The bridge to possible

Deploying Nexus Dashboard in your Organization

BRKDCN-2914

Matthias Wessendorf, Principal Engineer
@matteq4er

BRKDCN-2914

Cisco Webex App

Questions?

Use Cisco Webex App to chat
with the speaker after the session

How

1. Find this session in the Cisco Live Mobile App
2. Click “Join the Discussion”
3. Install the Webex App or go directly to the Webex space
4. Enter messages/questions in the Webex space

Webex spaces will be moderated
by the speaker until June 7, 2024.

https://ciscolive.ciscoevents.com/
ciscolivebot/#BRKDCN-2914

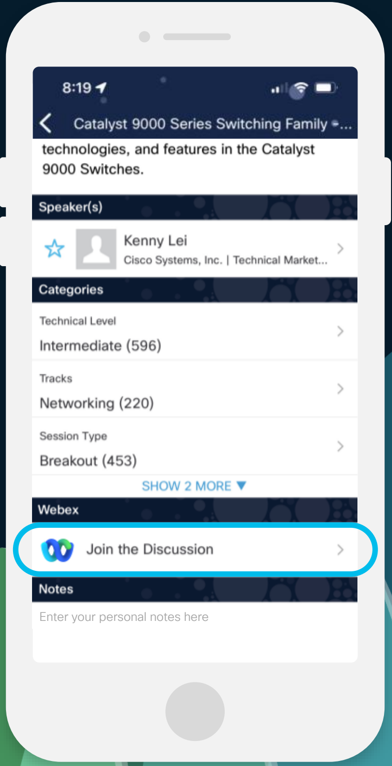

8:19

Catalyst 9000 Series Switching Family ...

technologies, and features in the Catalyst 9000 Switches.

Speaker(s)

Kenny Lei
Cisco Systems, Inc. | Technical Market...

Categories

Technical Level
Intermediate (596)

Tracks
Networking (220)

Session Type
Breakout (453)

SHOW 2 MORE ▼

Webex

Join the Discussion

Notes

Enter your personal notes here

cisco Live!

Agenda

- Introduction
- What is Nexus Dashboard? A view under the hood
- Deploying Nexus Dashboard
- Operating Nexus Dashboard
- Summary

At the end of the session you will ...

- Be able to define the requirements for deploying a Nexus Dashboard in your Organisation. By describing the
  - Deployment model, centralized vs. stretched
  - Network requirmenets and attachment to the network
  - Sizing a Nexus Dashboard for the different services.

Introduction

Nexus Dashboard

Deployment evolution

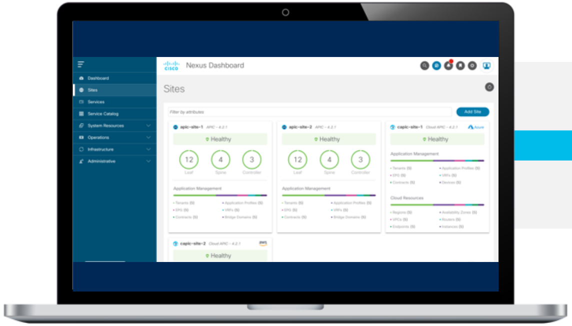

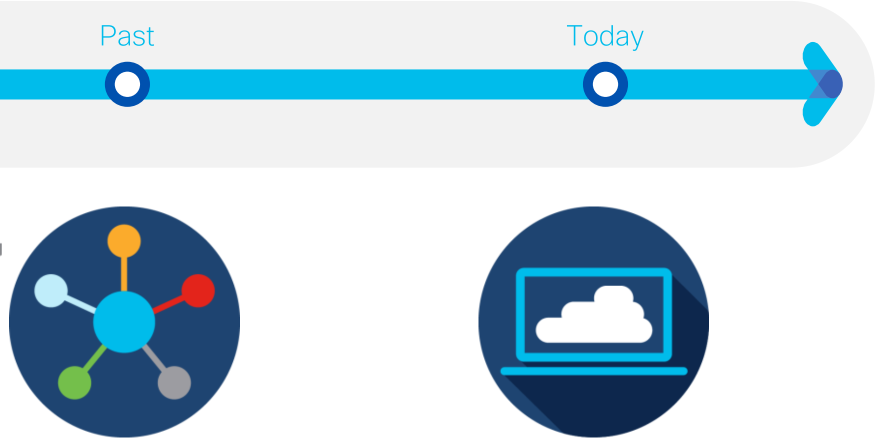

Physical Platform Cluster

Virtual/Cloud Platform

Powering automation
Unified agile platform

Nexus Dashboard

Simple to automate, simple to consume

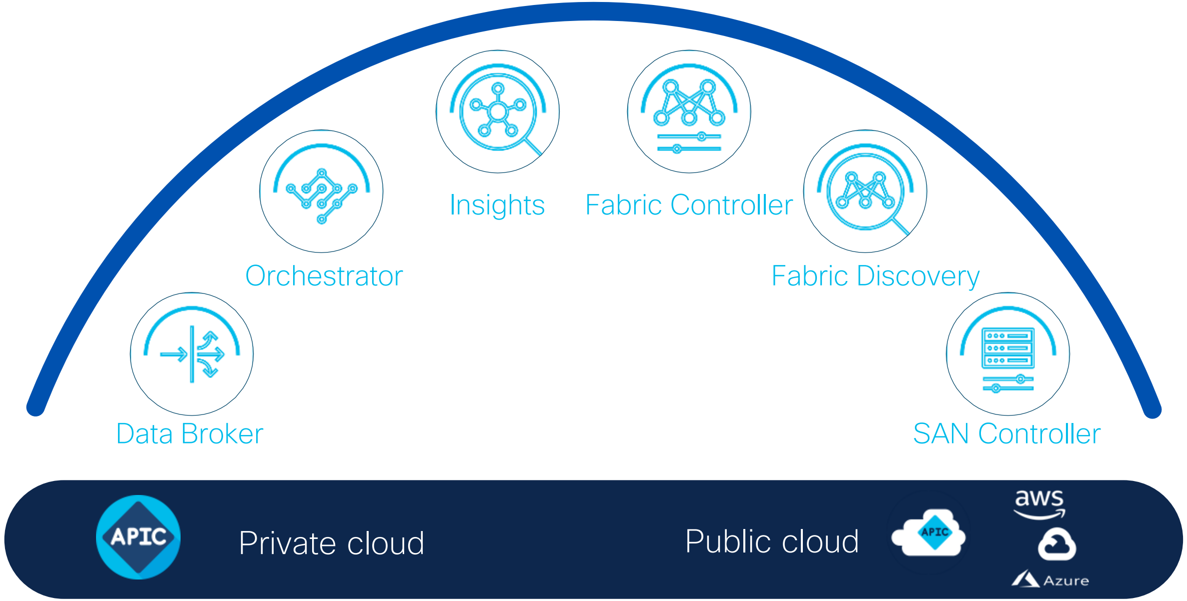

Consume all services in one place

Nexus Dashboard: One view

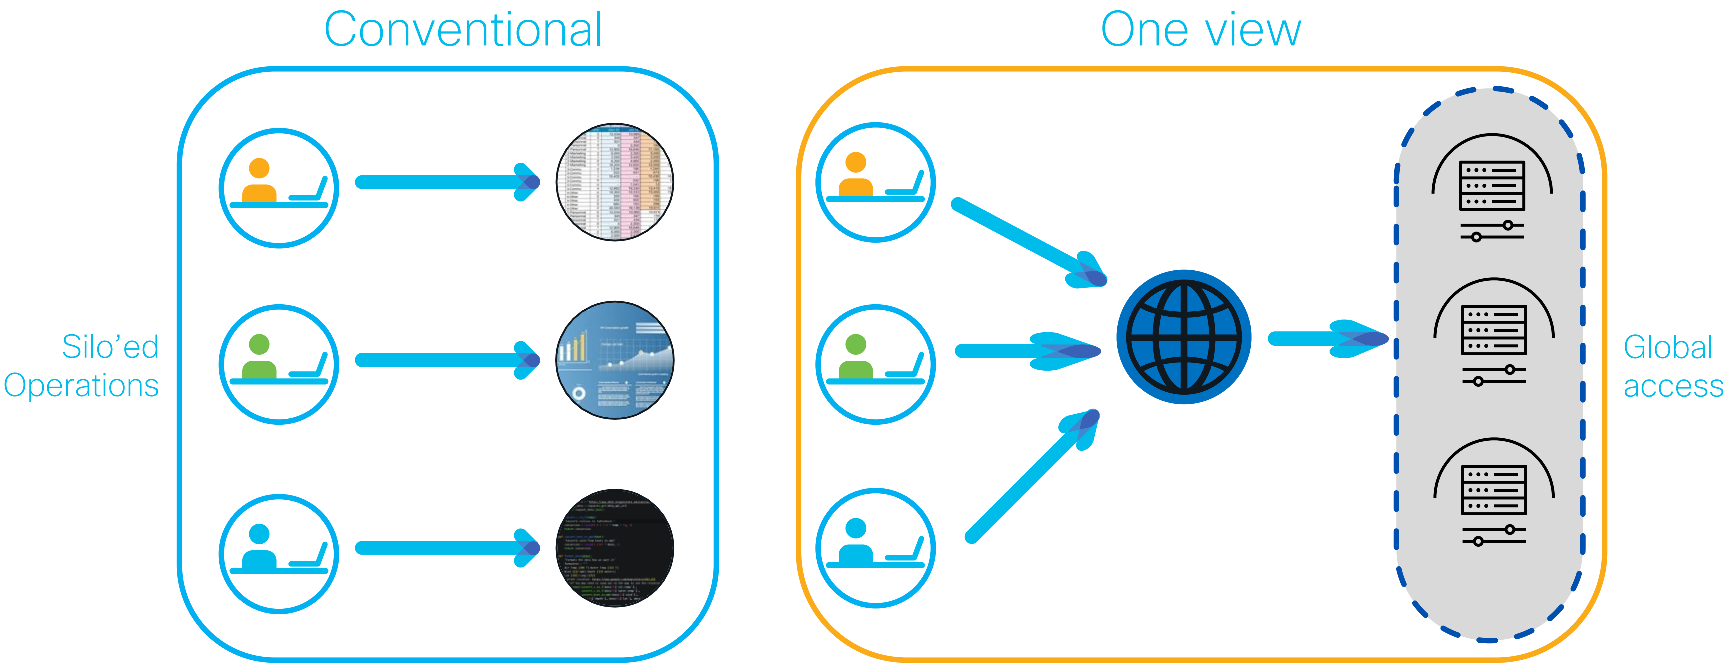

Conventional

One view

Silo’ed Operations

Global access

Cisco Nexus Dashboard Platform

Modern Scale-out application services stack to host data center operations applications

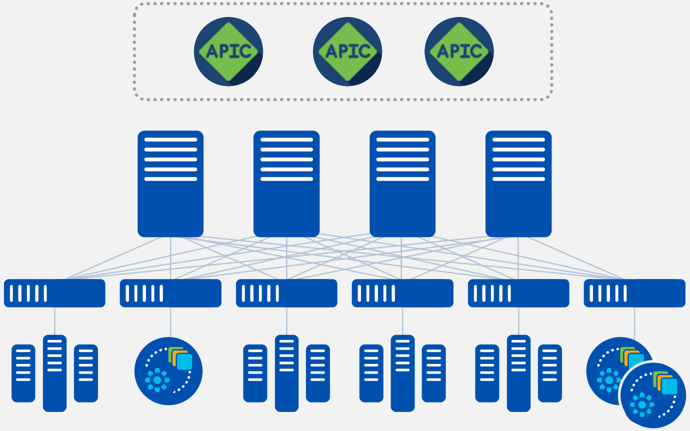

Nexus Dashboard Insights

3rd Party apps

NDO

Nexus Dashboard Orchestrator

2.2 GHz(Node-G2) or 2.8Ghz(Node-G4) CPU x 2

256 GB memory

2.4 TB x 4 HDD

10G/25G/40G connect

Network automation　Scale-out cluster　High Availability

Virtual Nexus Dashboard Platform

Virtual Platform to Support NDI ,NDO and NDFC in Production

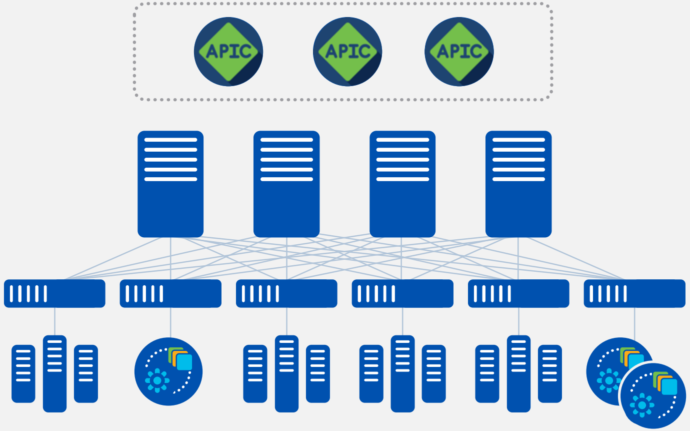

Nexus Dashboard Insights

3rd Party apps

NDO Nexus Dashboard Orchestrator

APP-Node	DATA-Node
64 GB memory	128 GB memory
550G/1536GB* SDD	3TB SSD/NVMe
16 vCPUs	32 vCPUs

Available for ESXi KVM

* For 3 APP node NDI installation

Nexus Dashboard: A Unified Agile Platform

The operator view

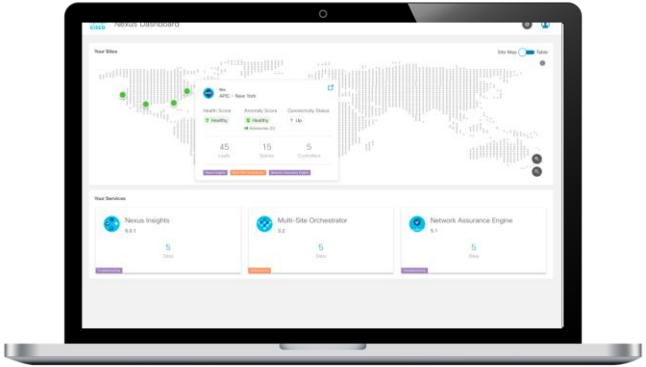

- Consume service(s) from single place
- Frictionless navigation across multiple services and sites
- Customize views and workflows

The admin view

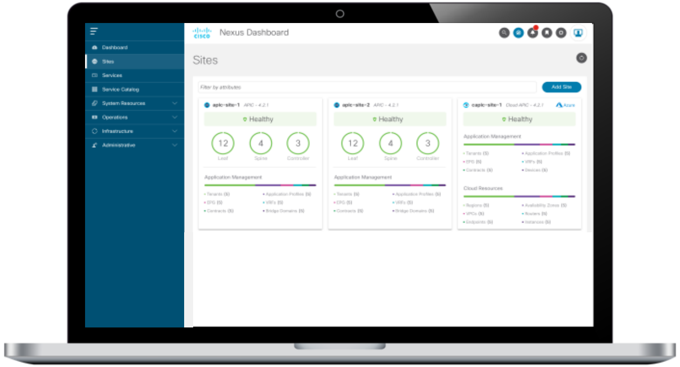

- Single dashboard for lifecycle management of services and Ops infra
- Consistent one-time onboarding of domains and services
- Consistent user management and access control

What is Nexus Dashboard? – a view under the hood –

Nexus Dashboard Platform—Under the Hood

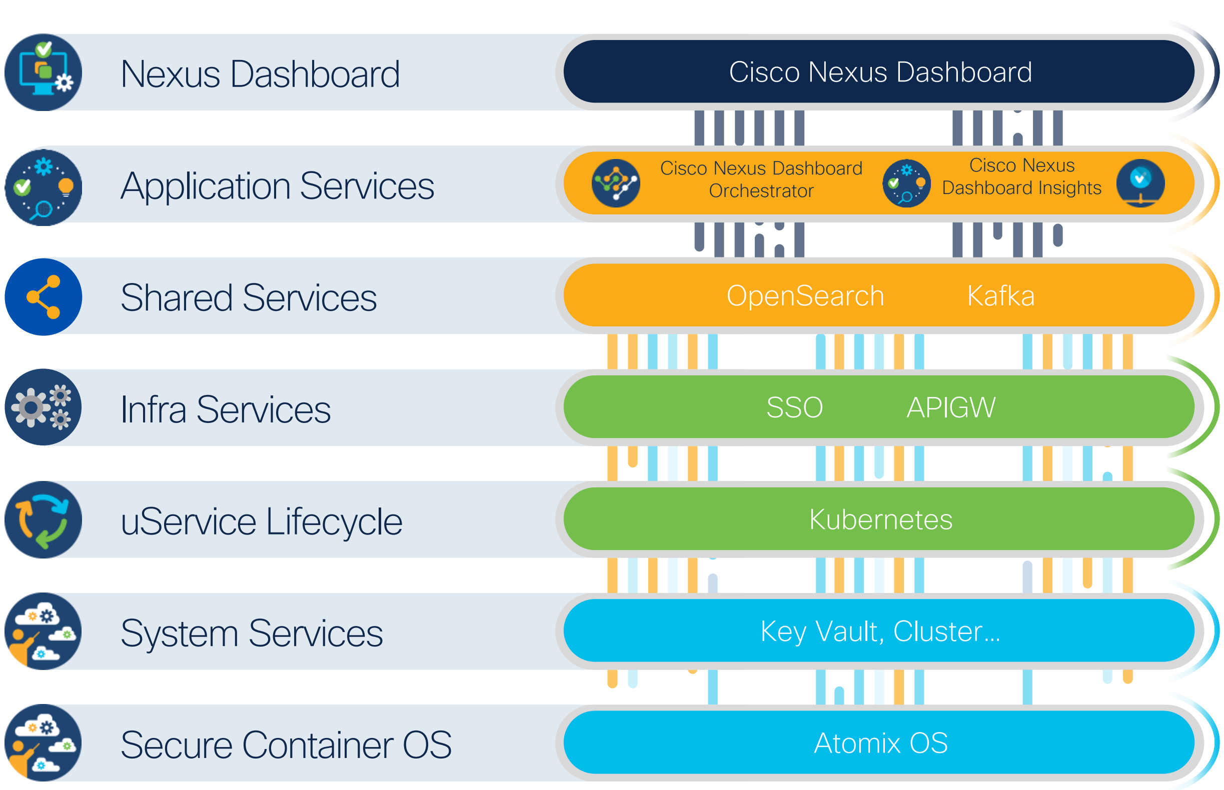

Nexus Dashboard — Cisco Nexus Dashboard

Application Services — Cisco Nexus Dashboard Orchestrator; Cisco Nexus Dashboard Insights

Shared Services — OpenSearch; Kafka

Infra Services — SSO; APIGW

uService Lifecycle — Kubernetes

System Services — Key Vault, Cluster…

Secure Container OS — Atomix OS

- Single Dashboard to view health
- Single UI to consume app services

- Proactive Notifications
- Configuration compliancy
- Realtime Telemetry Collection
- Faster resolution through correlation

- Frictionless access
- Standardized service access
- Industry standard container management

- Secure cluster and app management
- Infrastructure for secure K8 bringup
- Hardened secure container OS

ND Node Role

Primary	Secondary	Standby
• Hosting the infra-services of ND • Hosting workload-services • Can be deployed as 1 or 3 in a cluster • No dynamic adding afterwards	• Hosting workload-services • Used for scale-out-computing • Can be deployed as 1 or 3 in a cluster • Can be dynamically added	• Not hosting any services • Used for redundancy • Can be used to replace a failed primary node • Max 2 per cluster

Deployment Model

- Depending on the services (NDI/NDO) being deployed on top of vND the number of required nodes and which node type must be deployed as Primary is changing
- Scale numbers are documented in the ND cluster sizing tool

Deployed Services	NDI	NDO**	NDI	NDFC***
Total number of nodes needed	3	3	6	3
Type of Primary nodes	App	APP	DATA	APP
Total number of DATA nodes needed	0	0	3	0
Total number of APP nodes needed	3	3	3	3

** 1 APP node PoC setup for NDO with reduced scale is available

*** 1 APP node PoC setup for NDFC with reduced scale is available

Nexus Dashboard Connectivity

ND DATA Interface

- Used to communicate to Fabrics
    - Telemetry
    - SSH to Fabric
    - HTTPS to Fabric
    - KAFKA to Fabric

ND MGMT Interface

- Used to communicate for mgmt. purposes
    - AAA
    - Syslog
    - HTTPS
    - KAFKA
    - ND Federation

Pods on top of ND get interfaces assigned into MGMT and/or DATA Interface. This is defining the communication path.

ND to APIC Connectivity Considerations

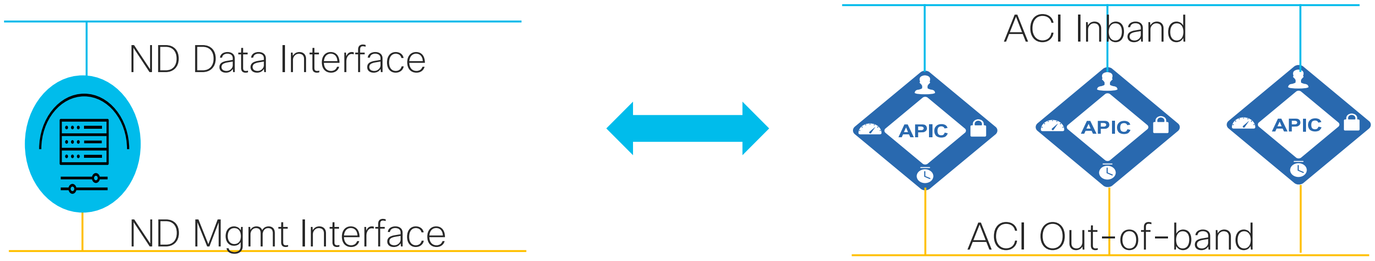

- An ACI fabric is onboarded on ND by specifying the IP address of one of the nodes of the APIC cluster
  - This can be either the APIC's IB or OOB address. In case of the usage of NDI it must be the APIC's IB address
- ND uses the Data Interface to establish the initial connection to that APIC's IP address
  - If the connection is successful, ND discovers all the OOB and IB IP addresses for the other nodes in the APIC cluster

ND to NDFC Connectivity Considerations

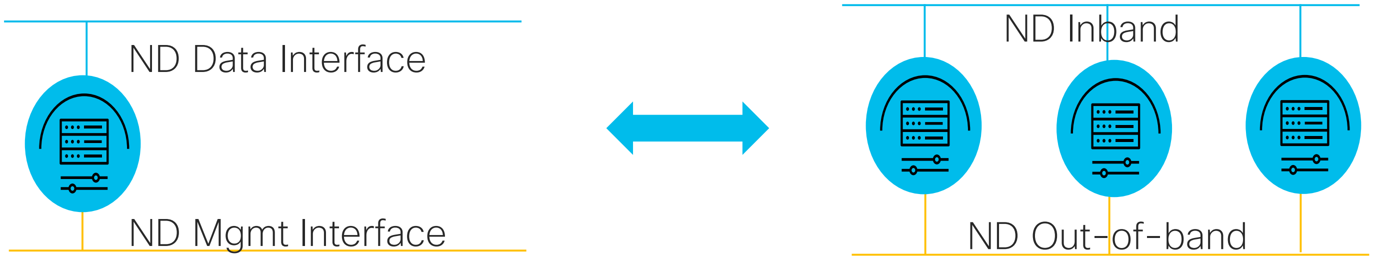

- An NDFC site is onboarded on ND by specifying the Inband IP address of the ND hosting the NDFC, no other IP is supported
- ND uses the Data Interface to establish the initial and ongoing connection to that ND Data IP address hosting NDFC

Persistent IPs and their usage

Usage of Persistent IPs

2. Switches are programmed to stream to this IP

1. Telemetry Pods get an IP assigned (Persistent IP) at application start

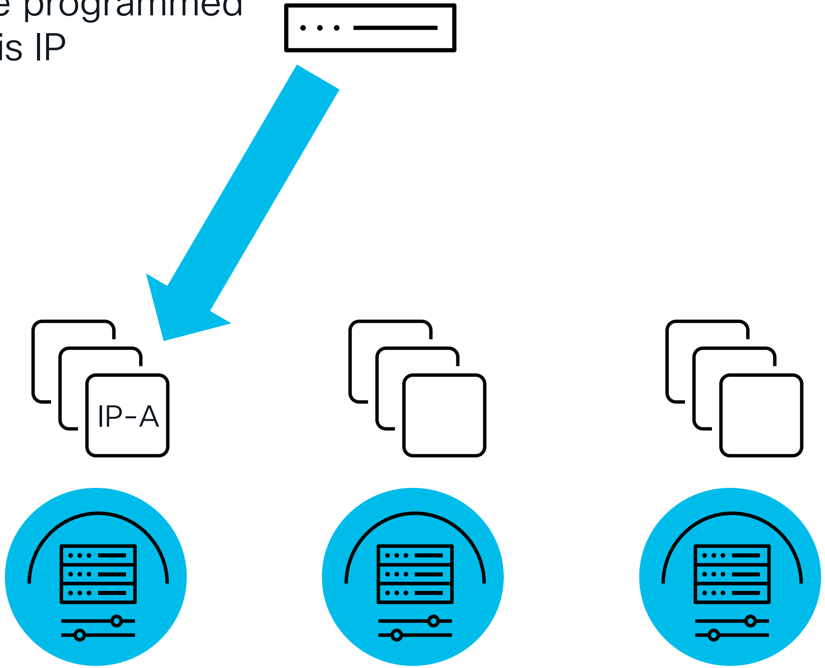

Usage of Persistent IPs

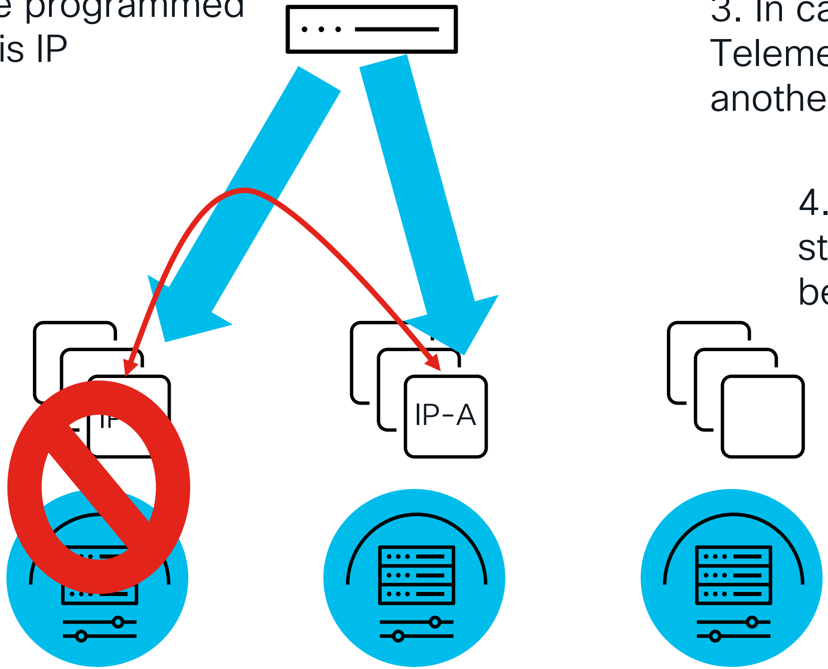

1. Telemetry Pods get an IP assigned (Persistent IP) at application start

2. Switches are programmed to stream to this IP

3. In case of a node failure, Telemetry Pod is moved to another Node

4. Switches continue to stream to this IP, without being reprogrammed

Persistent IP Pool 1/2

- Is needed to assign persistent IPs to Services/Apps
- These IPs are staying the same even the Service/App is moved to another ND Node
- Are entered as host IP addresses under Cluster Configuration->External Service Pools
- Used by NDI and NDFC

