Cisco TechClub webinář

Nastává čas ACI 6.0

Pepa Venzhöfer
Technical Solutions Specialist – CCIE DC#59794
20.6.2023

Agenda

- Životní cyklus ACI SW release
- Škálovatelnost a nový HW
- Konfigurace interface
- Možnosti nasazení APIC
- Automatizace ACI
- Symetrický PBR

Application Centric Infrastructure building blocks

Built on Cisco Nexus 9000

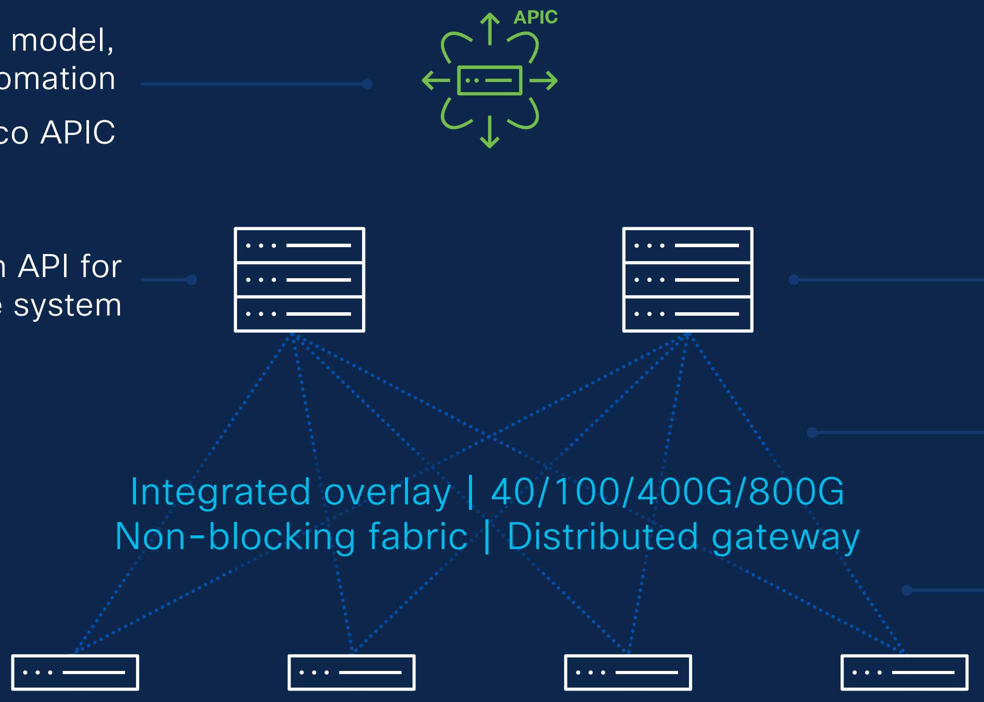

Centralized policy model, network automation

Cisco APIC

Single open API for entire system

Flexible – modular and fixed spine options

Non-blocking 40/100/400G fabric, CLOS fabric

Built-in distributed stateless firewall, multi-tenant security

Integrated overlay | 40/100/400G/800G
Non-blocking fabric | Distributed gateway

Physical, virtual and container workloads

WAN interconnect

IP storage

Network service appliances

Industry leading

- Price
- Performance
- Port density
- Programmability
- Power efficiency

Primer: ACI key market differentiators

Automation out-of-the-box; physical fabric and underlay

Application-aware service-chaining

Virtual Machine Manager (VMM)

True multi-tenancy with administrative Tenants

Hybrid cloud capability; public cloud-like networking constructs

Single API for 100s of switches and 1000s of workload connections

BU-supported Terraform capability

ACI is the on-prem anchor fabric in a hybrid cloud deployment model

Cisco ACI 6.0 themes

Operational simplicity

Network security

Network automation

Multicloud

Open APIs
Large partner ecosystem

Performance 400G to 800G evolution

Scaling

Flexible deployment models

Infrastructure as code

Container networking

Controller virtualization

Životní cyklus ACI SW release

ACI Software release cadence

Key objectives Predictable software release cadence | Reach maintenance mode quickly

