
cisco Live!
Let's go

The bridge to possible

Deploying Nexus Dashboard in your Organization

Matthias Wessendorf, Principal Engineer
@matteq4er

cisco Live!

BRKDCN-3914

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

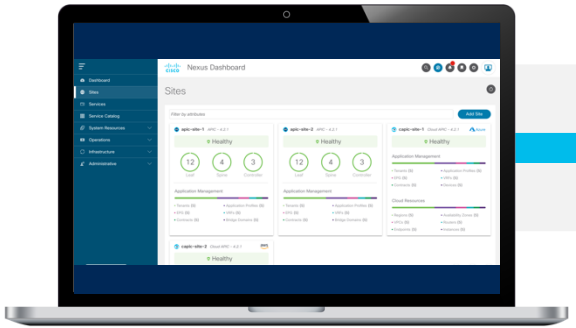

Past

Today

Physical Platform Cluster

Virtual/Cloud Platform

Nexus Dashboard

Simple to automate, simple to consume

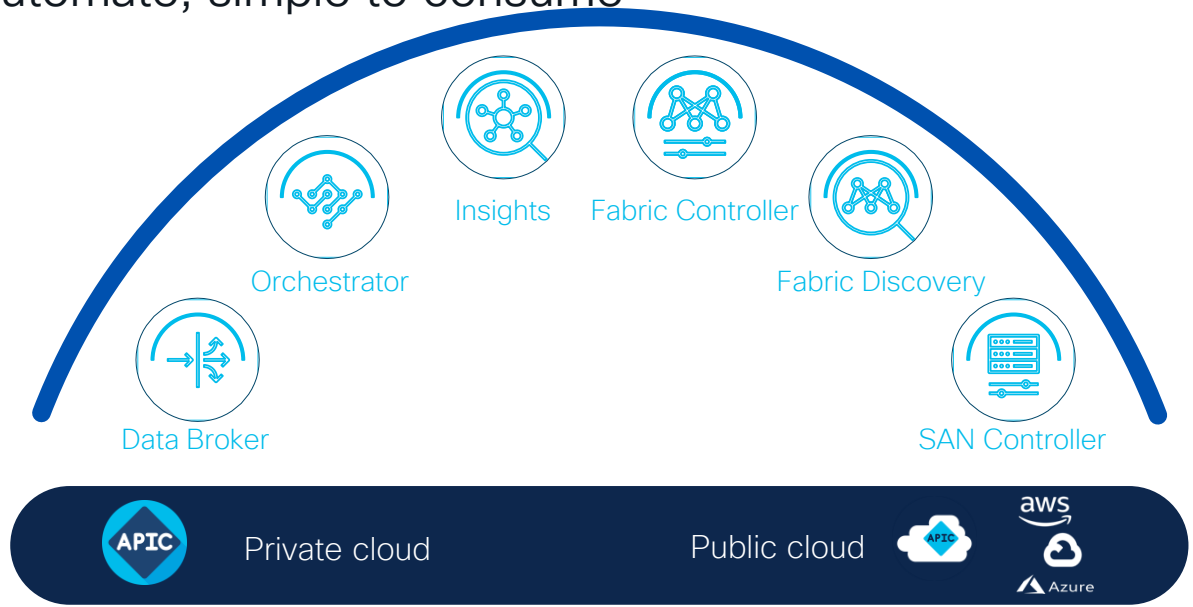

Consume all services in one place

Nexus Dashboard: One view

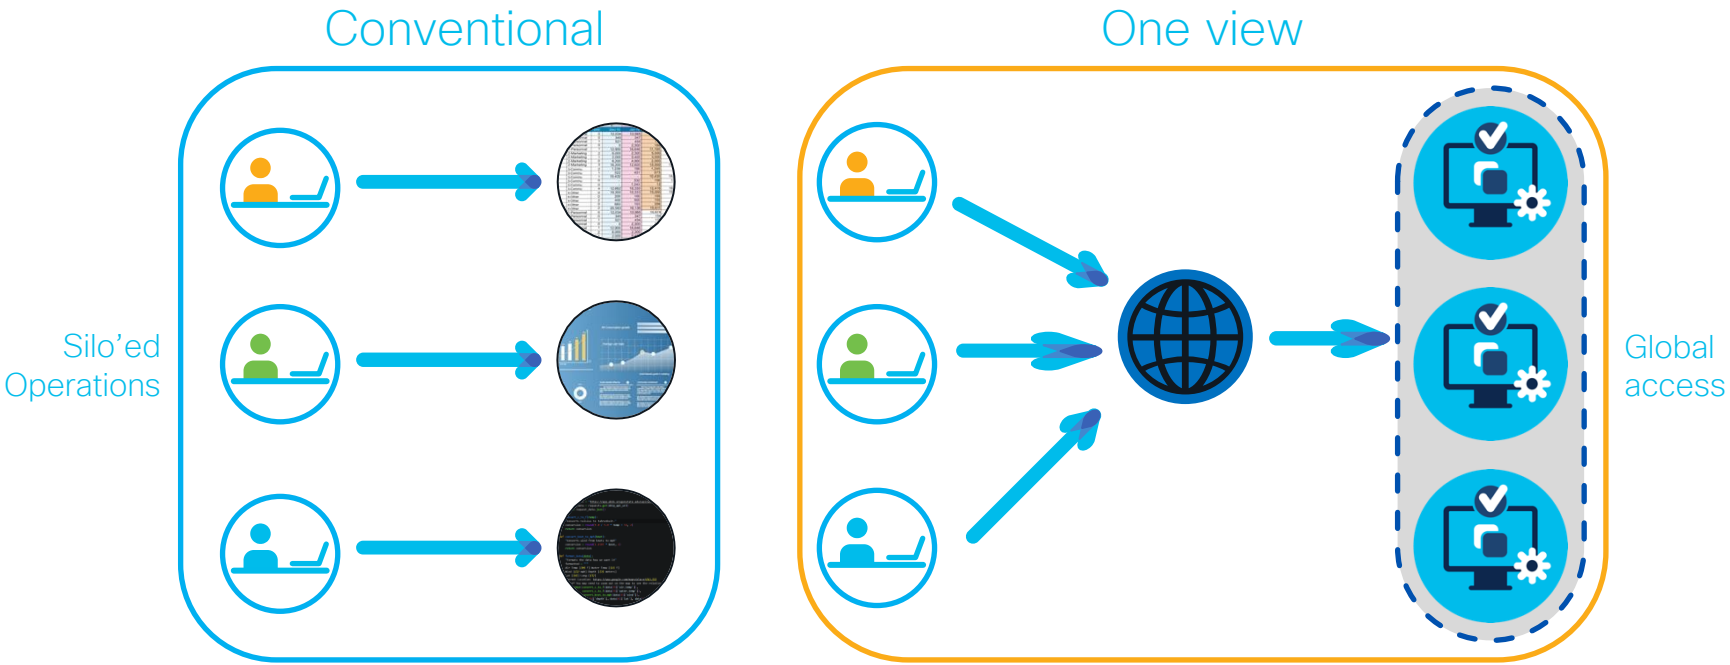

Cisco Nexus Dashboard Platform

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

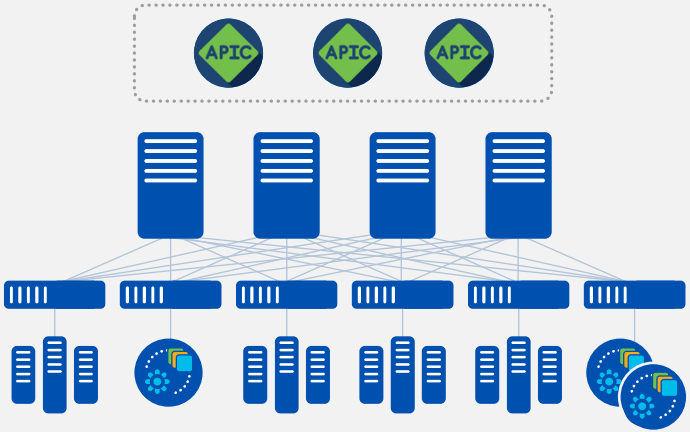

Nexus Dashboard Insights

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

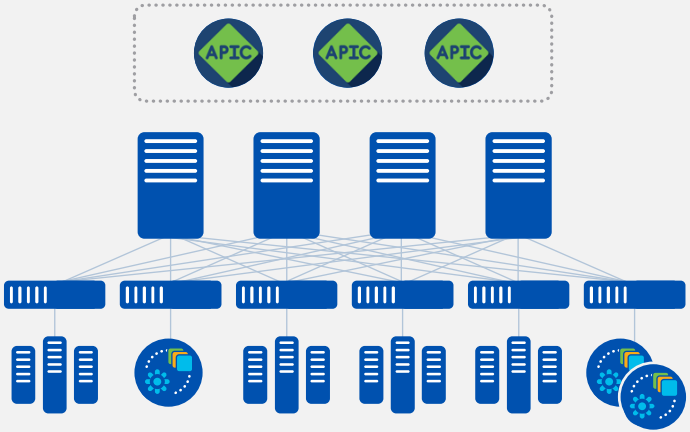

Nexus Dashboard Insights

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

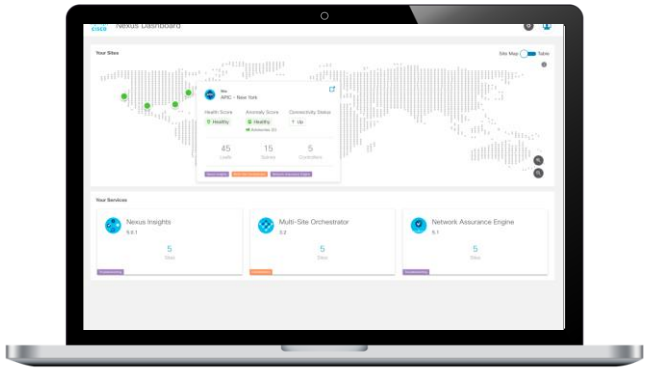

- Consume service(s) from single place
- Frictionless navigation across multiple services and sites
- Customize views and workflows

The admin view

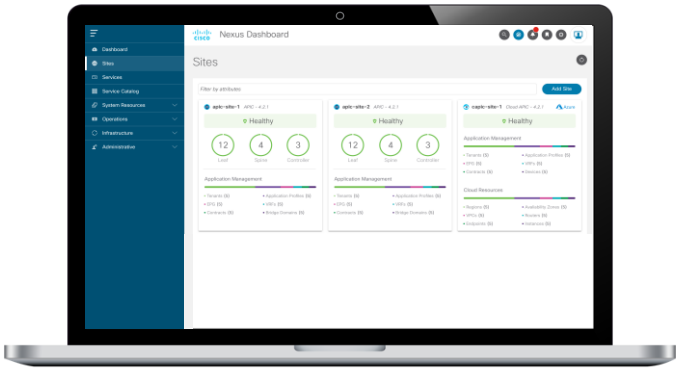

- Single dashboard for lifecycle management of services and Ops infra
- Consistent one-time onboarding of domains and services
- Consistent user management and access control

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

Nexus Dashboard Platform–Under the Hood

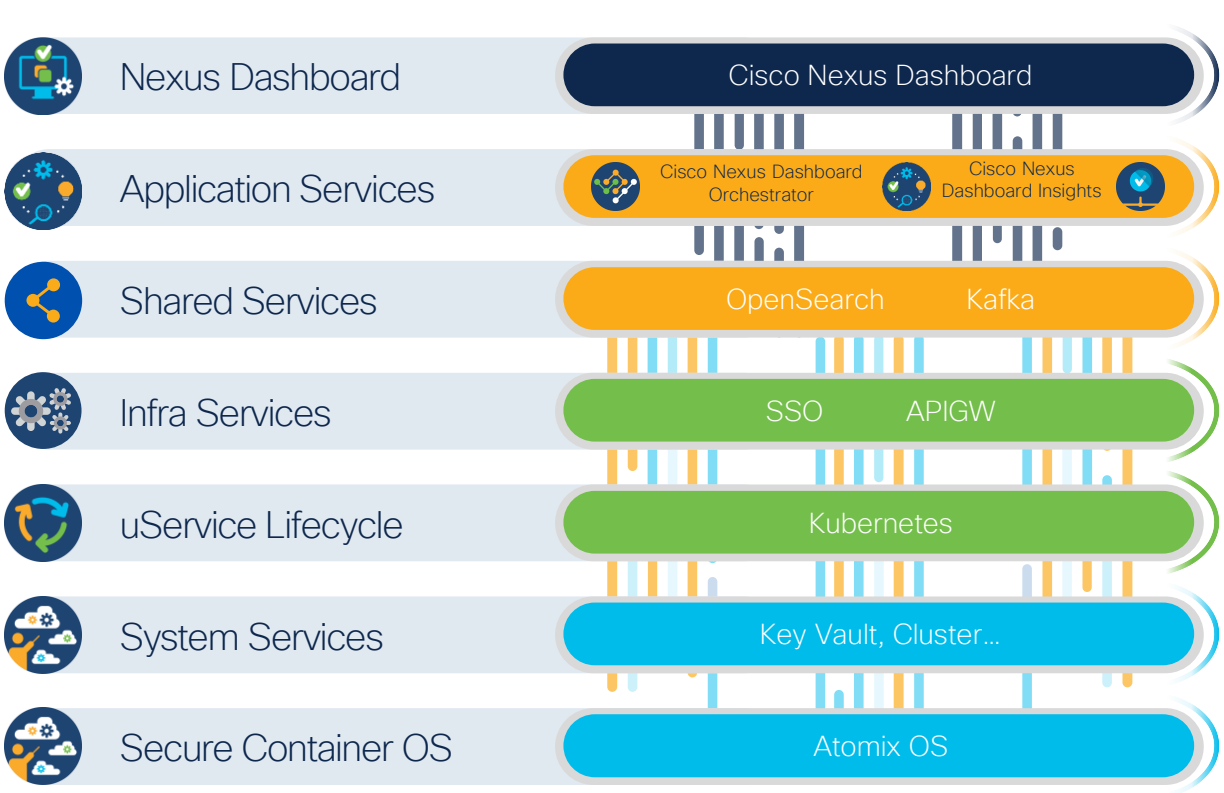

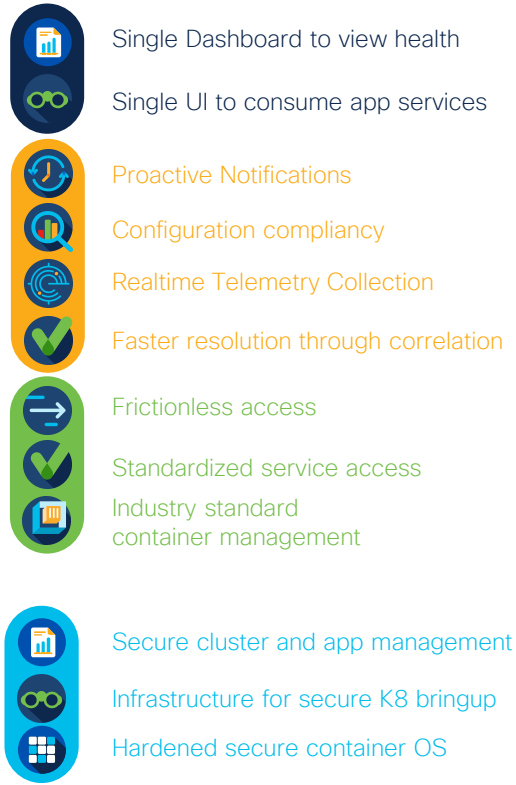

Deployment Model

- 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 master 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 master 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

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 DCNM Connectivity Considerations

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

ND to NDFC Connectivity Considerations

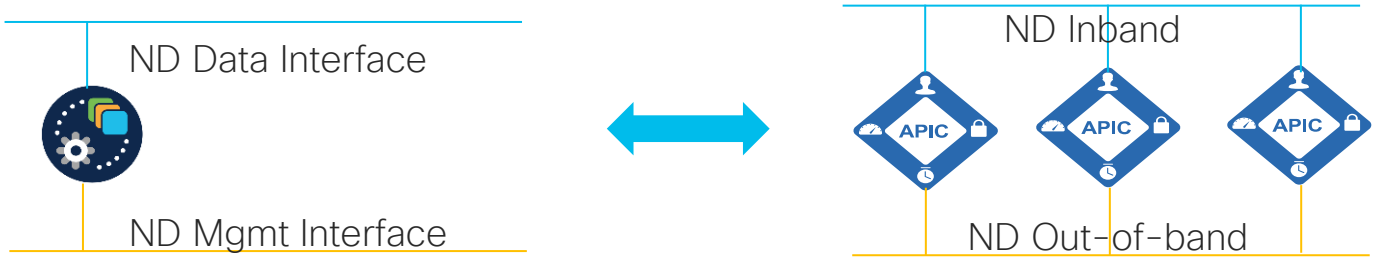

- An NDFC site is onboarded on ND by specifying the Inband IP address of the ND hosting the NDFC, no other IP is supported
- ND uses the Data Interface to establish the initial and ongoing connection to that ND Data IP address hosting NDFC

vND Considerations for ND 2.2 or earlier

Attaching vND to the Network (via UCS FI or equivalent or direct)

- If you plan to leverage Persistent IPs for NDI or NDFC
  - Port-Group and virtual Switch, where the vND is connected to has to be:
    - Connected via PC or vPC
    - Connected via a single link
    - A/A without PC or vPC is not supported
    - A/S at Hypervisor level without PC or vPC is not supported
    - Interface failover at UCS level (or equivalent) without PC or vPC is supported
  - In a nutshell the virtual switch has to have a single logical uplink.

- This is addressed in ND 2.3 and later release.

Attaching vND to Network (via UCS FI or equivalent)

Unsupported

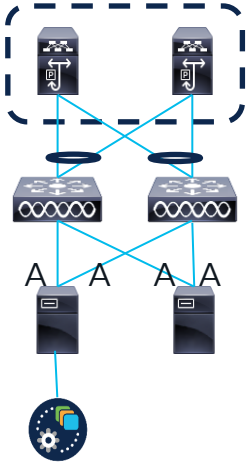

A/A uplinks of Port-Group/virtual Switch without PC or vPC

Unsupported

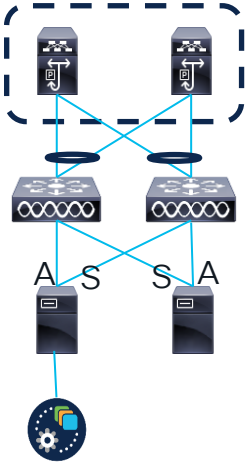

A/S uplinks of Port-Group /virtual Switch at Hypervisor level without PC or vPC

Supported

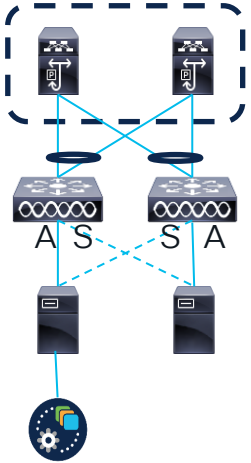

A/S uplinks of Port-Group /virtual Switch at UCS level (aka as Fabric Failover) without PC or vPC

Supported

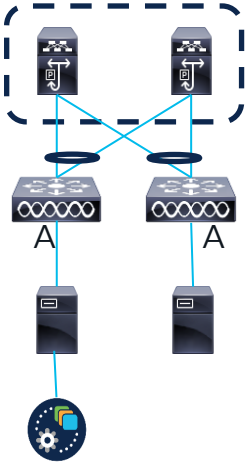

Single uplinks of Port-Group /virtual Switch

Attach vND to Network (directly)

Supported

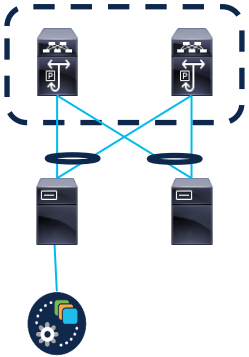

A/A uplinks of Port-Group / virtual Switch with PC or vPC

Unsupported

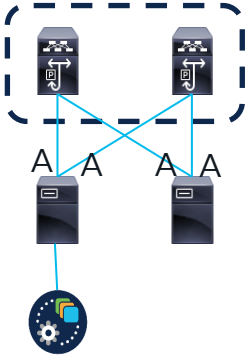

A/A uplinks of Port-Group / virtual Switch without PC or vPC

Unsupported

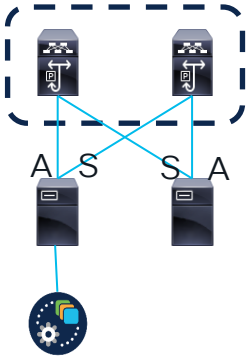

A/S uplinks of Port-Group /virtual Switch at Hypervisor level without PC or vPC

Supported

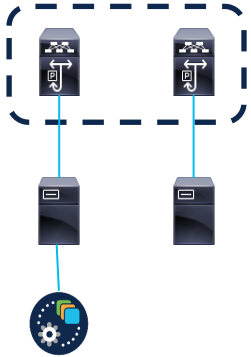

Single uplink of Port-Group /virtual Switch

Supported

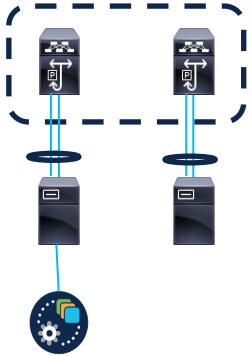

Port-Channel used as uplink of Port-Group /virtual Switch

Persistent IPs and their usage

Important Requirement for NDI 5.1 and later for DCNM/NDFC and for NetFlow/SFlow

- Nexus Dashboard Cluster Nodes need to be Layer-2 Adjacent on Data Interface
- IPv4 requirements:
  - You need to assign 6 IPs, out of the range of the Data Interface Subnet, Nexus Dashboard Cluster. 3 IP are needed for SW Telemetry receiver and 3 for HW Telemetry.
- IPv6 requirements:
  - You need to assign 7 IPs, out of the range of the Data Interface Subnet, Nexus Dashboard Cluster. 3 IP are needed for SW Telemetry receiver, 3 for HW Telemetry and 1 for Assurance Collector