Persistent IP Pool 2/2

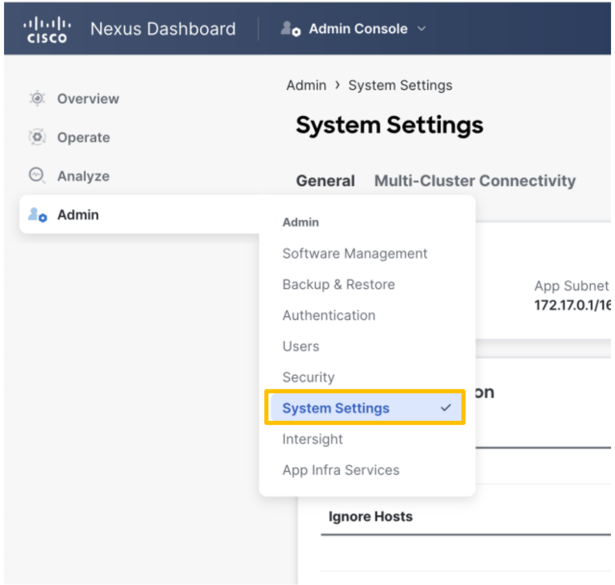

Admin › System Settings

System Settings

General | Multi-Cluster Connectivity

Admin
- Software Management
- Backup & Restore
- Authentication
- Users
- Security
- System Settings ✓
- Intersight
- App Infra Services

App Subnet
172.17.0.1/16

Ignore Hosts

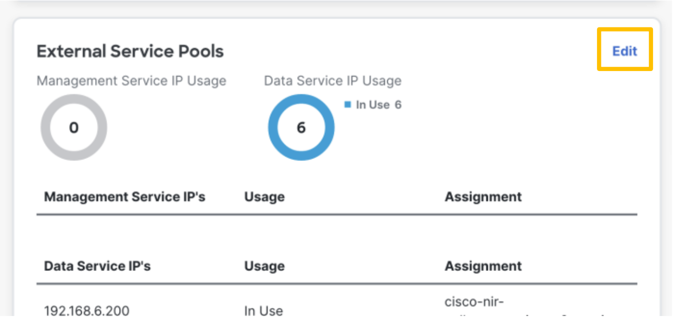

External Service Pools

Edit

Management Service IP Usage: 0

Data Service IP Usage: 6 (In Use 6)

Management Service IP's	Usage	Assignment

Data Service IP's	Usage	Assignment
192.168.6.200	In Use	cisco-nir-

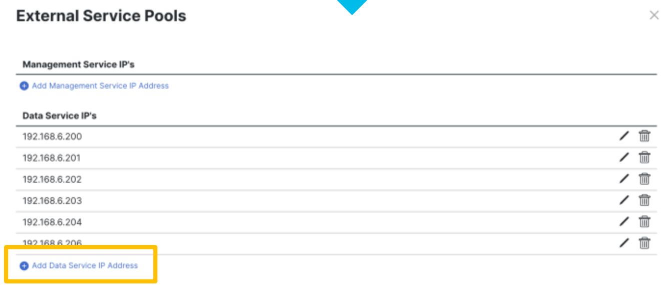

External Service Pools

Management Service IP's

Add Management Service IP Address

Data Service IP's

- 192.168.6.200
- 192.168.6.201
- 192.168.6.202
- 192.168.6.203
- 192.168.6.204
- 192.168.6.206

Add Data Service IP Address

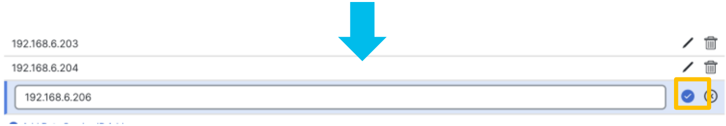

- 192.168.6.203
- 192.168.6.204
- 192.168.6.206

ND Persistent IP Connectivity Options

- For use of persistent IPs, there are now 2 choices:
  - 1. L2
    - All ND data interfaces are in the same subnet/L2 Domain and Persistent IPs are out of the same Network
  - 2. L3
    - All ND data interfaces can be in different subnets and have a BGP peering towards the network. Persistent IPs must not be out of any of these subnets.
    - ND nodes will only update the external peer with persistent IPs and not learn any prefixes. The local routing table will still be honored
    - Only supported on ND Data Interface

eBGP Peering with Network

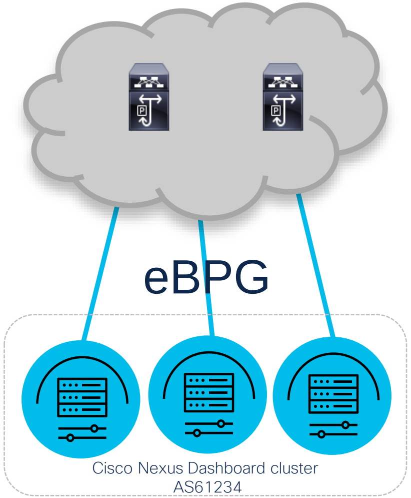

Reachability of Persitent IPs per ND Node

- Each ND node can be a separate AS or all in a single AS
- Multi-hop BGP peering is not supported
- Each ND node can peer to multiple Nodes (max 2) via IPv4 or IPv6
- Can be configured during bootstrap or added later
- Persitent IPs have to be out of an IP subnet not overlapping with any ND local IP.

Apps	Mgmt Interface	Data Interface	Persistent IPs	Support for Data and Mgmt in the same Subnet**
NDFC	L2 adjacent	L2 adjacent / L3 adjacent with L3 HA	2 IPs in mgmt network (for default settings) or 2 IPs data network (for POAP etc. via data network) + 1 IP per fabric for EPL in data network	no
NDI for DCNM based Sites	L3 adjacent	L2 adjacent	6 IPs in data network (+1 for IPv6)	no
NDI for ACI based Sites	L3 adjacent	L3 adjacent / L2 Adjacent	-/-	yes
NDI with SFLOW/Netflow function	L3 adjacent	L2 adjacent	6 IPs in data interface network*	no
NDO	L3 adjacent	L3 adjacent	-/-	yes

* if NDI is for DCNM no additional IPs are needed.

** supported but not recommended

Attaching ND to your Network

ND Cluster attached to any Networking Infra

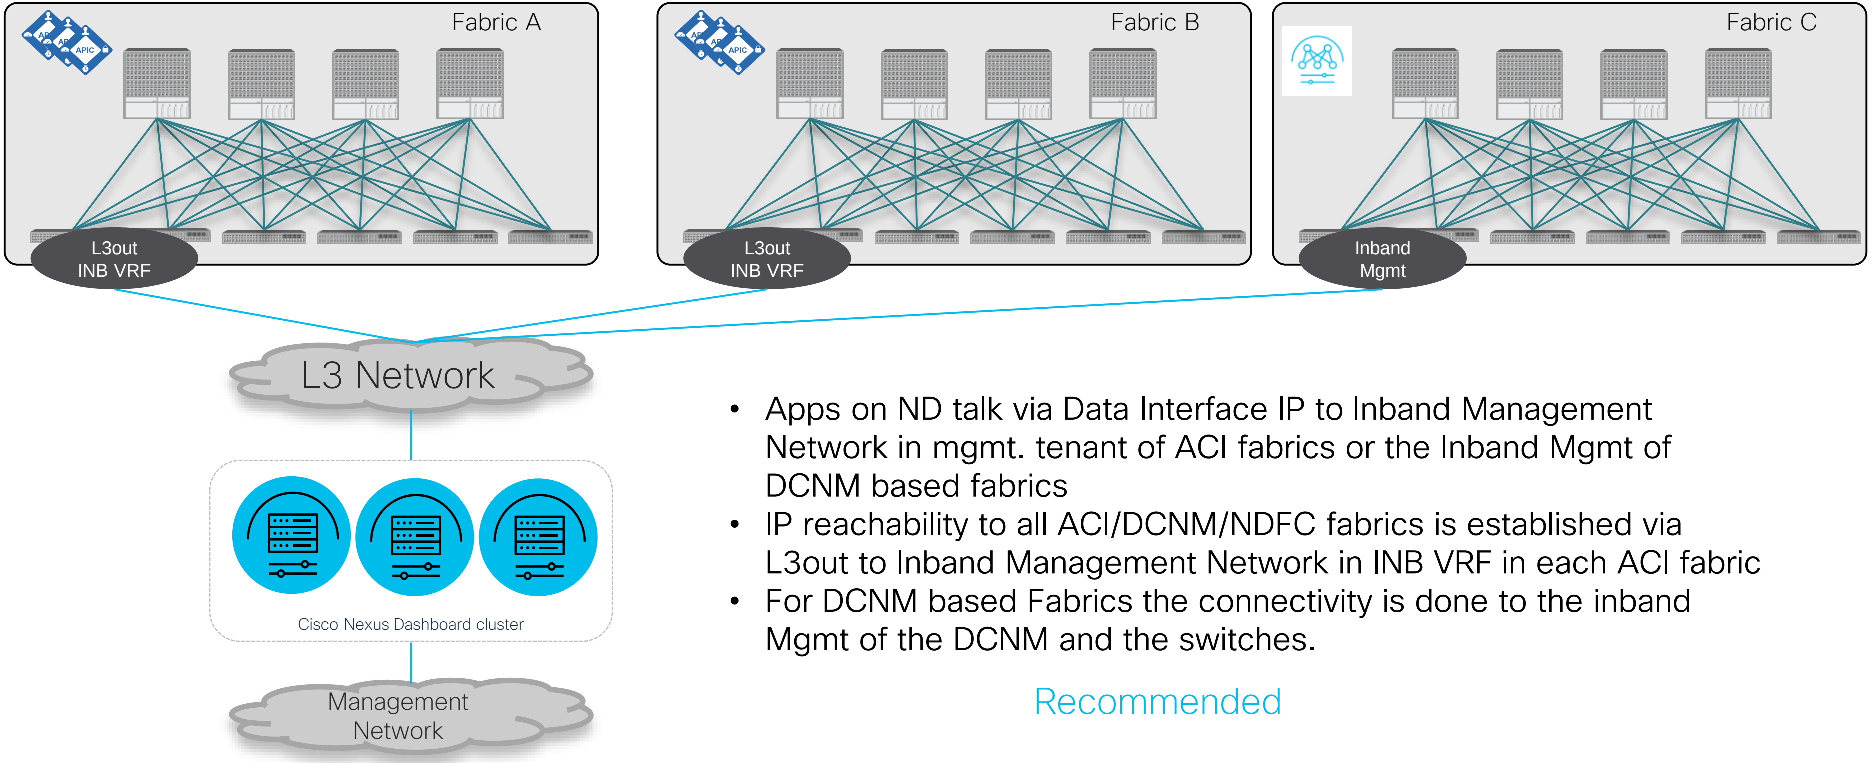

- Apps on ND talk via Data Interface IP to Inband Management Network in mgmt. tenant of ACI fabrics or the Inband Mgmt of DCNM based fabrics
- IP reachability to all ACI/DCNM/NDFC fabrics is established via L3out to Inband Management Network in INB VRF in each ACI fabric
- For DCNM based Fabrics the connectivity is done to the inband Mgmt of the DCNM and the switches.

Recommended

ND Cluster attached to DCNM/NDFC based Fabric

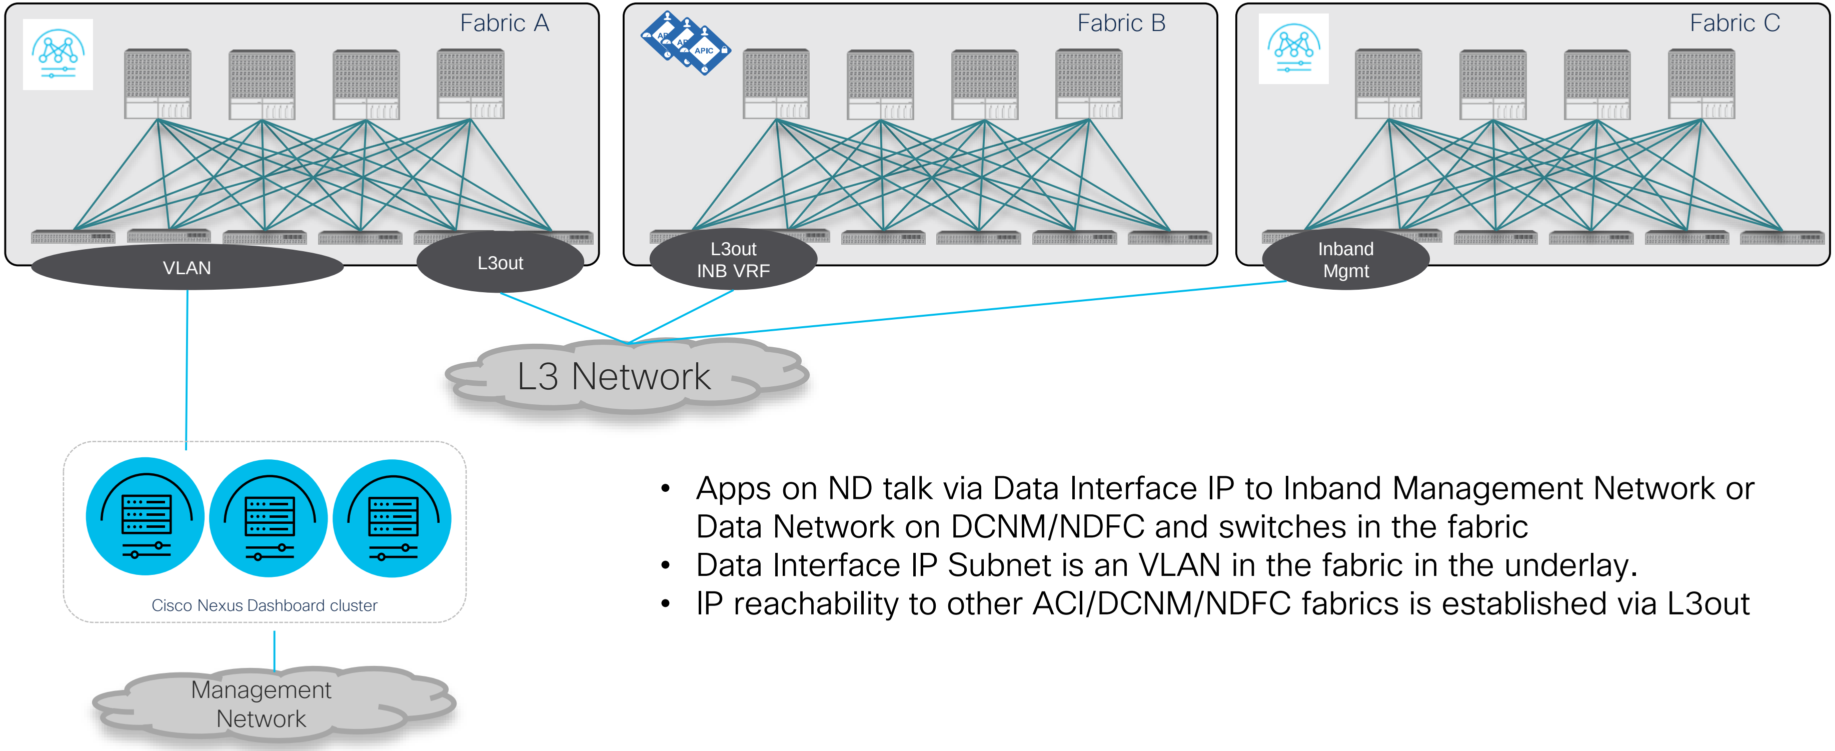

- Apps on ND talk via Data Interface IP to Inband Management Network or Data Network on DCNM/NDFC and switches in the fabric
- Data Interface IP Subnet is an VLAN in the fabric in the underlay.
- IP reachability to other ACI/DCNM/NDFC fabrics is established via L3out

ND Cluster attached to ACI Fabric

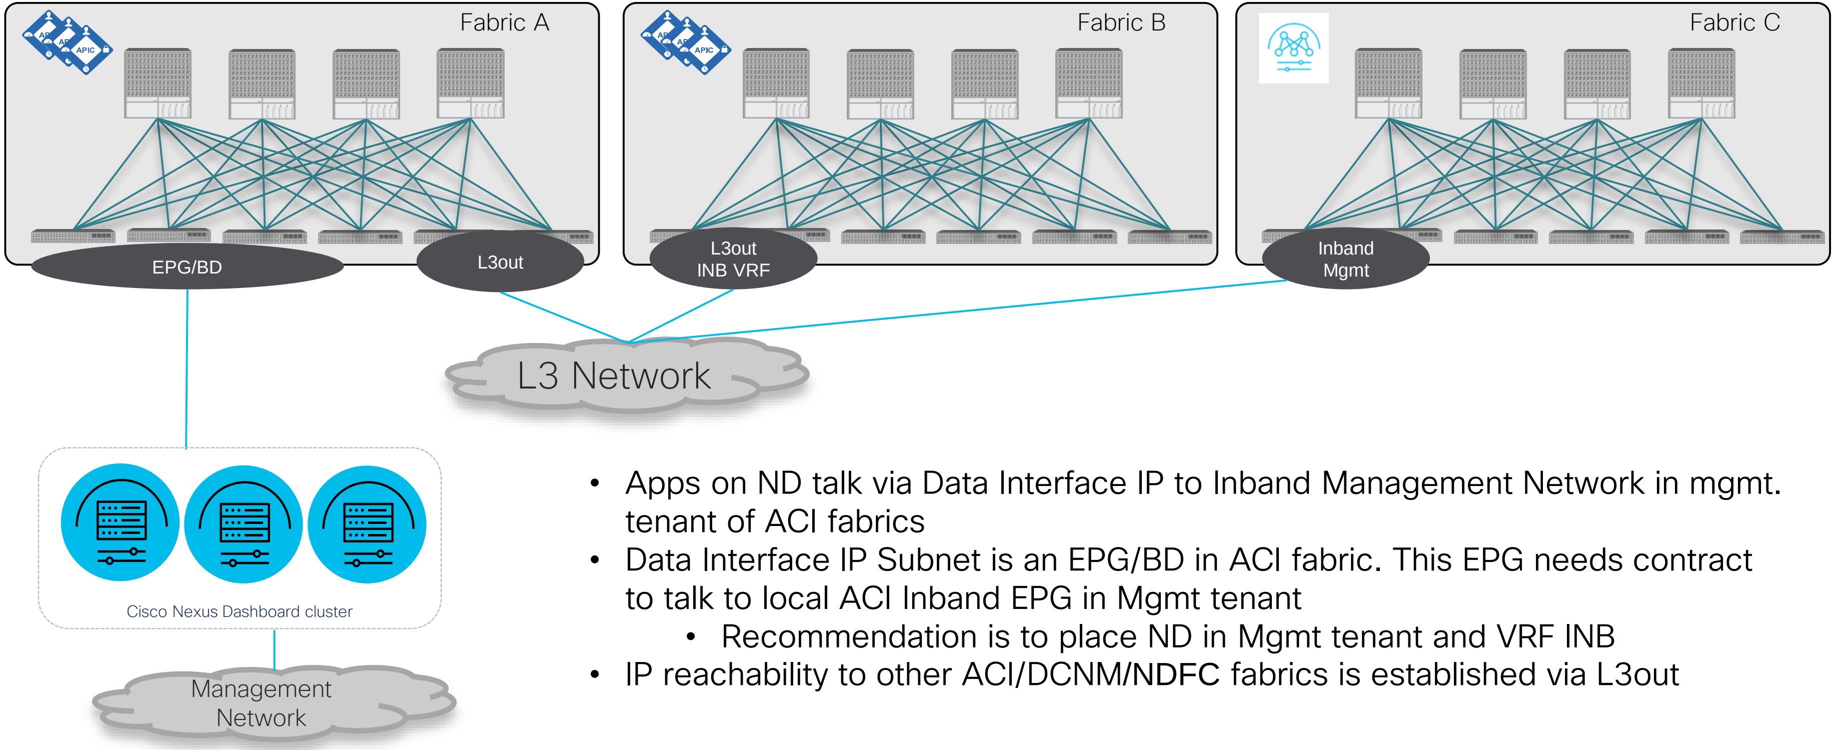

- Apps on ND talk via Data Interface IP to Inband Management Network in mgmt. tenant of ACI fabrics
- Data Interface IP Subnet is an EPG/BD in ACI fabric. This EPG needs contract to talk to local ACI Inband EPG in Mgmt tenant
  - Recommendation is to place ND in Mgmt tenant and VRF INB
- IP reachability to other ACI/DCNM/NDFC fabrics is established via L3out

Pro/Contra of connecting to an ACI/NDFC/DCNM fabric

Pro	Contra
- Easy connection between ND and Inband Management of ACI fabric	- ND cluster is tied to a single fabric - Reachability to other sites/fabrics has to go via L3out - ND cluster relies on single ACI fabric

Pro/Contra of connecting to any Networking Infra

Pro	Contra
- ND Cluster is not tied to any ACI Fabric - Same communication paths between all sites.	- All communications between ACI Apps on ND need to go via L3out

Recommendations/Best Practice

- Do not connect whenever possible to an ACI Fabric/DCNM based Fabric directly:
  - ND and Apps are relying on a functioning of the fabric, could be impacting during outages or maintenance
  - If you monitor multiple sites the ND cluster is not depend on a single site
- If a ND cluster is connected to a single fabric:
  - Fully supported/working BUT keep in mind
  - Issues in the fabric may impact the function of the ND cluster and the apps as they share fate.

Placement of Primary/Standby Nodes for Distribute/Stretched ND Clusters

(recommended for NDO)

Number of Sites	1	2	3	4	5
1	P1, P2, P3				
2	P1,P2	P3,S1			
3	P1	P2	P3		
4	P1	P2	P3	S1	
5	P1	P2	P3	S1	

P1, P2, P3 : ND Primary Nodes
S1 : ND Standby Node

When Centralized or Distributed/Stretched Cluster

Centralized	Distributed/Stretched
- With NDI/NDFC deployed	- For redundancy/DR for NDO
- NDI do not gain any better redundancy with distribute/stretched clusters. You more likely expose the cluster to interconnection failures with a distributed/stretched cluster	
- Synchronization traffic is kept between the ND nodes and only telemetry traffic is streamed via WAN	
- Same traffic path for reaching each site	
Recommended for NDI/NDFC	Recommended for NDO

Deployment Options for ND

Definition Terms and Assumptions/Requirements

- Site: geographical datacenter location with 1 or more fabrics
- RTT requirements for:
  - ND: between ND nodes <50ms
  - NDO : to APIC <500ms, to DCNM <50ms, between ND/NDO nodes <50ms
  - NDI: between ND/NDI nodes <50ms, to APIC/Fabric <50ms
  - NDFC: between ND/NDFC nodes <50ms, to Fabric <50ms (<200ms if no PoAP is used)
- Always select the lowest common denominator.
  - E.g. NDI and NDO co-hosted : between ND nodes <50ms, to APIC/Fabric <50ms

Deployment Requirements

- Customer has more than 1 Site
  - Number of ND clusters is driven by number of switches and combination of apps
    - Location of the ND clusters is driven by type of the apps:
    - NDO: cluster should be distributed for HA/DR reasons
    - NDI, NAE: cluster can be distributed, but should be placed close to source of telemetry data
    - Always keep virtual ND for NDO in consideration, to satisfy the HA/DR requirement
  - Please check the sizing calculator for ND for the supported apps and scale on CCO

Some Deployment Considerations

- In MPOD, ACI is taking care of the reachability, Keep in mind loosing IPN connectivity will e.g. break ND cluster
- In MSITE communication can not happen via ISN. It has to go via L3OUT in each site. Telemetry is sent via INB EPG in Mgmt Tenant, this is not managed by NDO!
- Data Interface IPs, have to be different from INB EPG/BD subnet of ACI, when ND cluster is connected to ACI fabric
- All communication of Apps hosted on ND is initiated via Data Interface IPs

HA/Redundancy with Stretched ND clusters

- 2 ND primary nodes are always needed to keep the ND cluster operational. If you deploy a stretched cluster across 2 sites, you SHOULD deploy in the site with a single ND primary node, a ND standby node.
- In case of a failure of 2 ND primary nodes, you have to manual promote the standby to Primary to replace a failed primary.
    - NDO/NDFC are the only apps surviving this.
    - App needs to be reinstalled
    - Backup of NDO/NDFC needs to be applied.
    - After the failed Primary comes back online, it needs to be wiped and re-added as standby node.