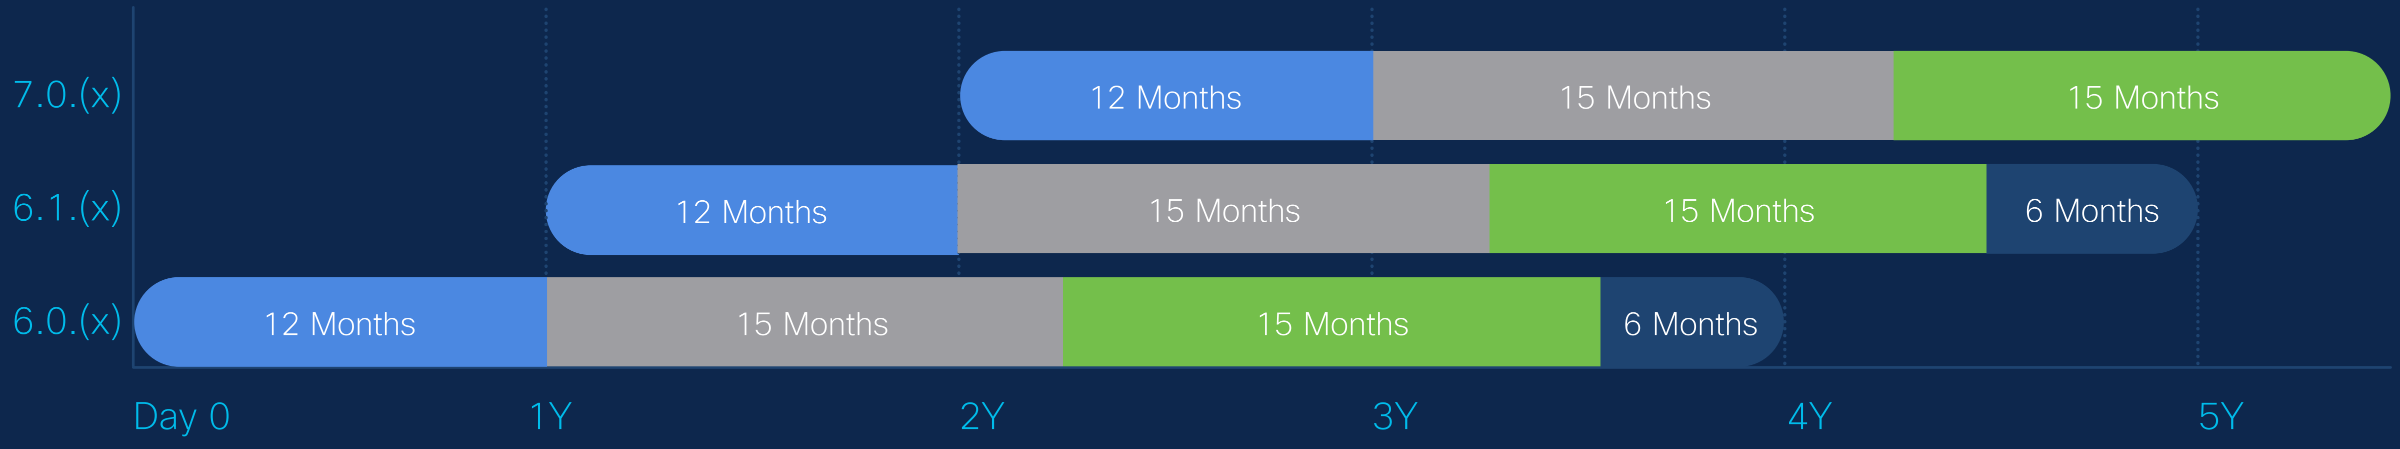

Legend: Development cycle | Maintenance cycle | Extended support with PSIRT fixes | TAC support

No short-lived and long-lived release tags

Two feature releases from FCS date, including FCS release

Hardware lifecycle is defined by multiple releases and not tied to a single release