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
- Currently used by NDI 6.0, when monitoring DCNM based Sites or Netflow/Sflow collection used for ACI/DCNM
  - Only required for the Data Subnet of ND

Persistent IP Pool 2/2

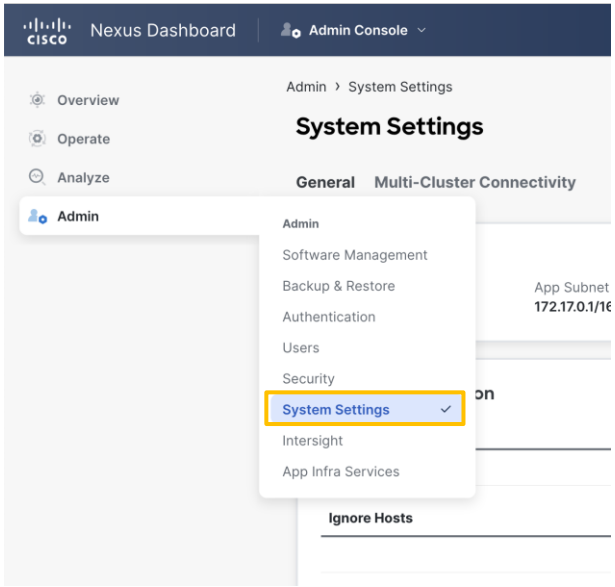

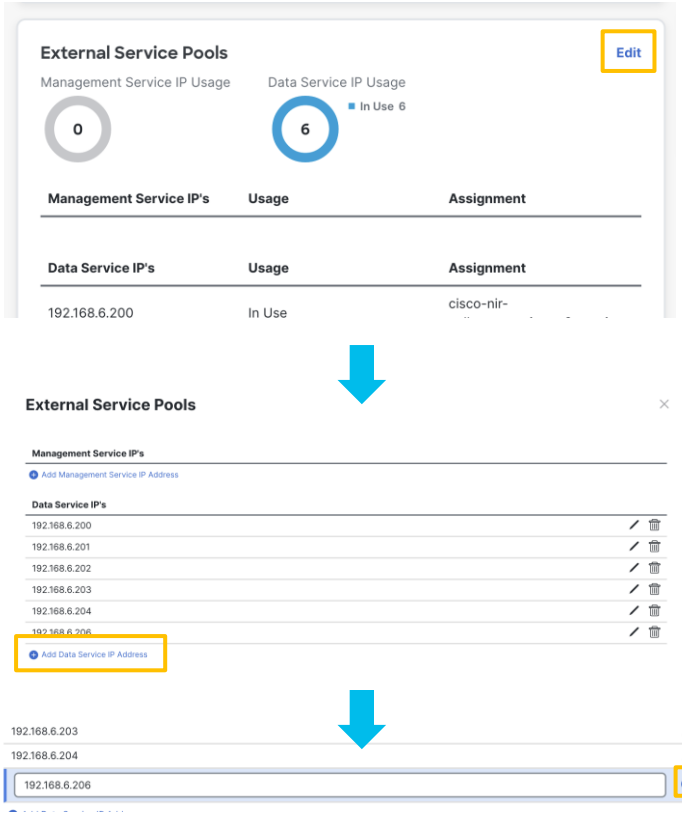

Apps	Mgmt Interface	Data Interface	Persistent IPs	Support for Data and Mgmt in the same Subnet**
NDFC	L2 adjacent	L2 adjacent / L3 adjacent with L3 HA	2 IPs in mgmt network (for default settings) or 2 IPs data network (for POAP etc. via data network) + 1 IP per fabric for EPL in data network	no
NDI for 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

ND L3 peering / L3 HA

- 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

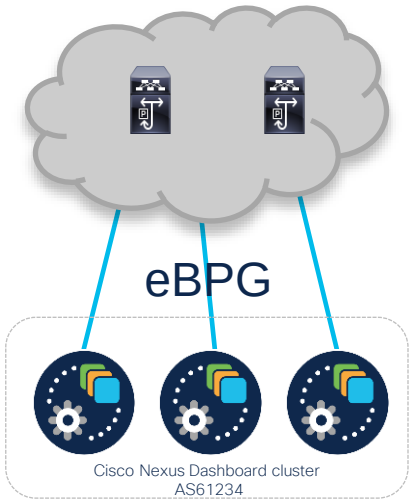

Reachability of Persitent IPs per ND Node

- 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.

Attaching ND to your Network

ND Cluster attached to any Networking Infra

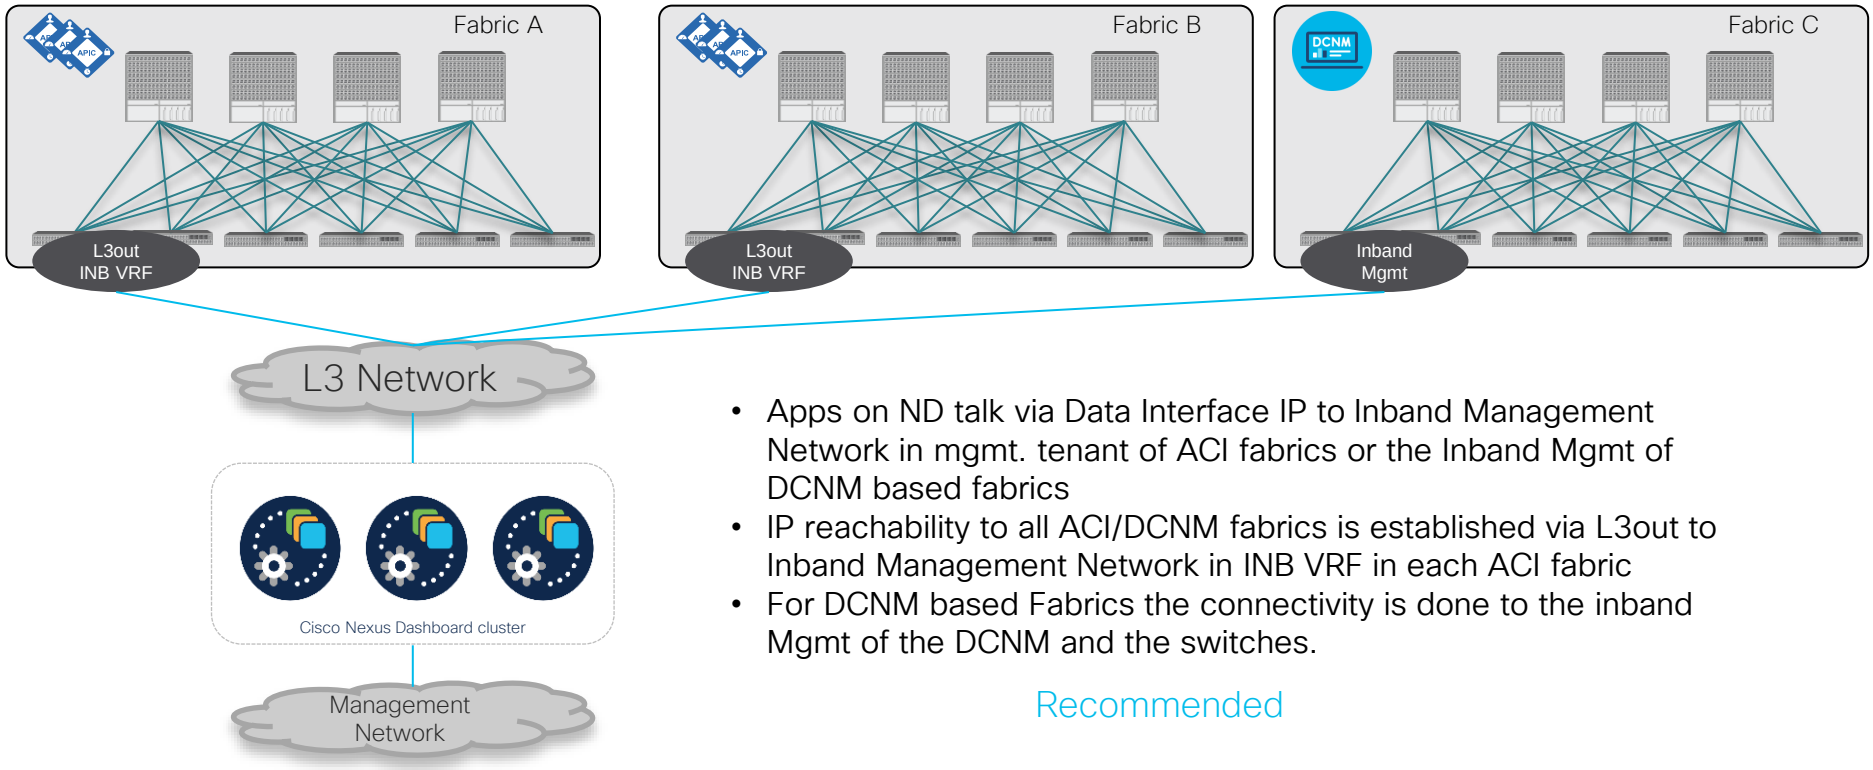

- Apps on ND talk via Data Interface IP to Inband Management Network in mgmt. tenant of ACI fabrics or the Inband Mgmt of DCNM based fabrics
- IP reachability to all ACI/DCNM fabrics is established via L3out to Inband Management Network in INB VRF in each ACI fabric
- For DCNM based Fabrics the connectivity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

Recommended

ND Cluster attached to DCNM/NDFC based Fabric

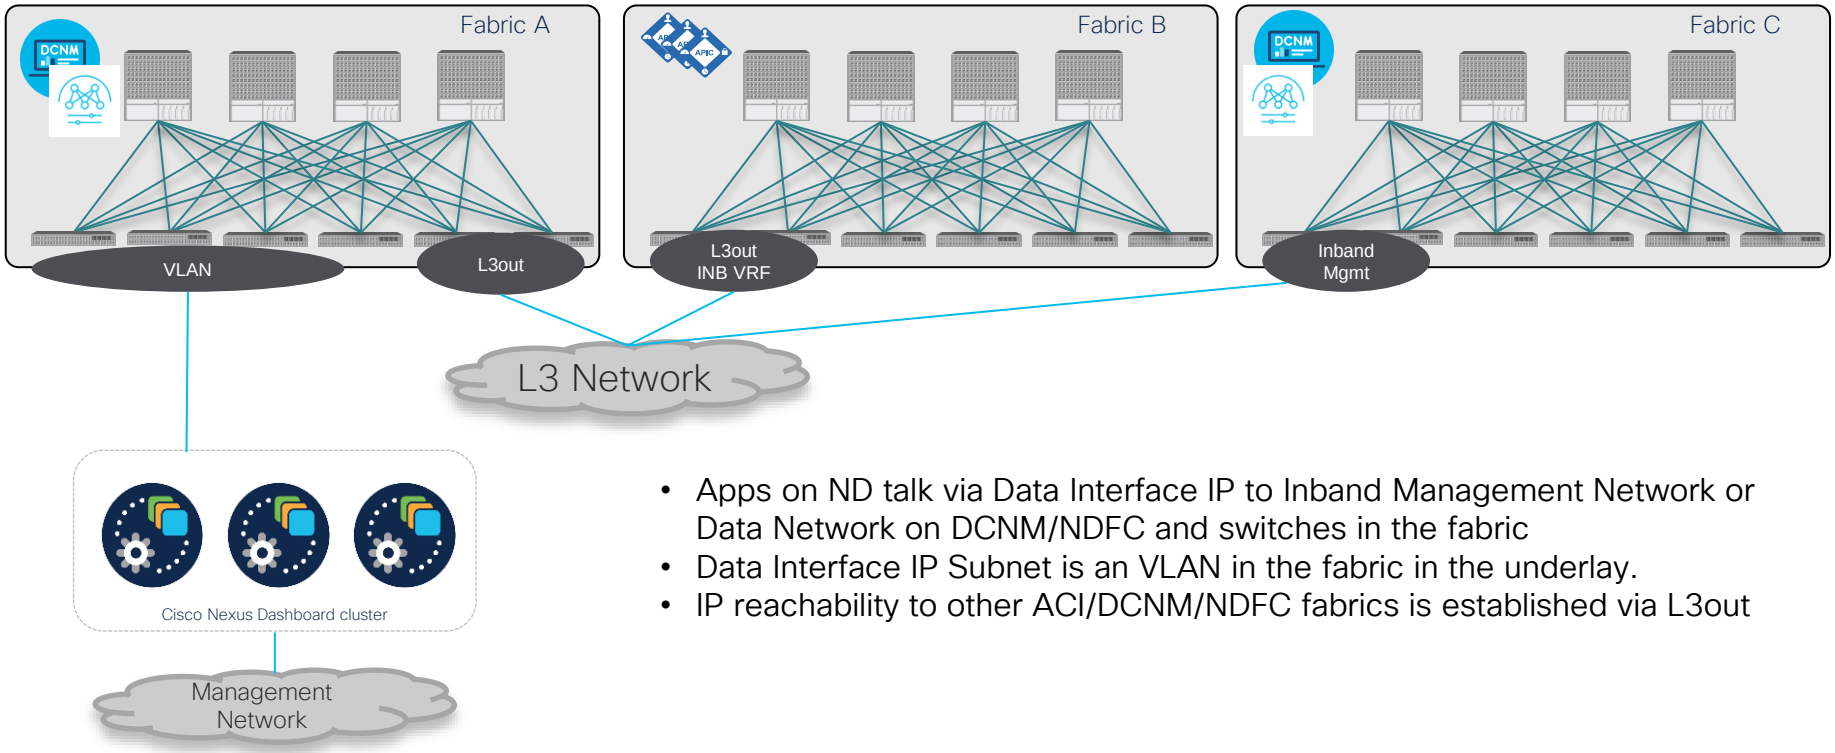

- Apps on ND talk via Data Interface IP to Inband Management Network or Data Network on DCNM/NDFC and switches in the fabric
- Data Interface IP Subnet is an VLAN in the fabric in the underlay.
- IP reachability to other ACI/DCNM/NDFC fabrics is established via L3out

ND Cluster attached to ACI Fabric

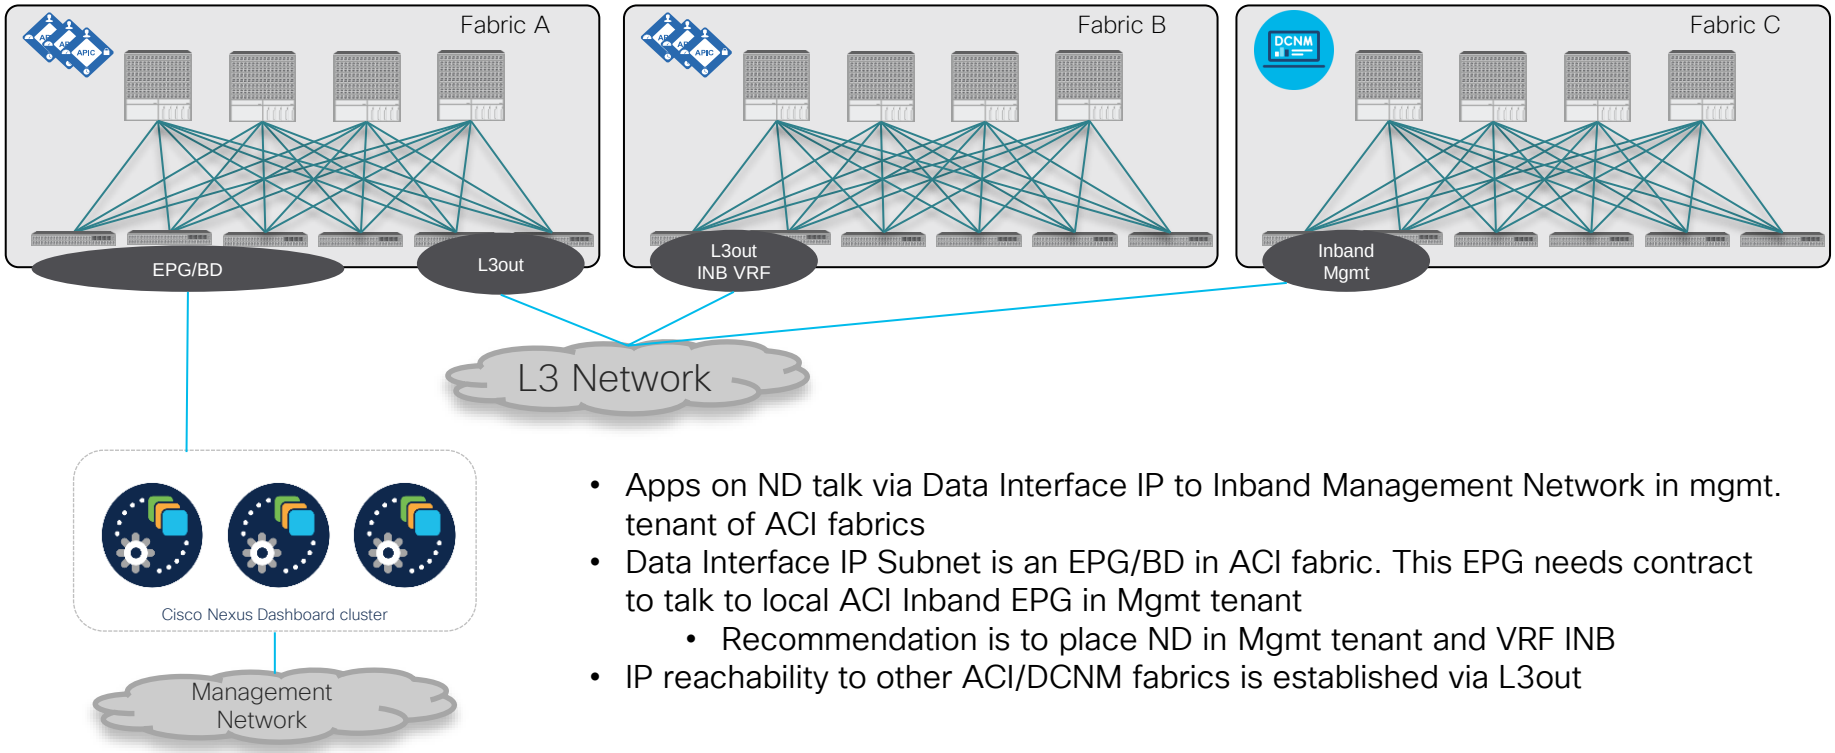

- Apps on ND talk via Data Interface IP to Inband Management Network in mgmt. tenant of ACI fabrics
- Data Interface IP Subnet is an EPG/BD in ACI fabric. This EPG needs contract to talk to local ACI Inband EPG in Mgmt tenant
  - Recommendation is to place ND in Mgmt tenant and VRF INB
- IP reachability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

(recommended for NDO)

Number of Sites	1	2	3	4	5
1	M1, M2, M3				
2	M1,M2	M3,S1			
3	M1	M2	M3		
4	M1	M2	M3	S1	
5	M1	M2	M3	S1	

M1, M2, M3 : ND Master 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 1/2

- Try to keep the potential points of failure for reachability between the ND nodes as low as possible.
- When distributing a ND cluster
  - ND Data and Mgmt interface of ND nodes can be in different subnets. Only IP connectivity is needed. (Please allow ports listed in documentation)!
  - For NDI being hosted on ND2.1 or later for DCNM/NDFC based fabrics, you need to have the Data Interfaces of the ND nodes L2 adjacent or eBGP enabled and provide persistant IPs!
  - For NDI being hosted on ND2.1 or later leveraging Netflow/Sflow, you need to have the Data Interfaces of the ND nodes L2 adjacent and provide persistant IPs!
  - When deploying NDFC on ND2.1 or later the Management Interfaces of ND nodes have to be L2 adjacent. Also Data Interfaces of the ND nodes have to be L2 adjacent.

Some Deployment Considerations 2/2

- In MPOD, ACI is taking care of the reachability, Keep in mind loosing IPN connectivity will e.g. break NDI
- 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 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 master 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 master node, a ND standby node.
- In case of a failure of 2 ND master nodes, you have to manual promote the standby to master to replace a failed master.
  - NDO/NDFC are the only apps surviving this.
  - After the failed master comes back online, it needs to be wiped and re-added as standby node.