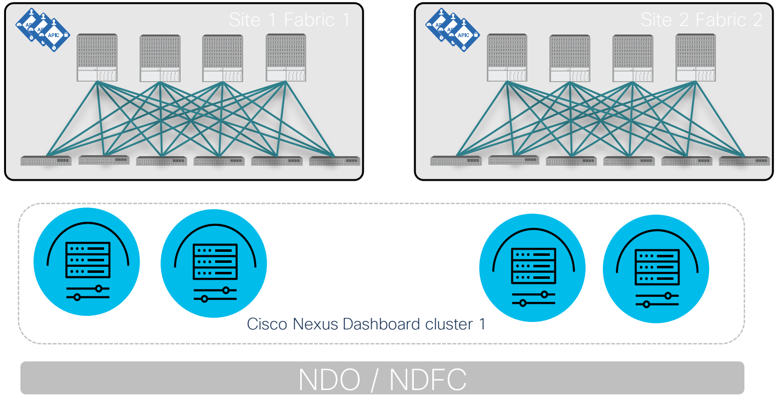

Option 1: 1 Site/Fabric (below 500 nodes) NDI

- Single cluster (x number of nodes, cluster connected to either ACI fabric or legacy infra with IP reachability)

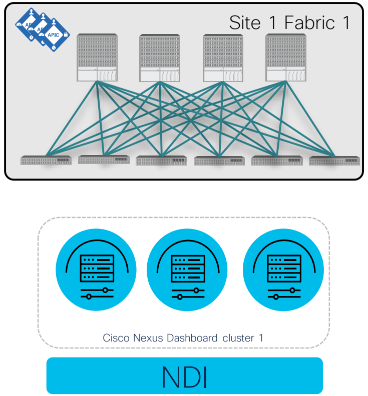

Option 2: 1+ Site (below 500 nodes) NDI

- Single cluster (x number of nodes, cluster connected to either ACI fabric or legacy infra with IP reachability, Cluster can be stretched or local to a site)

Recommended

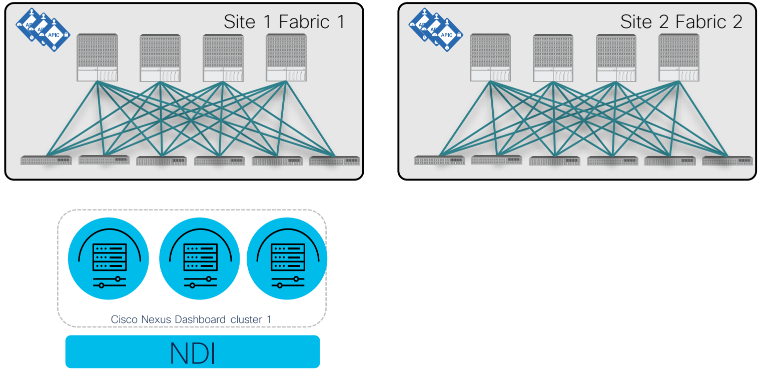

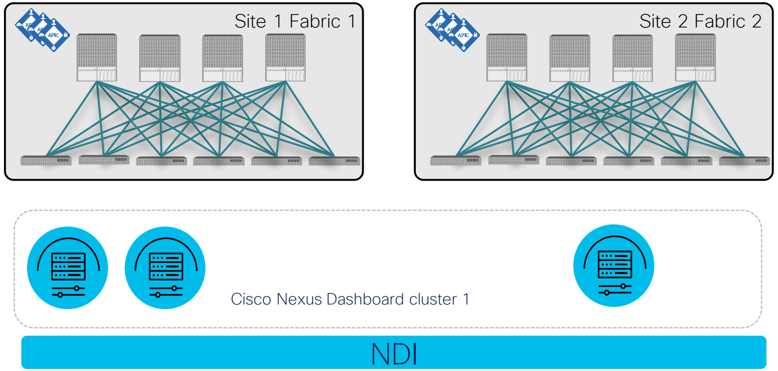

Option 3a: 1+ Site (below 500 nodes) NDI and NDO

- Single ND cluster for NDI (x number of nodes, cluster connected to either ACI fabric or legacy infra with IP reachability)
- Single additional virtual ND cluster for NDO to meet HA/DR requirements

Recommended

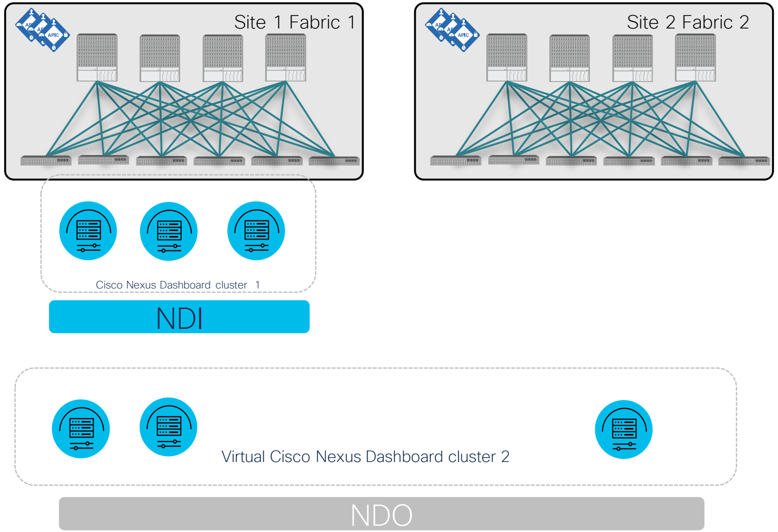
Site 1 Fabric 1
Site 2 Fabric 2
APIC
Cisco Nexus Dashboard cluster 1
NDI
Virtual Cisco Nexus Dashboard cluster 2
NDO

Option 3b: 1+ Site (below 500 nodes) NDI and NDO

- Single ND cluster (x number of nodes, cluster connected to either ACI fabric or legacy infra with IP reachability)

Not recommended as NDO is not distributed

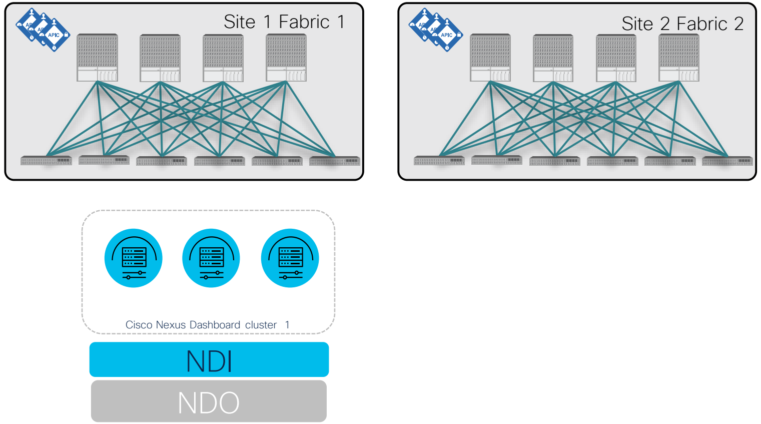

Site 1 Fabric 1

Site 2 Fabric 2

Cisco Nexus Dashboard cluster 1

NDI

NDO

Not recommended as NDI is distributed, consider vND for NDO (Option 3a)

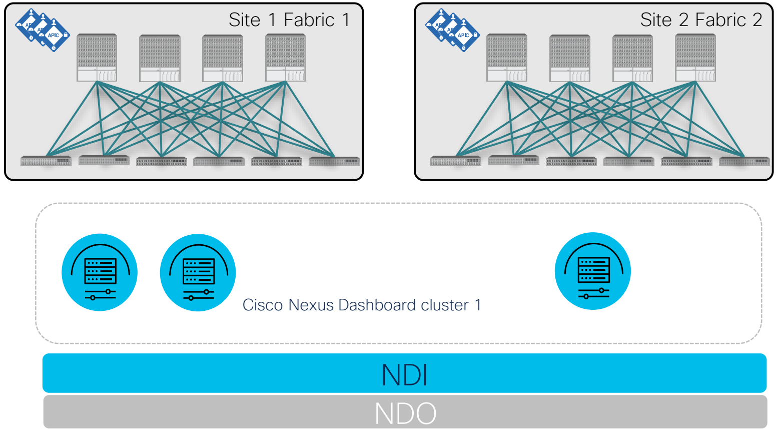

Site 1 Fabric 1

Site 2 Fabric 2

Cisco Nexus Dashboard cluster 1

NDI

NDO

Option 4: 1+ Site (above 500 nodes) NDI and NDO

- Multiple ND cluster (x number of nodes, cluster connected to either ACI fabric or legacy infra with IP reachability) and ND federation

Recommended

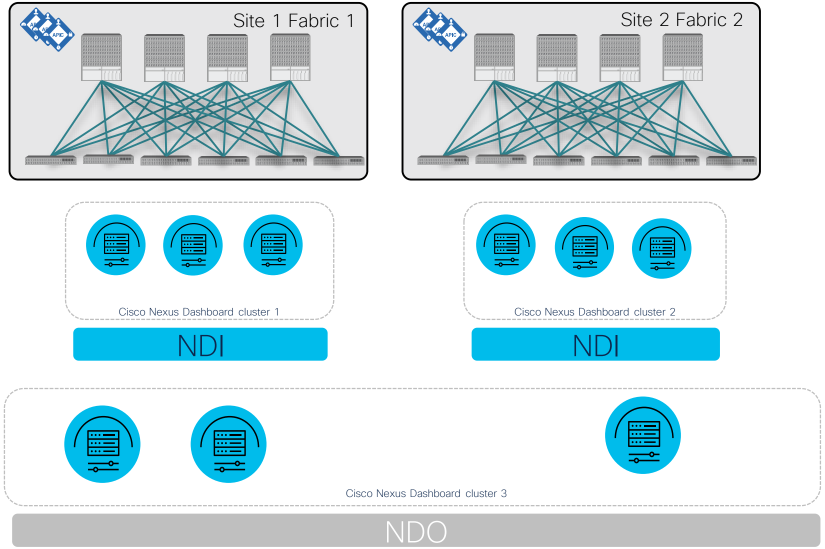
Site 1 Fabric 1
Site 2 Fabric 2
Cisco Nexus Dashboard cluster 1
NDI
Cisco Nexus Dashboard cluster 2
NDI
Cisco Nexus Dashboard cluster 3
NDO

Installing Nexus Dashboard with Unified Image

Choosing deployment mode during install

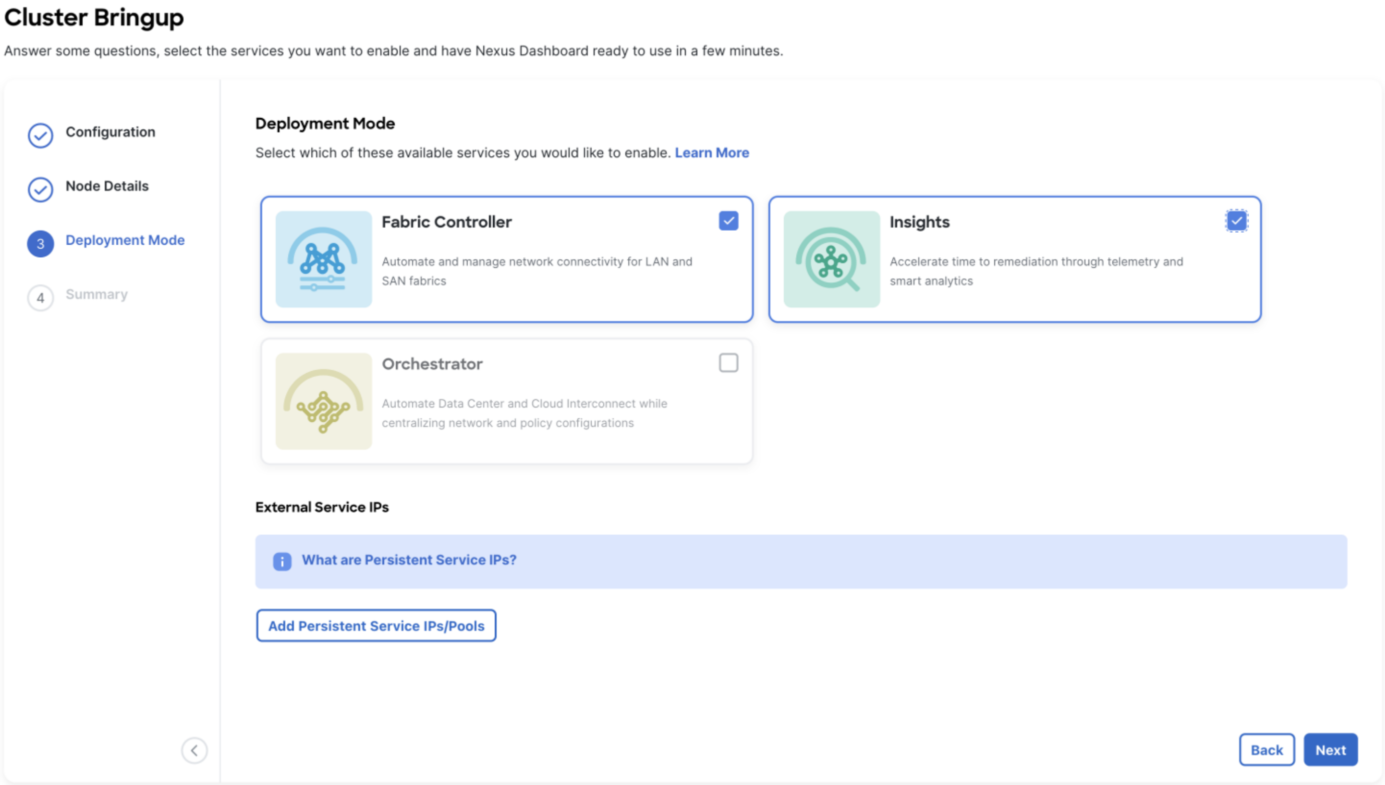

Cluster Bringup

Answer some questions, select the services you want to enable and have Nexus Dashboard ready to use in a few minutes.

- ✓ Configuration
- ✓ Node Details
- 3 Deployment Mode
- 4 Summary

Deployment Mode

Select which of these available services you would like to enable. Learn More

☑ **Fabric Controller**
Automate and manage network connectivity for LAN and SAN fabrics

☑ **Insights**
Accelerate time to remediation through telemetry and smart analytics

□ **Orchestrator**
Automate Data Center and Cloud Interconnect while centralizing network and policy configurations

External Service IPs

What are Persistent Service IPs?

ND Install Workflow

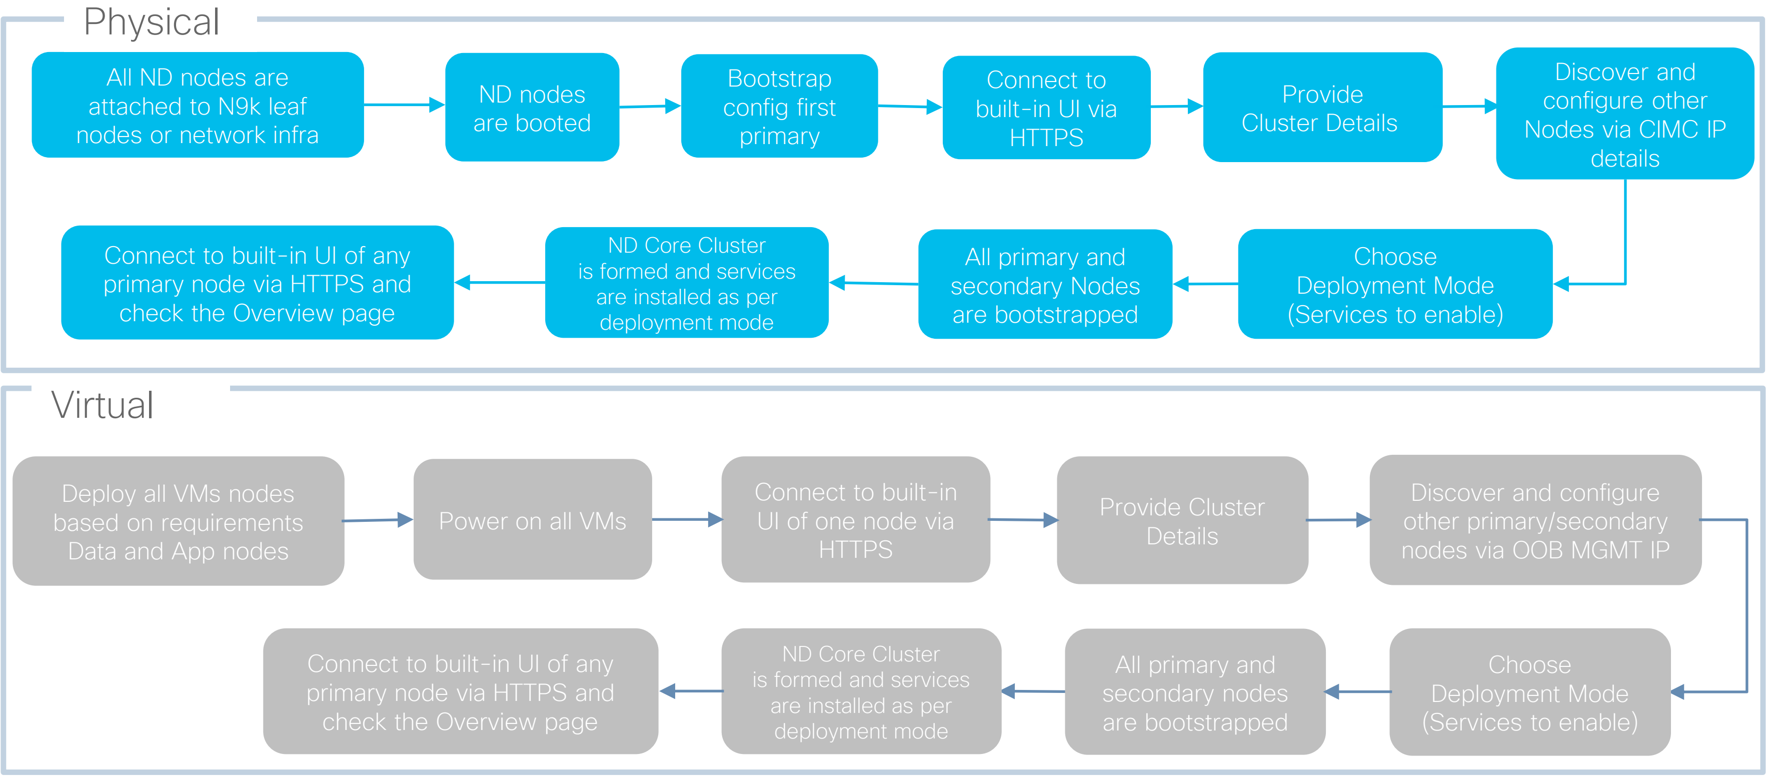

Physical

1. All ND nodes are attached to N9k leaf nodes or network infra
2. ND nodes are booted
3. Bootstrap config first primary
4. Connect to built-in UI via HTTPS
5. Provide Cluster Details
6. Discover and configure other Nodes via CIMC IP details
7. Choose Deployment Mode (Services to enable)
8. All primary and secondary Nodes are bootstrapped
9. ND Core Cluster is formed and services are installed as per deployment mode
10. Connect to built-in UI of any primary node via HTTPS and check the Overview page

Virtual

1. Deploy all VMs nodes based on requirements Data and App nodes
2. Power on all VMs
3. Connect to built-in UI of one node via HTTPS
4. Provide Cluster Details
5. Discover and configure other primary/secondary nodes via OOB MGMT IP
6. Choose Deployment Mode (Services to enable)
7. All primary and secondary nodes are bootstrapped
8. ND Core Cluster is formed and services are installed as per deployment mode
9. Connect to built-in UI of any primary node via HTTPS and check the Overview page

ND Unified Image/Installation Process

ND 3.0 and before

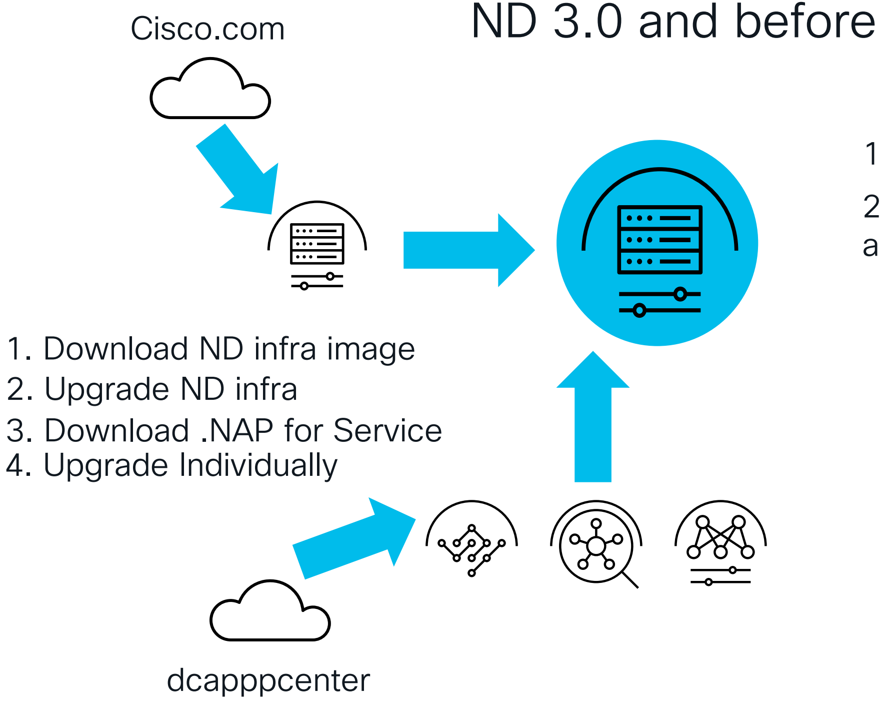

1. Download ND infra image
2. Upgrade ND infra
3. Download .NAP for Service
4. Upgrade Individually

ND 3.1 and later

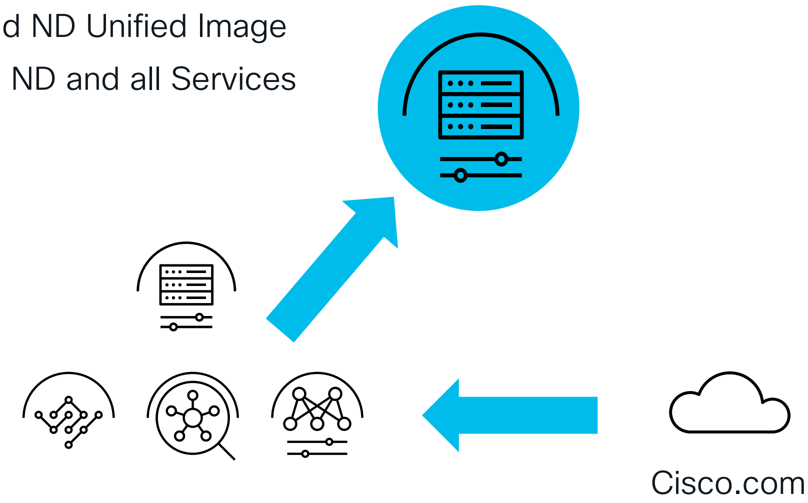

1. Download ND Unified Image
2. Upgrade ND and all Services at once*

* Based on Deployment Mode

Operating Nexus Dashboard

OneView aka as ND Federation

Overview

- ND Federation is an association of several ND clusters that allows working across with them as if they were a single entity and simplify the consumption of their resources
- ND clusters onboard other ND clusters creating a trusted environment which allows to learn about those clusters and to communicate and share information with each other
- Information shared between clusters is visible on each cluster being part of that federation. Also this data is accessible from each cluster.
- Apps can query for information related to other clusters in the federation for purposes such as onboarding (for eg NDI/Sites) or grouping
- Remote User is required to setup and use ND Federation

Federation Architecture

- User configures an ND cluster as Federation manager (FM) and connects it to other ND clusters
- FM manages the federation keeping track of member cluster reachability, node status, sites. etc.

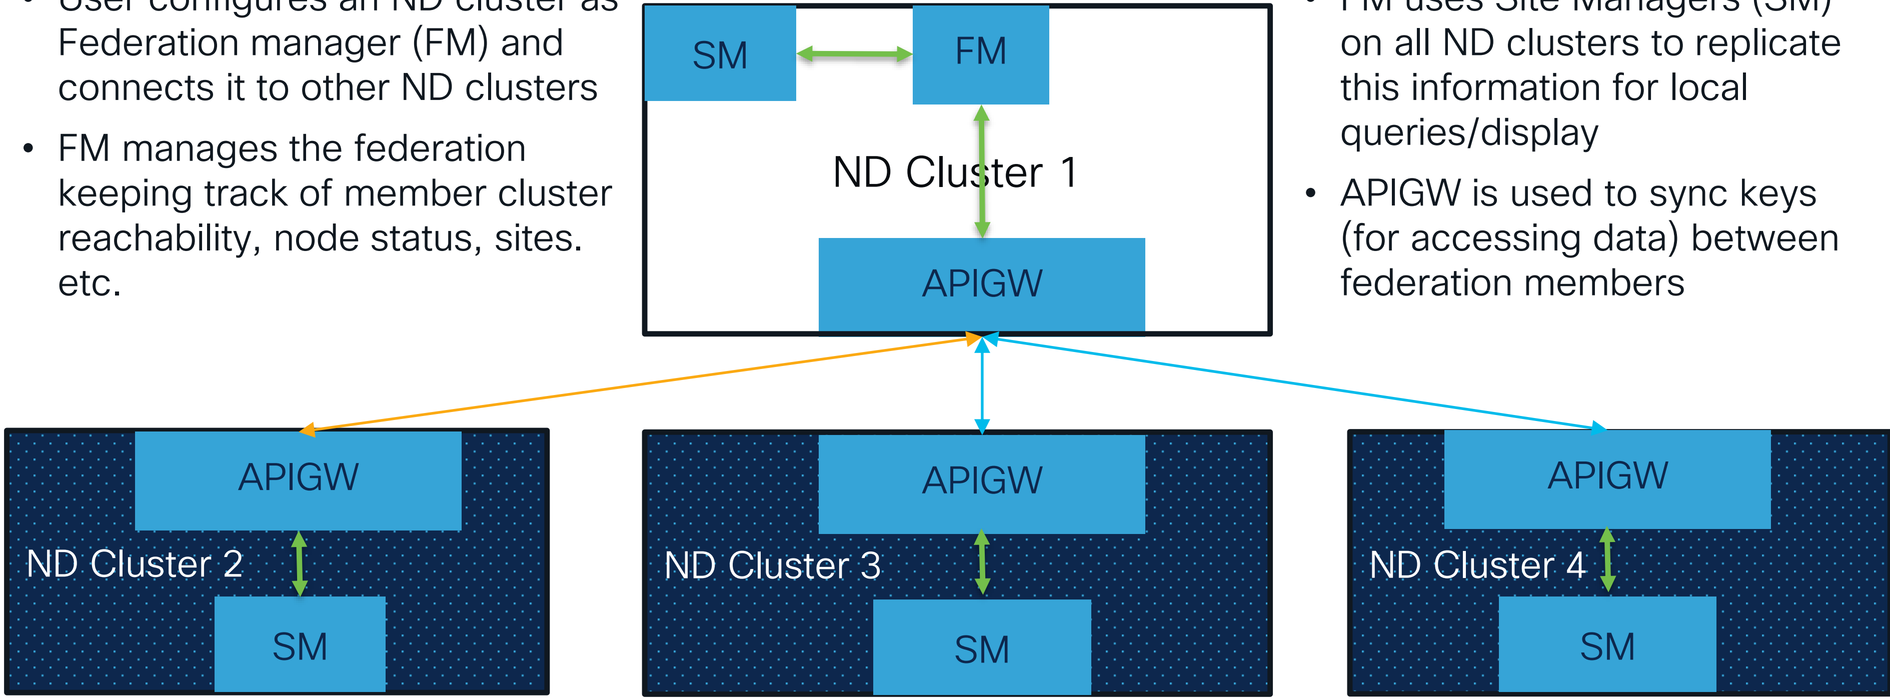

- FM uses Site Managers (SM) on all ND clusters to replicate this information for local queries/display
- APIGW is used to sync keys (for accessing data) between federation members

Onboard Clusters (Federation Configuration)

- Expand the Infrastructure menu
- Select Cluster Configuration
- Go to the Multi Cluster Connectivity tab
- Click “Connect Cluster”

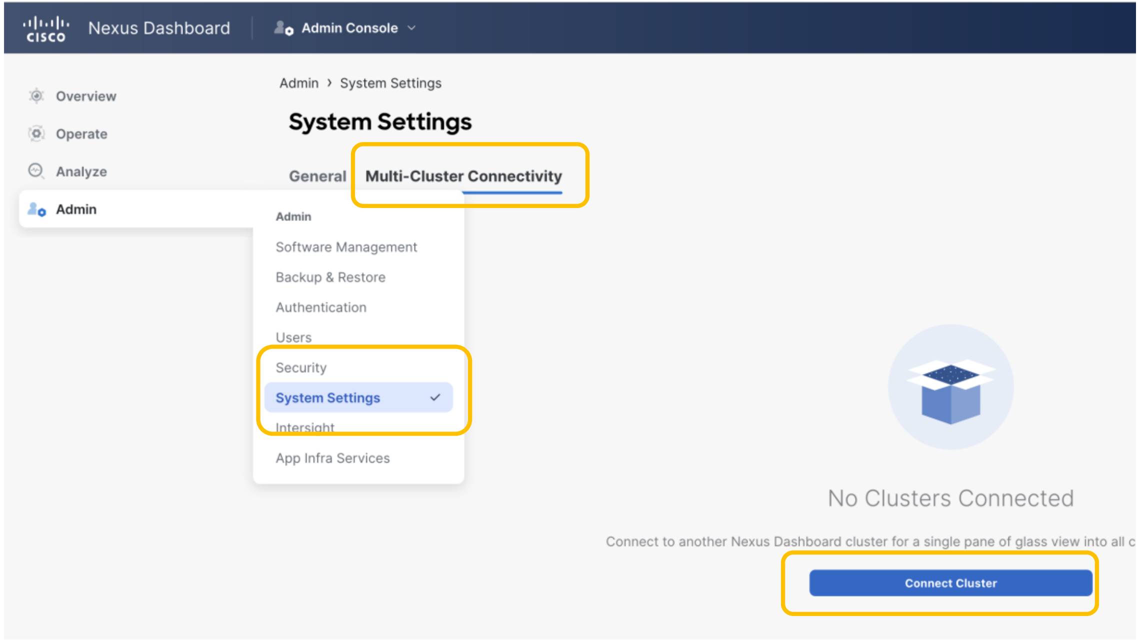

System Settings

No Clusters Connected

Connect to another Nexus Dashboard cluster for a single pane of glass view into all c

Onboard Clusters (Federation Configuration)

- Complete the target cluster information (IP of Mgmt Interface of remote cluster)
- Click save

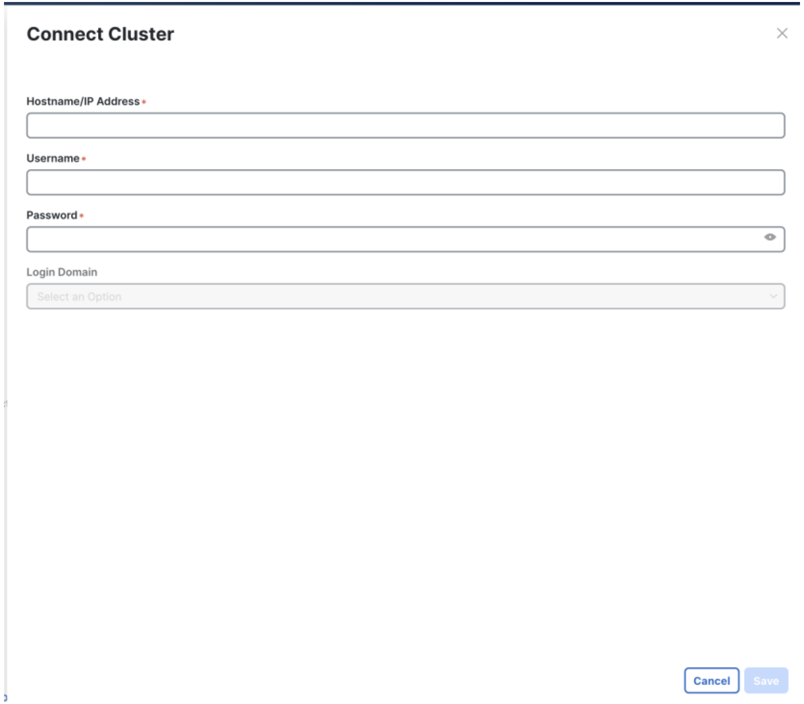

Connect Cluster

Hostname/IP Address *

Username *

Password *

Login Domain

Select an Option

Viewing Connected Clusters’ Information

- After connecting a cluster, it will show up on the Multi Cluster Connectivity table
- User would be able to connect more clusters or disconnect clusters from the table
- The cluster name on the header bar becomes a link to selecta specific cluster
- Central Dashboard is added to the header bar
- Local cluster and FM are marked in the list