Total release lifecycle of four years

ACI Release Timeline

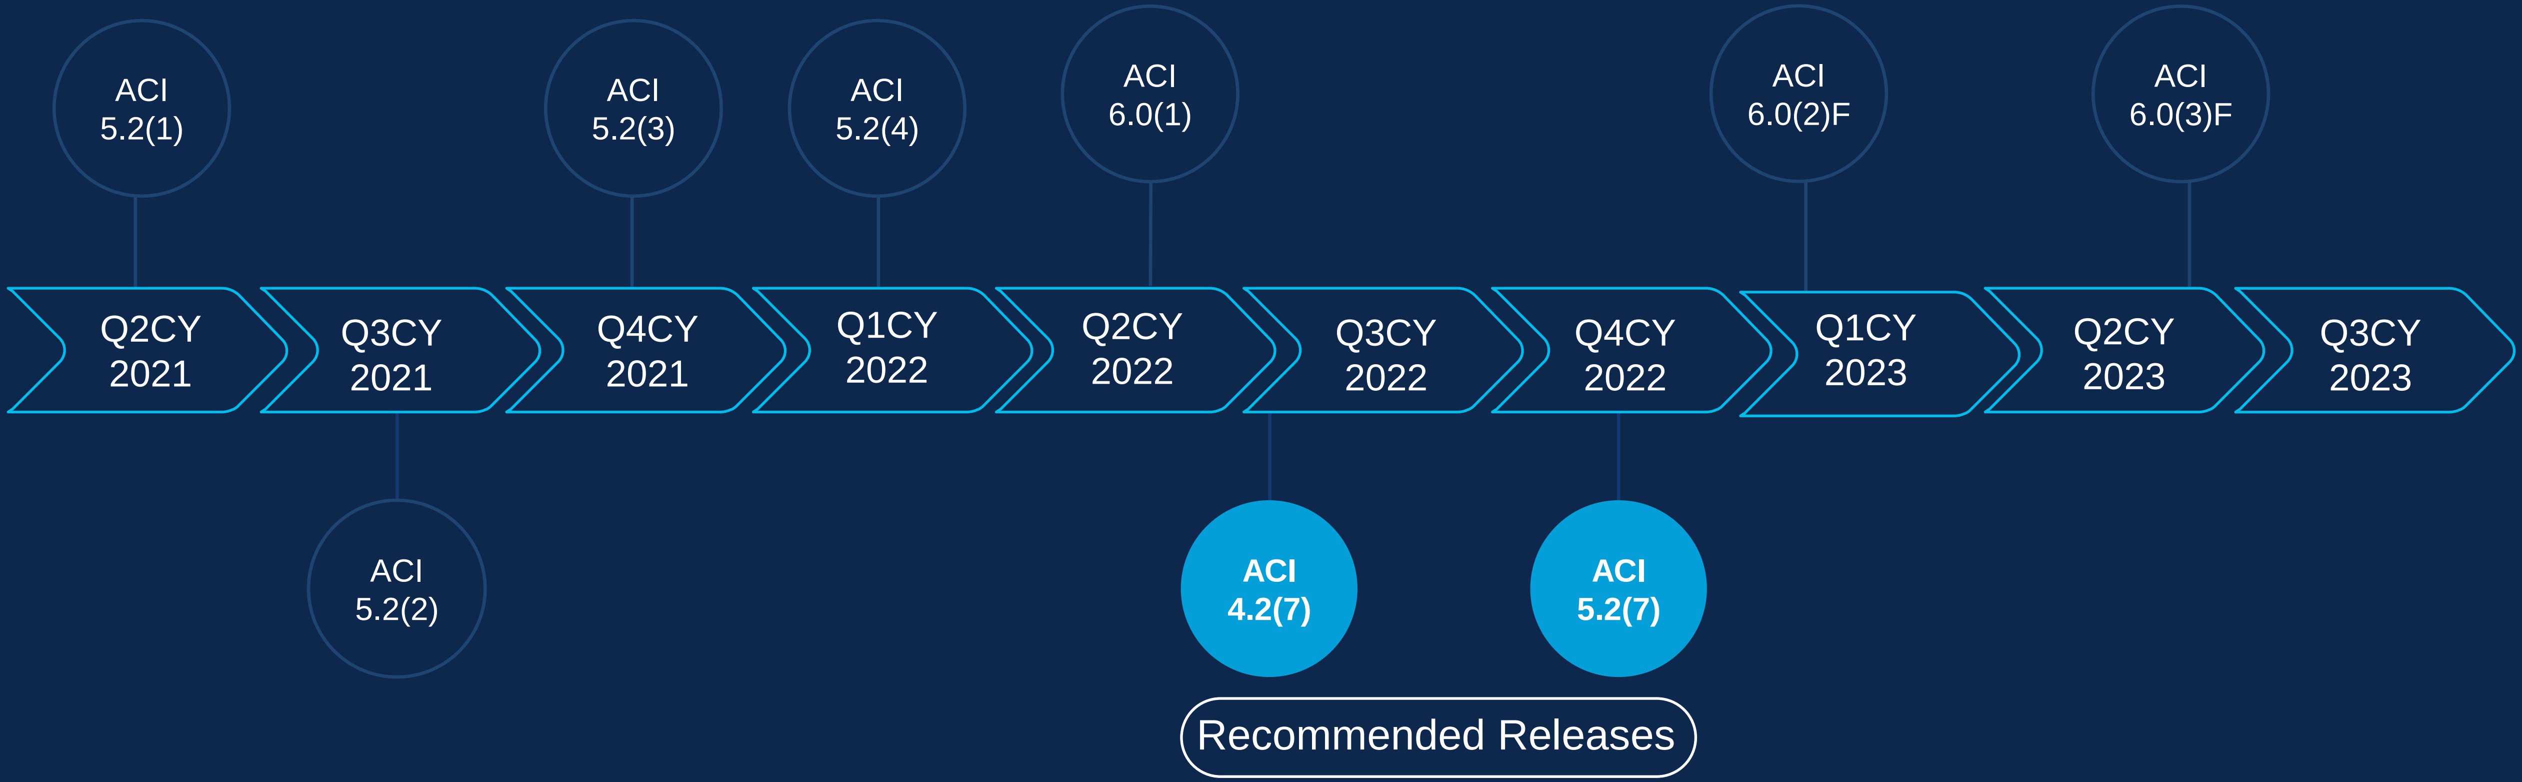

Note: ‘F’ indicates feature release

- No short-lived and long-lived release tags
- Three feature releases from FCS date, including FCS release 6.0 (1), 6.0(2)F, 6.0(3)F
- Fourth release is a maintenance release (MR), target for golden star 6.0(4)M
- Hardware lifecycle is defined by multiple release and not tied to a single release
- Total release lifecycle of four years

Škálovatelnost a nový HW

Increased scalability

Increasing number of pods in fabric from 12 to 25

Increasing from 3K to 10k VRF per fabric

Increased number of routed ports per leaf from 16 to 48
Increasing from 1K to 2k sub-interfaces with 2k BGP/OSPF/Static/2k BFD

Increasing from 2K to12k MCP (Port x VLAN) per leaf
Increasing from 256 to 2k VLANs for MCP per interface

Expand ACI fabrics to 400G and beyond

400G Leaf or spine

Nexus 9408

(N9K-C9408 ACI)

8 slot modular

Pay as you grow up to 64 port at 400G plus breakout

Leaf switch

N9K-C93108TC-FX3H (24P 10g Copper)

N9K-C93180YC-FX3H a (24P 10g/25g SFP+)

Flexible deployments with 400G optics

QDD-400G-SR4-BD/4X100G-SRBD1.2

Spine to leaf at 4x100 breakout

Leaf to server at 4x100 breakout

QDD-400G-ZRP-S

Connect fabrics at greater distances

Konfigurace interface

Access policies prior to 6.0(1)