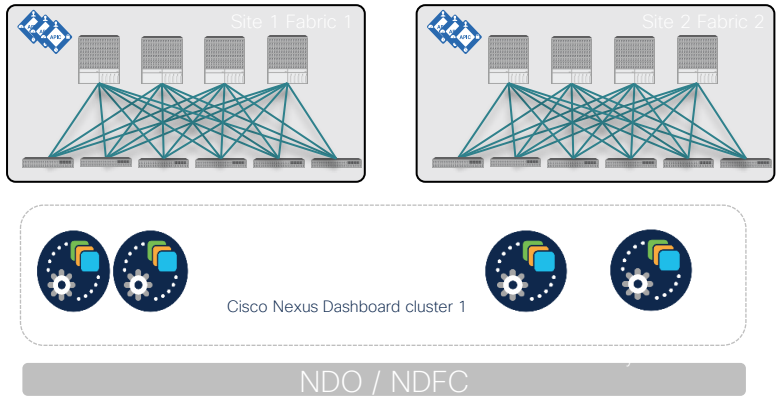

Option 1: 1 Site/Fabric (below 500 nodes) NDI

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

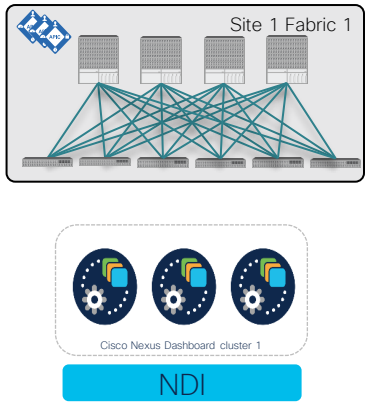

Option 2: 1+ Site (below 500 nodes) NDI

- 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

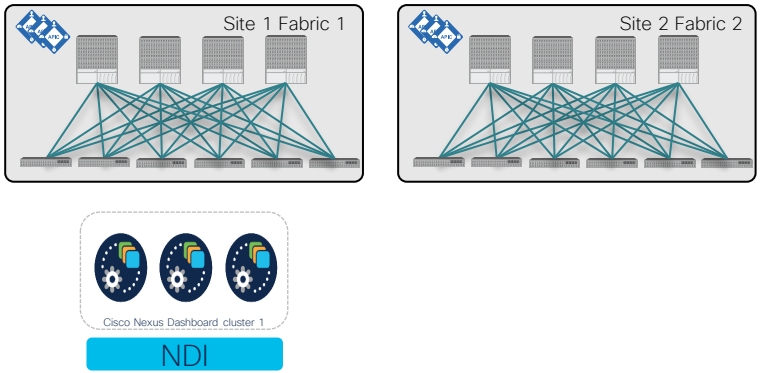

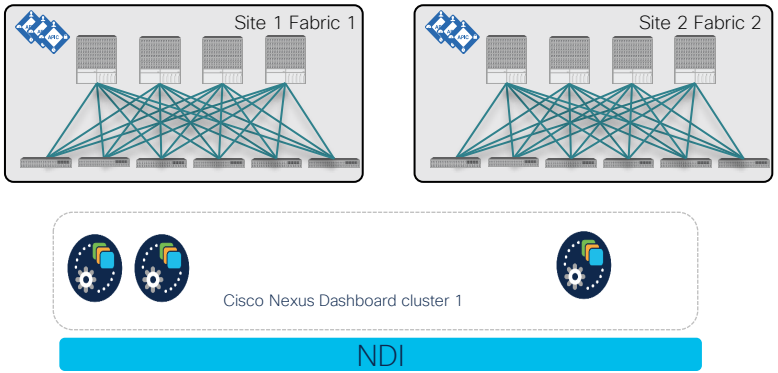

Option 3a: 1+ Site (below 500 nodes) NDI and NDO

- 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

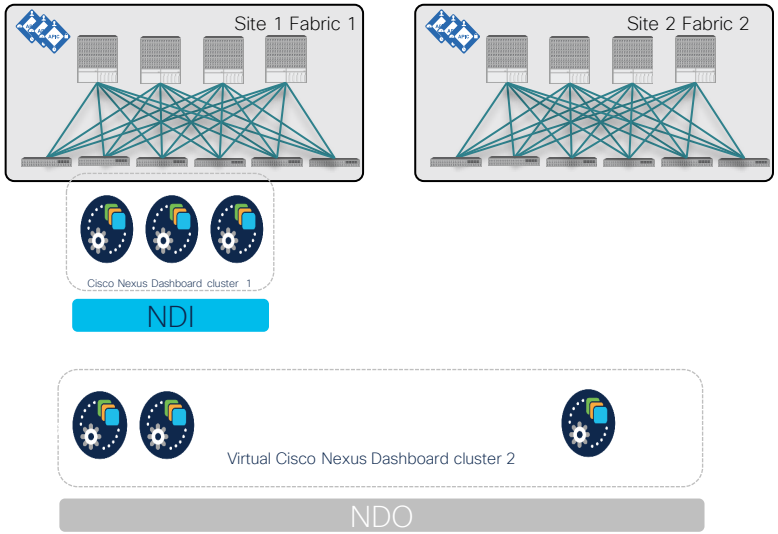

Option 3b: 1+ Site (below 500 nodes) NDI and NDO

- 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

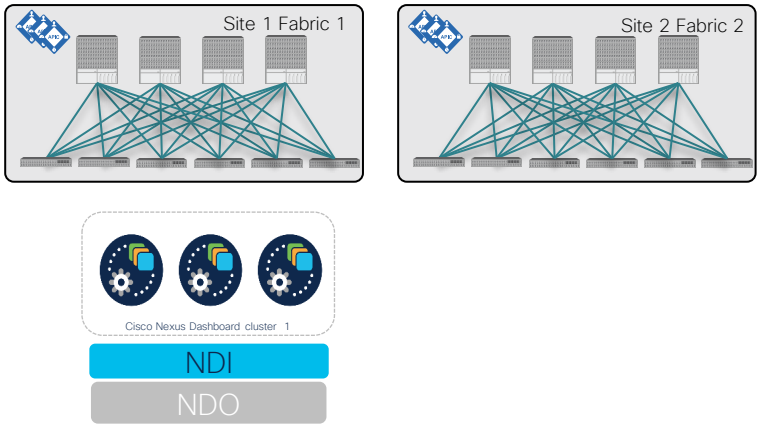

Not recommended as NDI is distributed, consider vND for NDO (Option 3a)

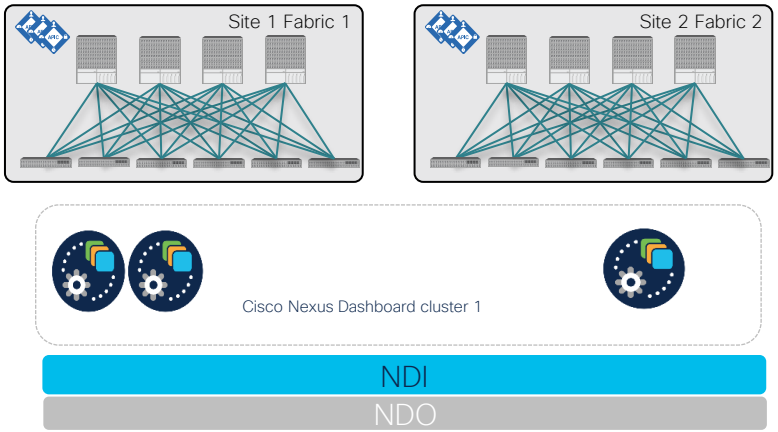

Option 4: 1+ Site (above 500 nodes) NDI and NDO

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

Recommended

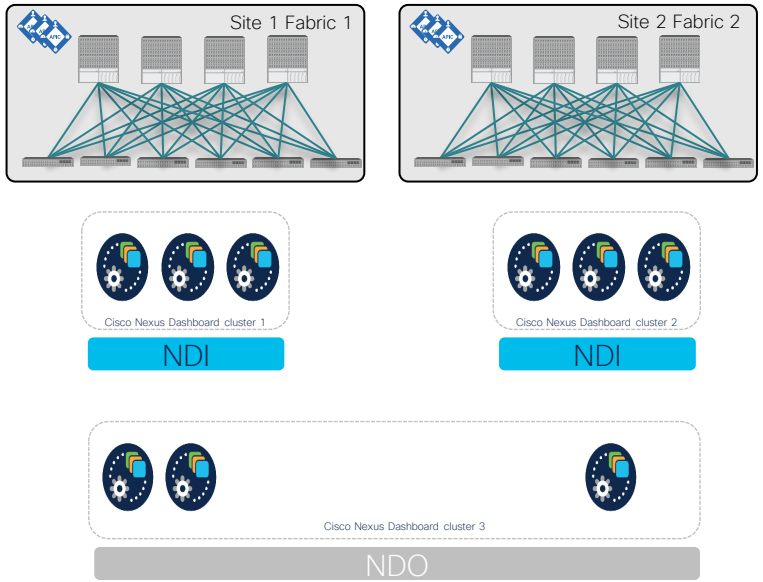

Operating Nexus Dashboard

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.

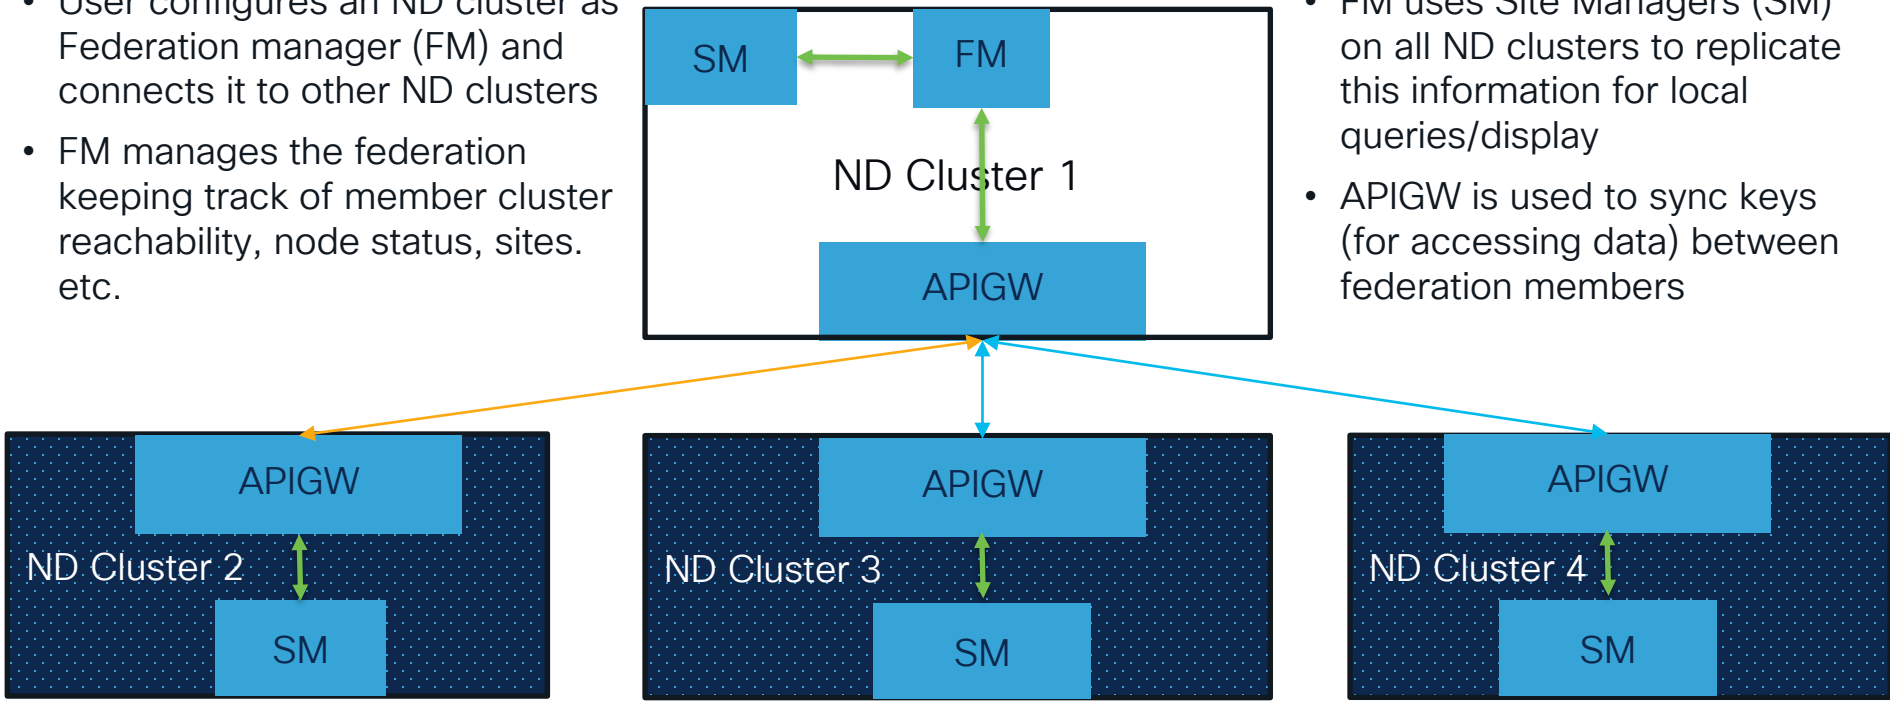

- FM uses Site Managers (SM) on all ND clusters to replicate this information for local queries/display
- APIGW is used to sync keys (for accessing data) between federation members

Onboard Clusters (Federation Configuration)

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

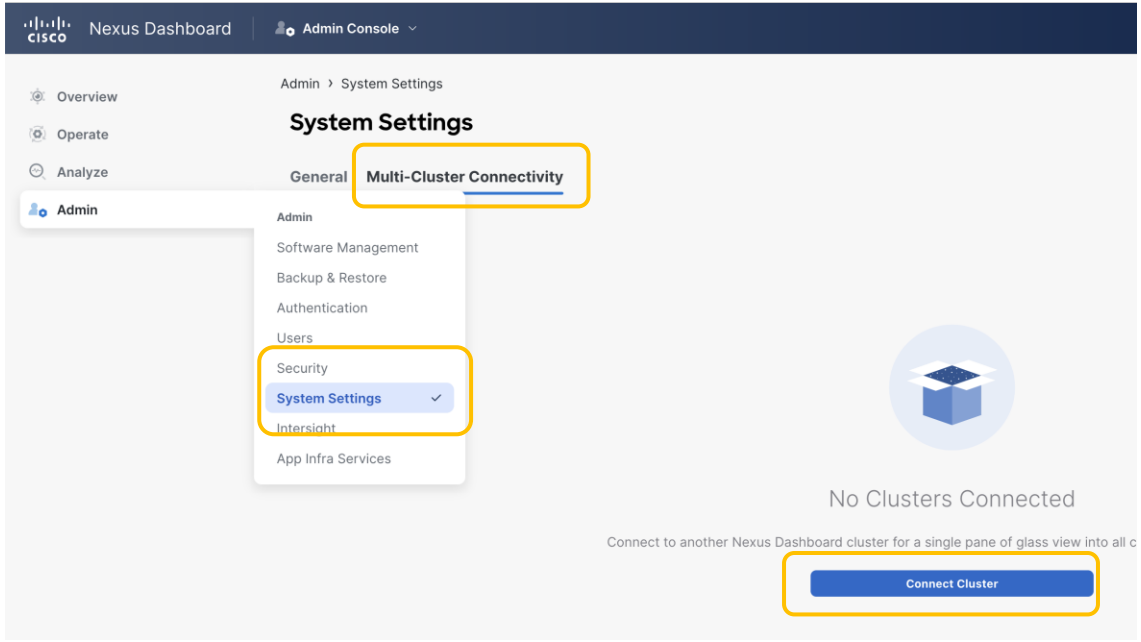

Onboard Clusters (Federation Configuration)