Central Dashboard

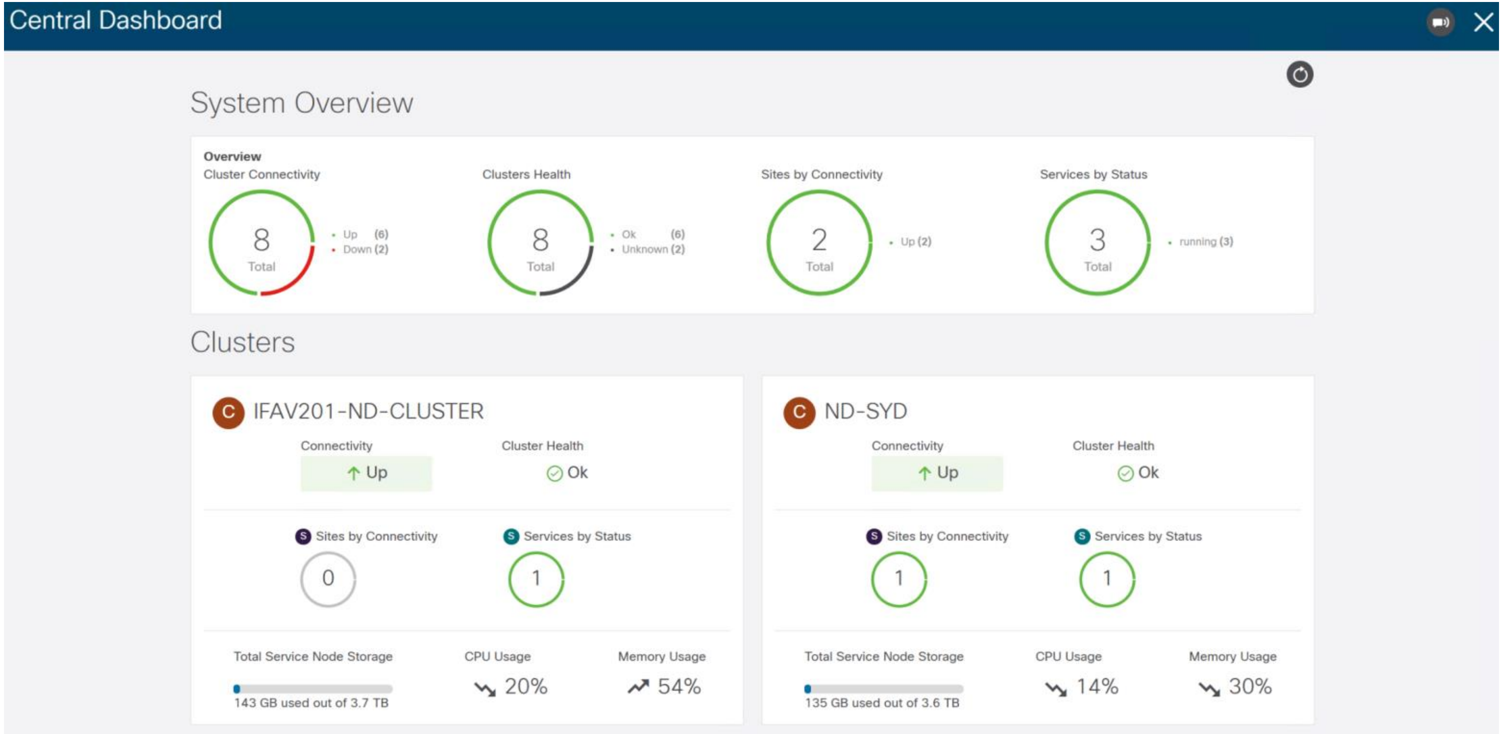

Central Dashboard

System Overview

Overview

Cluster Connectivity: 8 Total — Up (6), Down (2)

Clusters Health: 8 Total — Ok (6), Unknown (2)

Sites by Connectivity: 2 Total — Up (2)

Services by Status: 3 Total — running (3)

Clusters

C IFAV201-ND-CLUSTER

Connectivity: Up

Cluster Health: Ok

Sites by Connectivity: 0

Services by Status: 1

Total Service Node Storage: 143 GB used out of 3.7 TB

CPU Usage: 20%

Memory Usage: 54%

C ND-SYD

Connectivity: Up

Cluster Health: Ok

Sites by Connectivity: 1

Services by Status: 1

Total Service Node Storage: 135 GB used out of 3.6 TB

CPU Usage: 14%

Memory Usage: 30%

Public API

Overview

- API publicly available
- Swagger built-in
- Apps onboarded to ND populate their APIs there as well (e.g. NDI)

API UI

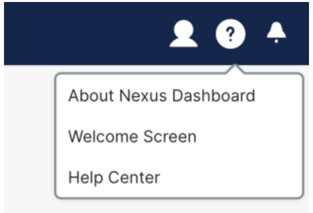

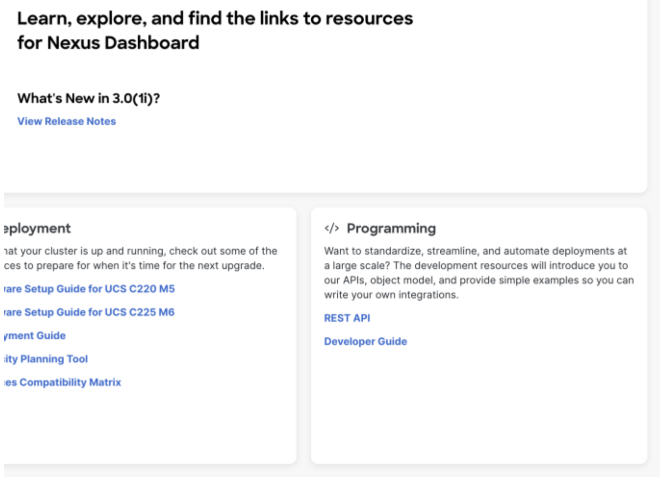

Learn, explore, and find the links to resources for Nexus Dashboard

What's New in 3.0(1i)?

View Release Notes

eployment

ıat your cluster is up and running, check out some of the ces to prepare for when it's time for the next upgrade.

ıare Setup Guide for UCS C220 M5

ıare Setup Guide for UCS C225 M6

ʏment Guide

ity Planning Tool

es Compatibility Matrix

</> Programming

Want to standardize, streamline, and automate deployments at a large scale? The development resources will introduce you to our APIs, object model, and provide simple examples so you can write your own integrations.

REST API

Developer Guide

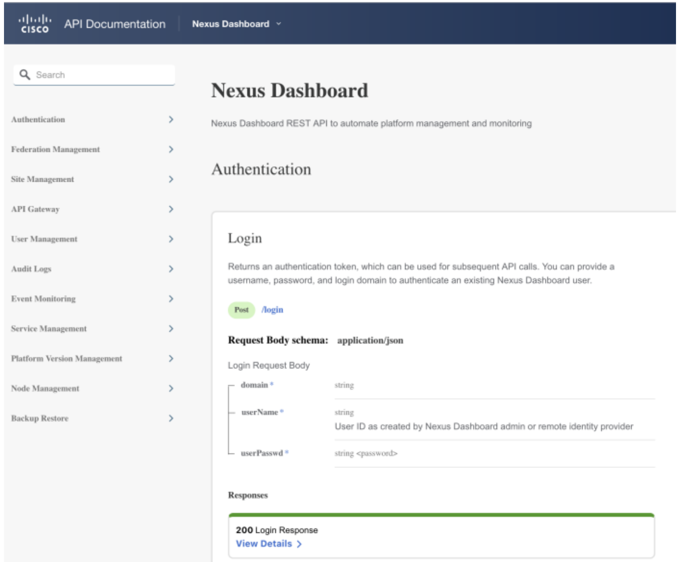

Nexus Dashboard

Nexus Dashboard REST API to automate platform management and monitoring

Authentication

Login

Returns an authentication token, which can be used for subsequent API calls. You can provide a username, password, and login domain to authenticate an existing Nexus Dashboard user.

Post /login

Request Body schema: application/json

Login Request Body

domain *	string
userName *	string User ID as created by Nexus Dashboard admin or remote identity provider
userPasswd *	string <password>

Responses

200 Login Response
View Details >

Registering Nodes to existing Cluster and Standby Node

Register new Nodes and Standby Primary

- New nodes are discovered via CIMC and bootstrapped
- During registration Role is selected (Worker or Standby)
- Worker Node is for horizontal Scaling
- Standby Node is increasing HA as it can replace a failed Primary
- Difference between Replace and Standby is, that Replace is a RMA workflow where the new node is installed and brought up. Standby is replacing a failed Primary with an already bootstrapped node
- Workers can only be replaced by delete and re-add

Lifecycle of non-Primary Nodes

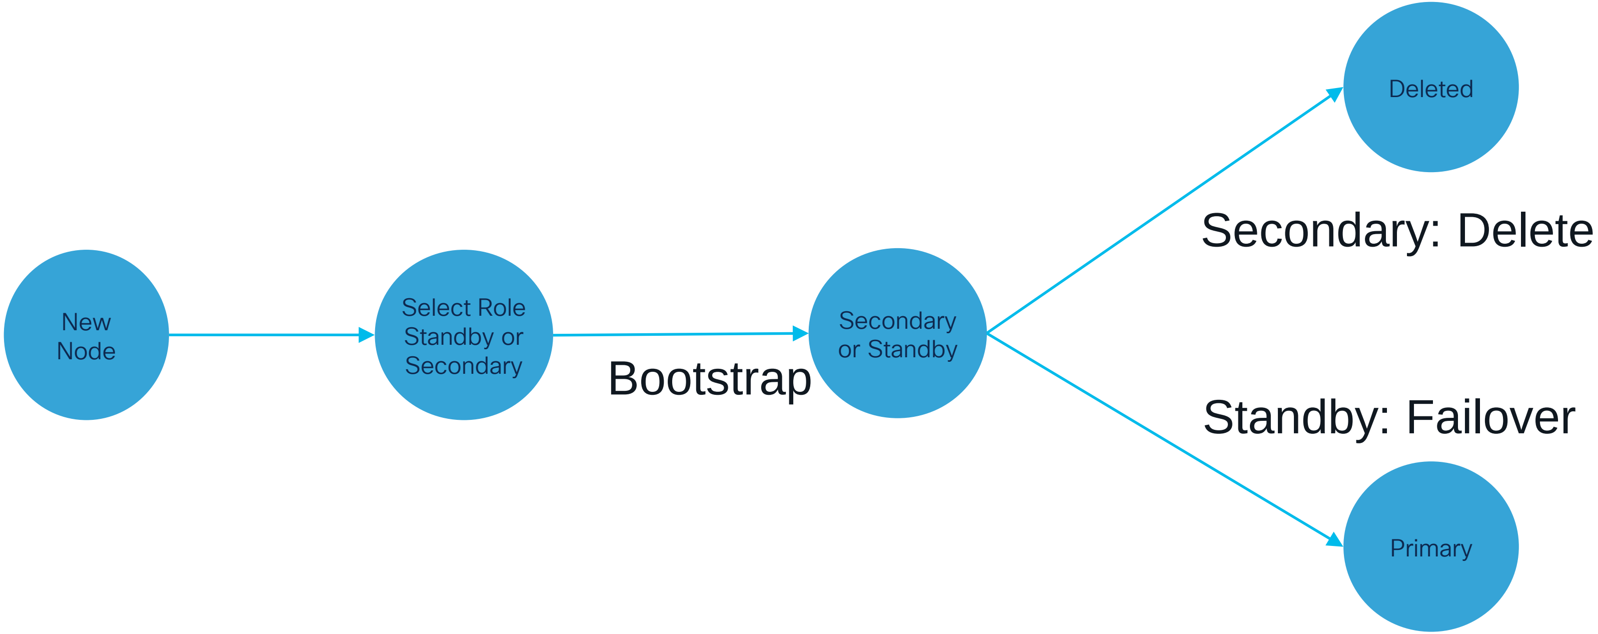

Adding a new Node

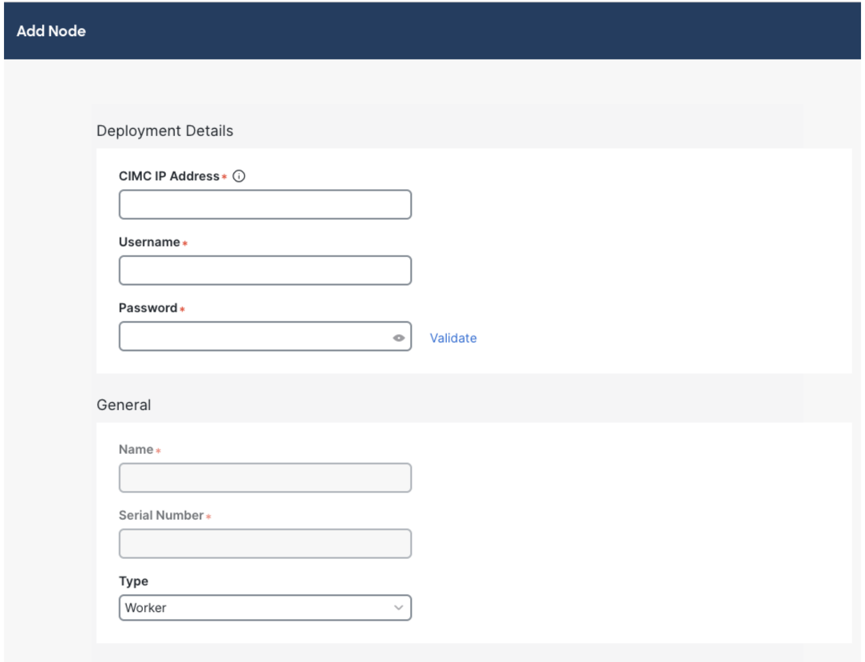

Add Node

Deployment Details

CIMC IP Address *

Username *

Password *

Validate

General

Name *

Serial Number *

Type

Worker

1. Provide CIMC details to discover node
2. Fill in node details
3. Node is bootstrapped and registered
4. Node status will change from "unregistered" to "discovering" to "active"

Replace a failed Primary with Standby Node

Primary is failed

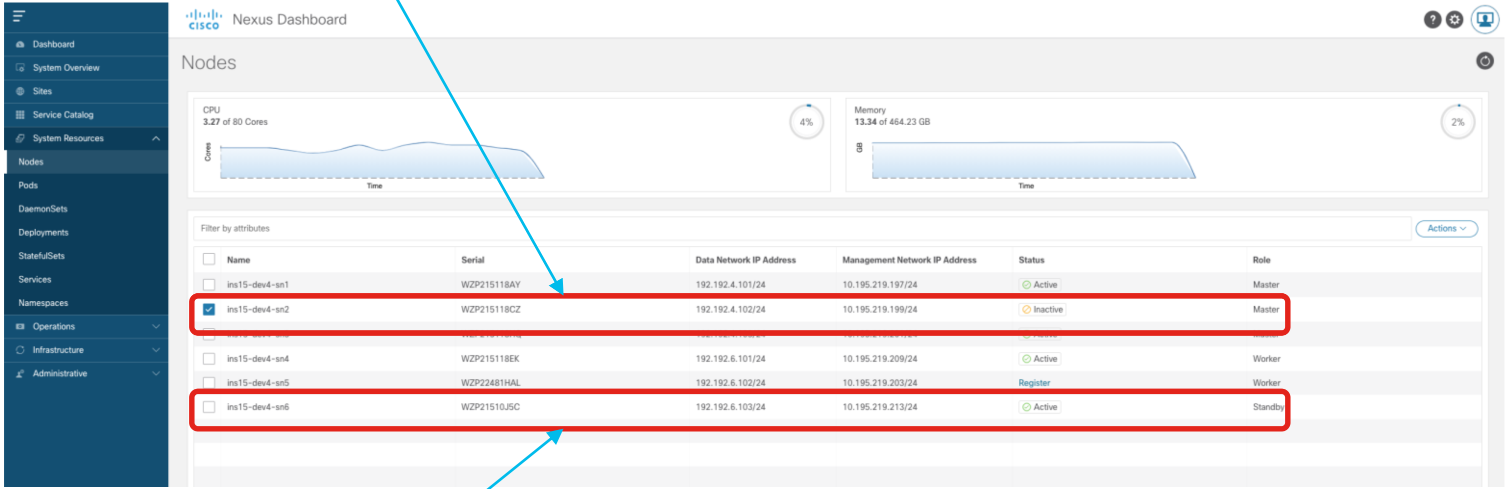

Nodes

☐	Name	Serial	Data Network IP Address	Management Network IP Address	Status	Role
☐	ins15-dev4-sn1	WZP215118AY	192.192.4.101/24	10.195.219.197/24	Active	Master
☑	ins15-dev4-sn2	WZP215118CZ	192.192.4.102/24	10.195.219.199/24	Inactive	Master
☐	ins15-dev4-sn3	[illegible]	192.192.4.103/24	[illegible]	Active	Master
☐	ins15-dev4-sn4	WZP215118EK	192.192.6.101/24	10.195.219.209/24	Active	Worker
☐	ins15-dev4-sn5	WZP22481HAL	192.192.6.102/24	10.195.219.203/24	Register	Worker
☐	ins15-dev4-sn6	WZP21510J5C	192.192.6.103/24	10.195.219.213/24	Active	Standby

Standby Node is part of Cluster

Failover to Standby

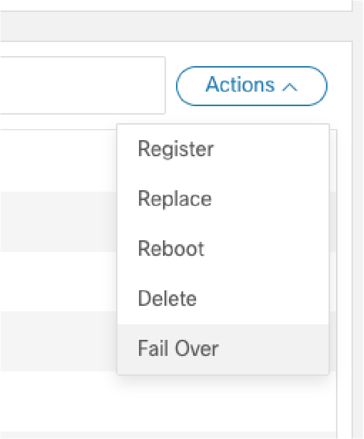

Select failed Primary and click Fail Over

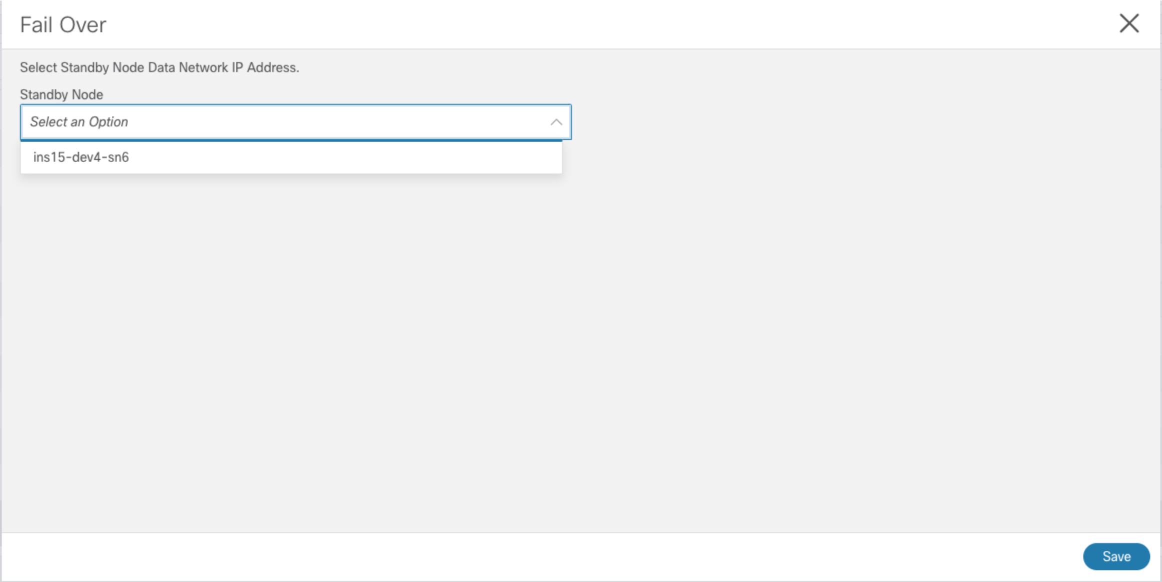

Select Standby Node Data Network IP Address.

Standby Node: *Select an Option*

ins15-dev4-sn6

Select Standby to replace failed Primary

If you receive a replacement for the failed node, you can register it as a Standby node

Manual Recovery of 2 failed Primaries

Recovery Process if 2 Primaries are down 1/3

- 2 Primary Nodes are failed
- 1 Standby Nodes are required to get the system back online
- Log in to the remaining primary
  - Run "`acs failover`" command to failover one of failed primary to standby

```
acs failover --failedIP <Primary-to-failover> \
             --failedIP <other-failed-Primary> \
             --standbyIP <standby-ip>
Note: Use inband ipaddress for above parameters
```

Recovery Process if 2 Primaries are down 2/3

- *acs cluster masters* will show 1 Active Primary and 2 Inactive Primaries

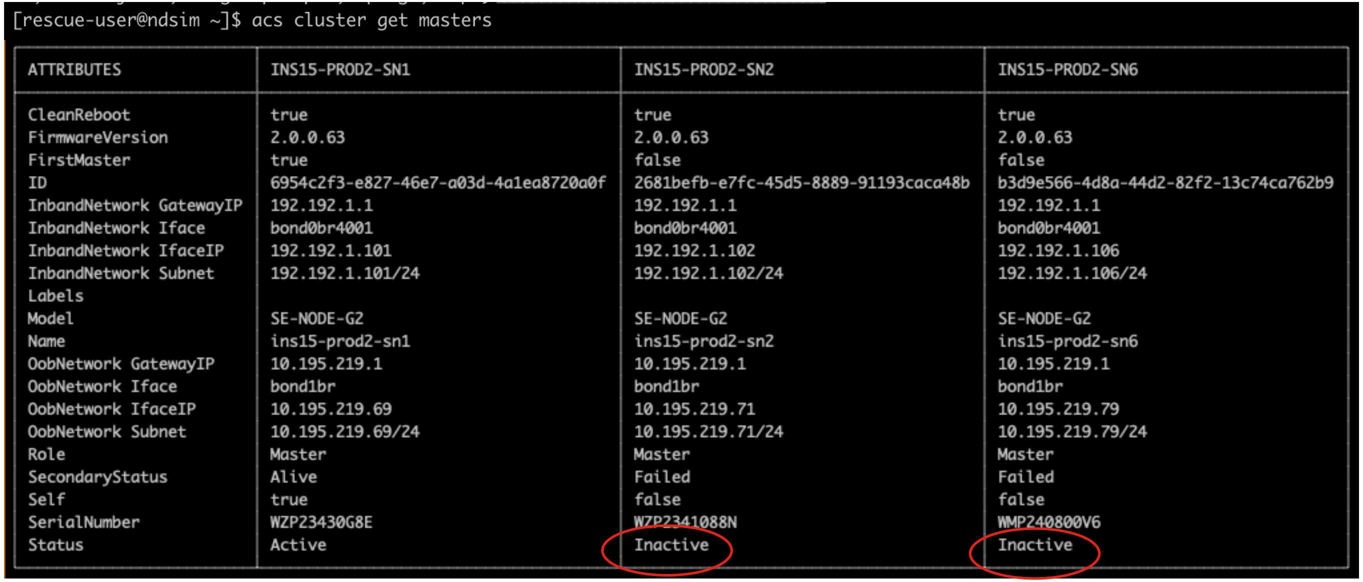

```
[rescue-user@ndsim ~]$ acs cluster get masters

 ATTRIBUTES             INS15-PROD2-SN1                        INS15-PROD2-SN2                        INS15-PROD2-SN6

 CleanReboot            true                                   true                                   true
 FirmwareVersion        2.0.0.63                               2.0.0.63                               2.0.0.63
 FirstMaster            true                                   false                                  false
 ID                     6954c2f3-e827-46e7-a03d-4a1ea8720a0f   2681befb-e7fc-45d5-8889-91193caca48b   b3d9e566-4d8a-44d2-82f2-13c74ca762b9
 InbandNetwork GatewayIP 192.192.1.1                           192.192.1.1                            192.192.1.1
 InbandNetwork Iface    bond0br4001                            bond0br4001                            bond0br4001
 InbandNetwork IfaceIP  192.192.1.101                          192.192.1.102                          192.192.1.106
 InbandNetwork Subnet   192.192.1.101/24                       192.192.1.102/24                       192.192.1.106/24
 Labels
 Model                  SE-NODE-G2                             SE-NODE-G2                             SE-NODE-G2
 Name                   ins15-prod2-sn1                        ins15-prod2-sn2                        ins15-prod2-sn6
 OobNetwork GatewayIP   10.195.219.1                           10.195.219.1                           10.195.219.1
 OobNetwork Iface       bond1br                                bond1br                                bond1br
 OobNetwork IfaceIP     10.195.219.69                          10.195.219.71                          10.195.219.79
 OobNetwork Subnet      10.195.219.69/24                       10.195.219.71/24                       10.195.219.79/24
 Role                   Master                                 Master                                 Master
 SecondaryStatus        Alive                                  Failed                                 Failed
 Self                   true                                   false                                  false
 SerialNumber           WZP23430G8E                            WZP2341088N                            WMP240800V6
 Status                 Active                                 Inactive                               Inactive
```

Recovery Process if 2 Primaries are down 3/3

- Command (both failed Primaries needs to be entered):

  acs failover --failedIP 192.192.1.102
  --failedIP 192.192.1.106
  --standbyIP 192.192.1.105

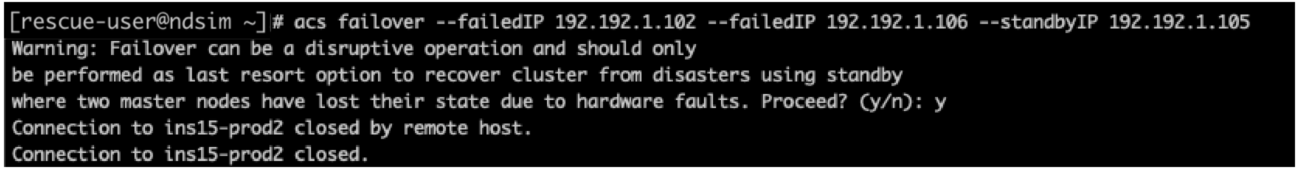

```
[rescue-user@ndsim ~]# acs failover --failedIP 192.192.1.102 --failedIP 192.192.1.106 --standbyIP 192.192.1.105
Warning: Failover can be a disruptive operation and should only
be performed as last resort option to recover cluster from disasters using standby
where two master nodes have lost their state due to hardware faults. Proceed? (y/n): y
Connection to ins15-prod2 closed by remote host.
Connection to ins15-prod2 closed.
```

- State will be copied from remaining Primary to Standby node
- Both nodes will reboot
- Standby node will reboot and come up as Primary

Recovery Process of a virtual ND

Recovery Process of a virtual ND

- Ensure that the failed node’s VM is powered down.
- Ensure new VM is deployed and powered on.
- Use the Replace workflow for the inactive node.

