I/O status of the hosts disks, and so on. Note If you see a License Status tab, it indicates a licensing issue.
Start	Starts the services on the node.
Stop	Stops the services on the node.
View Details	Restarts the services on the node.
Restart	Deletes node from the cluster.
Delete Node to Bare Metal	The node is removed from the cluster and disassociated from the service profile. The node becomes a Bare Metal server.

Step 6 Select an account and click **View Details**.

You can start, stop, or restart the Splunk cluster.

Step 7 Click the **Performance** tab.

Step 8 Click **Run Test**.

The Performance tab displays a default Big Data Metrics Report. This report shows the statistics collected for each host before the Splunk cluster creation and the reports post Splunk cluster creation only when you check the **Memory Test**, **Network Test**, and **Disk Test** check boxes in the **Pre Cluster Performance Tests** section of the **Management** tab. If you enable the precluster disk test, it impacts Splunk cluster creation.

Step 9 Click **Submit**, and then click **OK**.

For the following actions, choose the performance report:

Name	Description
View	Displays the metrics in the Big Data Metrics Report.
Compare	Compares and displays the metrics in the Big Data Metrics Report.
View Graph Report	Displays graphically the following reports from the Summary tab: <ul style="list-style-type: none"> • Average TRIAD Rate (MB/Sec) • Average Network Bandwidth (MB/Sec)
Delete	Deletes the Big Data Metrics Report.
More Reports	Displays the metrics on an hourly, daily, weekly, or monthly basis.

Step 10 Click **Monitoring**.

Every time an inventory collection cycle is triggered, an entry listing the aggregate CPU, network bandwidth, and disk utilization metrics appears on the Monitoring Page.

Step 11 Select the entry you want to analyze and click **View Details**.

Step 12 Click **Back** to return to the **Monitoring** page.