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

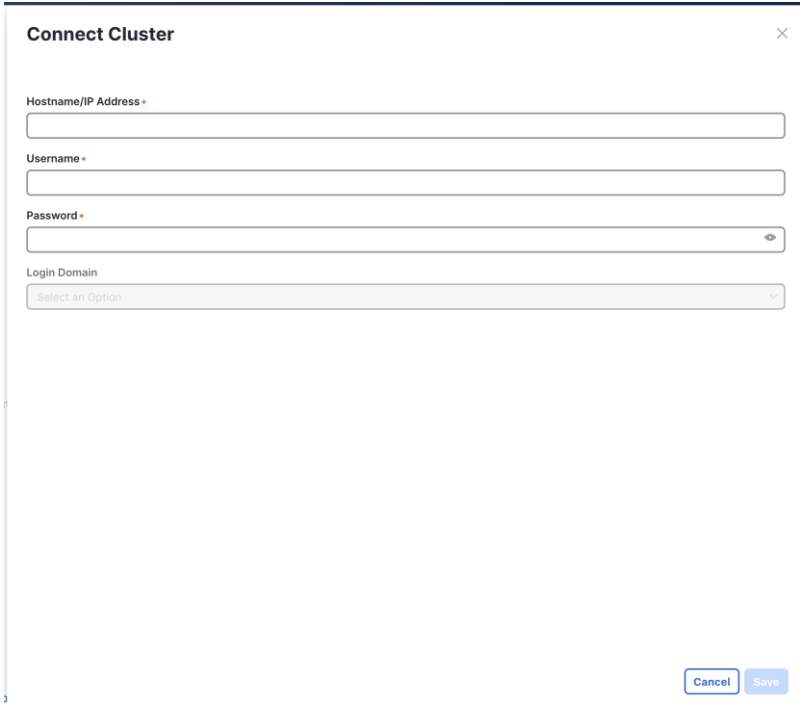

Viewing Connected Clusters' Information

- 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

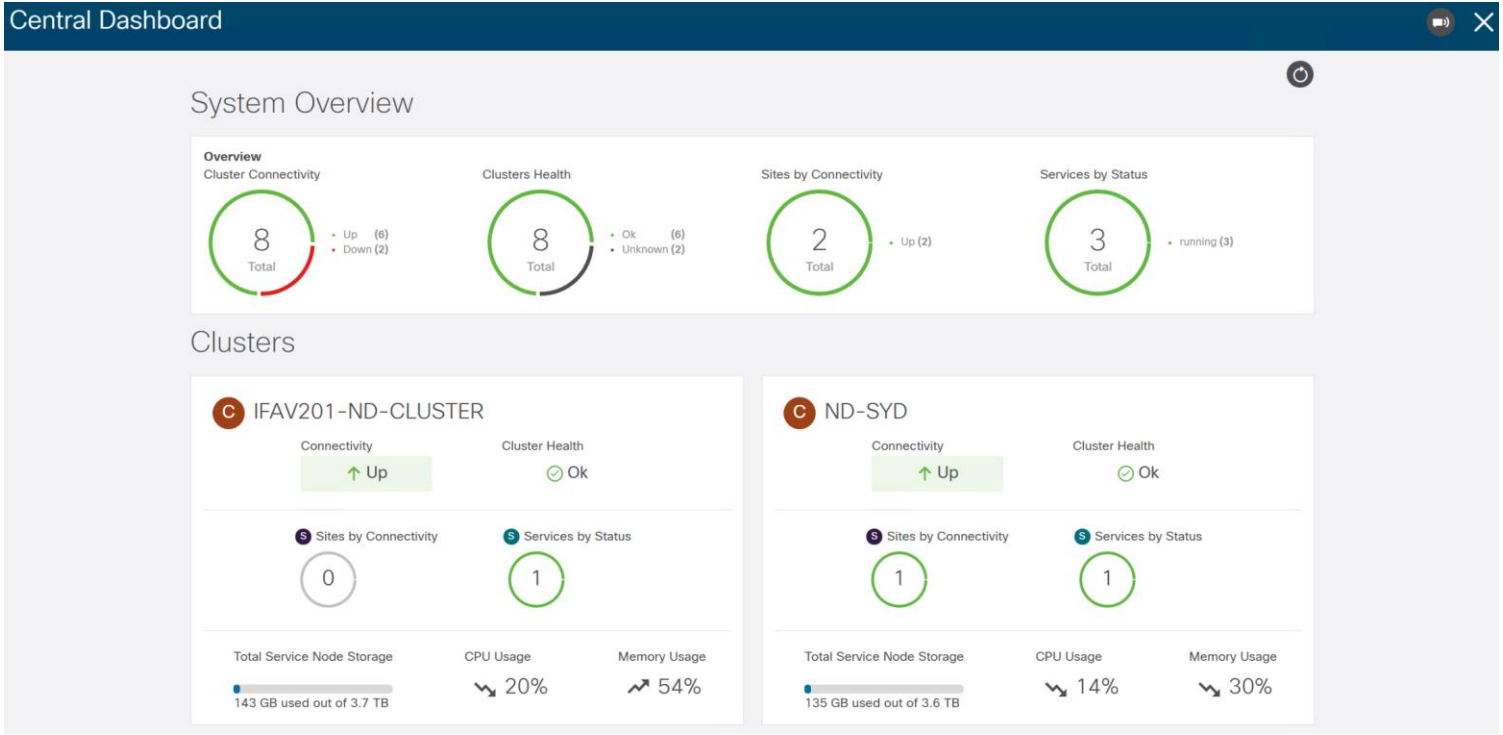

Public API

Overview

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

API UI

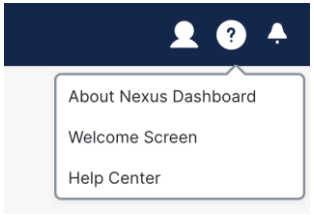

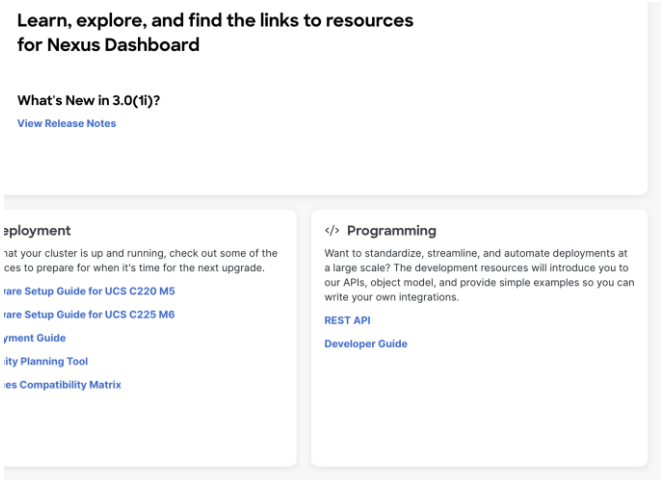

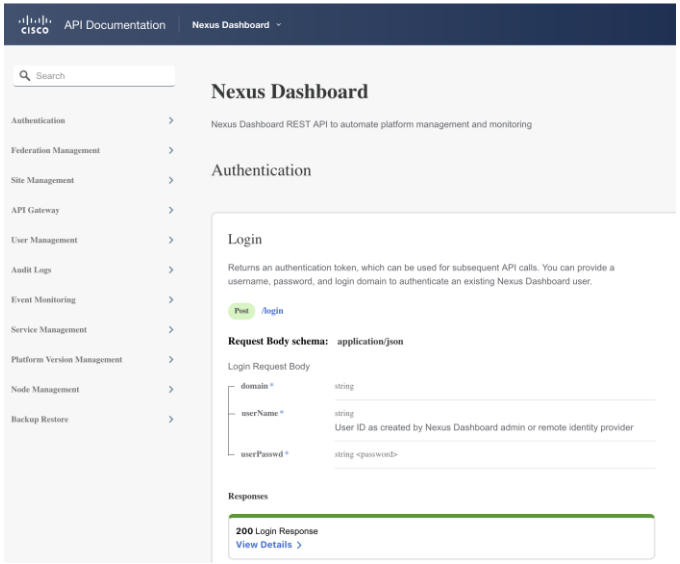

Registering Nodes to existing Cluster and Standby Node

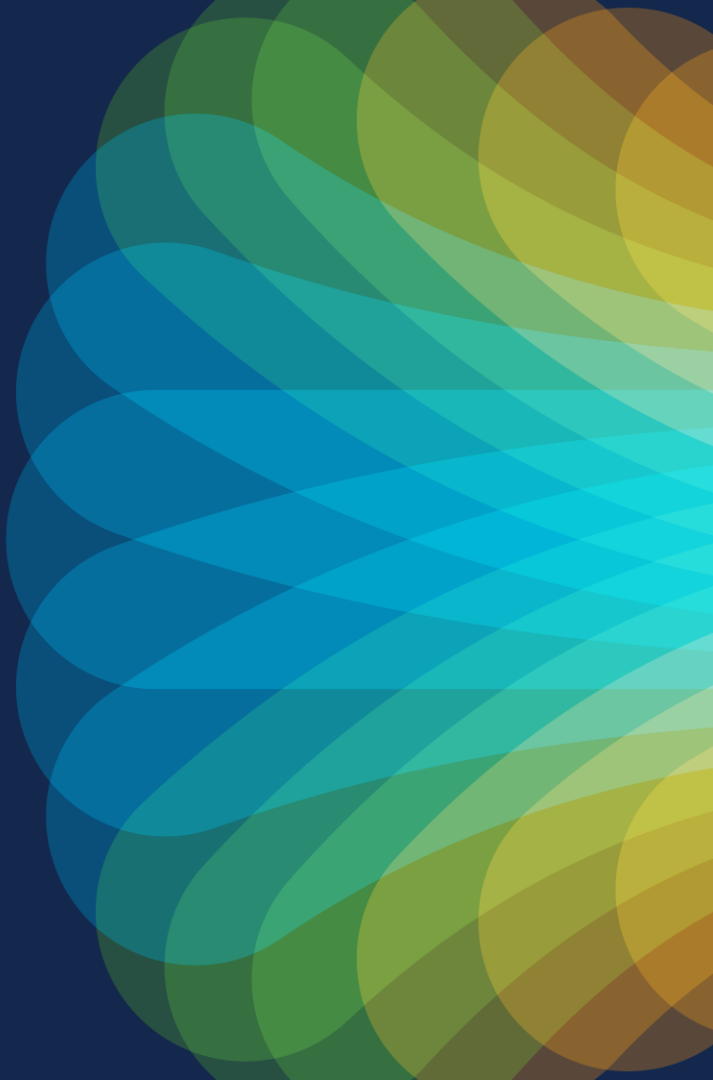

Register new Nodes and Standby Master

- 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 Master
- 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 master with an already bootstrapped node
- Workers can only be replaced by delete and re-add

Lifecycle of non-Master Nodes

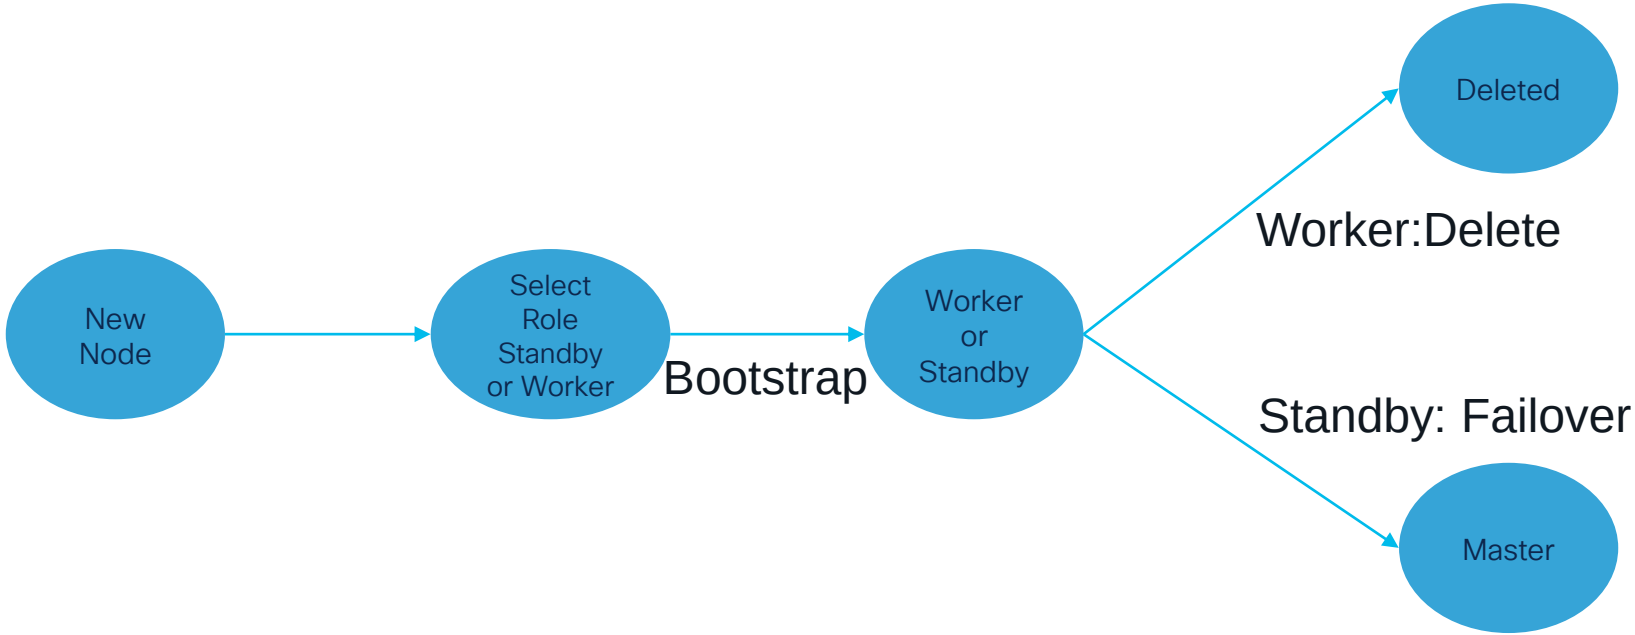

Adding a new Node

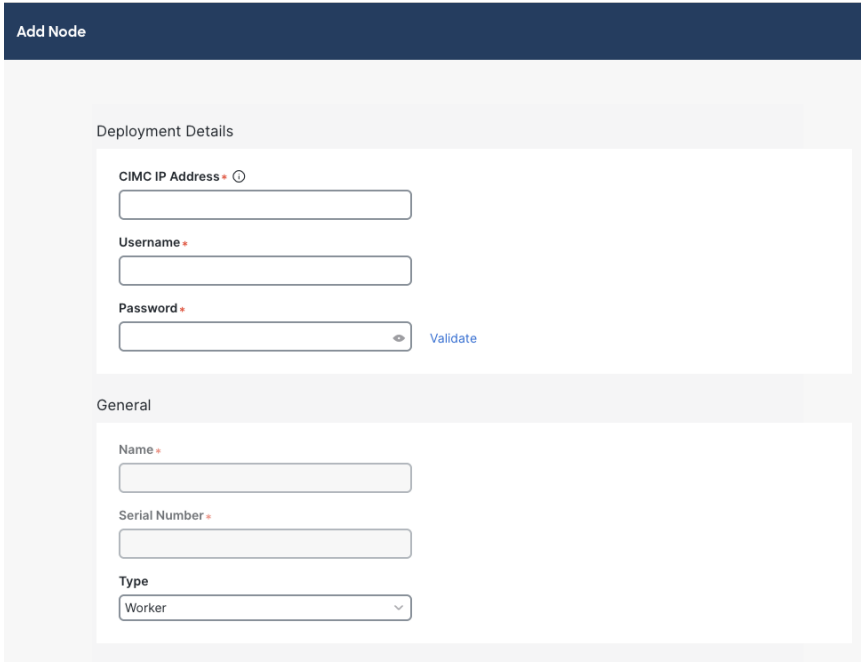

1. Provide CIMC details to discover node
2. Fill in node details
3. Node is bootstrapped and registered
4. Node status will change from “unregistered” to “discovering” to “active”