Firmware Upgrade

Firmware Upload

Click in Images first to upload a firmware image

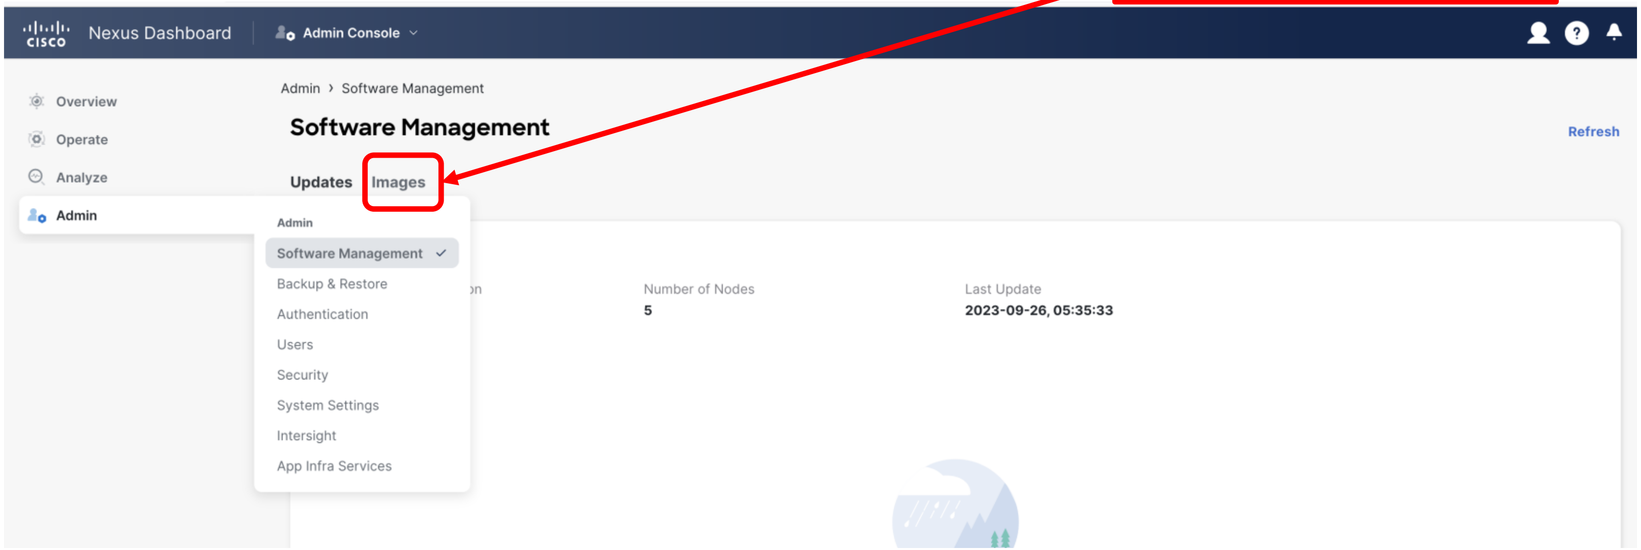

Admin › Software Management

Software Management

Number of Nodes
5

Last Update
2023-09-26, 05:35:33

Firmware Upload

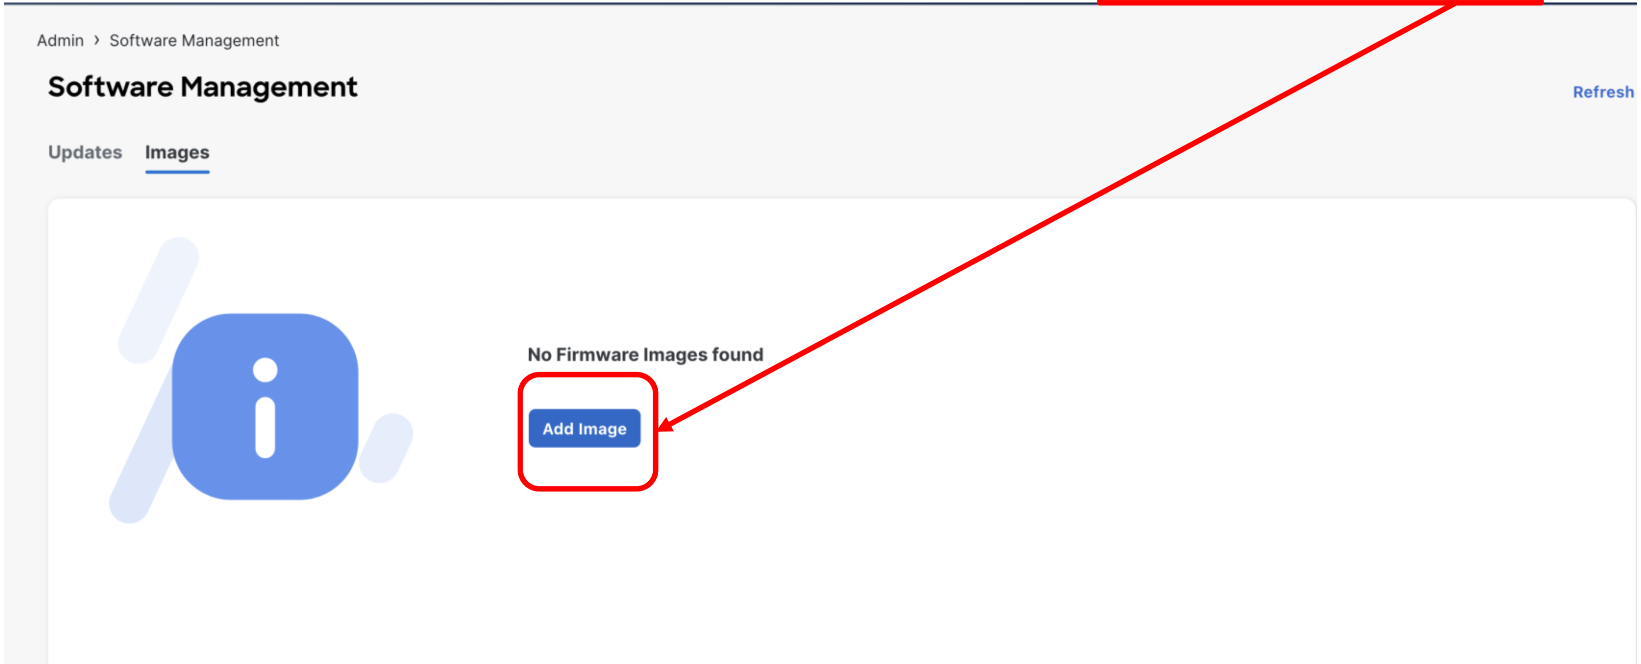

Admin › Software Management

Software Management

Refresh

Updates Images

No Firmware Images found

Add Image

Click Add Image

Firmware Upload

- 2 Options supported either via remote (WEB server) or local
- Remote upload is recommended

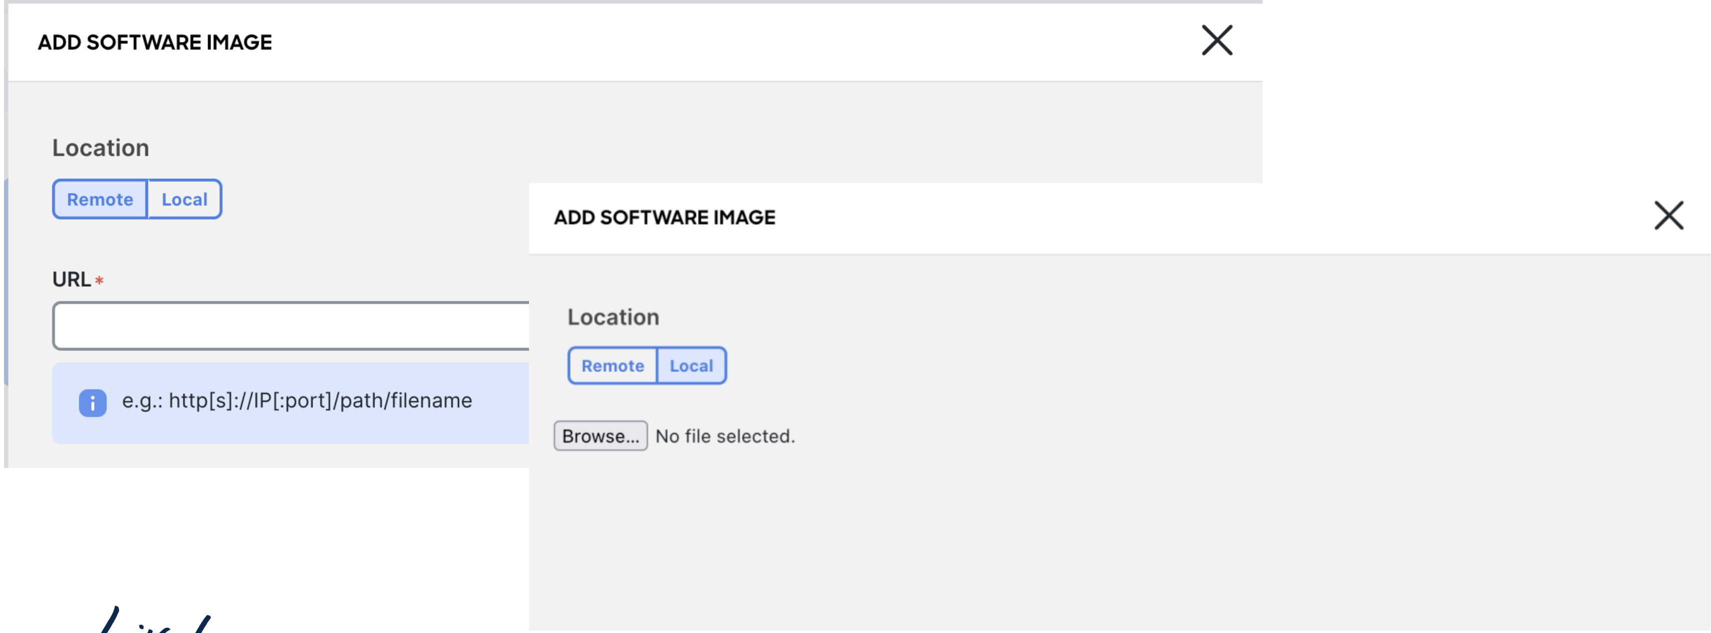

ADD SOFTWARE IMAGE

Location

Remote | Local

URL *

e.g.: http[s]://IP[:port]/path/filename

ADD SOFTWARE IMAGE

Location

Remote | Local

Browse... No file selected.

Firmware Upload

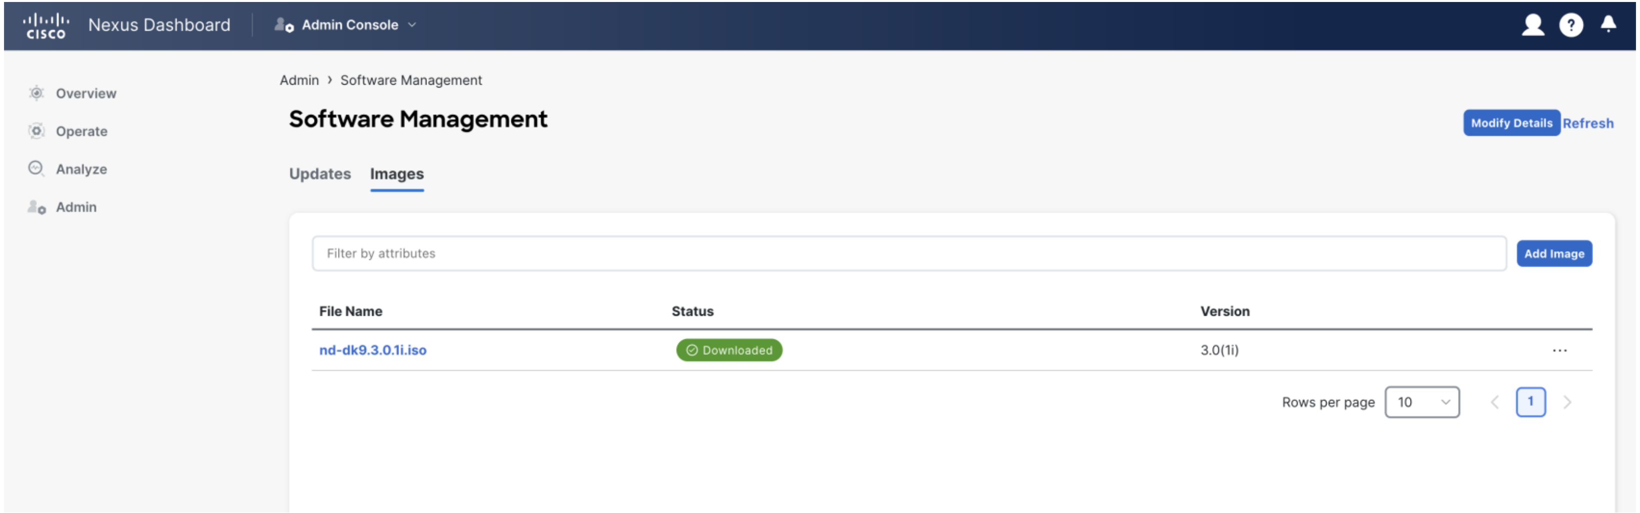

Software Management

File Name	Status	Version
nd-dk9.3.0.1i.iso	Downloaded	3.0(1i)

Setup Firmware Upgrade

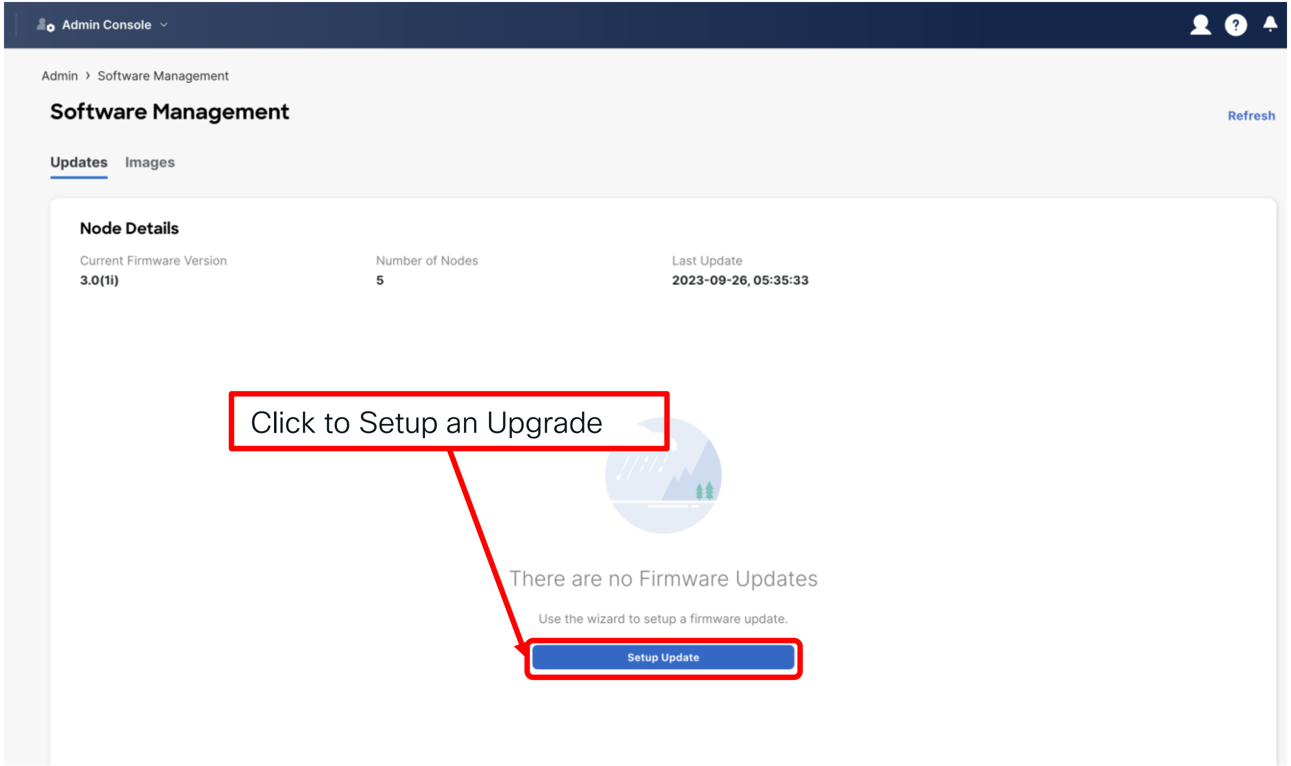

Admin Console

Admin › Software Management

Software Management

Refresh

Updates Images

Node Details

Current Firmware Version	Number of Nodes	Last Update
3.0(1i)	5	2023-09-26, 05:35:33

Click to Setup an Upgrade

There are no Firmware Updates

Use the wizard to setup a firmware update.

Select Firmware

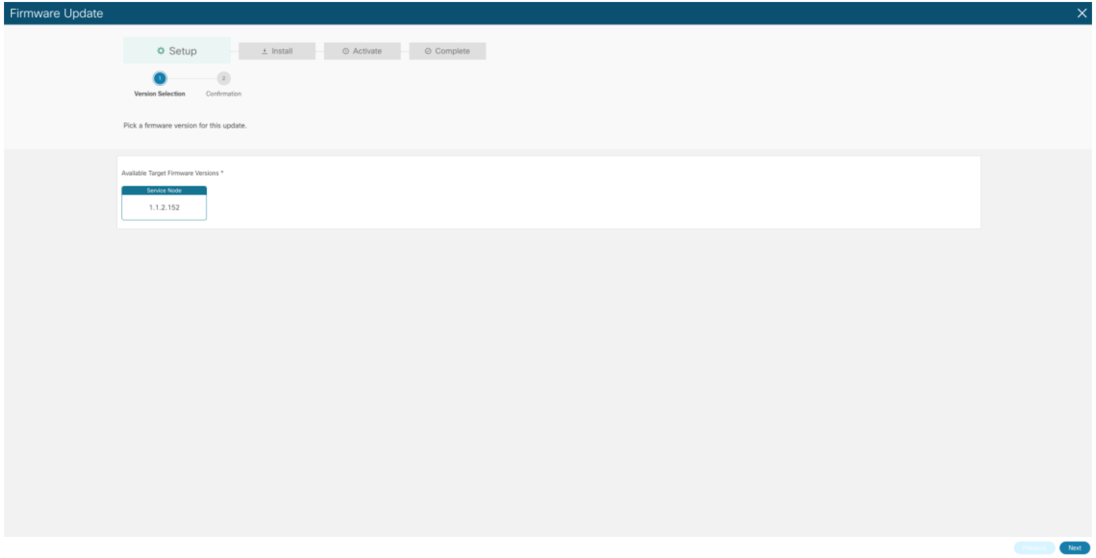

Pick a firmware version for this update.

Available Target Firmware Versions *

Service Node: 1.1.2.152

Current Cluster Setup is validated

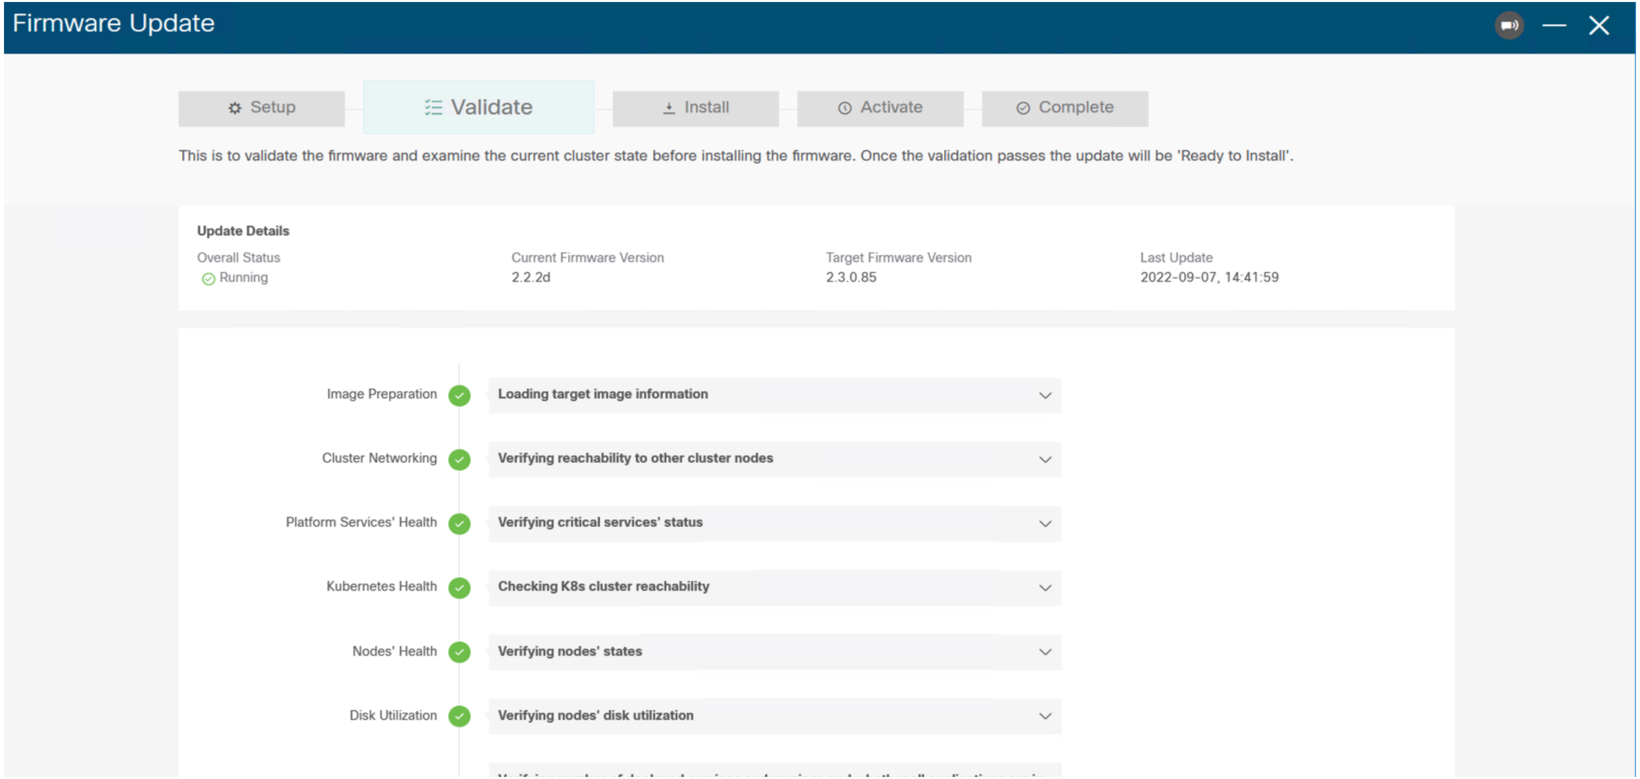

This is to validate the firmware and examine the current cluster state before installing the firmware. Once the validation passes the update will be 'Ready to Install'.

Update Details

Overall Status	Current Firmware Version	Target Firmware Version	Last Update
✓ Running	2.2.2d	2.3.0.85	2022-09-07, 14:41:59

Image Preparation	✓	**Loading target image information**
Cluster Networking	✓	**Verifying reachability to other cluster nodes**
Platform Services' Health	✓	**Verifying critical services' status**
Kubernetes Health	✓	**Checking K8s cluster reachability**
Nodes' Health	✓	**Verifying nodes' states**
Disk Utilization	✓	**Verifying nodes' disk utilization**
		Verifying number of deployed services and versions and whether all applications are in

Install Firmware to Nodes

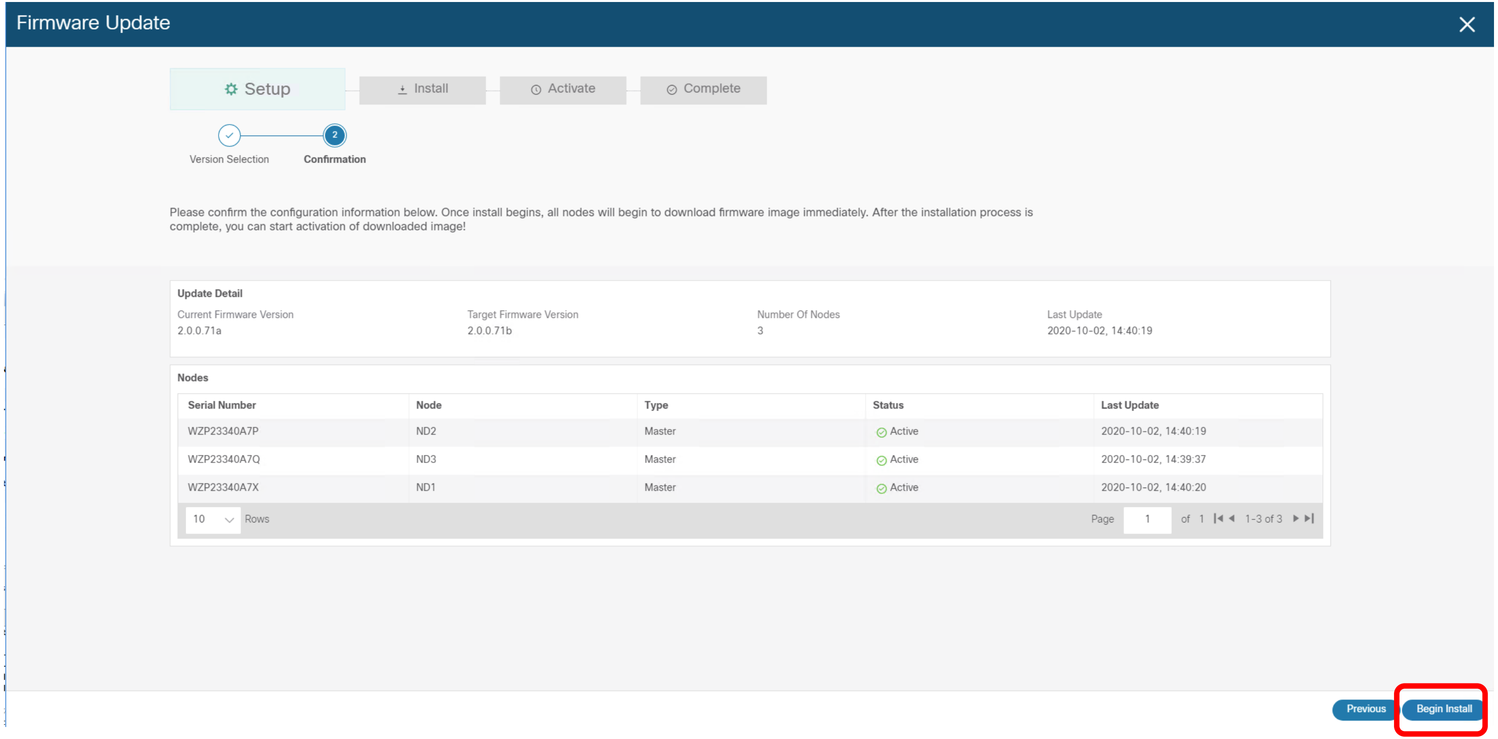

Please confirm the configuration information below. Once install begins, all nodes will begin to download firmware image immediately. After the installation process is complete, you can start activation of downloaded image!