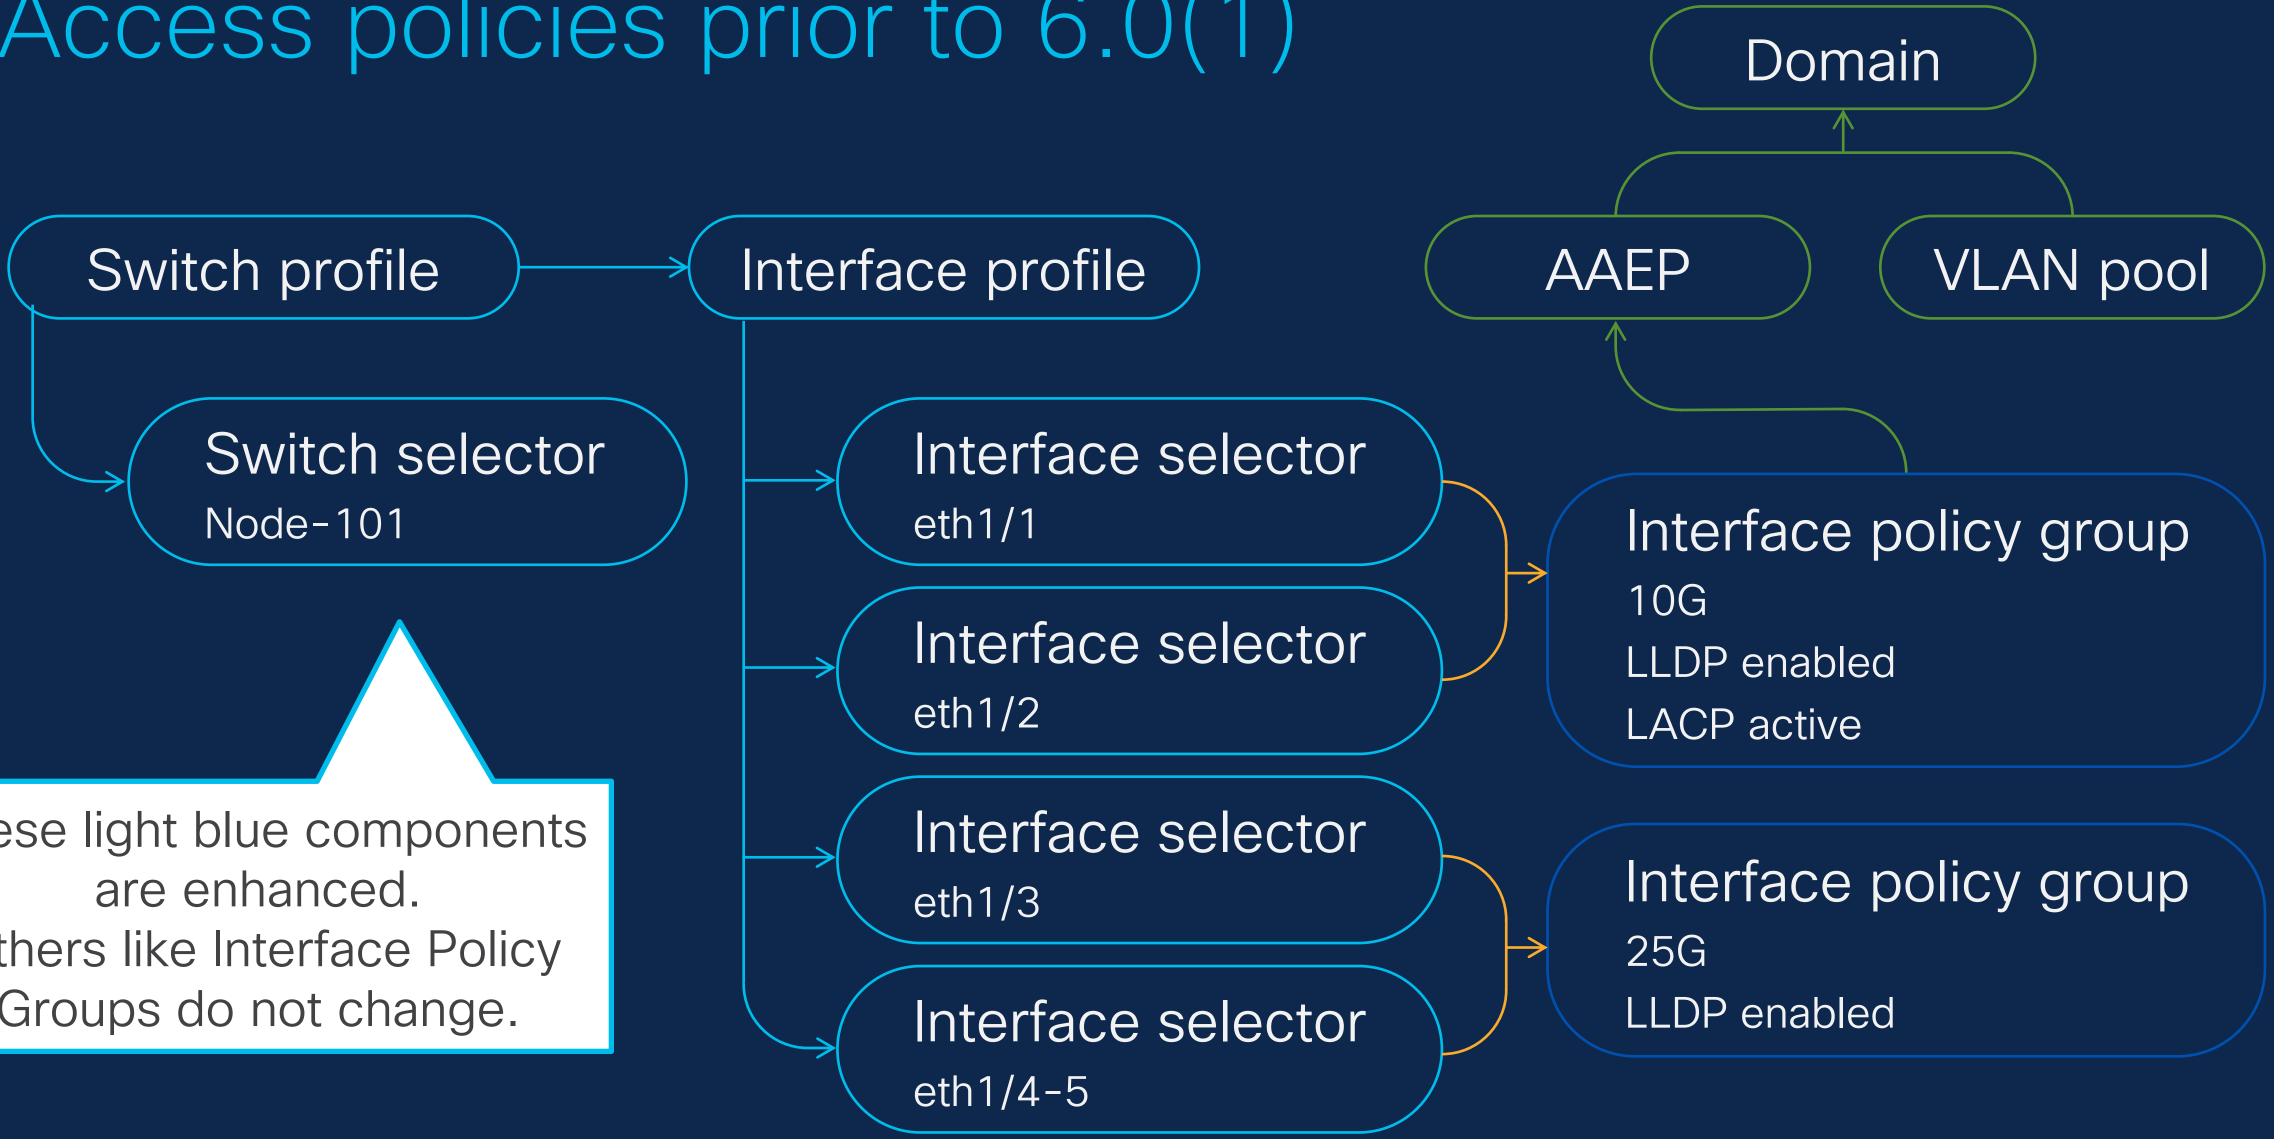

These light blue components are enhanced.
Others like Interface Policy Groups do not change.