Replace a failed Master with Standby Node

Master is failed

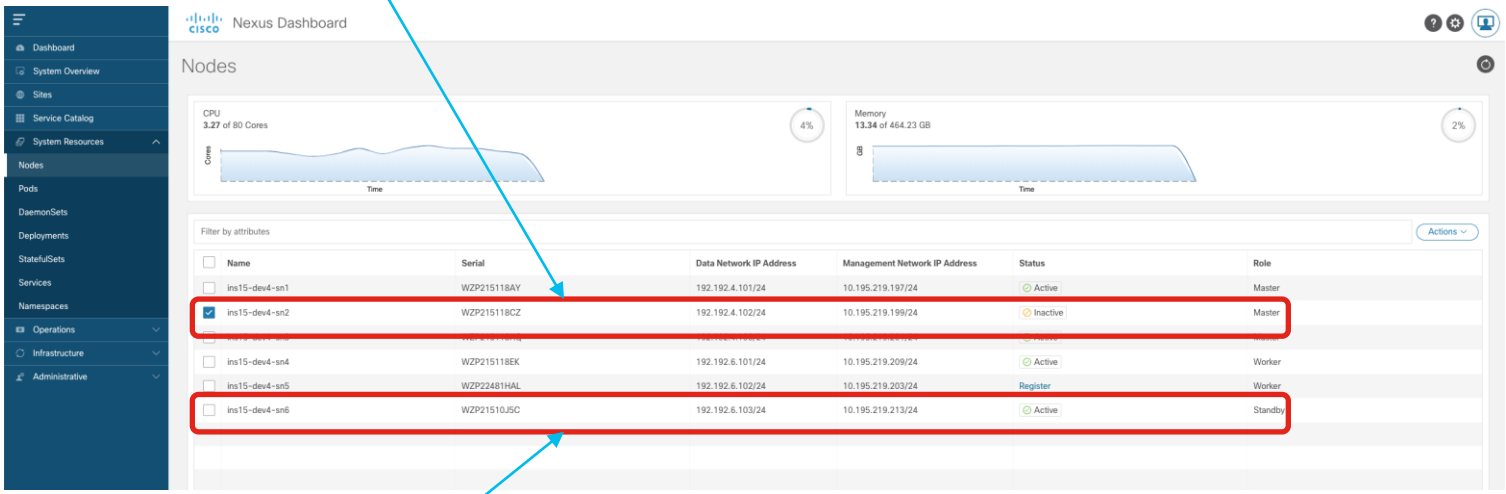

Standby Node is part of Cluster

Failover to Standby

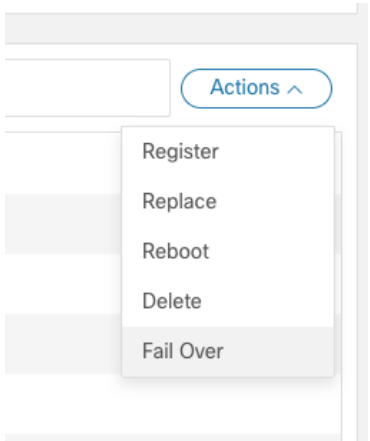

Select failed Master and click Fail Over

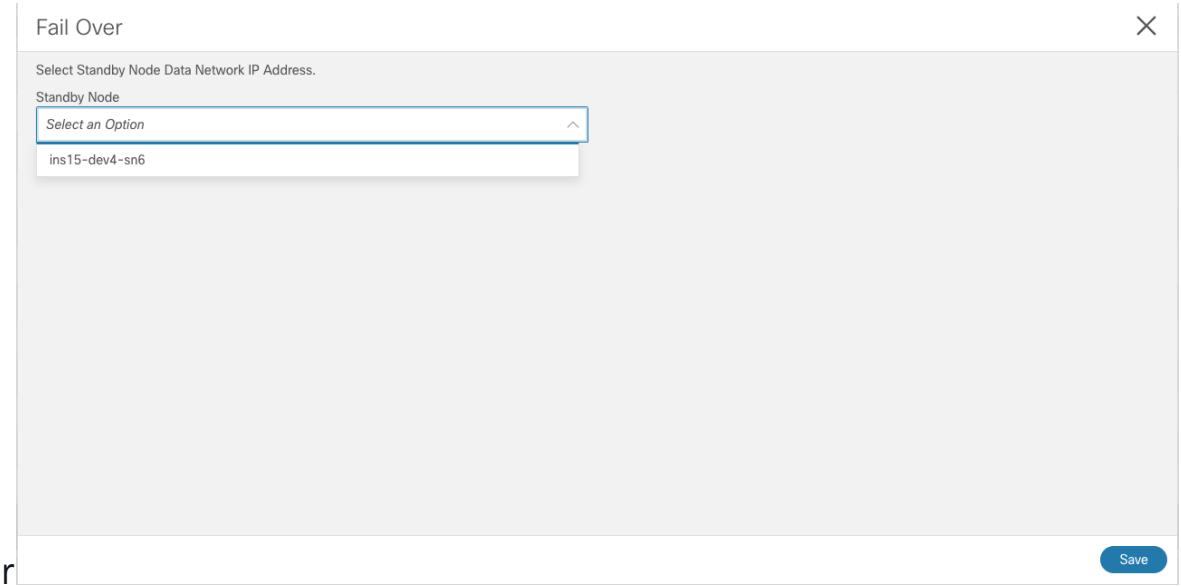

Select Standby to replace failed Master

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

Manual Recovery of 2 failed Masters

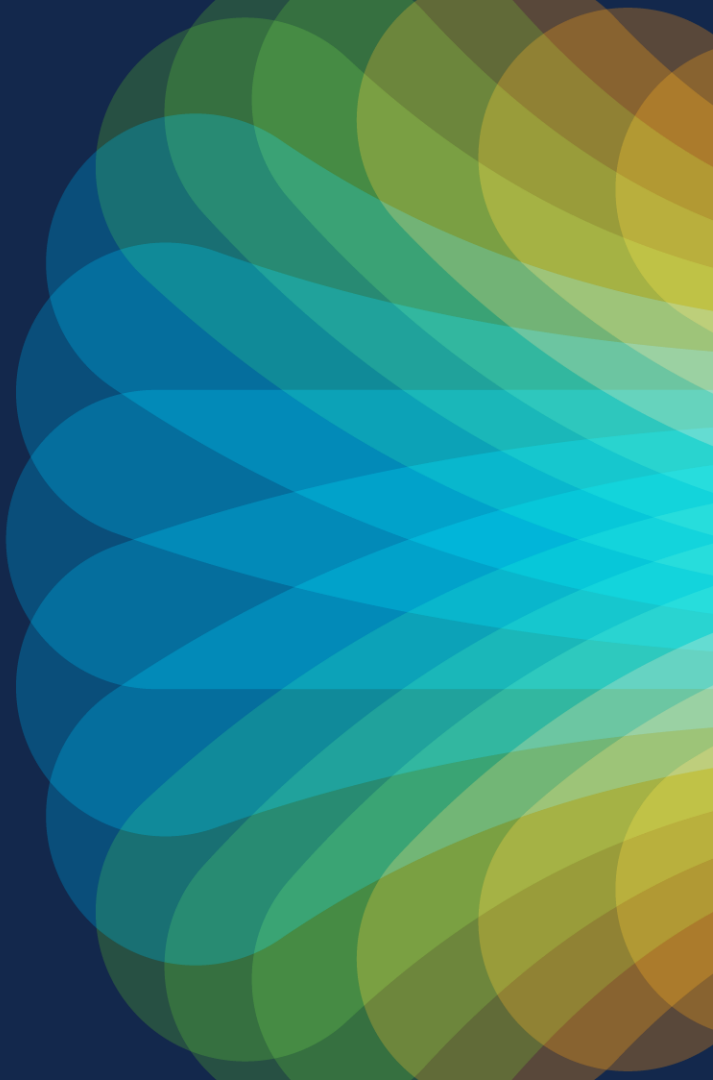

Recovery Process if 2 Masters are down 1/3

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

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

Recovery Process if 2 Masters are down 2/3

- *acs cluster masters* will show 1 Active Master and 2 Inactive Masters

```
[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 Masters are down 3/3

- Command (both failed Masters 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 Master to Standby node
- Both nodes will reboot
- Standby node will reboot and come up as Master

Recovery Process of a virtual ND

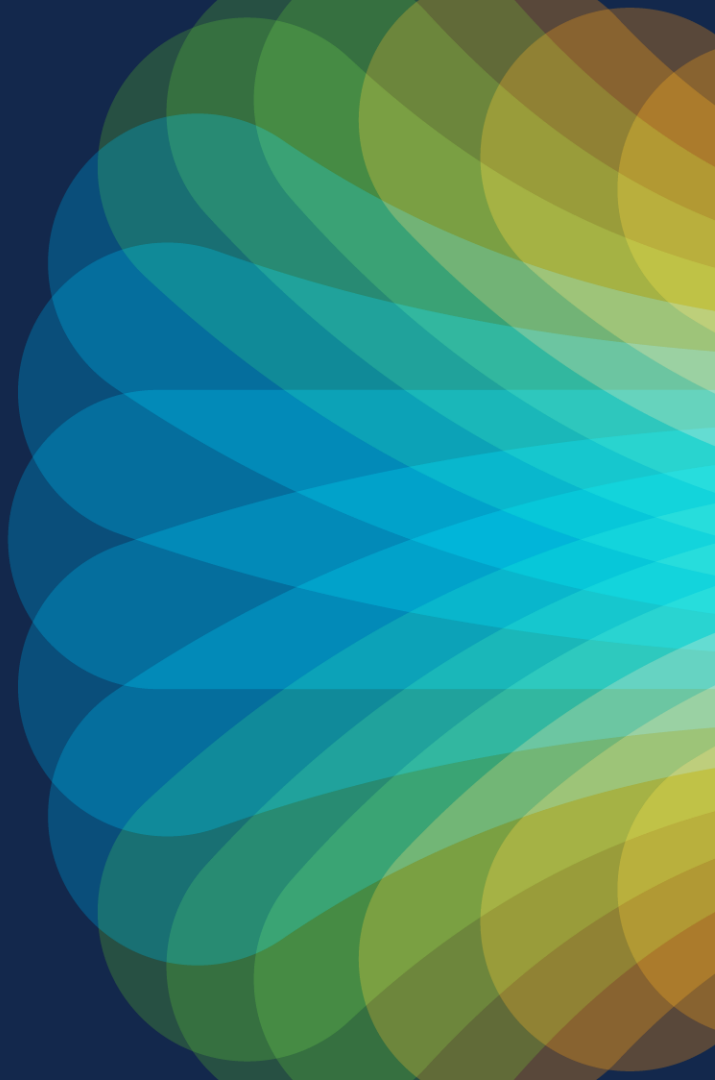

Recovery Process of a virtual ND

- 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

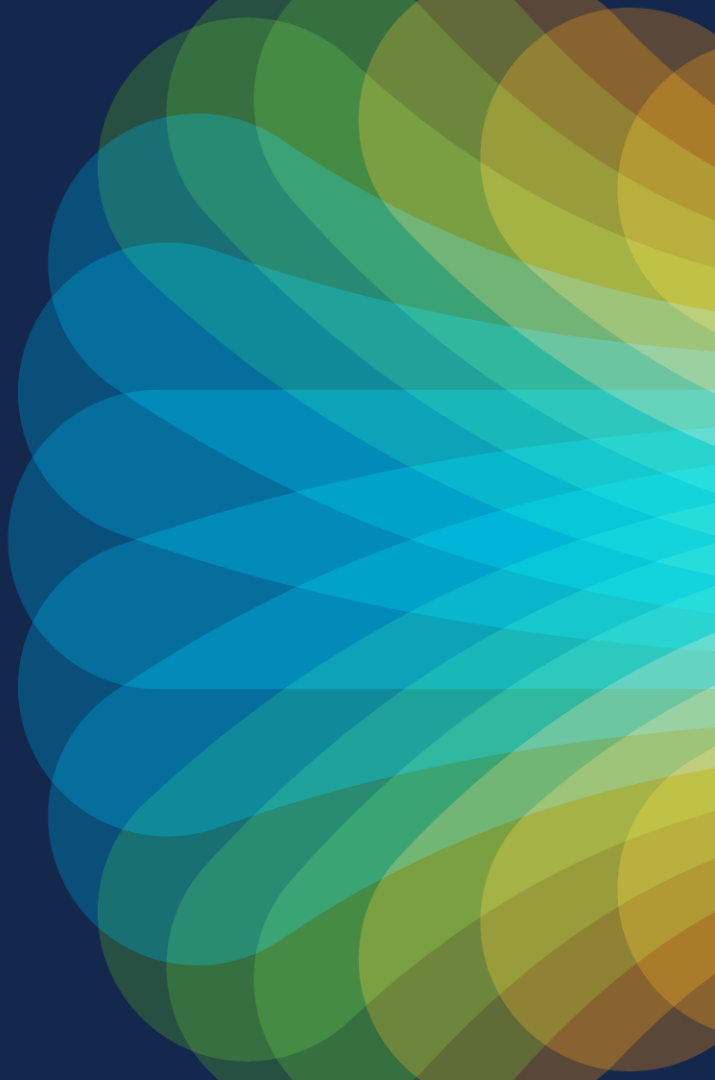