Update Detail

Current Firmware Version	Target Firmware Version	Number Of Nodes	Last Update
2.0.0.71a	2.0.0.71b	3	2020-10-02, 14:40:19

Nodes

Serial Number	Node	Type	Status	Last Update
WZP23340A7P	ND2	Master	✓ Active	2020-10-02, 14:40:19
WZP23340A7Q	ND3	Master	✓ Active	2020-10-02, 14:39:37
WZP23340A7X	ND1	Master	✓ Active	2020-10-02, 14:40:20

Installing Firmware to Nodes

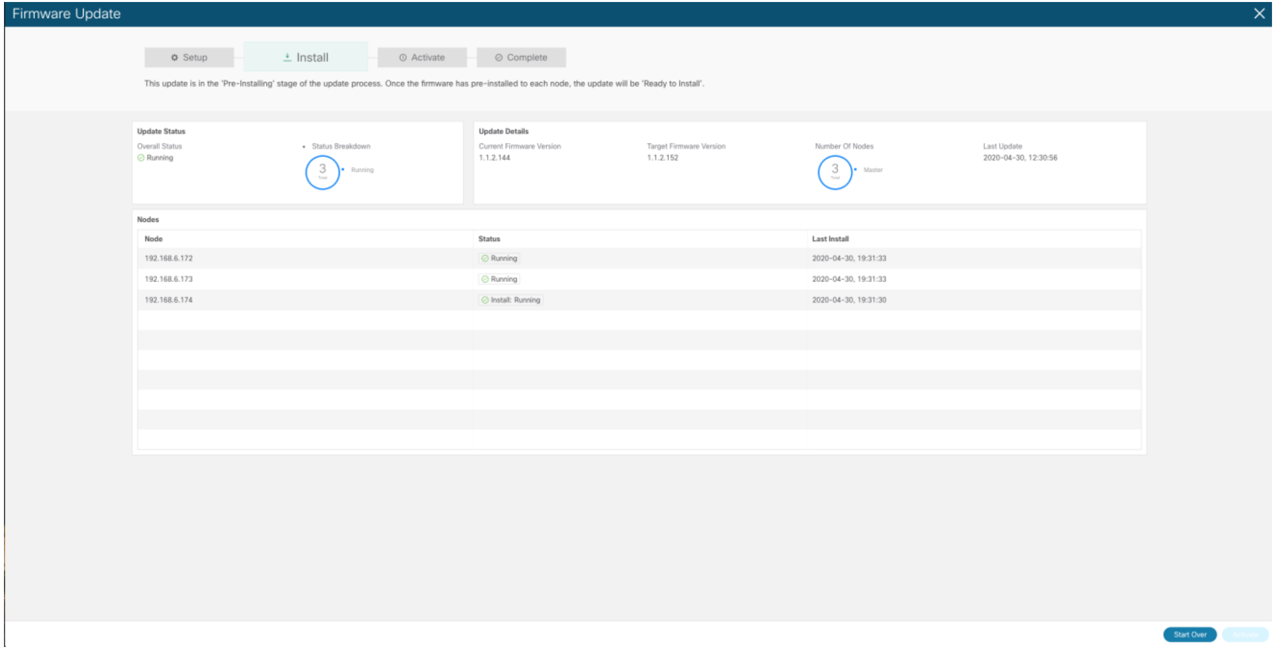

This update is in the 'Pre-Installing' stage of the update process. Once the firmware has pre-installed to each node, the update will be 'Ready to Install'.

Update Status

Overall Status
Running

Status Breakdown

Update Details

Current Firmware Version	Target Firmware Version	Number Of Nodes	Last Update
1.1.2.144	1.1.2.152	3 Total · Master	2020-04-30, 12:30:56

Nodes

Node	Status	Last Install
192.168.6.172	Running	2020-04-30, 19:31:33
192.168.6.173	Running	2020-04-30, 19:31:33
192.168.6.174	Install: Running	2020-04-30, 19:31:30

Once Install is done Click Activate

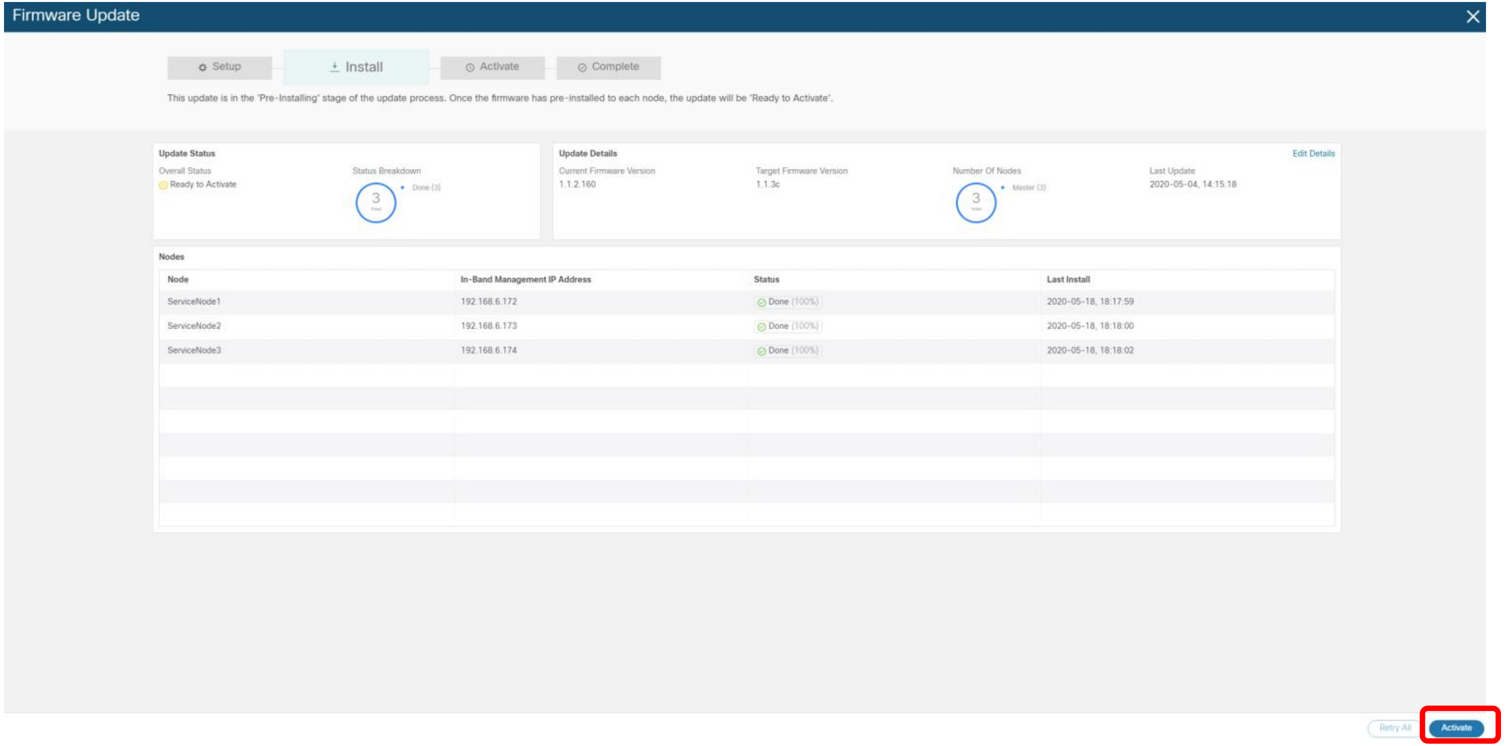

This update is in the 'Pre-Installing' stage of the update process. Once the firmware has pre-installed to each node, the update will be 'Ready to Activate'.

Update Status

Overall Status: Ready to Activate

Status Breakdown: 3 Total · Done (3)

Update Details

Current Firmware Version: 1.1.2.160

Target Firmware Version: 1.1.3c

Number Of Nodes: 3 Total · Master (3)

Last Update: 2020-05-04, 14:15:18

Nodes

Node	In-Band Management IP Address	Status	Last Install
ServiceNode1	192.168.6.172	Done (100%)	2020-05-18, 18:17:59
ServiceNode2	192.168.6.173	Done (100%)	2020-05-18, 18:18:00
ServiceNode3	192.168.6.174	Done (100%)	2020-05-18, 18:18:02

Activation Progress

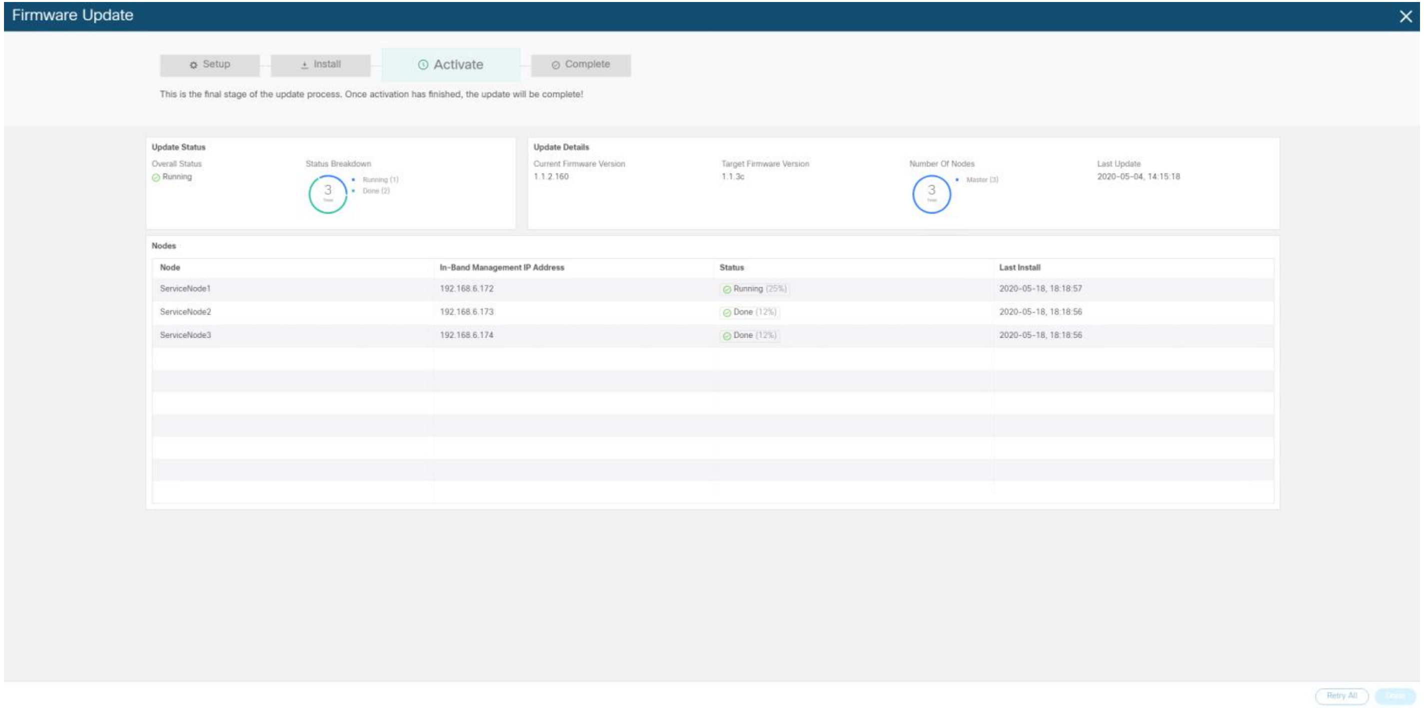

This is the final stage of the update process. Once activation has finished, the update will be complete!

Update Status

Overall Status: ✓ Running

Status Breakdown: 3 Total
- Running (1)
- Done (2)

Update Details

Current Firmware Version	Target Firmware Version	Number Of Nodes	Last Update
1.1.2.160	1.1.3c	3 Total Master (3)	2020-05-04, 14:15:18

Nodes

Node	In-Band Management IP Address	Status	Last Install
ServiceNode1	192.168.6.172	✓ Running (25%)	2020-05-18, 18:18:57
ServiceNode2	192.168.6.173	✓ Done (12%)	2020-05-18, 18:18:56
ServiceNode3	192.168.6.174	✓ Done (12%)	2020-05-18, 18:18:56

Monitoring Firmware Upgrade

- When the node you are connected to is activating, it will disconnect you. Please connect to another SE node. Check status via:

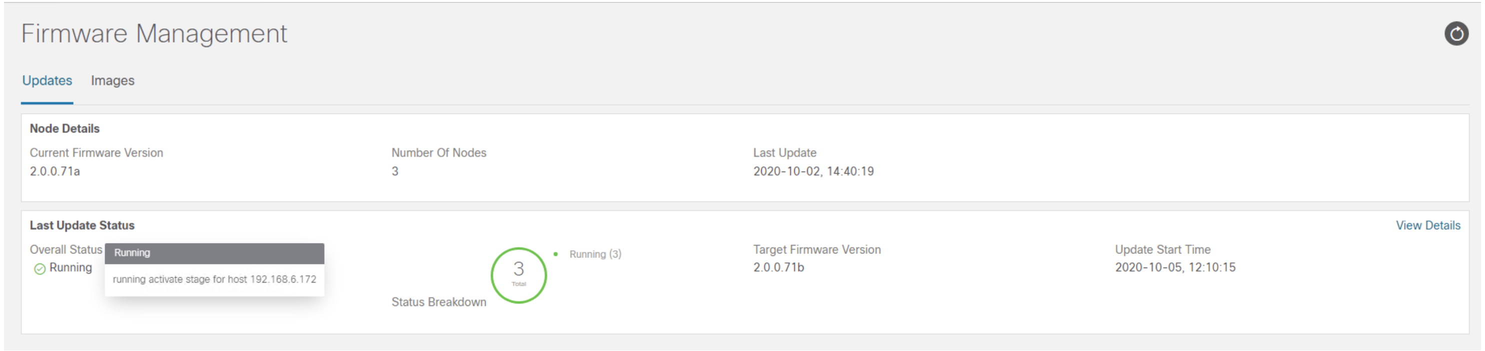

Firmware Management

Node Details

Current Firmware Version: 2.0.0.71a

Number Of Nodes: 3

Last Update: 2020-10-02, 14:40:19

Last Update Status

Overall Status: Running

Running

running activate stage for host 192.168.6.172

Status Breakdown: 3 Total — Running (3)

Target Firmware Version: 2.0.0.71b

Update Start Time: 2020-10-05, 12:10:15

- Node going through an update will display:

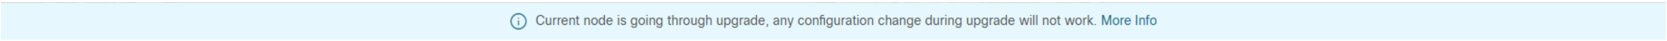

ⓘ Current node is going through upgrade, any configuration change during upgrade will not work. More Info

Remote Authentication

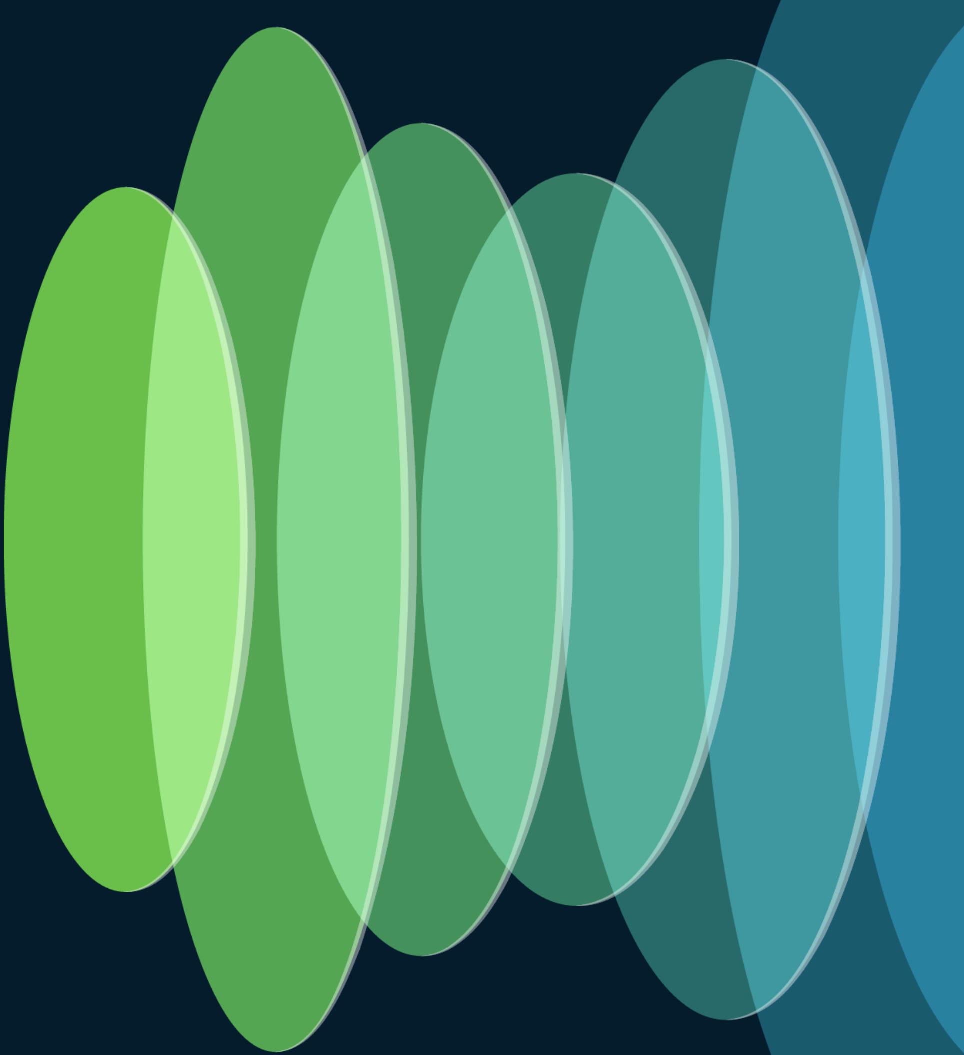

Remote Authentication

- ND adds support for following authentication providers
  - LDAP
  - TACACS
  - RADIUS
- RBAC is supported via cisco-avpair
- Is used for SSO, if the remote user has access rights to APIC, the user is automatically signed into APIC UI (4.2.6, 5.1 and later) and DCNM 11.5, when cross launching the UI. This is assuming the same auth. domain is used.

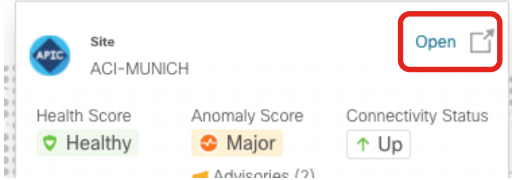
Site
ACI-MUNICH
Open
Health Score
Healthy
Anomaly Score
Major
Connectivity Status
Up
Advisories (2)

Login without and with enabled Login Domain

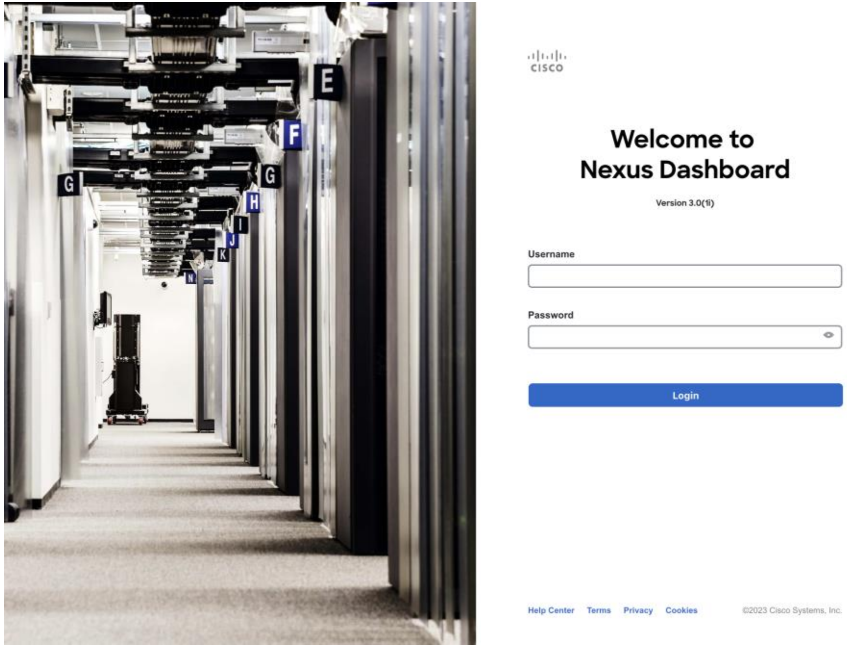

Welcome to
Nexus Dashboard

Version 3.0(1i)

Username

Password

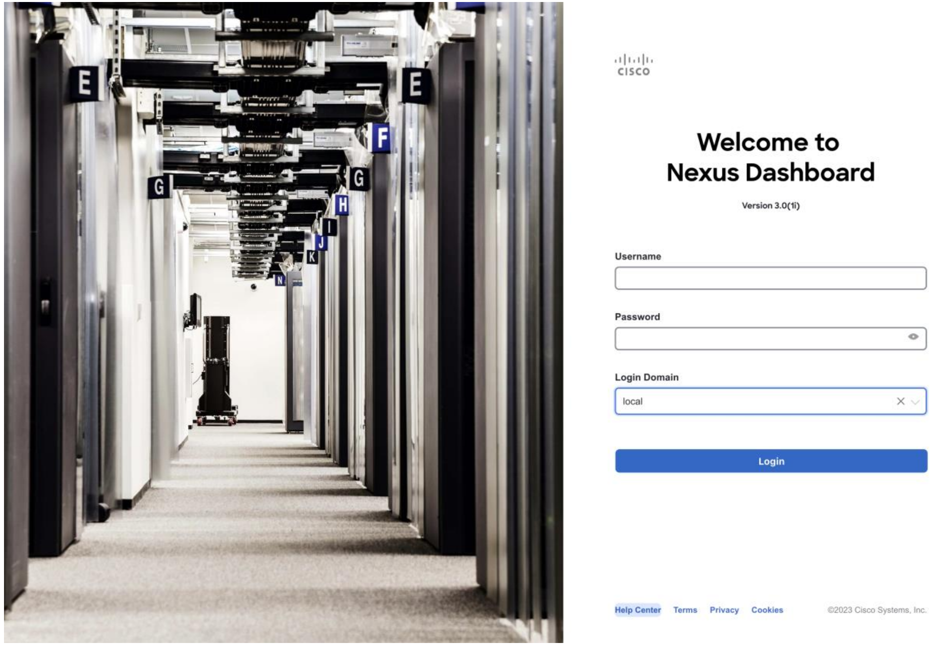

Welcome to
Nexus Dashboard

Version 3.0(1i)

Username

Password

Login Domain: local

Create a Login Domain

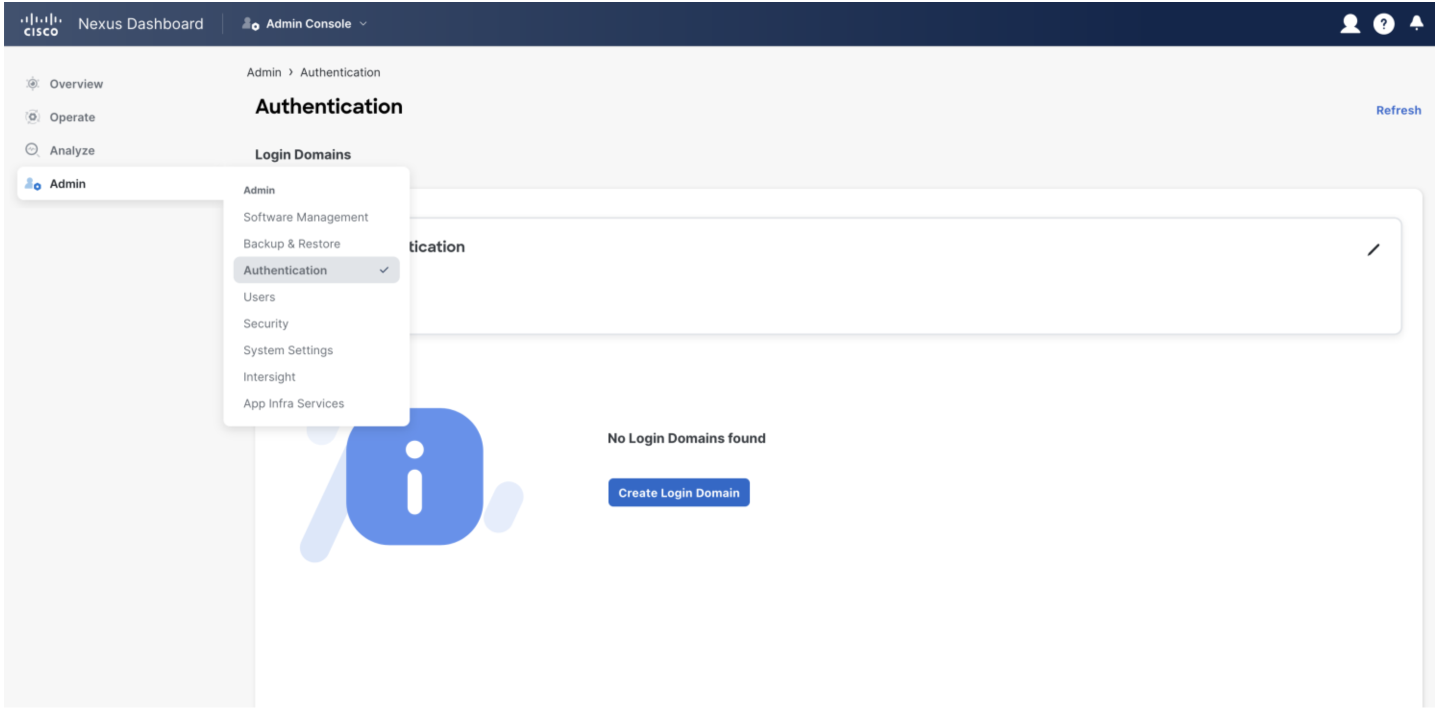

Create a Login Domain

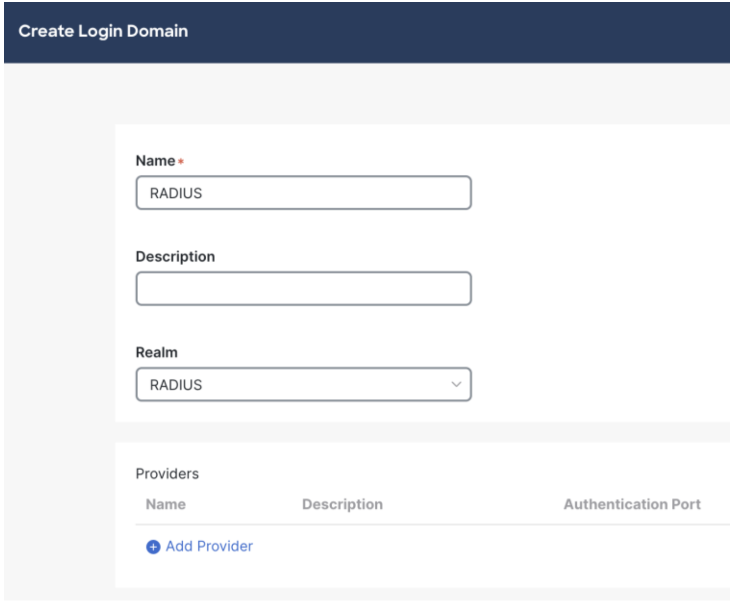

Create Login Domain

Name*

RADIUS

Description

Realm

RADIUS

Providers

Name	Description	Authentication Port

Add Provider

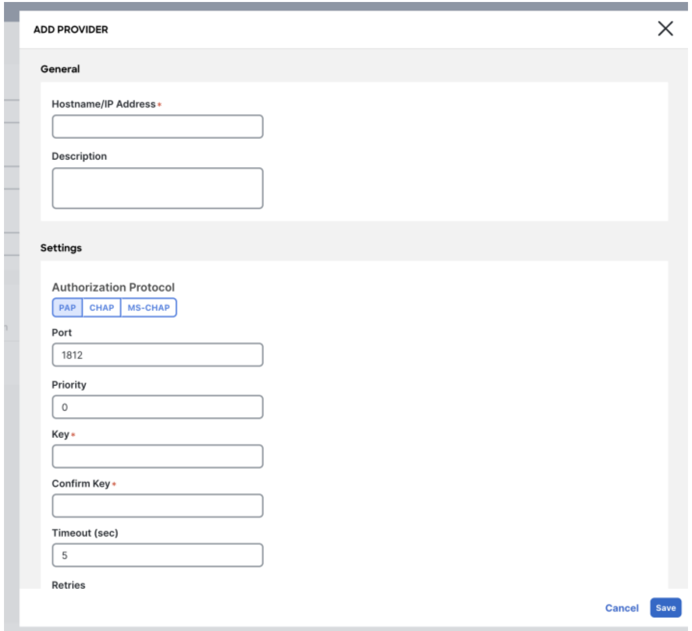

ADD PROVIDER

General

Hostname/IP Address*

Description

Settings

Authorization Protocol

PAP CHAP MS-CHAP

Port

1812

Priority

0

Key*

Confirm Key*

Timeout (sec)

5

Retries

Cancel Save

Need to have a valid remote user to add provider – backend will query the remote auth server with provider info and user/pass before it can be added.

Change Default Authentication for Login

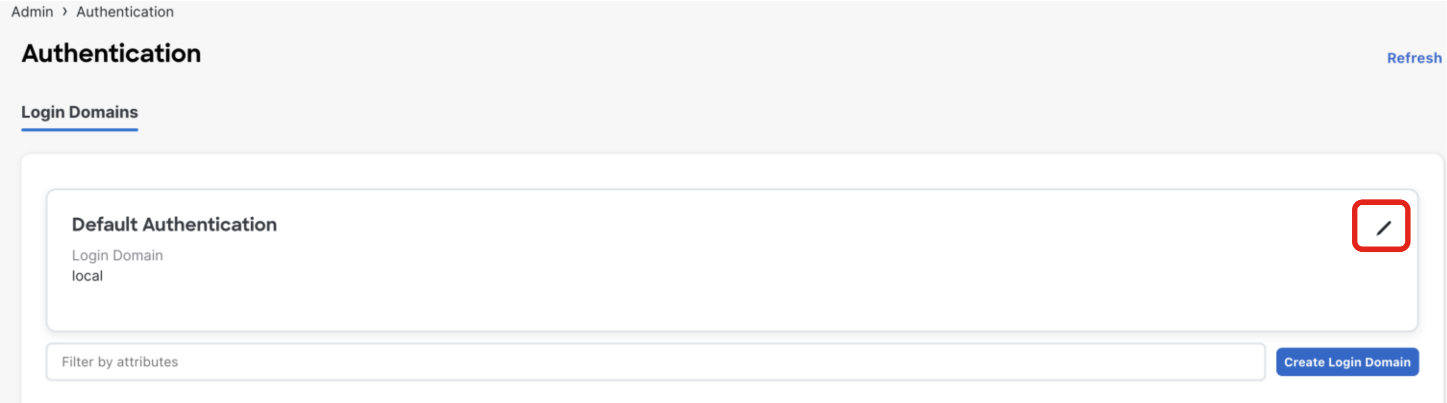

Admin › Authentication

Authentication

Refresh

Login Domains

Default Authentication

Login Domain
local

Filter by attributes

Create Login Domain

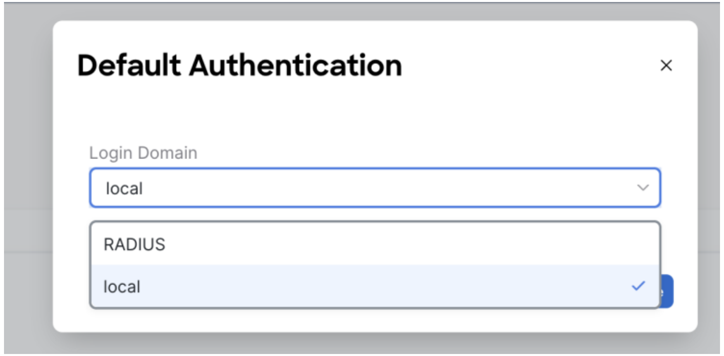

Default Authentication

Login Domain
local

RADIUS
local ✓

Login Screen with Login Domain

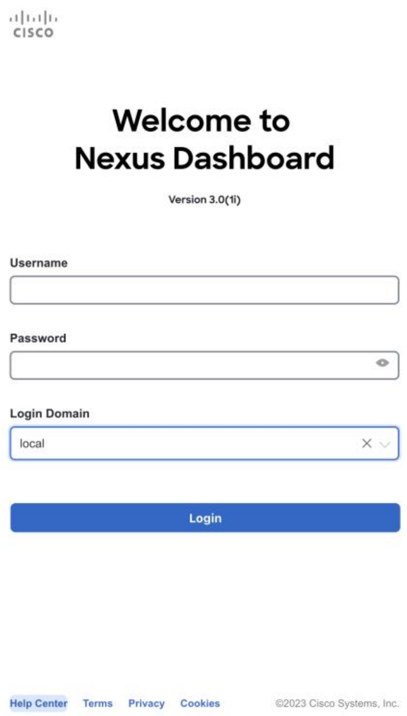

Welcome to
Nexus Dashboard

Version 3.0(1i)

Username

Password

Login Domain

local

©2023 Cisco Systems, Inc.

RBAC and User Roles 1/2

- Administrator – allows access to all objects and configurations. (Dashboard role)
  - AV Pair Value: admin
- User Manager – allows access to users and authentication configurations. (Dashboard role)
  - AV Pair Value: aaa
- Dashboard User – allows access only to the Dashboard view and launching applications; does not allow any changes to the Nexus Dashboard configurations. (Dashboard role)
  - AV Pair Value: app-user
- Site Administrator – allows access to configurations related to the sites on-boarding and configuration. (Dashboard role)
  - AV Pair Value: site-admin
- Site Manager – allows application user to manage the sites used by that application. (NDO App role)
  - AV Pair Value: config-manager
- Policy Manager – allows application user to view policy objects. (NDO App role)
  - AV Pair Value: site-policy
- Tenant Manager – allows application user to view tenants (NDO App role)
  - AV Pair Value: tenant-policy

RBAC and User Roles 2/2

- Cisco-avpair is used for RBAC via remote Auth
- AVPAIR format
  - shell:domains=<domain>/<writerole>|<writerole2>/<readrole>|<readrole2>
  - Example
    - All admin access: shell:domains=all/admin/
    - Tenant Mgr, Site Mgr and readonly AAA: shell:domains=all/tenant-policy|site-admin/aaa
- Local Users can be assigned to User roles as well while creating the User

User Roles for Local Users

Add Security Domain and Roles

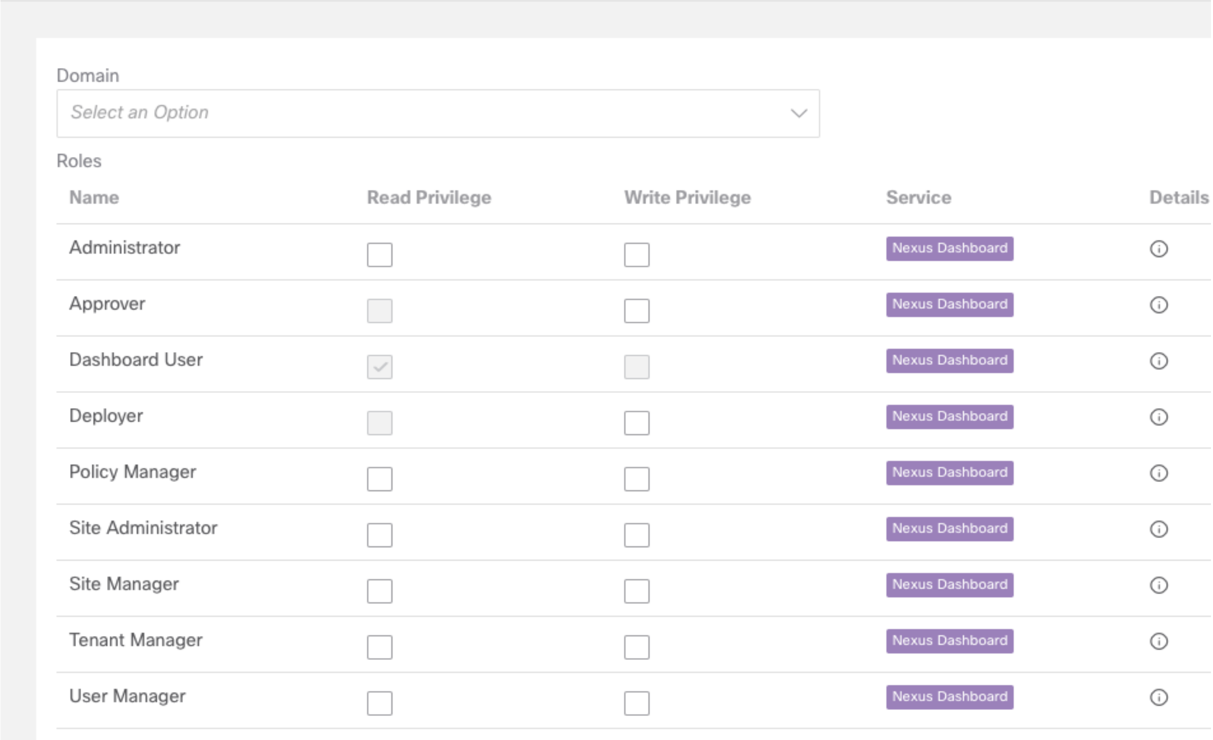

Domain

Select an Option

Roles

Name	Read Privilege	Write Privilege	Service	Details
Administrator	□	□	Nexus Dashboard	ⓘ
Approver	□	□	Nexus Dashboard	ⓘ
Dashboard User	☑	□	Nexus Dashboard	ⓘ
Deployer	□	□	Nexus Dashboard	ⓘ
Policy Manager	□	□	Nexus Dashboard	ⓘ
Site Administrator	□	□	Nexus Dashboard	ⓘ
Site Manager	□	□	Nexus Dashboard	ⓘ
Tenant Manager	□	□	Nexus Dashboard	ⓘ
User Manager	□	□	Nexus Dashboard	ⓘ

Configurable Security Settings

Configurable Security Settings

- Idle and Session Timeout is configurable
- Custom Certificates can be used
  - User needs to provide valid cert chain – backend does the validation before applying custom certs.
- Also with ND 2.3 and later you can have ND verify the Certificates of the onboarded Site-Controller before onboarding

Configure Security Settings

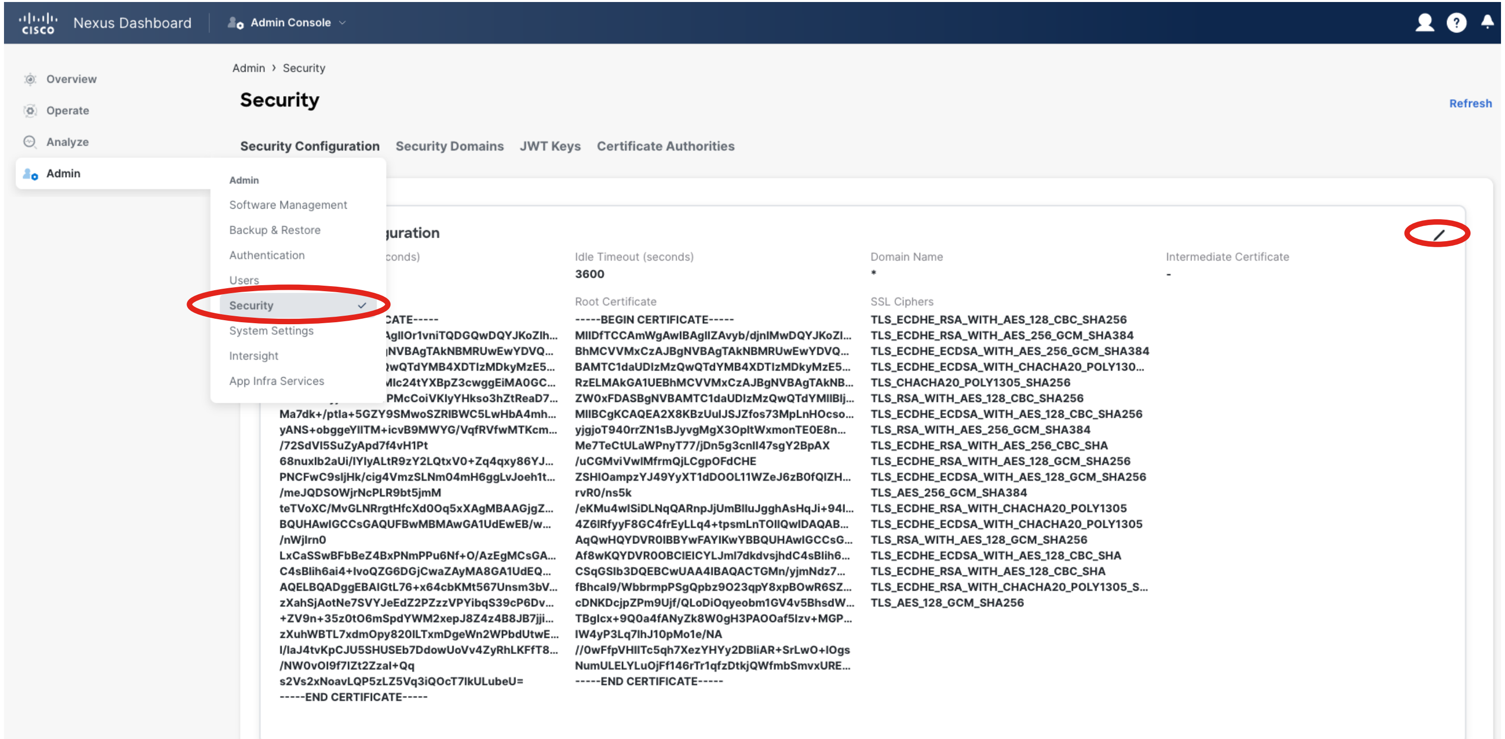

Admin › Security

Security

Security Configuration | Security Domains | JWT Keys | Certificate Authorities

…uration

…conds)

Idle Timeout (seconds)
3600

Domain Name
*

Intermediate Certificate
-

Root Certificate

SSL Ciphers

```
…CATE-----
…AgIIOr1vniTQDGQwDQYJKoZIh…
…NVBAgTAkNBMRUwEwYDVQ…
…wQTdYMB4XDTIzMDkyMzE5…
…MIc24tYXBpZ3cwggEiMA0GC…
…PMcCoiVKlyYHkso3hZtReaD7…
Ma7dk+/ptla+5GZY9SMwoSZRIBWC5LwHbA4mh…
yANS+obggeYIITM+icvB9MWYG/VqfRVfwMTKcm…
/72SdVI5SuZyApd7f4vH1Pt
68nuxIb2aUi/IYIyALtR9zY2LQtxV0+Zq4qxy86YJ…
PNCFwC9sljHk/cig4VmzSLNm04mH6ggLvJoeh1t…
/meJQDSOWjrNcPLR9bt5jmM
teTVoXC/MvGLNRrgtHfcXd0Oq5xXAgMBAAGjgZ…
BQUHAwIGCCsGAQUFBwMBMAwGA1UdEwEB/w…
/nWjIrn0
LxCaSSwBFbBeZ4BxPNmPPu6Nf+O/AzEgMCsGA…
C4sBlih6ai4+IvoQZG6DGjCwaZAyMA8GA1UdEQ…
AQELBQADggEBAIGtL76+x64cbKMt567Unsm3bV…
zXahSjAotNe7SVYJeEdZ2PZzzVPYibqS39cP6Dv…
+ZV9n+35z0tO6mSpdYWM2xepJ8Z4z4B8JB7jji…
zXuhWBTL7xdmOpy820ILTxmDgeWn2WPbdUtwE…
l/laJ4tvKpCJU5SHUSEb7DdowUoVv4ZyRhLKFfT8…
/NW0vOI9f7IZt2ZzaI+Qq
s2Vs2xNoavLQP5zLZ5Vq3iQOcT7IkULubeU=
-----END CERTIFICATE-----
```

```
-----BEGIN CERTIFICATE-----
MIIDfTCCAmWgAwIBAgIIZAvyb/djnIMwDQYJKoZI…
BhMCVVMxCzAJBgNVBAgTAkNBMRUwEwYDVQ…
BAMTC1daUDIzMzQwQTdYMB4XDTIzMDkyMzE5…
RzELMAkGA1UEBhMCVVMxCzAJBgNVBAgTAkNB…
ZW0xFDASBgNVBAMTC1daUDIzMzQwQTdYMIIBIj…
MIIBCgKCAQEA2X8KBzUuIJSJZfos73MpLnHOcso…
yjgjoT940rrZN1sBJyvgMgX3OpltWxmonTE0E8n…
Me7TeCtULaWPnyT77/jDn5g3cnII47sgY2BpAX
/uCGMviVwIMfrmQjLCgpOFdCHE
ZSHIOampzYJ49YyXT1dDOOL11WZeJ6zB0fQIZH…
rvR0/ns5k
/eKMu4wISiDLNqQARnpJjUmBIluJgghAsHqJi+94I…
4Z6IRfyyF8GC4frEyLLq4+tpsmLnTOIIQwIDAQAB…
AqQwHQYDVR0IBBYwFAYIKwYBBQUHAwIGCCsG…
Af8wKQYDVR0OBCIEICYLJmI7dkdvsjhdC4sBlih6…
CSqGSIb3DQEBCwUAA4IBAQACTGMn/yjmNdz7…
fBhcaI9/WbbrmpPSgQpbz9O23qpY8xpBOwR6SZ…
cDNKDcjpZPm9Ujf/QLoDiOqyeobm1GV4v5BhsdW…
TBgIcx+9Q0a4fANyZk8W0gH3PAOOaf5Izv+MGP…
IW4yP3Lq7lhJ10pMo1e/NA
//0wFfpVHIITc5qh7XezYHYy2DBIiAR+SrLwO+IOgs
NumULELYLuOjFf146rTr1qfzDtkjQWfmbSmvxURE…
-----END CERTIFICATE-----
```

```
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384
TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY130…
TLS_CHACHA20_POLY1305_SHA256
TLS_RSA_WITH_AES_128_CBC_SHA256
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256
TLS_RSA_WITH_AES_256_GCM_SHA384
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
TLS_AES_256_GCM_SHA384
TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305
TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY1305
TLS_RSA_WITH_AES_128_GCM_SHA256
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA
TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_S…
TLS_AES_128_GCM_SHA256
```

Configure Security Settings

Session and Idle Timeout in Seconds

Customer Certificate and Root Certificate, enabled SSL Ciphers etc.

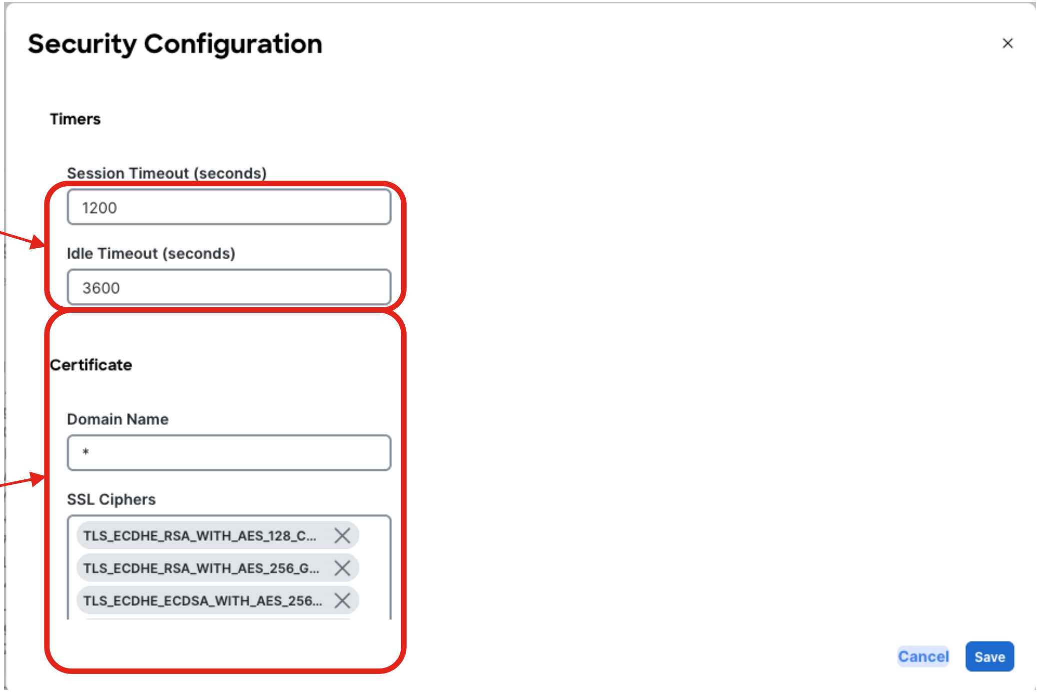

Security Configuration

Timers

Session Timeout (seconds)

1200

Idle Timeout (seconds)

3600

Certificate

Domain Name

*

SSL Ciphers

TLS_ECDHE_RSA_WITH_AES_128_C...

TLS_ECDHE_RSA_WITH_AES_256_G...

TLS_ECDHE_ECDSA_WITH_AES_256...

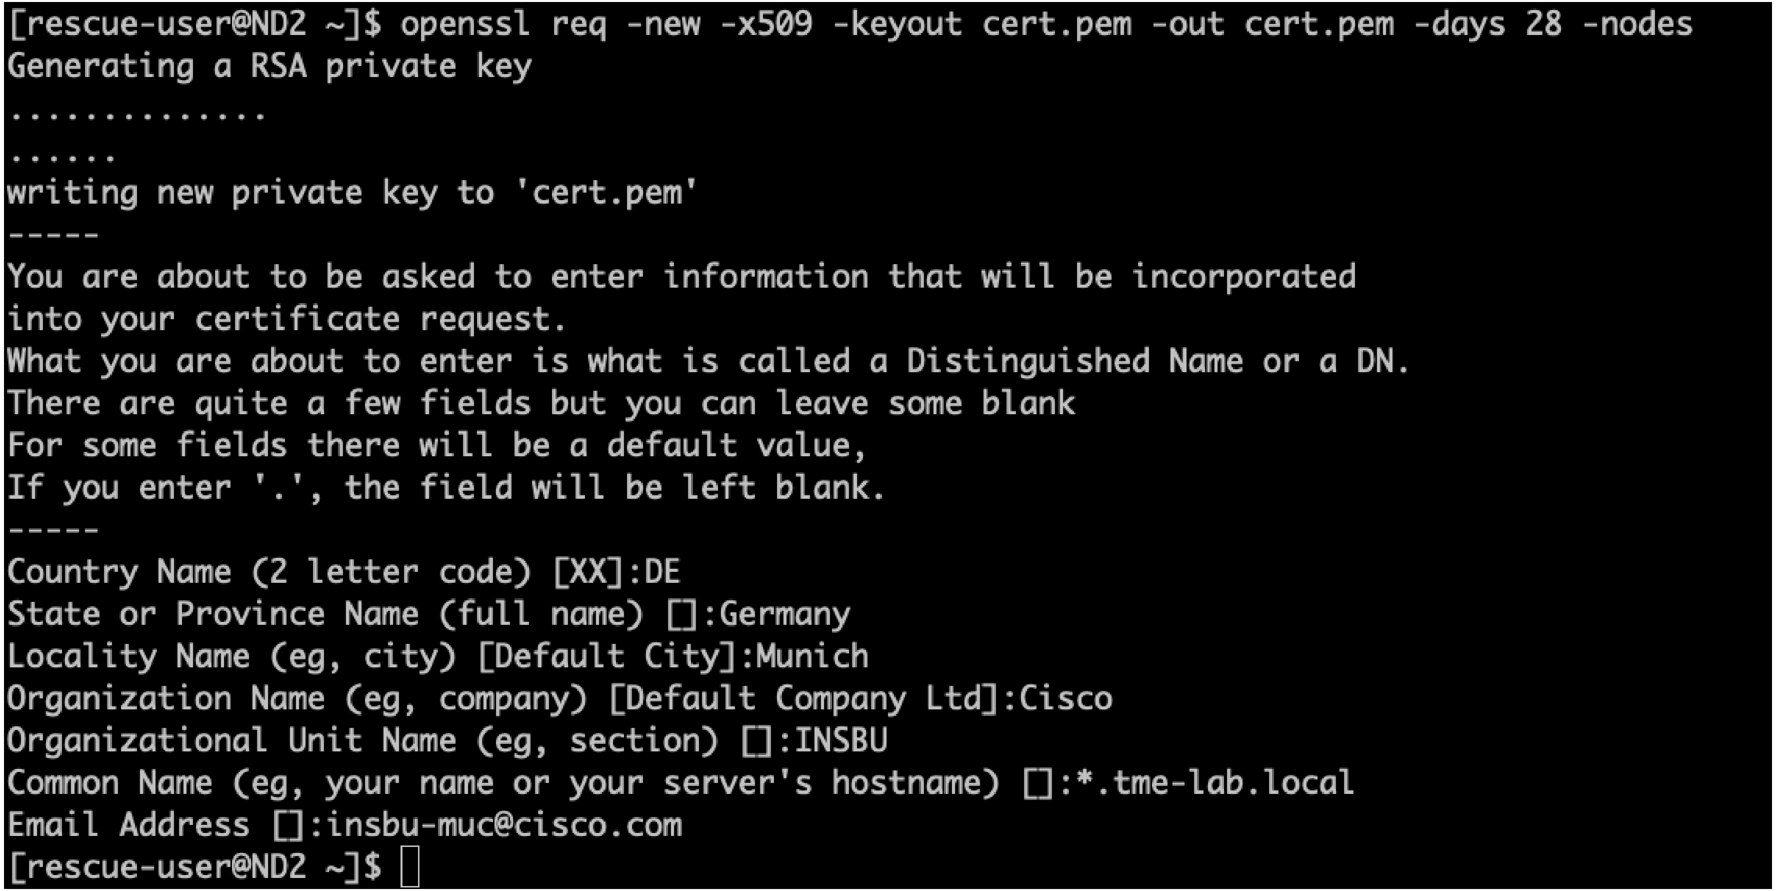

```
[rescue-user@ND2 ~]$ openssl req -new -x509 -keyout cert.pem -out cert.pem -days 28 -nodes
Generating a RSA private key
..............
......
writing new private key to 'cert.pem'
-----
You are about to be asked to enter information that will be incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) [XX]:DE
State or Province Name (full name) []:Germany
Locality Name (eg, city) [Default City]:Munich
Organization Name (eg, company) [Default Company Ltd]:Cisco
Organizational Unit Name (eg, section) []:INSBU
Common Name (eg, your name or your server's hostname) []:*.tme-lab.local
Email Address []:insbu-muc@cisco.com
[rescue-user@ND2 ~]$
```

Resource Monitoring

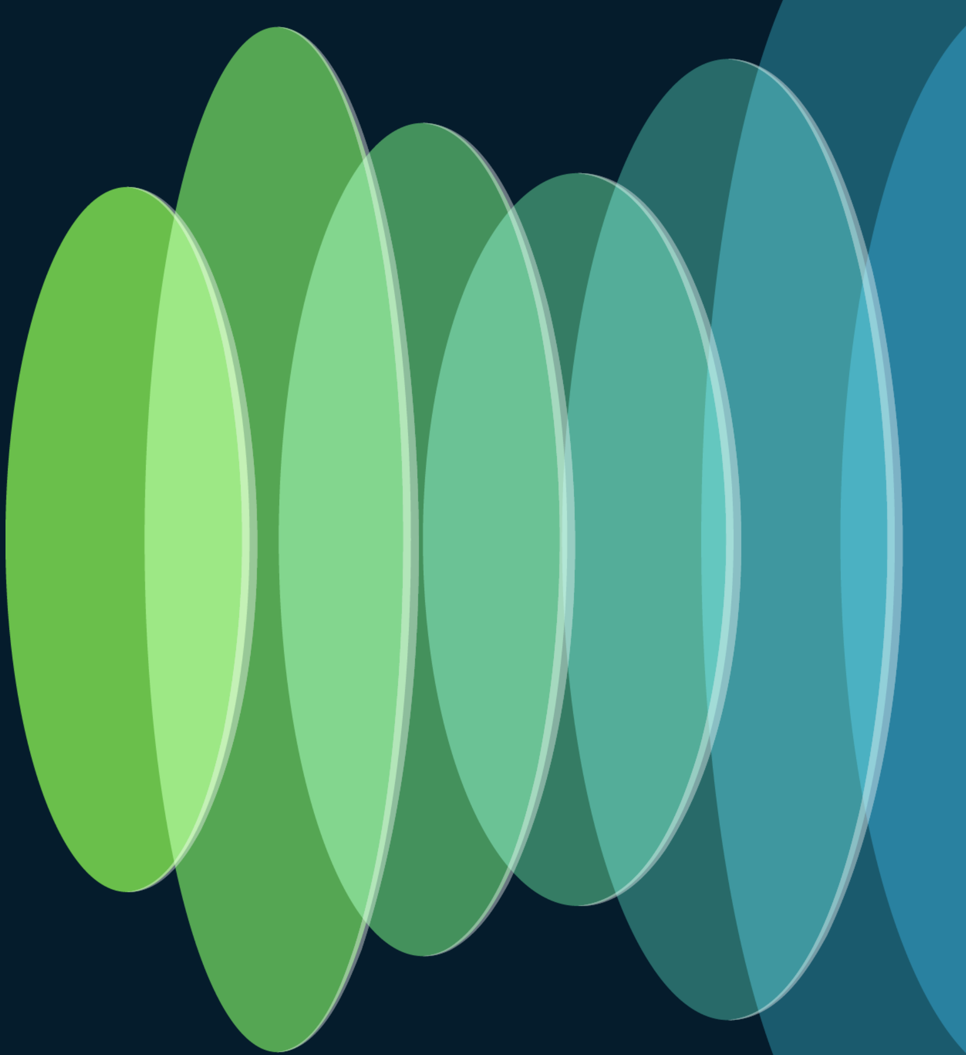

Resource Monitoring

- Provides Monitoring on
  - CPU
  - RAM
  - I/O Disk
  - I/O Network
- Node or Cluster level View
- Namespaces View