Firmware Upload

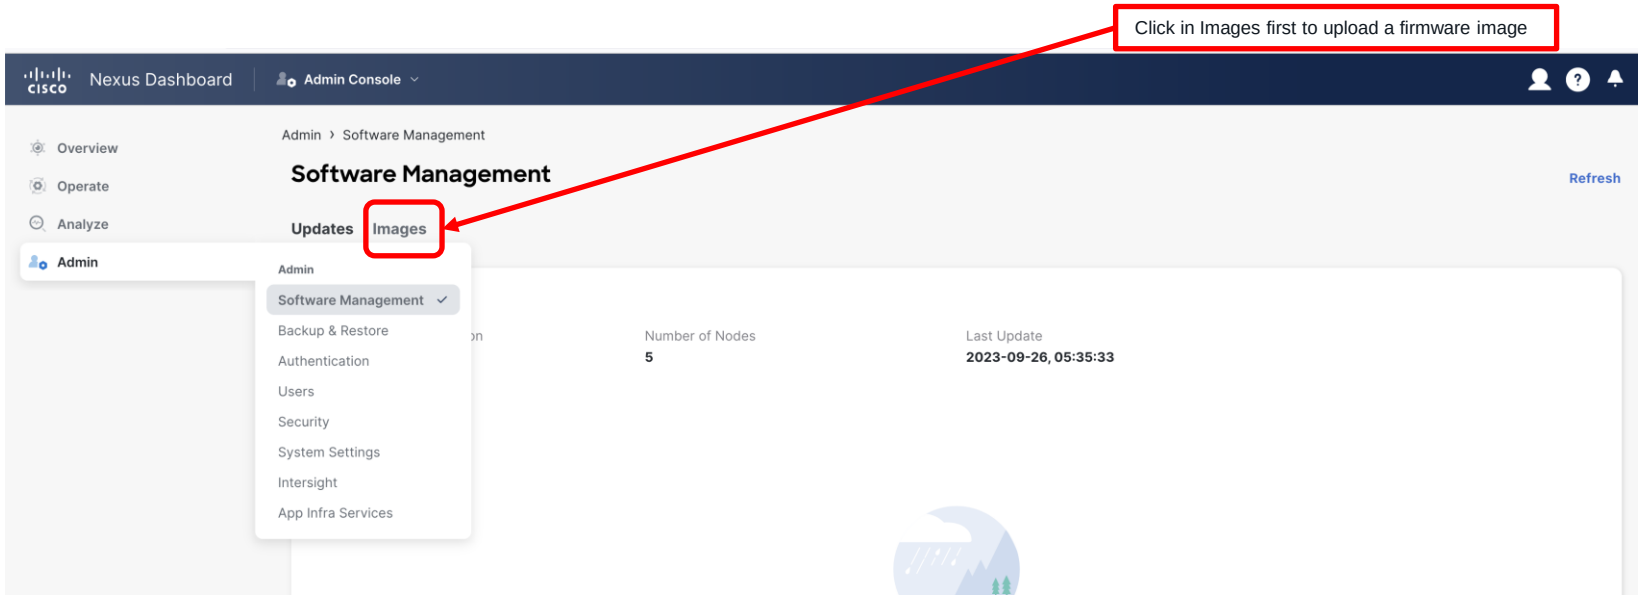

Firmware Upload

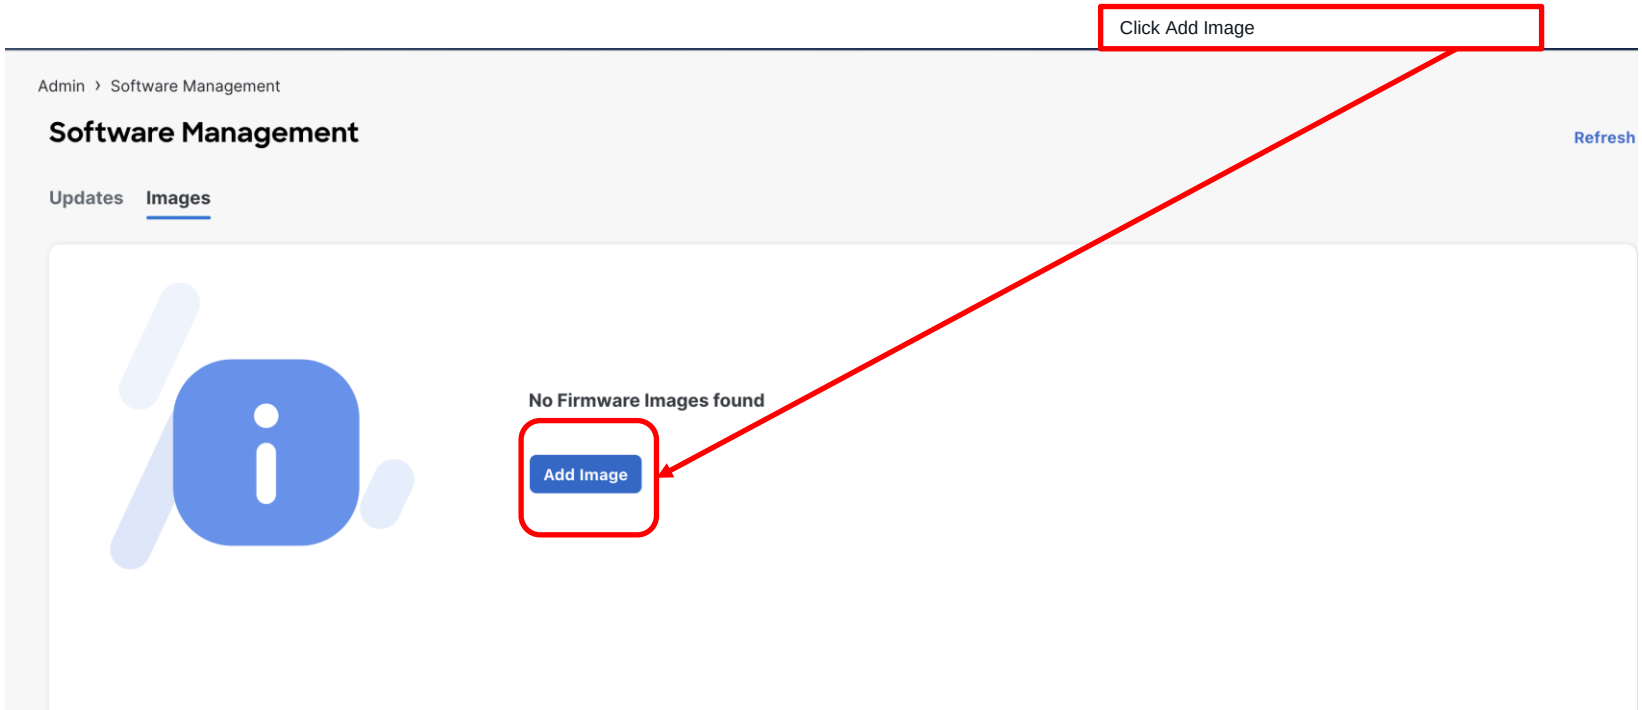

Firmware Upload

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

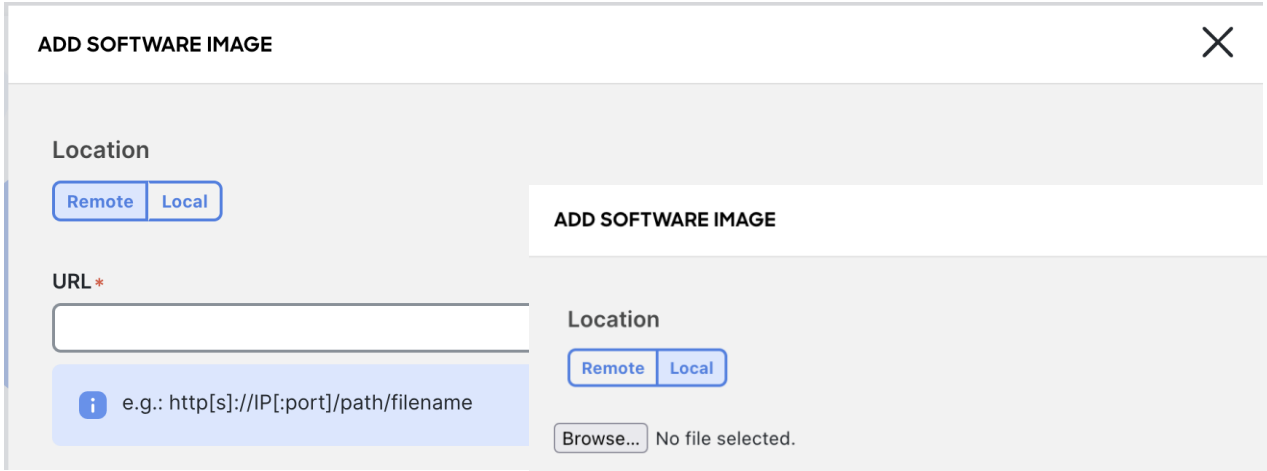

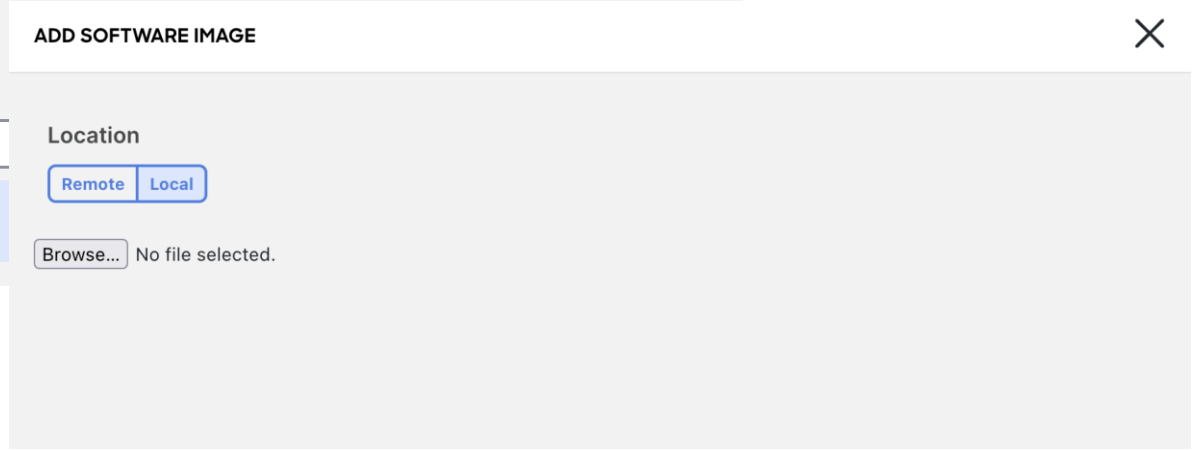

Firmware Upload

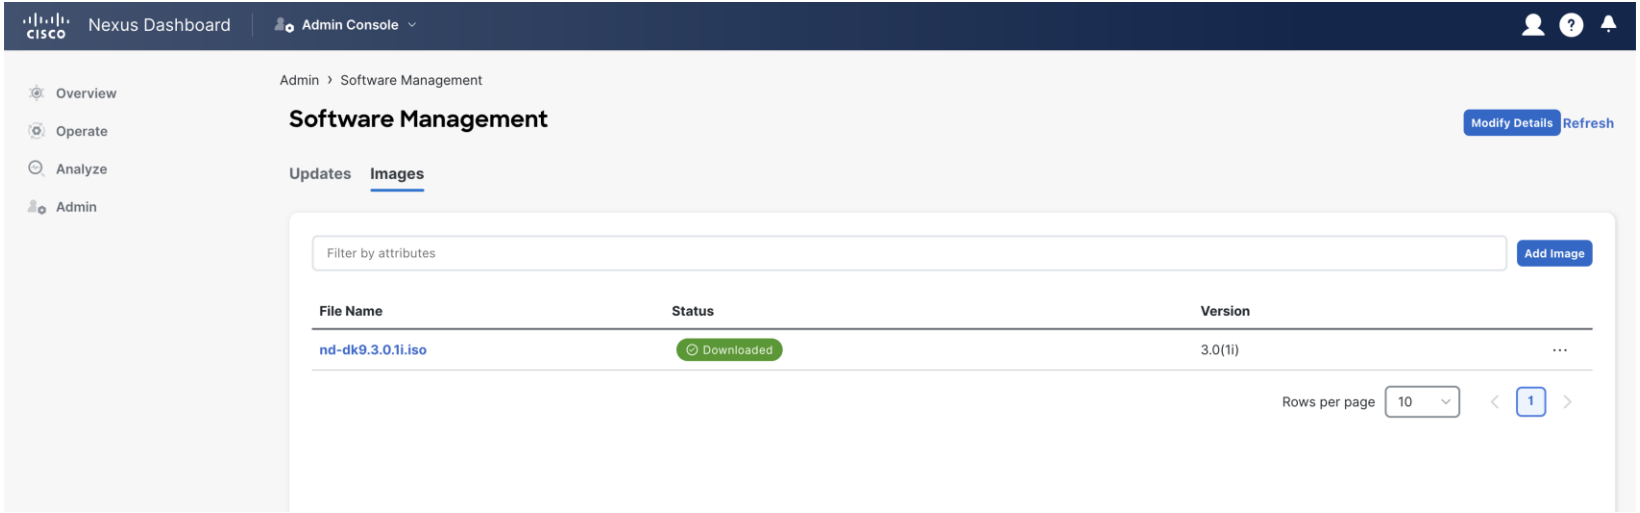

Setup Firmware Upgrade

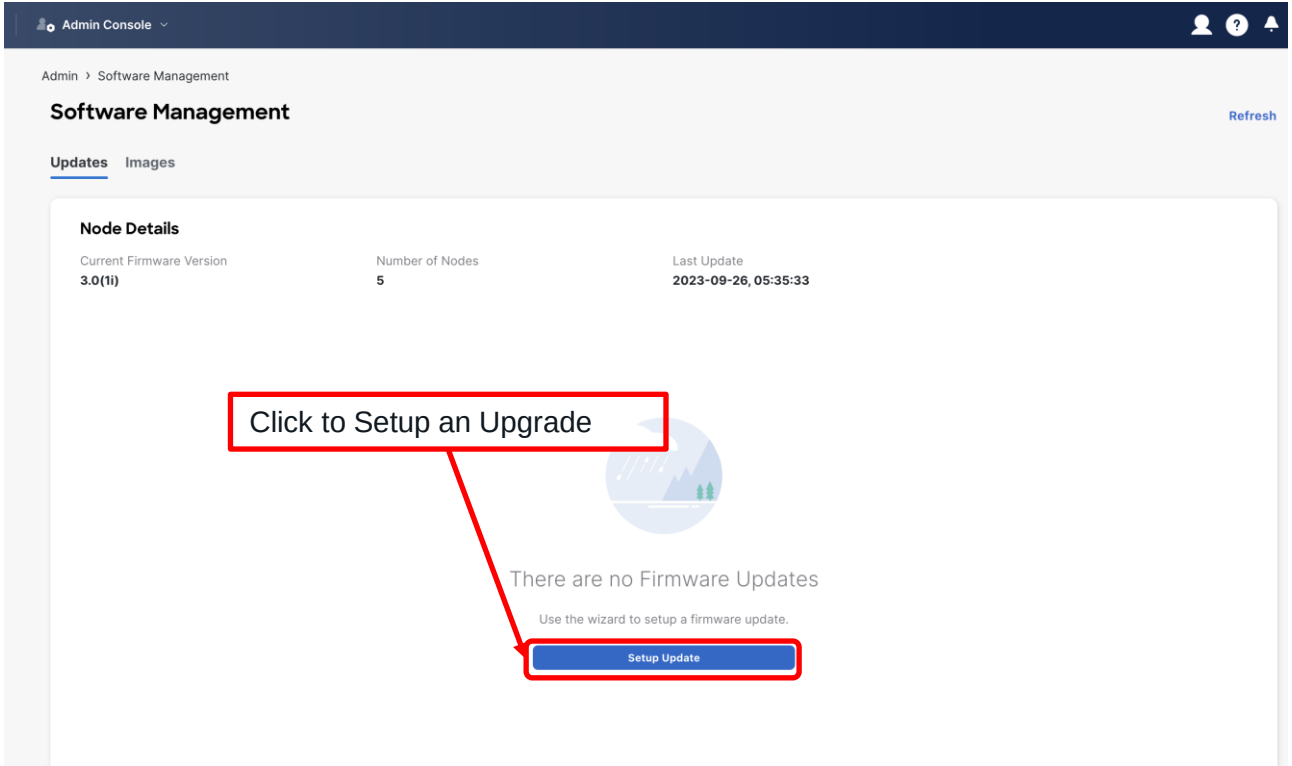

Select Firmware

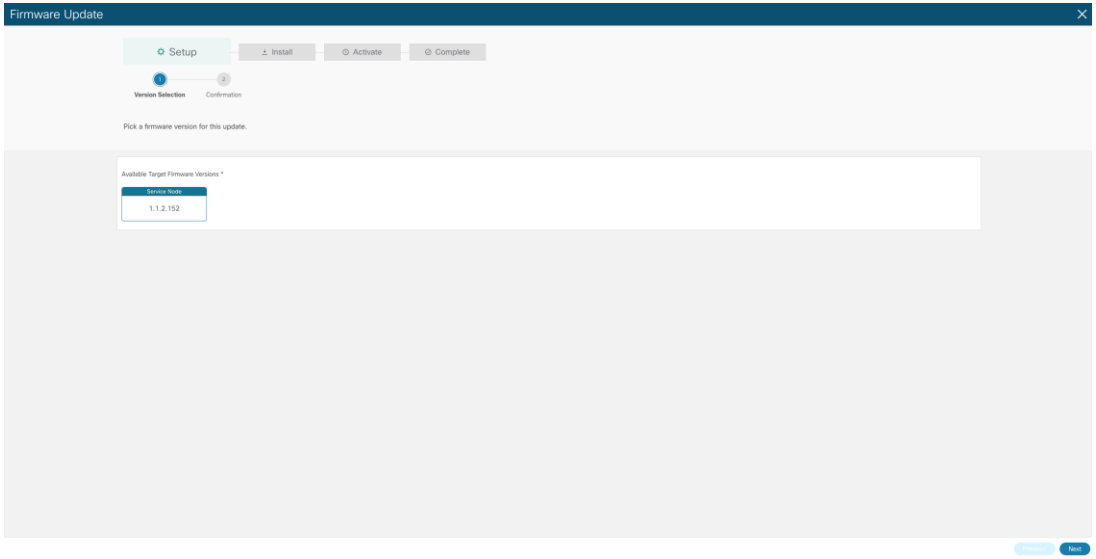

Current Cluster Setup is validated

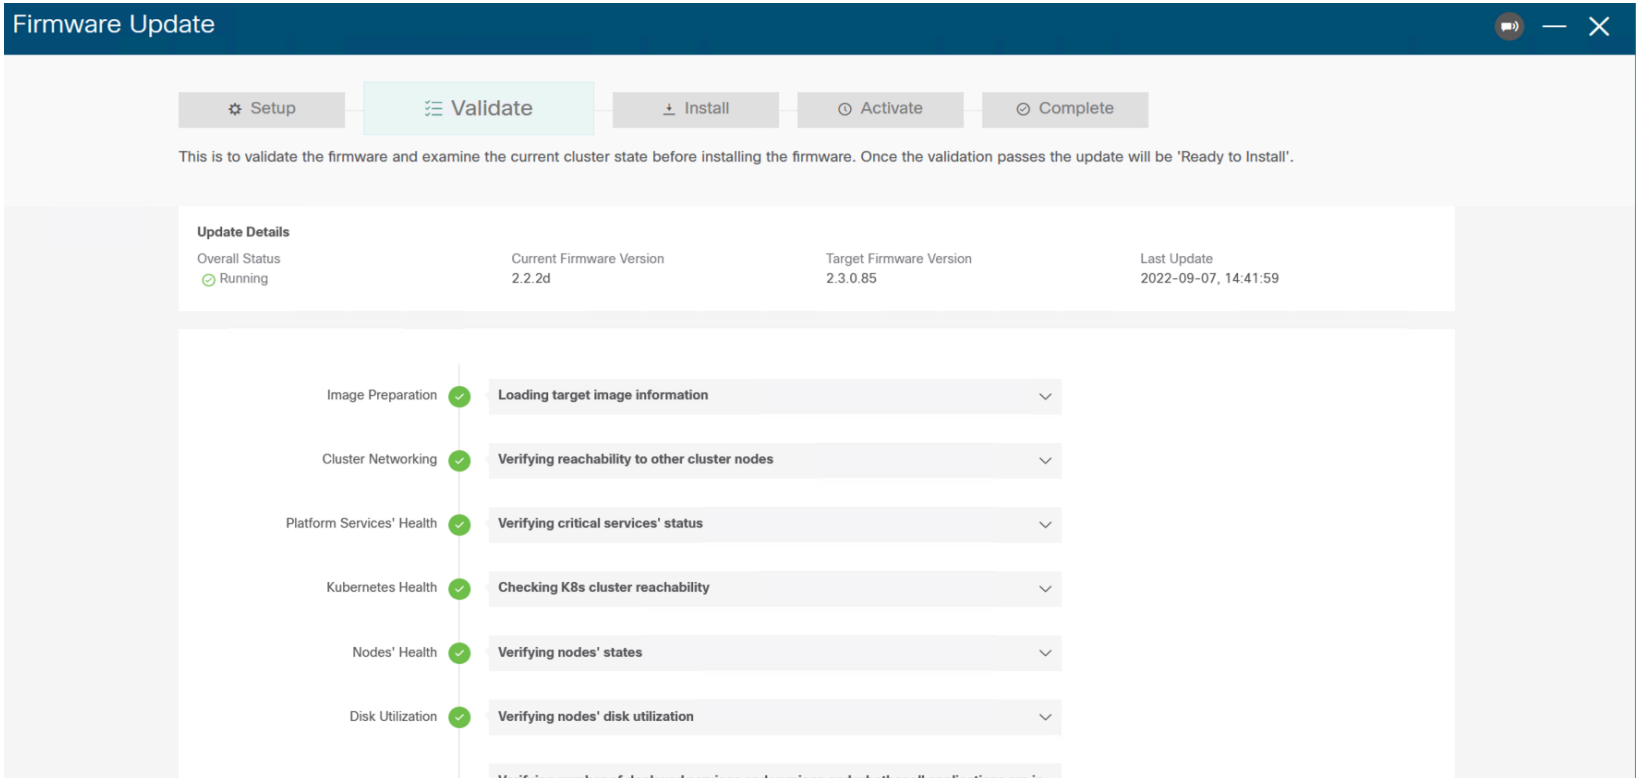

Install Firmware to Nodes

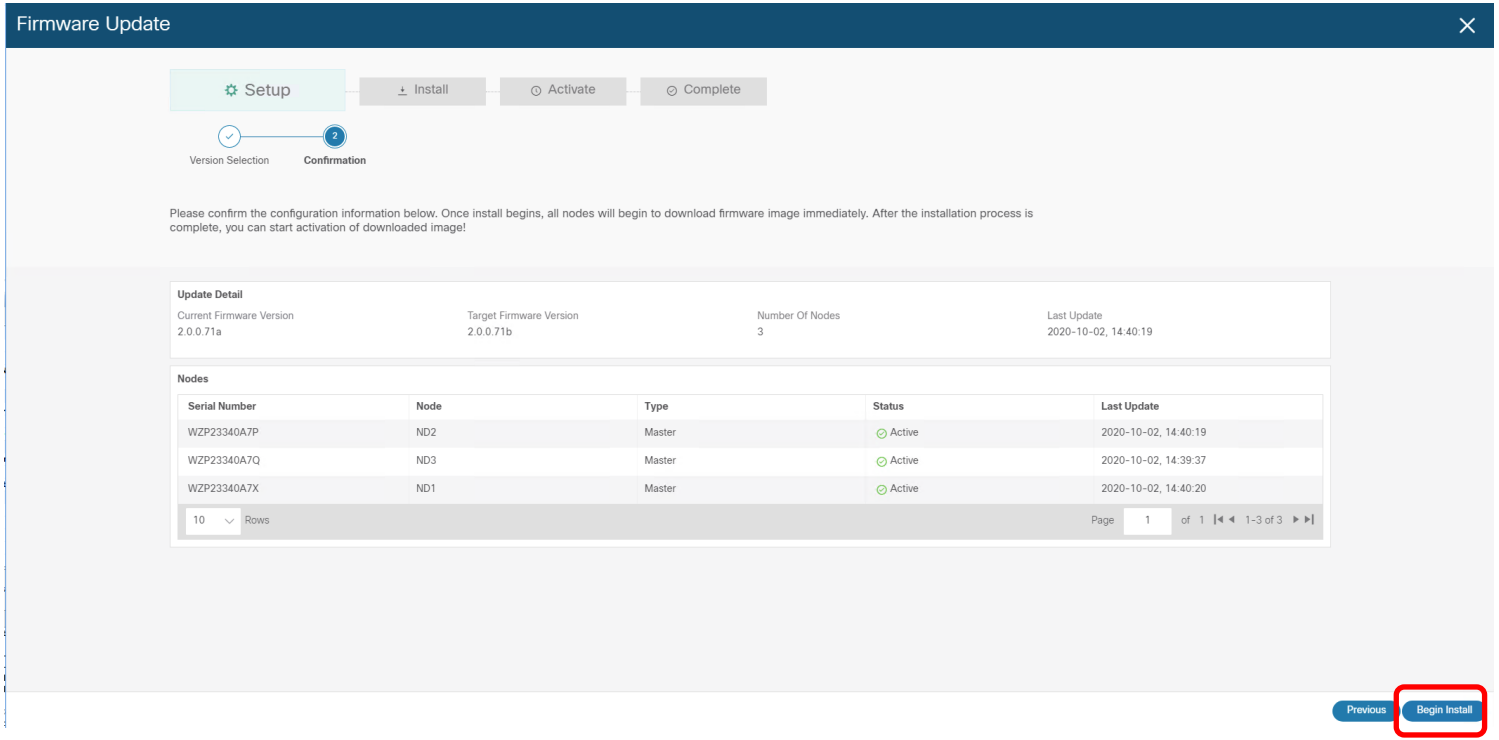

Installing Firmware to Nodes

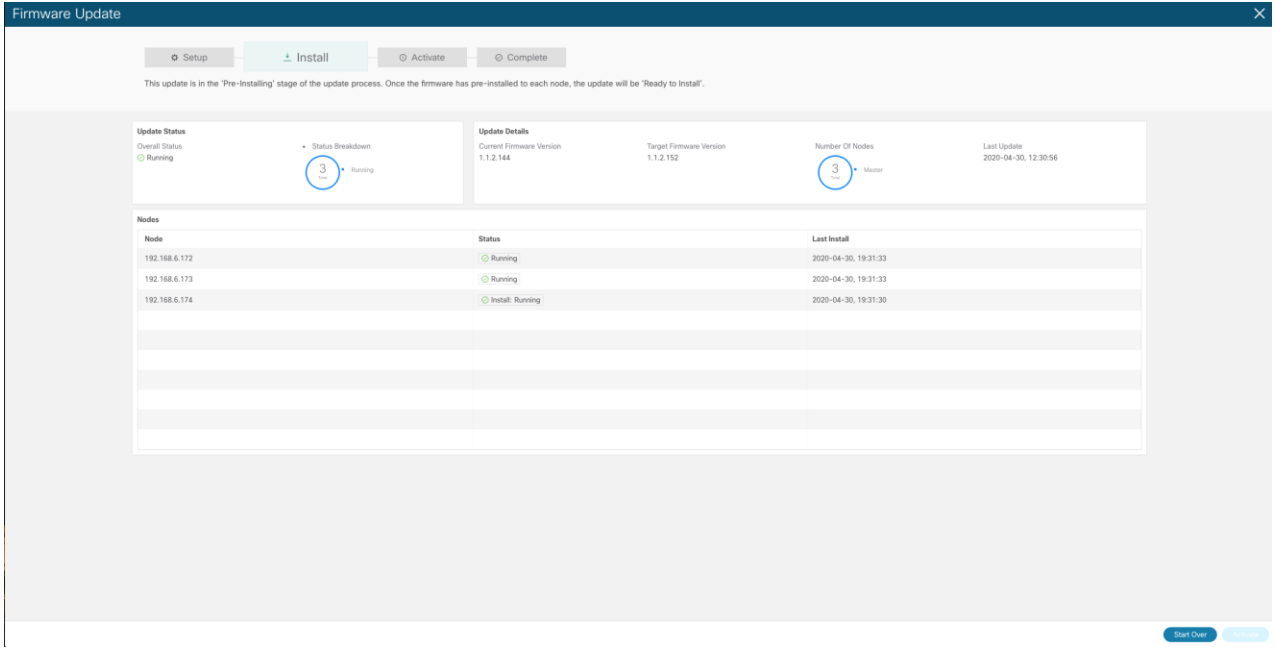

Once Install is done Click Activate

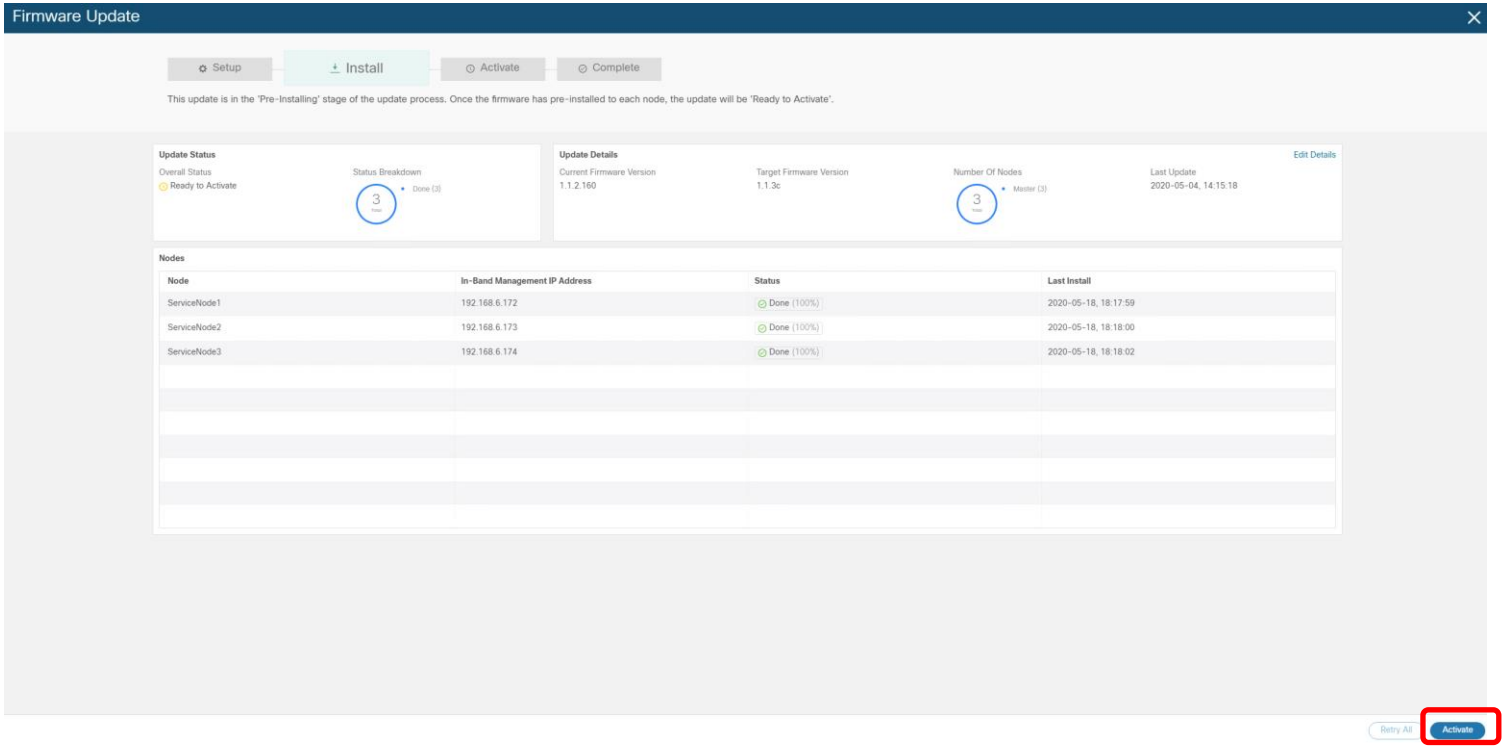

Activation Progress

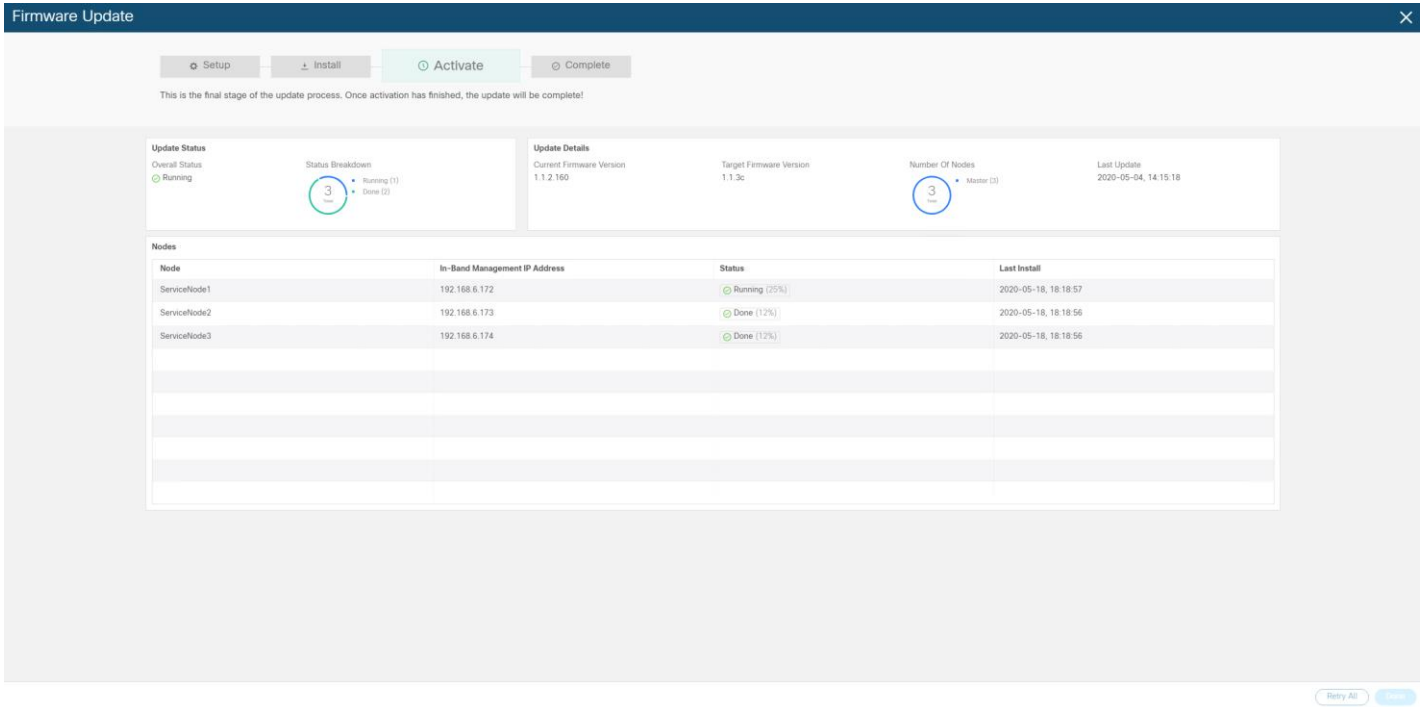

Monitoring Firmware Upgrade

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

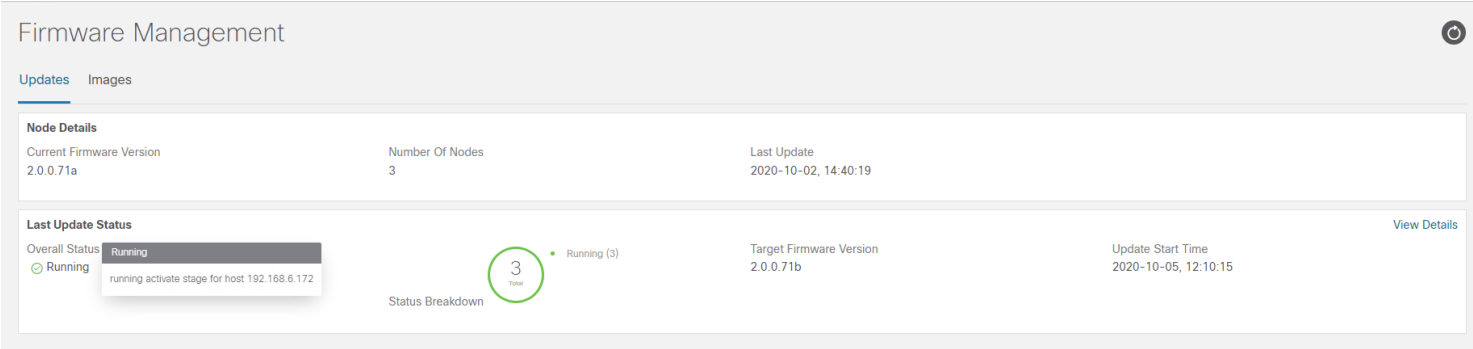

- Node going through an update will display:

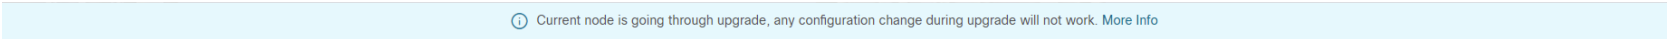

Remote Authentication

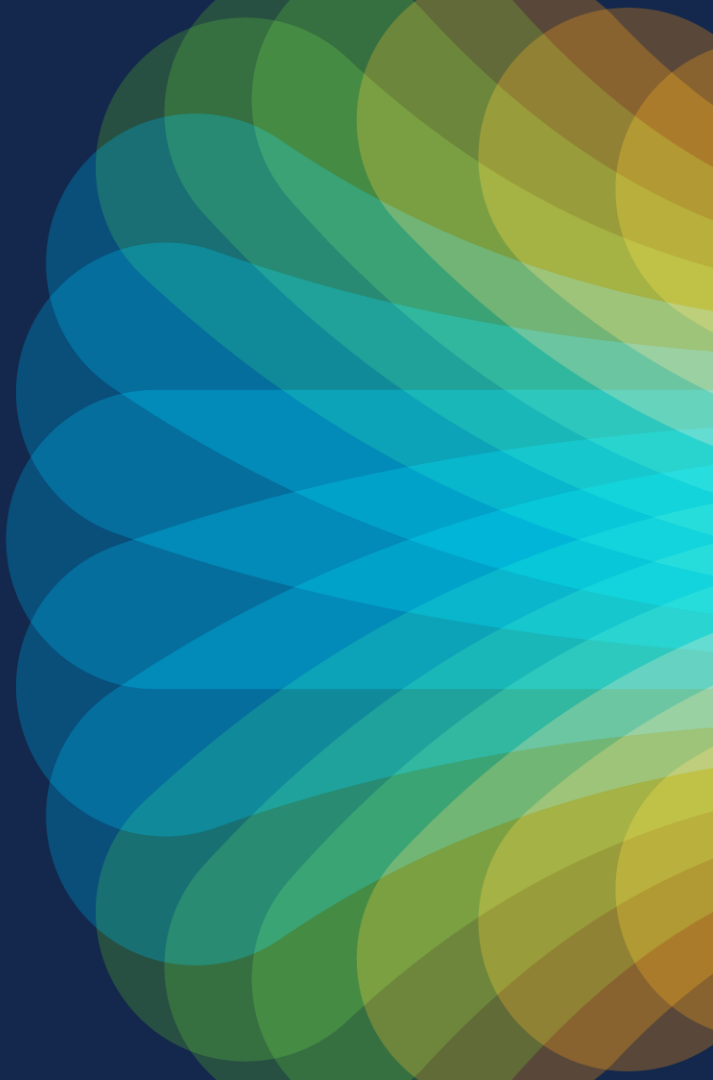

Remote Authentication

- 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.

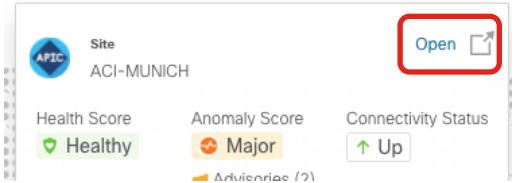

Login without and with enabled Login Domain

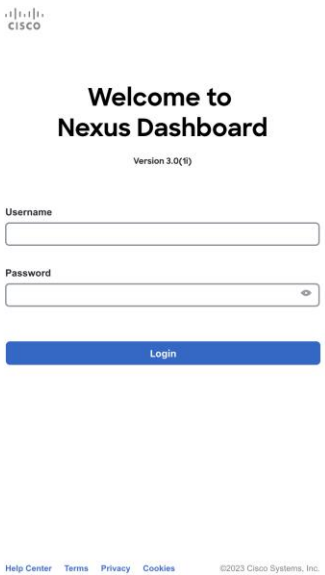

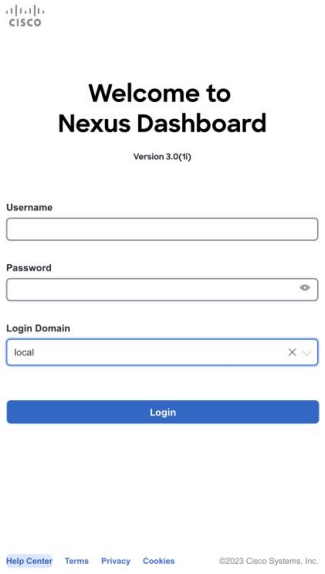

Create a Login Domain

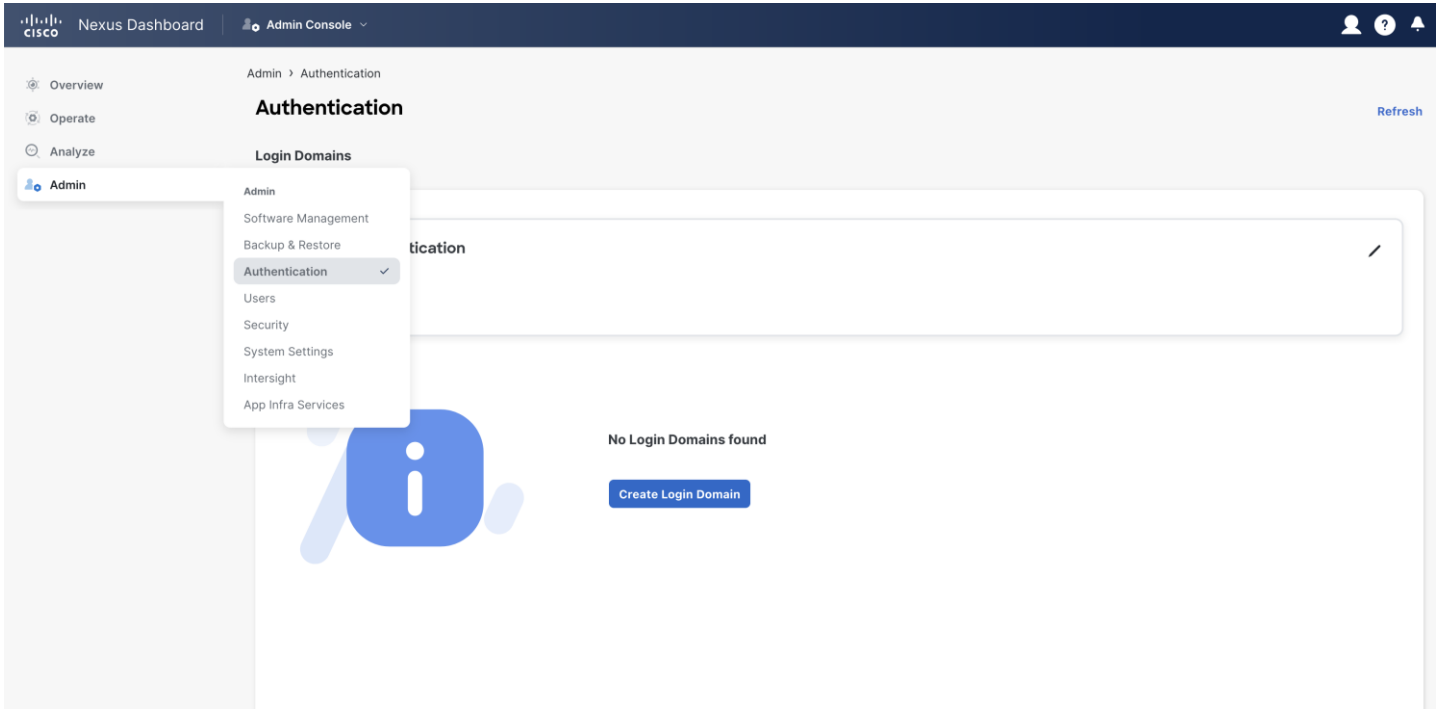

Create a Login Domain

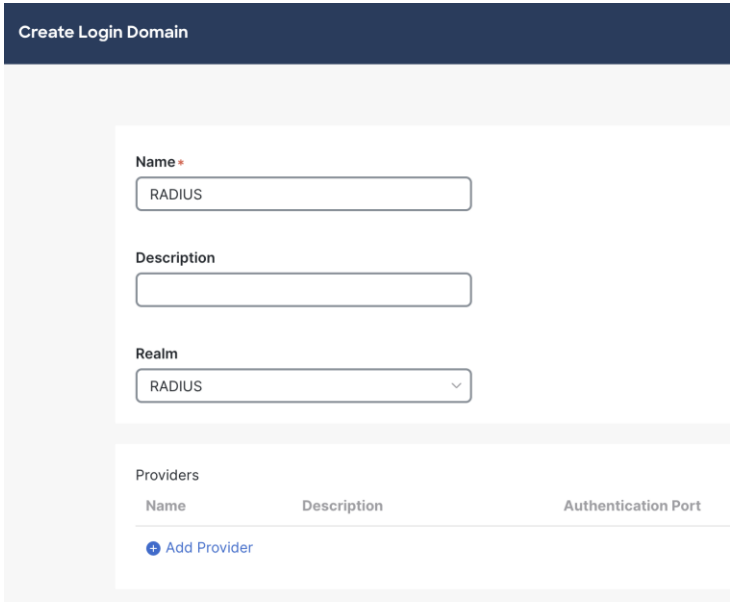

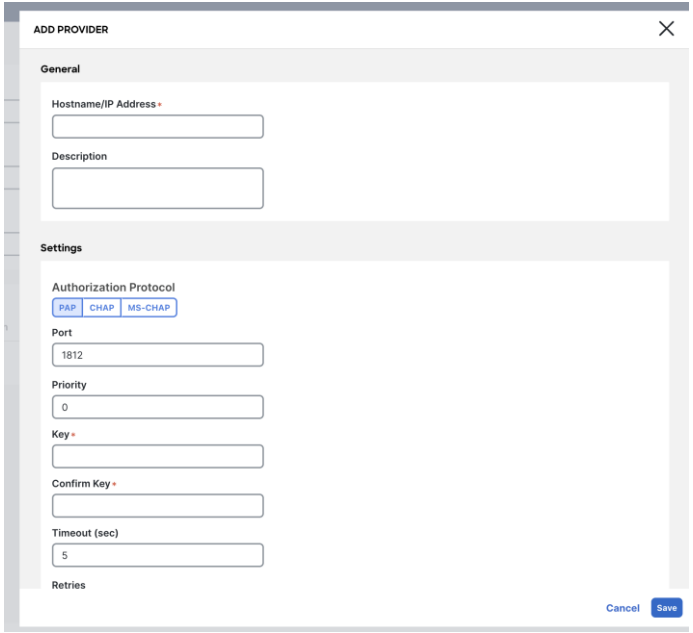

Need to have a valid remote user to add provider – backend will query the remote auth server with provider info and user/pass before it can be added.

Change Default Authentication for Login

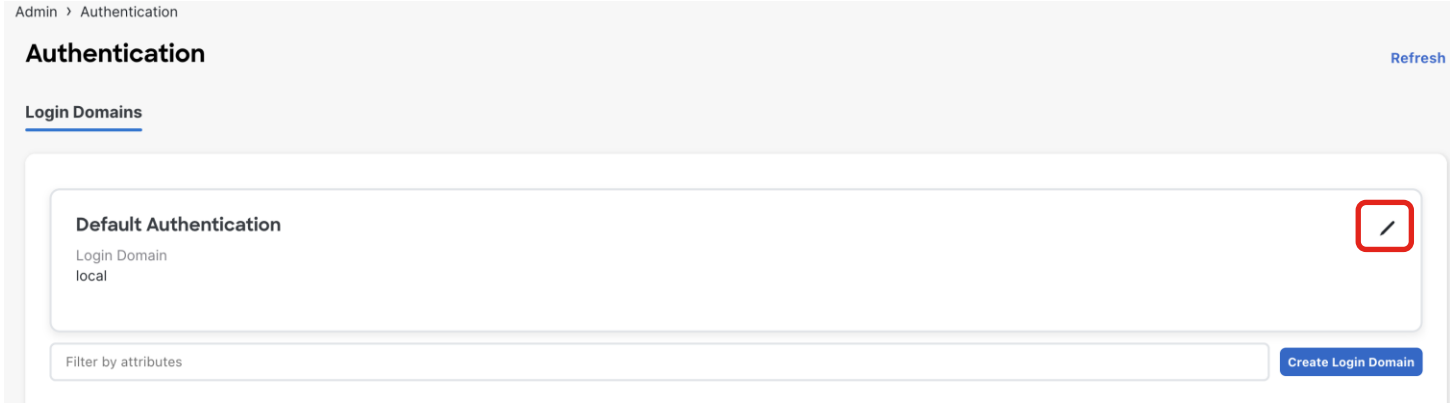

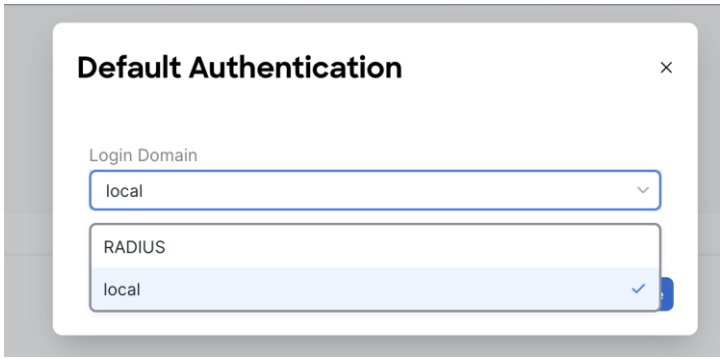

Login Screen with Login Domain

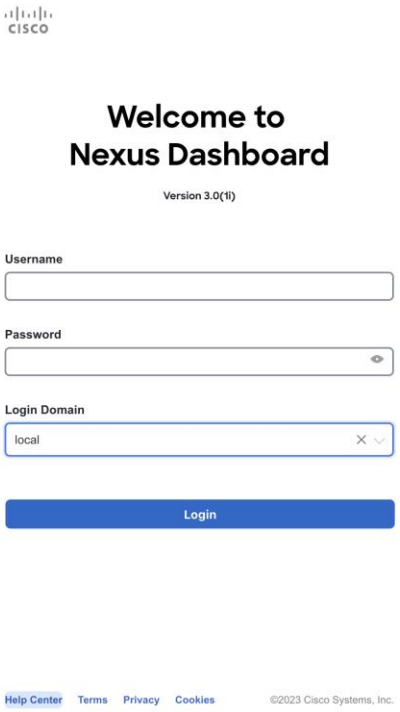

RBAC and User Roles 1/2

- 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

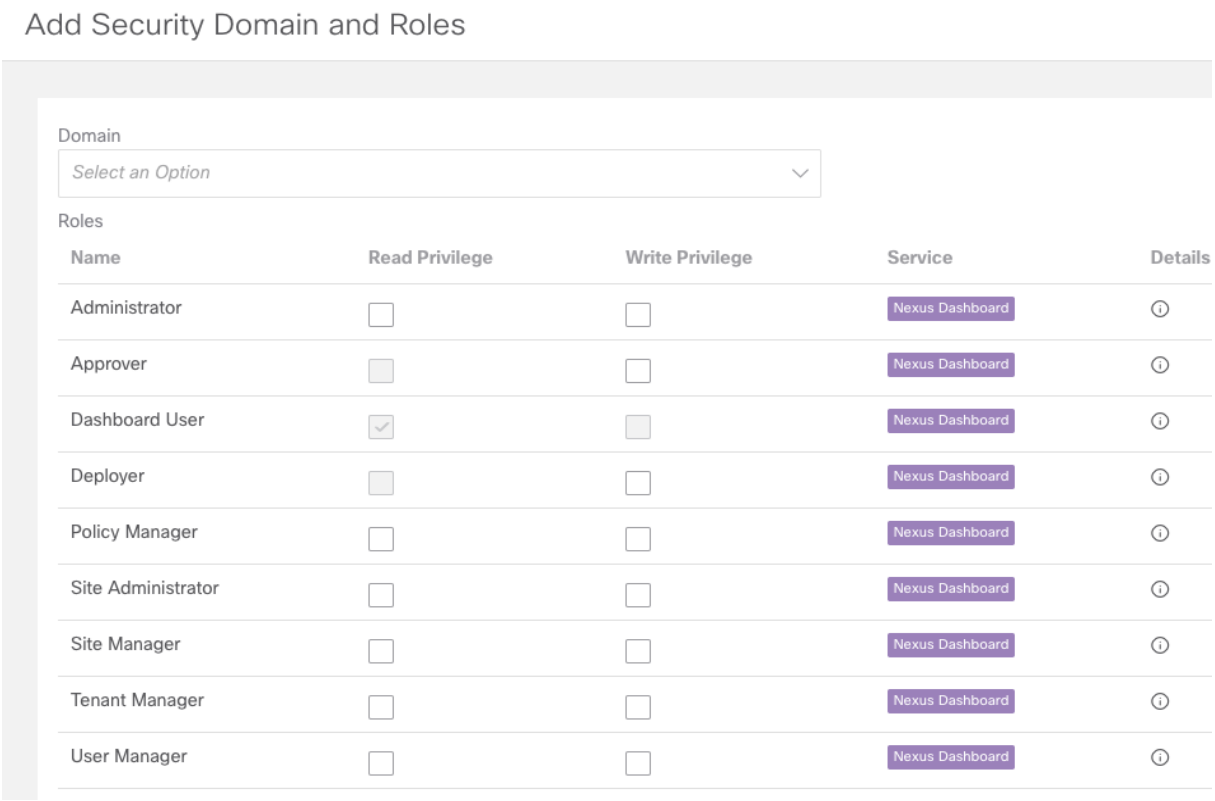

Configurable Security Settings

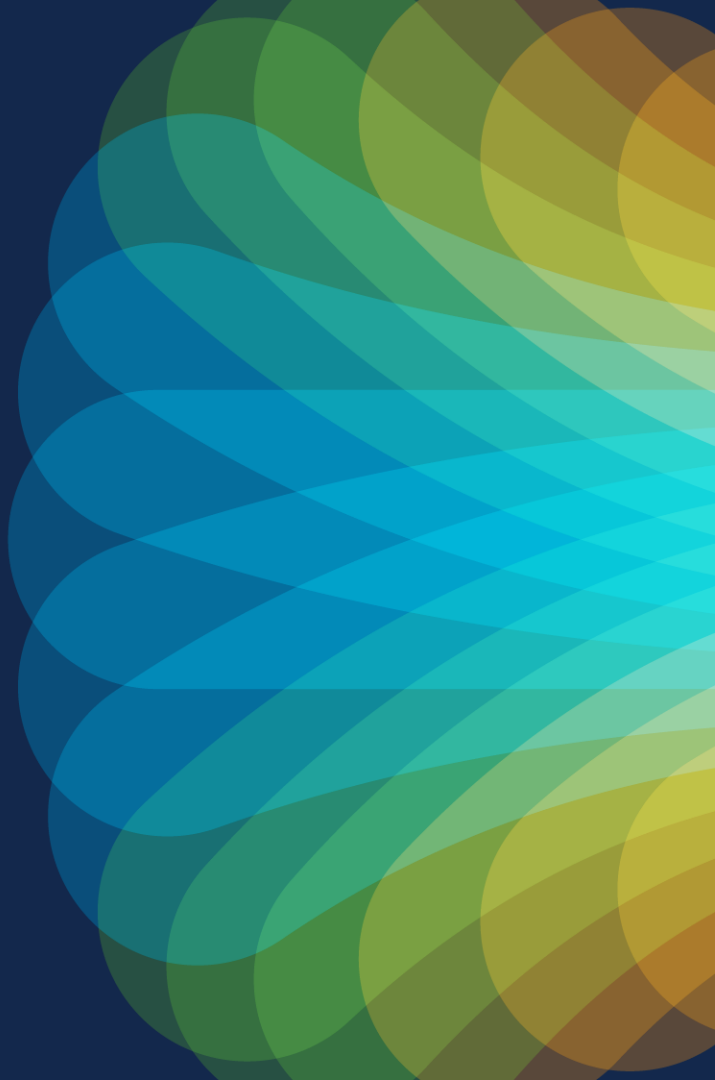

Configurable Security Settings

- 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

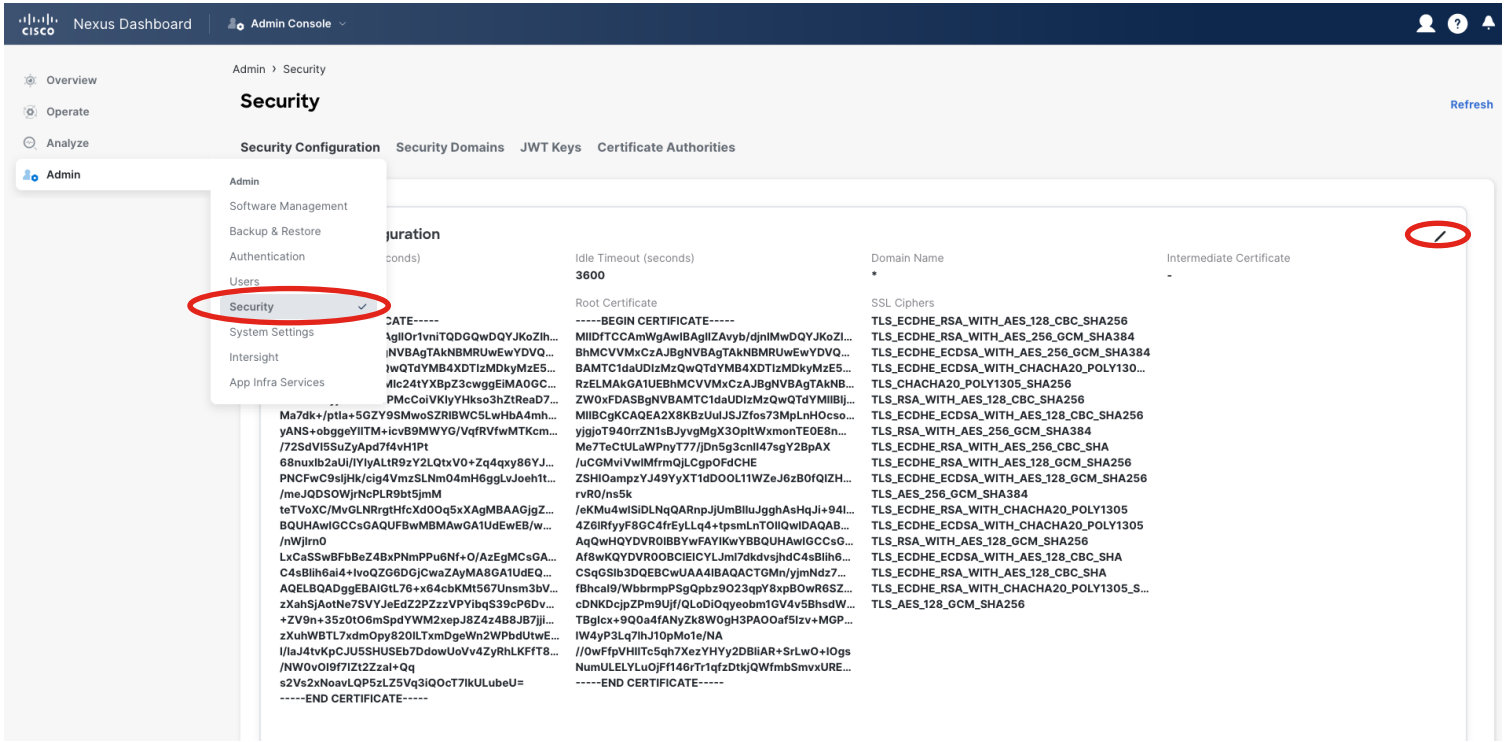

Configure Security Settings

Session and Idle Timeout in Seconds

Customer Certificate and Root Certificate, enabled SSL Chiphers etc.

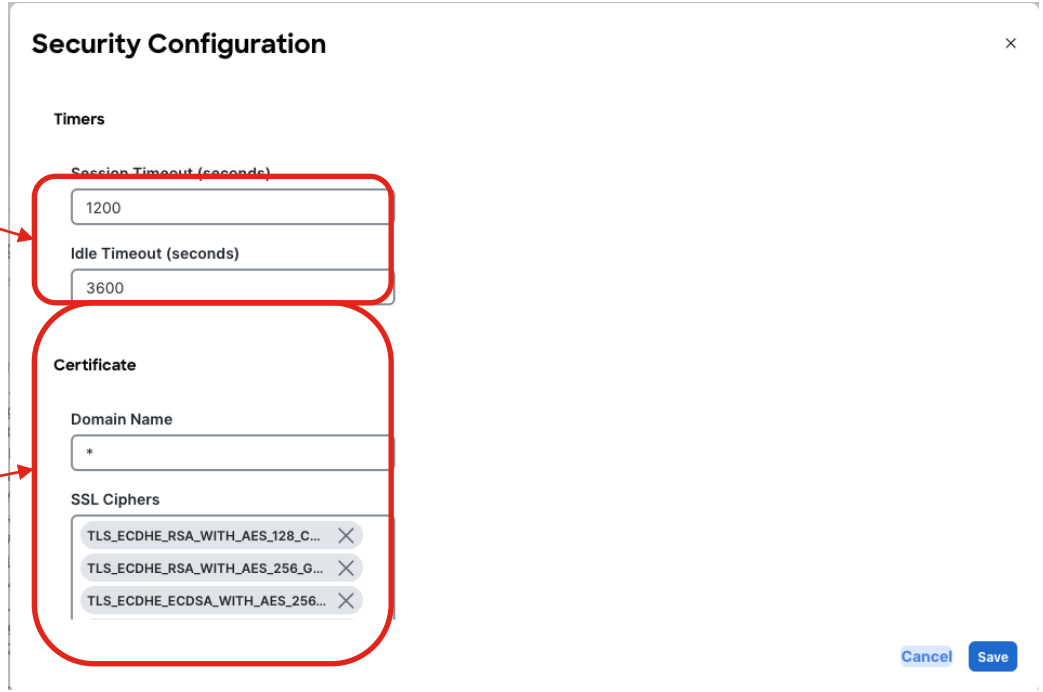