Resource Monitoring on Node and Cluster Level

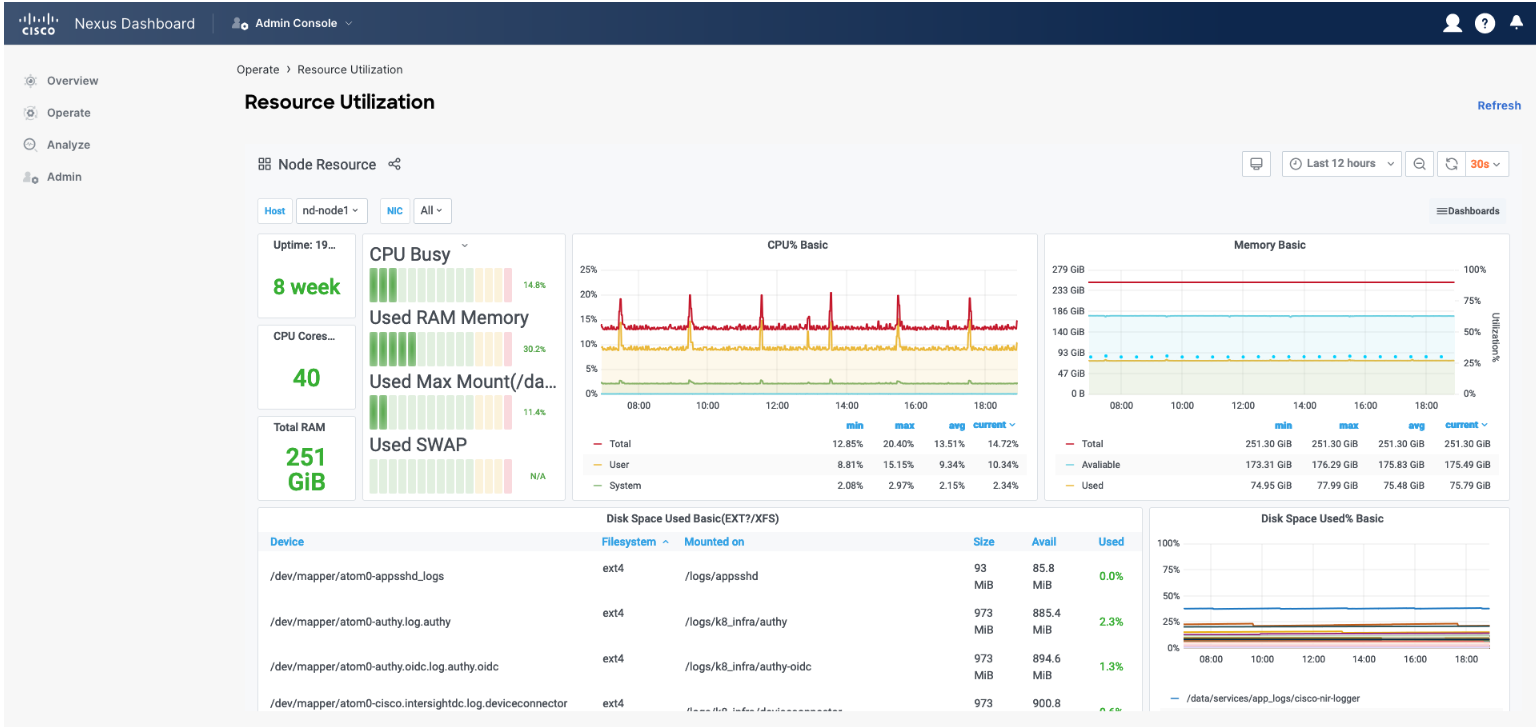

Resource Utilization

Uptime: 19...
8 week

CPU Cores...
40

Total RAM
251 GiB

CPU Busy 14.8%

Used RAM Memory 30.2%

Used Max Mount(/da... 11.4%

Used SWAP N/A

CPU% Basic

	min	max	avg	current
Total	12.85%	20.40%	13.51%	14.72%
User	8.81%	15.15%	9.34%	10.34%
System	2.08%	2.97%	2.15%	2.34%

Memory Basic

	min	max	avg	current
Total	251.30 GiB	251.30 GiB	251.30 GiB	251.30 GiB
Avaliable	173.31 GiB	176.29 GiB	175.83 GiB	175.49 GiB
Used	74.95 GiB	77.99 GiB	75.48 GiB	75.79 GiB

Disk Space Used Basic(EXT?/XFS)

Device	Filesystem	Mounted on	Size	Avail	Used
/dev/mapper/atom0-appsshd_logs	ext4	/logs/appsshd	93 MiB	85.8 MiB	0.0%
/dev/mapper/atom0-authy.log.authy	ext4	/logs/k8_infra/authy	973 MiB	885.4 MiB	2.3%
/dev/mapper/atom0-authy.oidc.log.authy.oidc	ext4	/logs/k8_infra/authy-oidc	973 MiB	894.6 MiB	1.3%
/dev/mapper/atom0-cisco.intersightdc.log.deviceconnector	ext4	/logs/k8_infra/deviceconnector	973	900.8	0.6%

Event Analytic

Event Analytic

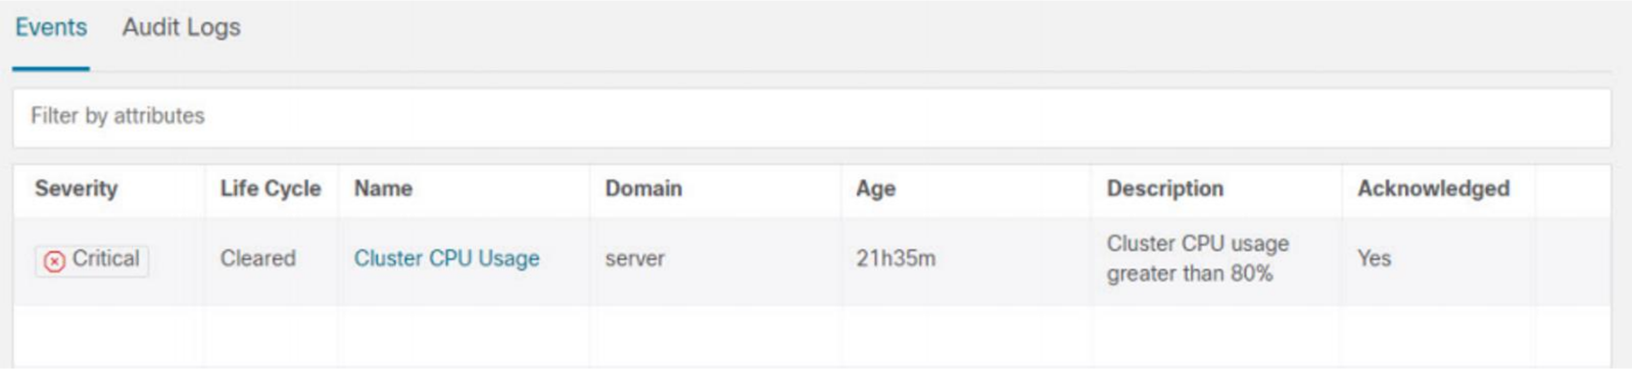

Severity	Life Cycle	Name	Domain	Age	Description	Acknowledged
⊗ Critical	Cleared	Cluster CPU Usage	server	21h35m	Cluster CPU usage greater than 80%	Yes

Event Analytics enables easy access your Nexus Dashboard's events and audit logs. In addition to viewing the events and logs directly in the Nexus Dashboard GUI, you can also configure the cluster to stream the events to an external syslog server (TCP/UDP)

Events

- Node CPU exceeding threshold (80%)
- Node storage exceeding threshold (80%)
- Node memory exceeding threshold (80%)
- Cluster node is unreachable
- Cluster node is rebooted
- All audit events
- NTP is not synchronized
- BGP peers are not reachable

Configuring Syslog Servers 1/2

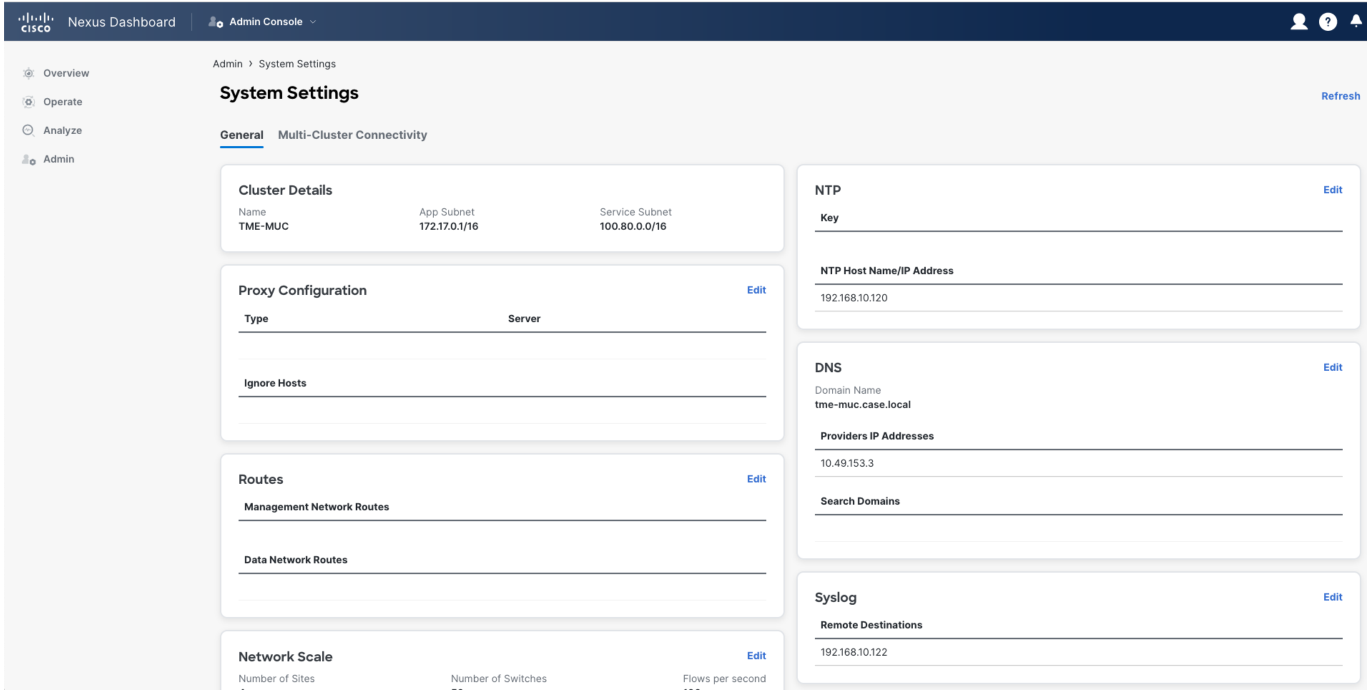

System Settings

Cluster Details

Name	App Subnet	Service Subnet
TME-MUC	172.17.0.1/16	100.80.0.0/16

Proxy Configuration

Type	Server

Ignore Hosts

Routes

Management Network Routes

Data Network Routes

Network Scale

Number of Sites | Number of Switches | Flows per second

NTP

Key

NTP Host Name/IP Address

192.168.10.120

DNS

Domain Name
tme-muc.case.local

Providers IP Addresses

10.49.153.3

Search Domains

Syslog

Remote Destinations

192.168.10.122

Configuring Syslog Servers 2/2

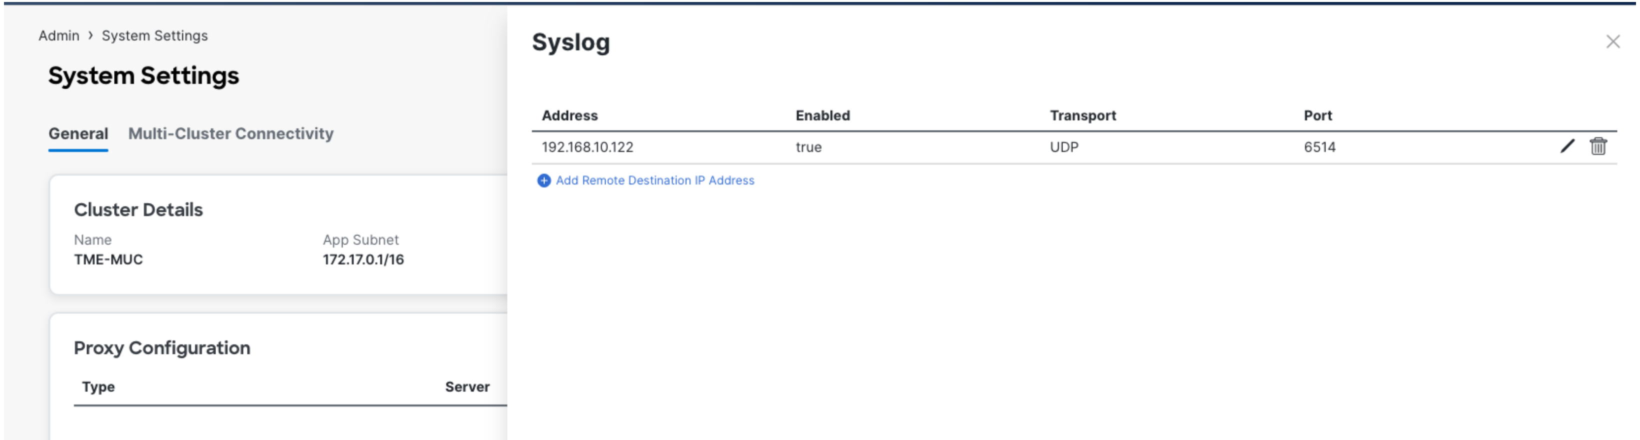

Admin › System Settings

System Settings

General   Multi-Cluster Connectivity

Cluster Details

Name: TME-MUC

App Subnet: 172.17.0.1/16

Proxy Configuration

Type	Server

Syslog

Address	Enabled	Transport	Port
192.168.10.122	true	UDP	6514

Add Remote Destination IP Address

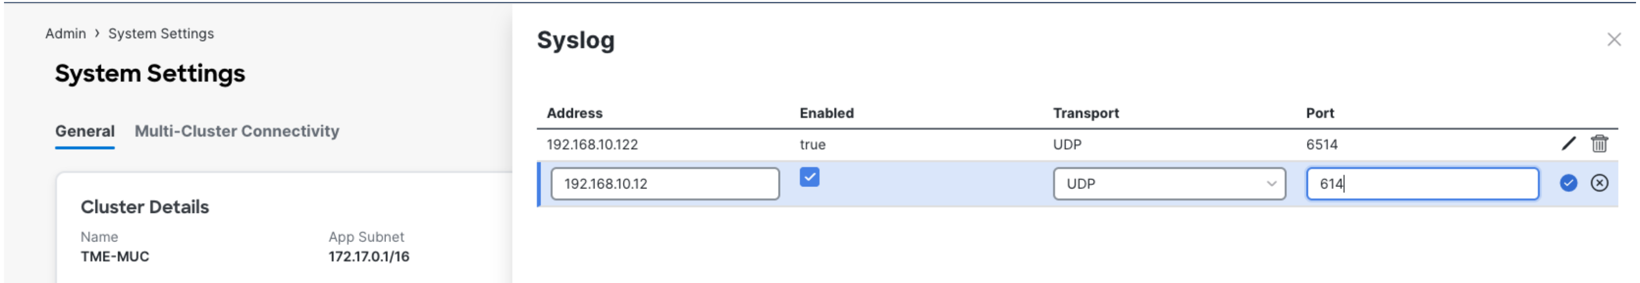

Admin › System Settings

System Settings

General   Multi-Cluster Connectivity

Cluster Details

Name: TME-MUC

App Subnet: 172.17.0.1/16

Syslog

Address	Enabled	Transport	Port
192.168.10.122	true	UDP	6514
192.168.10.12	☑	UDP	614

Hardware Monitoring of ND via CIMC

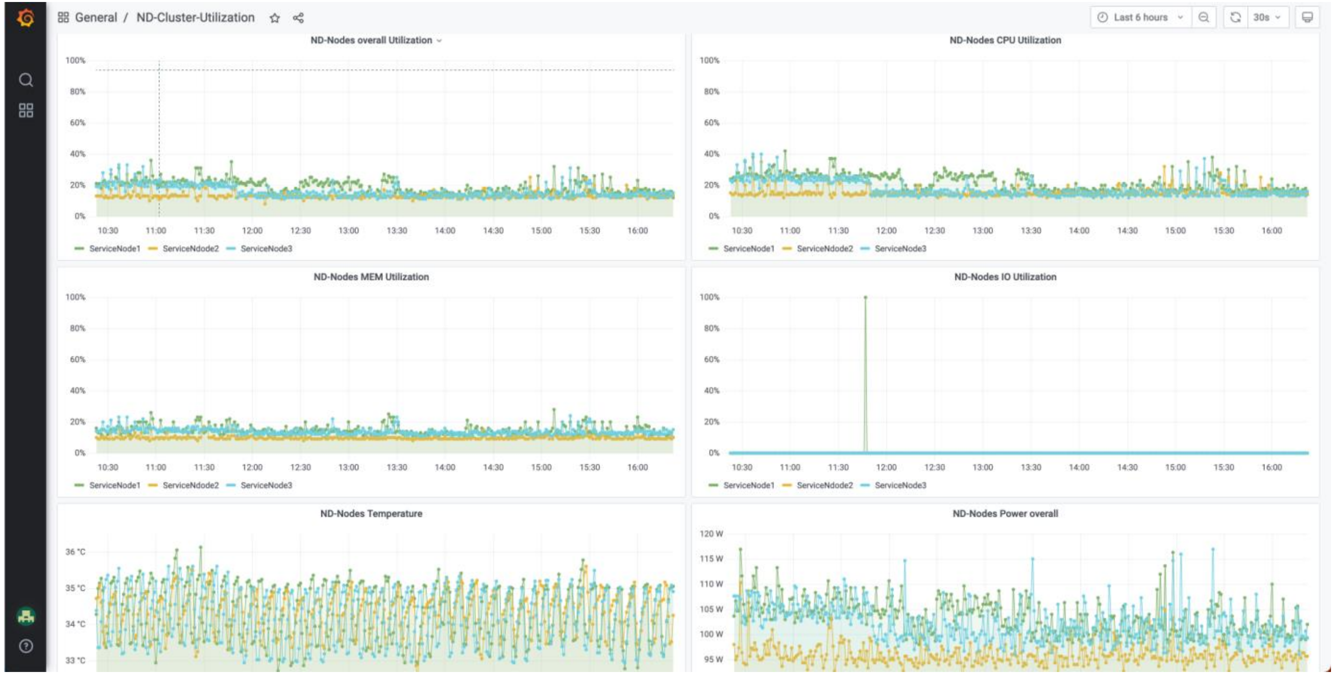

Hardware Monitoring of ND via CIMC

- Leveraging REST-API of CIMC to get:
  - Power draw
  - Temperature
  - CPU, I/O and RAM Utilization
- Querying the following dns:
  - CPU, I/O and RAM : dn="sys/rack-unit-1/utilization”
  - Temperature: dn=="sys/rack-unit-1/temperature”
  - Power: dn=“sys/rack-unit-1/pwrmonitor-Platform”

SW Stack Example

Telegraf calling a Python script to collect periodically data from CIMC

Telegraf storing data as timeseries in InfluxDB. Grafana visualizes the data

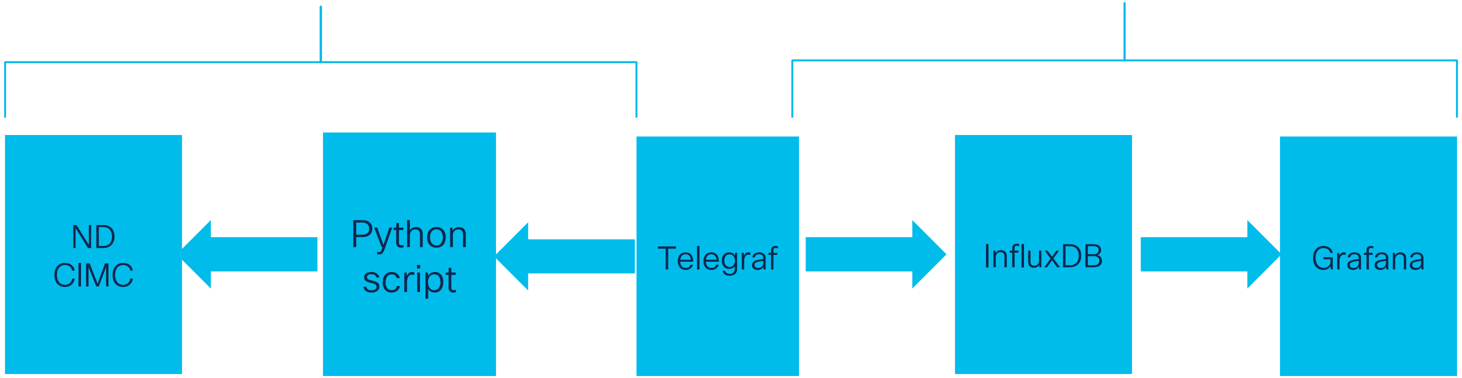

Basic Troubleshooting

Basic Troubleshooting

- Accessing ND Console, only via “rescue-user” with “admin” password
- Usage of ACS

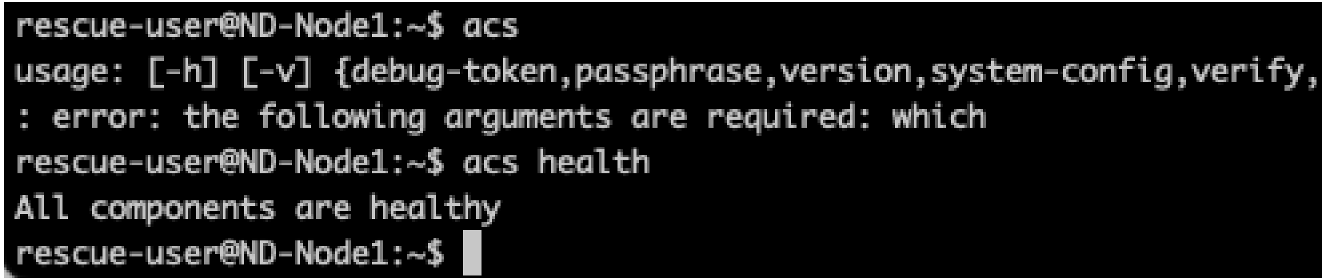

```
rescue-user@ND-Node1:~$ acs
usage: [-h] [-v] {debug-token,passphrase,version,system-config,verify,
: error: the following arguments are required: which
rescue-user@ND-Node1:~$ acs health
All components are healthy
rescue-user@ND-Node1:~$
```

Basic Troubleshooting

Usage of Kubectl to get information of the K8S

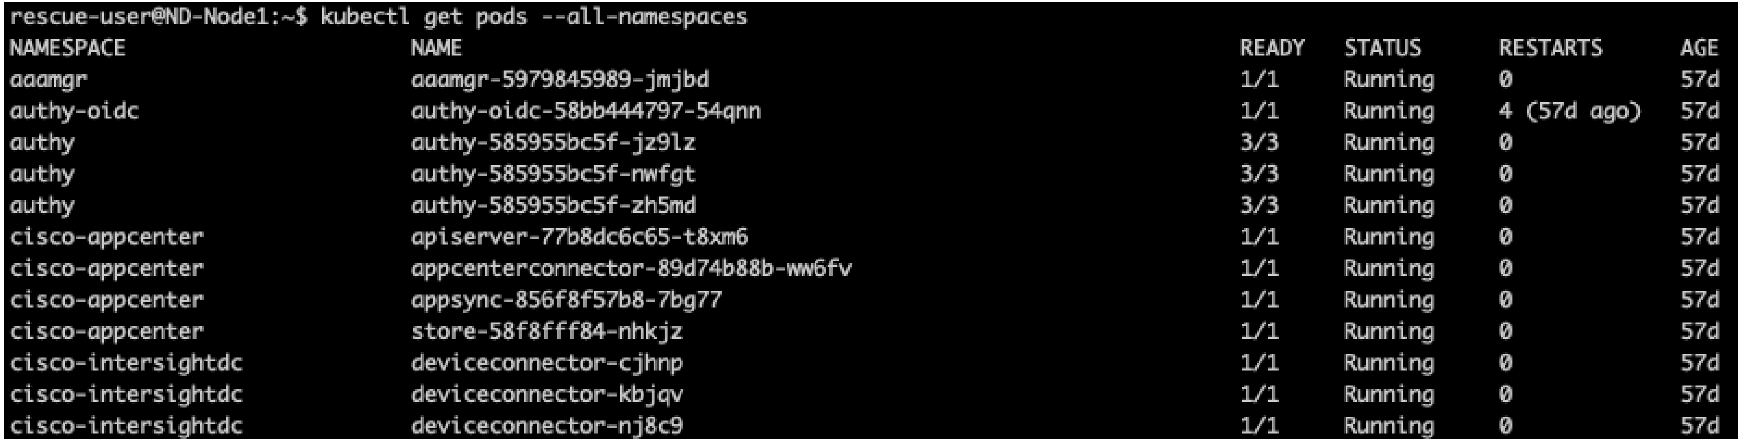

```
rescue-user@ND-Node1:~$ kubectl get pods --all-namespaces
NAMESPACE               NAME                                    READY   STATUS    RESTARTS       AGE
aaamgr                  aaamgr-5979845989-jmjbd                 1/1     Running   0              57d
authy-oidc              authy-oidc-58bb444797-54qnn             1/1     Running   4 (57d ago)    57d
authy                   authy-585955bc5f-jz9lz                  3/3     Running   0              57d
authy                   authy-585955bc5f-nwfgt                  3/3     Running   0              57d
authy                   authy-585955bc5f-zh5md                  3/3     Running   0              57d
cisco-appcenter         apiserver-77b8dc6c65-t8xm6              1/1     Running   0              57d
cisco-appcenter         appcenterconnector-89d74b88b-ww6fv      1/1     Running   0              57d
cisco-appcenter         appsync-856f8f57b8-7bg77                1/1     Running   0              57d
cisco-appcenter         store-58f8fff84-nhkjz                   1/1     Running   0              57d
cisco-intersightdc      deviceconnector-cjhnp                   1/1     Running   0              57d
cisco-intersightdc      deviceconnector-kbjqv                   1/1     Running   0              57d
cisco-intersightdc      deviceconnector-nj8c9                   1/1     Running   0              57d
```

Conclusion

Take Away

- Attachment of ND Cluster depends on use case
  - E.g. NDI, NDO, NDFC
- Unified Image made upgrading ND cluster more easy
- ND provides all tools to operate ND cluster

Complete Your Session Evaluations

Complete a minimum of 4 session surveys and the Overall Event Survey to be entered in a drawing to **win 1 of 5 full conference passes** to Cisco Live 2025.

Earn 100 points per survey completed and compete on the Cisco Live Challenge leaderboard.

Level up and earn **exclusive prizes!**

Complete your surveys in the **Cisco Live mobile app.**

Continue your education

- Visit the Cisco Showcase for related demos
- Book your one-on-one Meet the Engineer meeting
- Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs
- Visit the On-Demand Library for more sessions at www.CiscoLive.com/on-demand

The bridge to possible

Thank you