```
[rescue-user@ND2 ~]$ openssl req -new -x509 -keyout cert.pem -out cert.pem -days 28 -nodes
Generating a RSA private key
..............
......
writing new private key to 'cert.pem'
-----
You are about to be asked to enter information that will be incorporated
into your certificate request.
What you 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

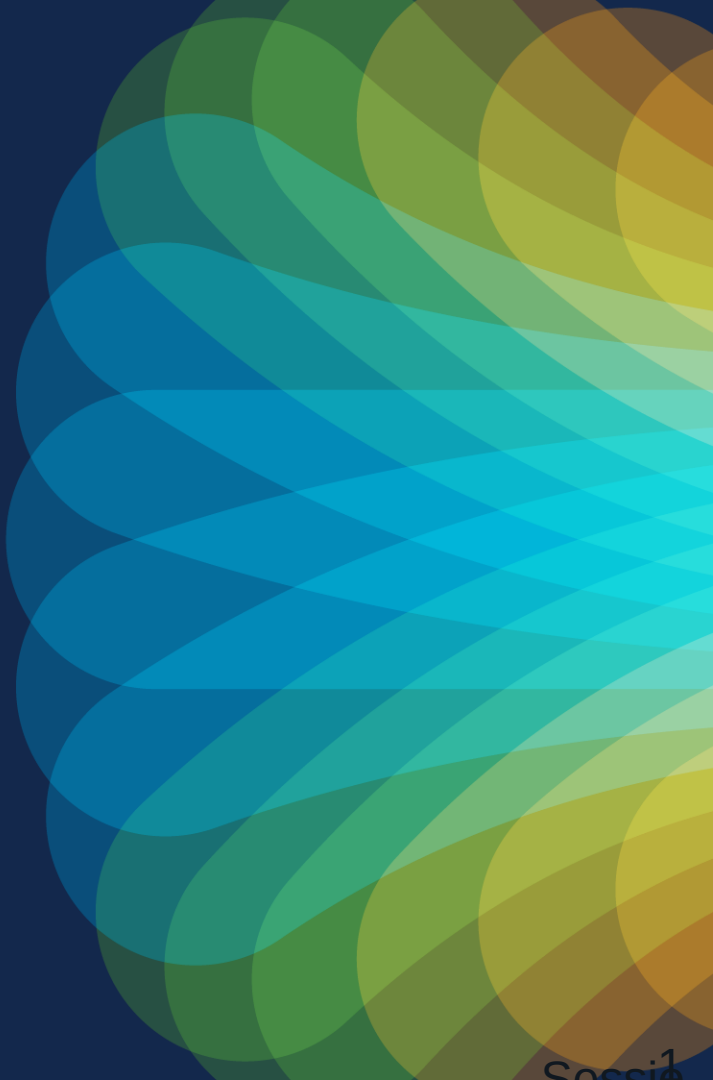

Resource Monitoring

- 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

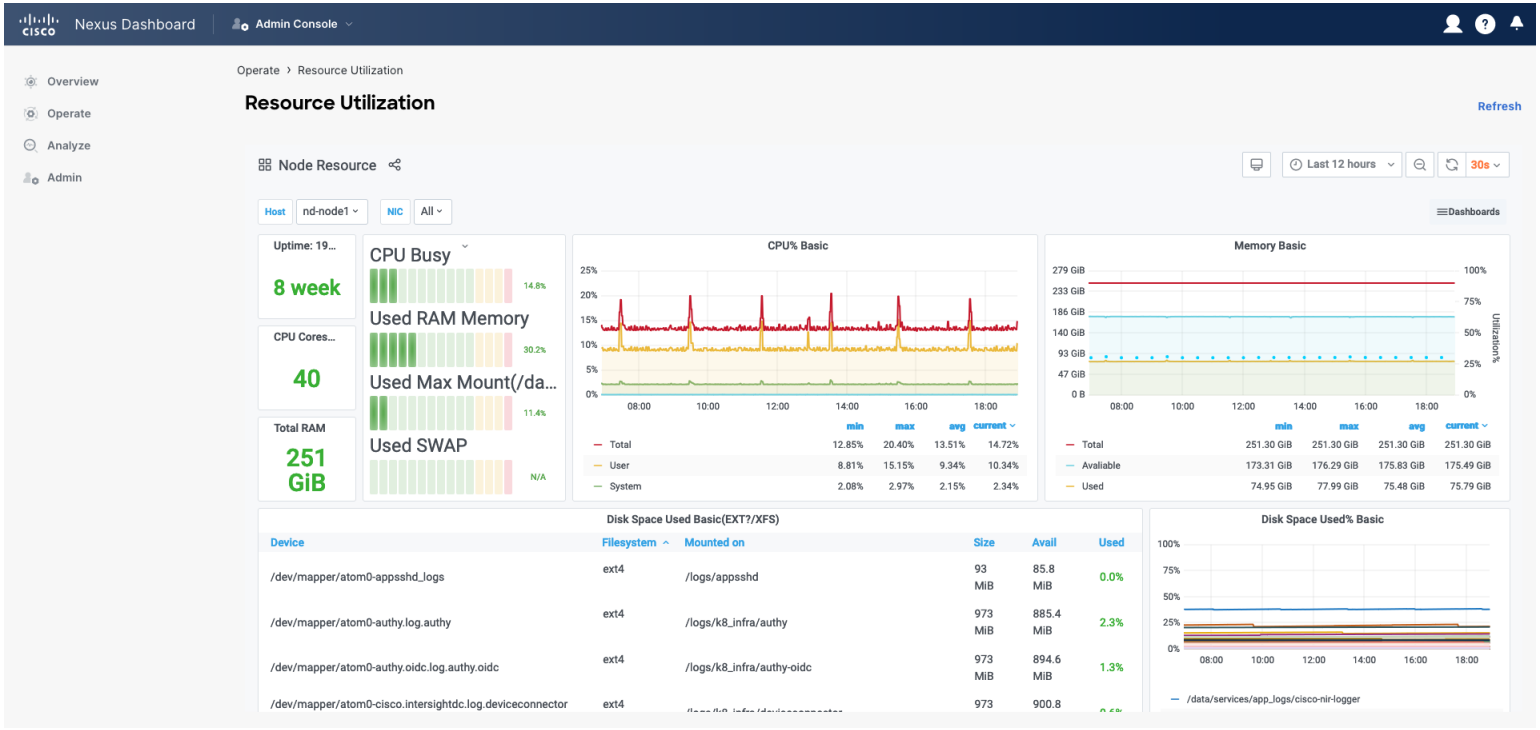

Event Analytic

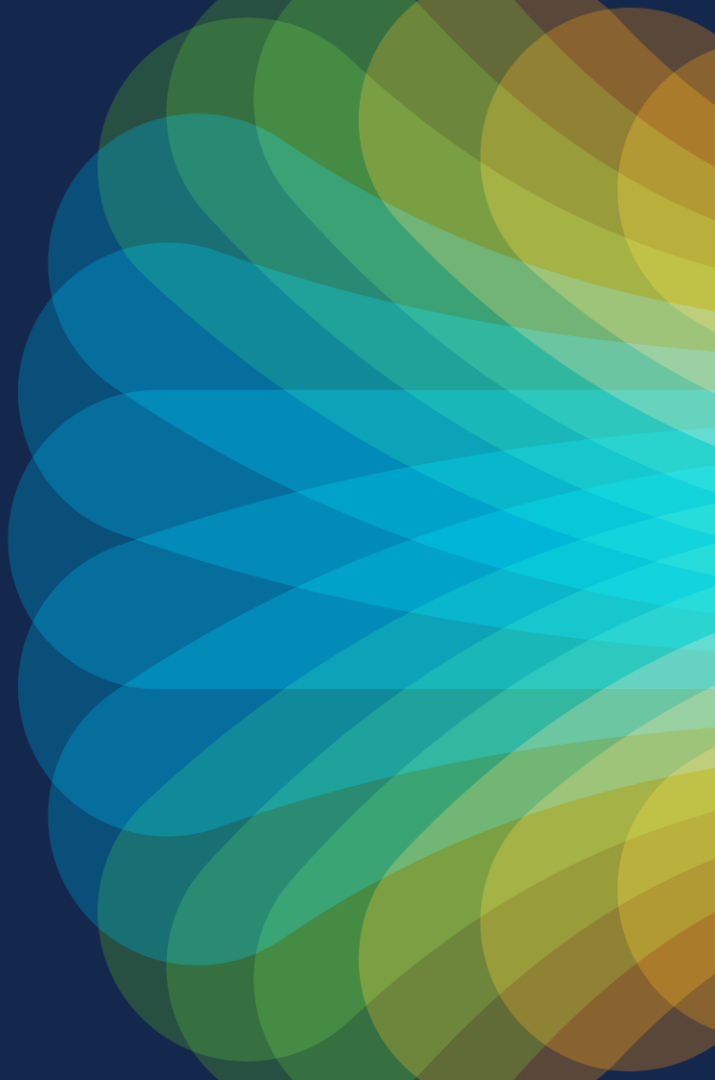

Event Analytic

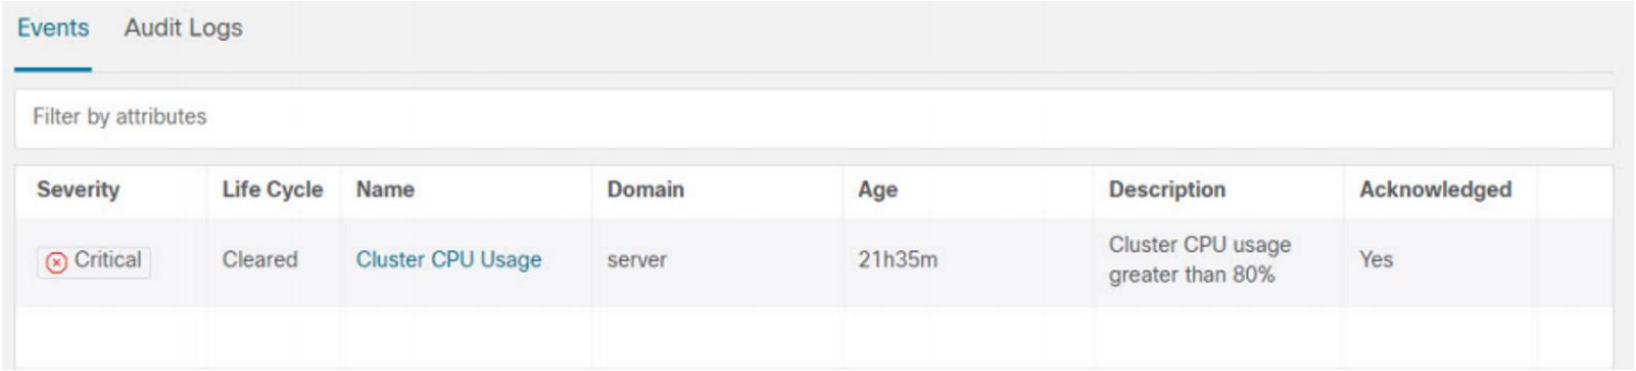

Event Analytics enables easy access your Nexus Dashboard's events and audit logs. In addition to viewing the events and logs directly in the Nexus Dashboard GUI, you can also configure the cluster to stream the events to an external syslog server (TCP/UDP)

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

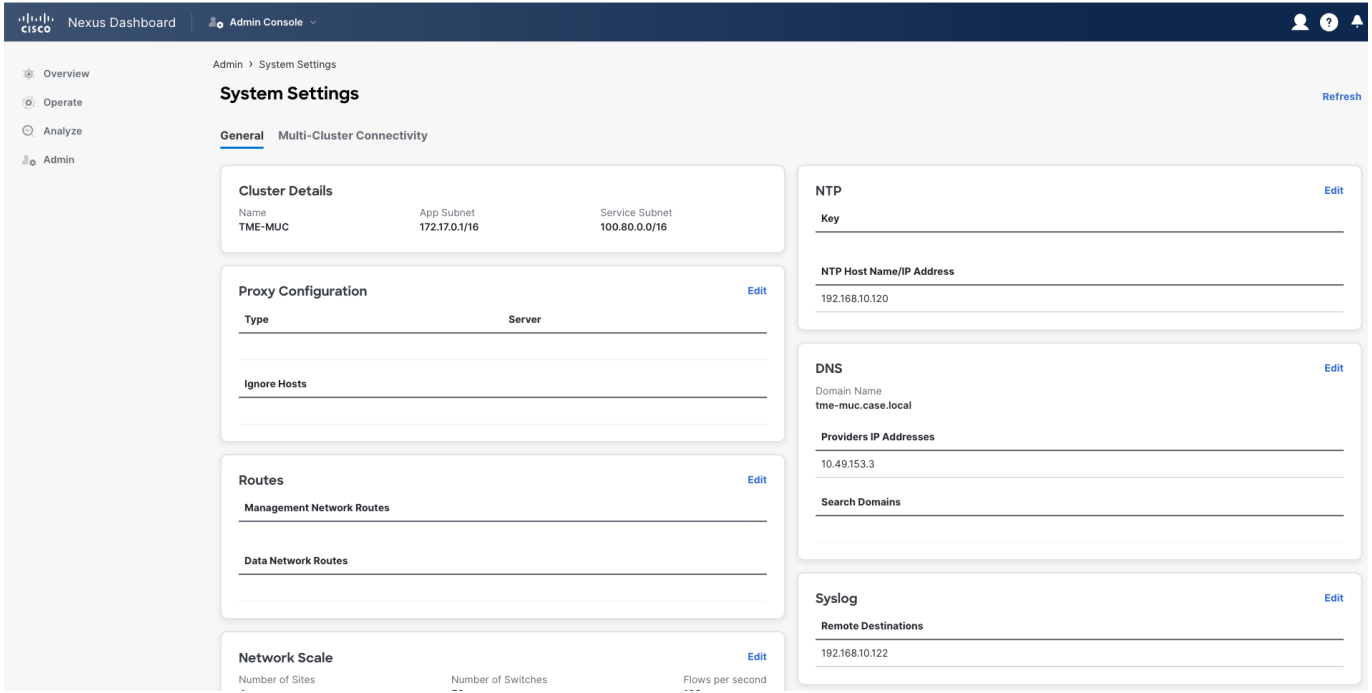

Configuring Syslog Servers 2/2

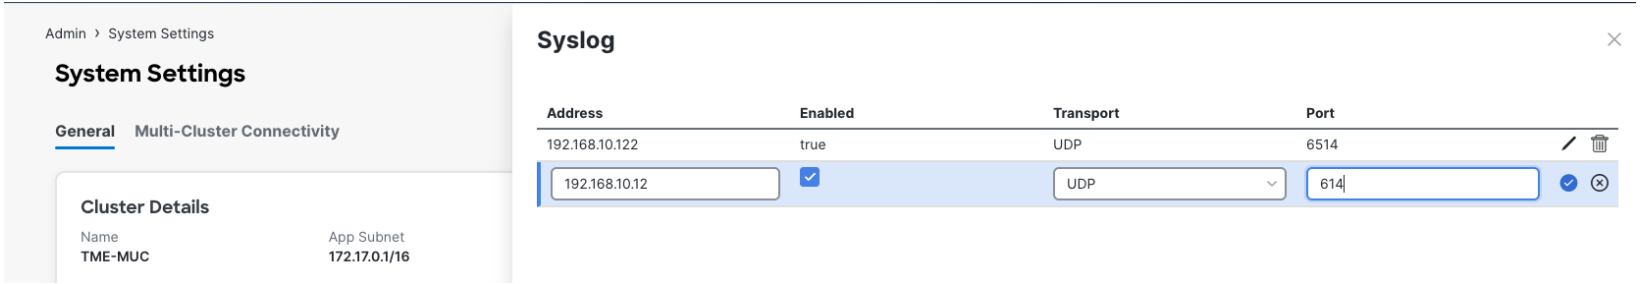

Hardware Monitoring of ND via CIMC

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

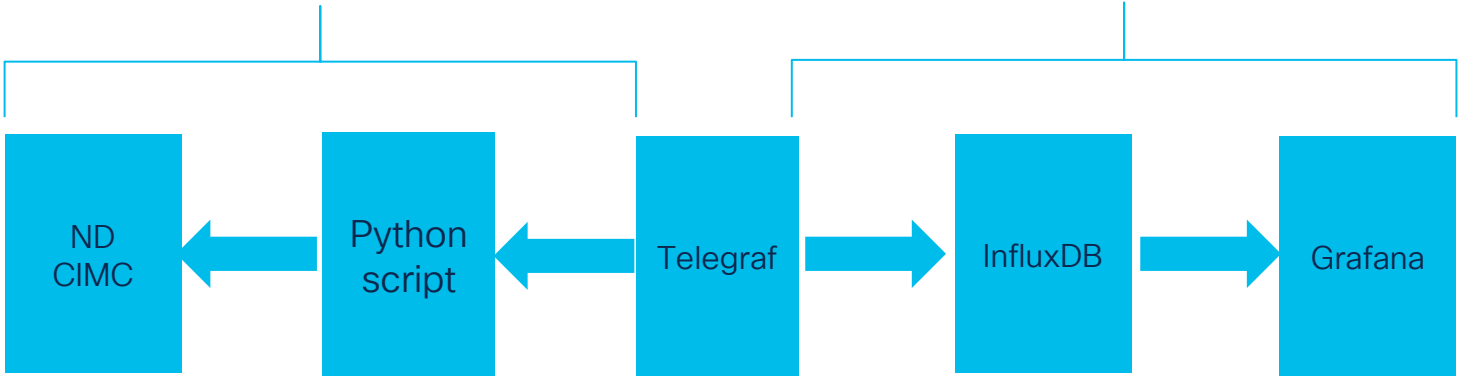

Basic Troubelshooting

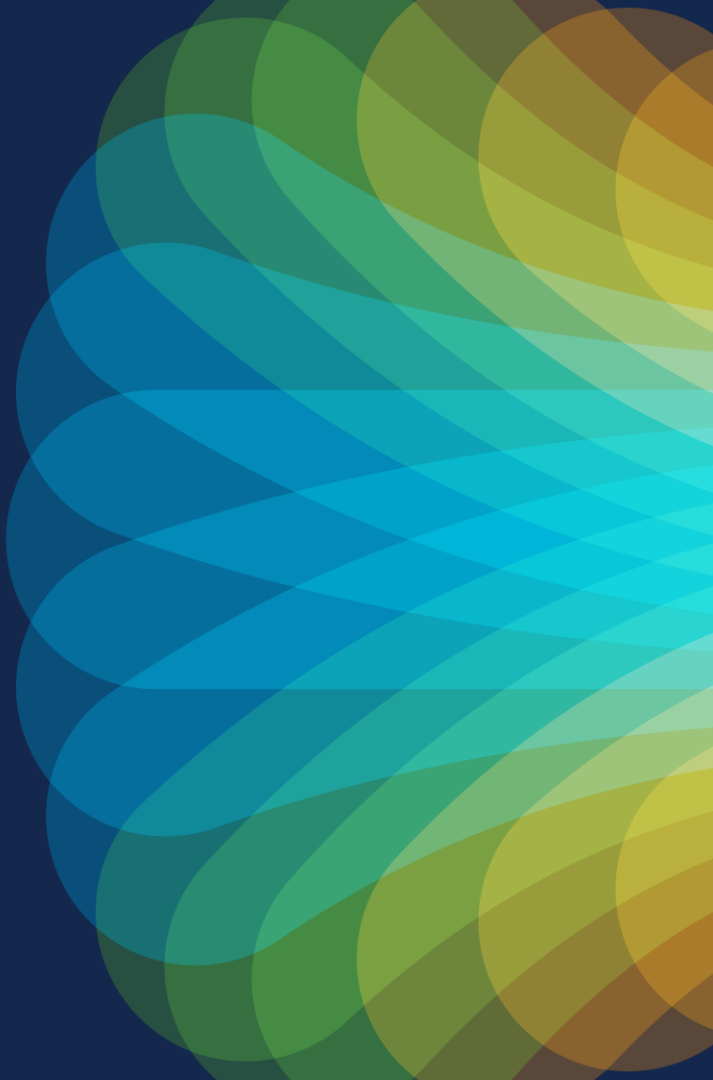

Basic Troubleshooting

- 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

- Better visibility with real time analysis
- Meaningful, actionable anomalies
- Root Cause is a few clicks away
- Assurance for your configuration intent


cisco
The bridge to possible
Thank you
cisco Live!


cisco Live!
Let's go