The new option for access and fabric policies

4 step configuration

1 step configuration

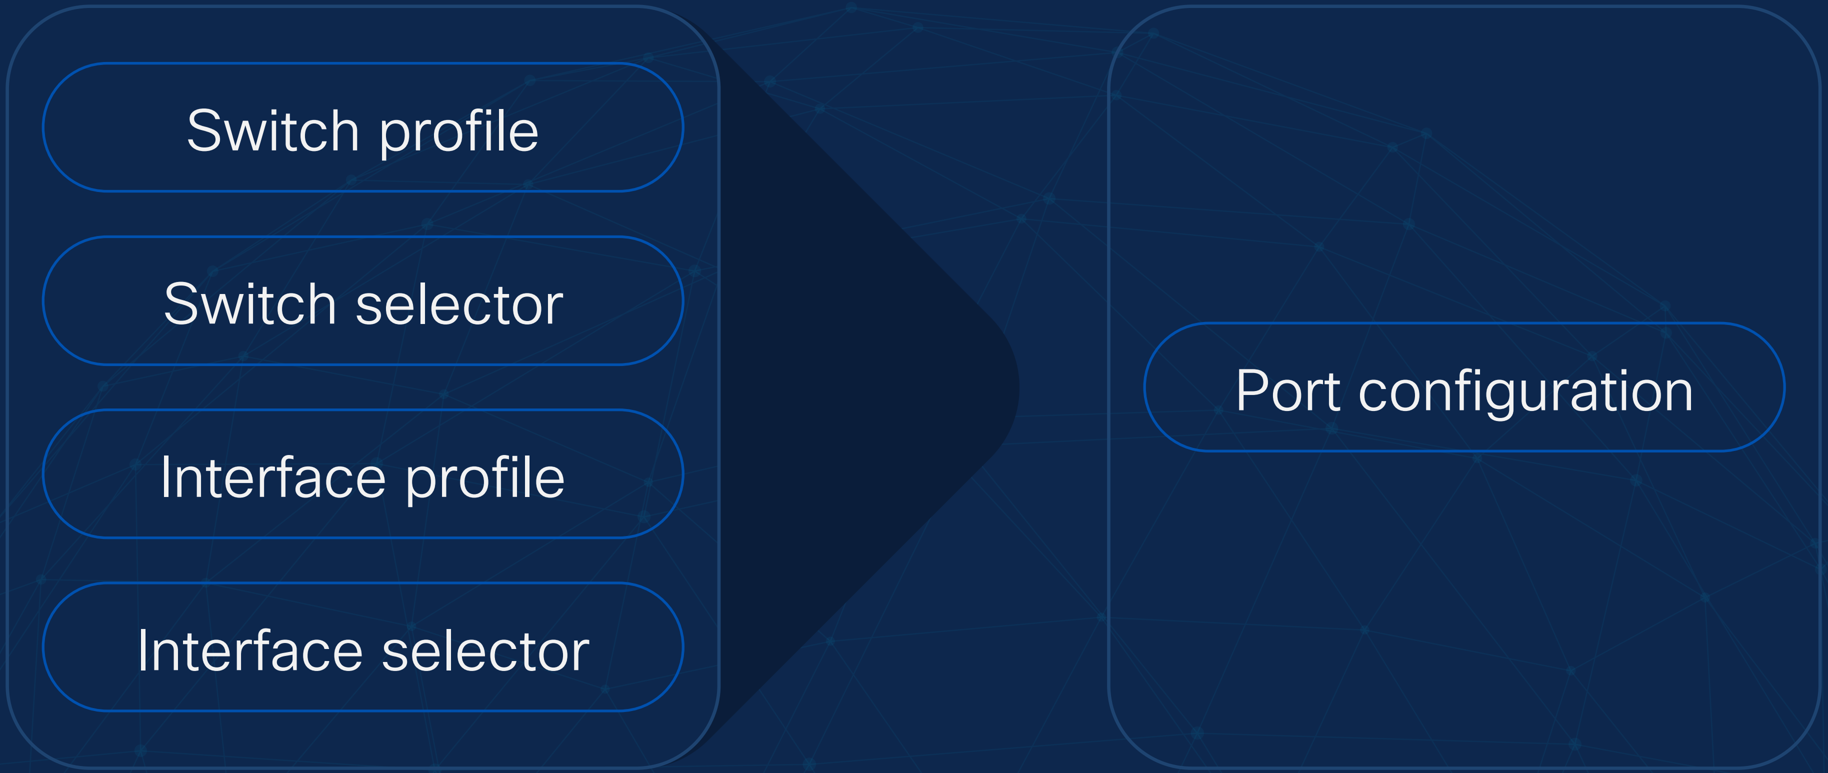

Switch profile

Switch selector

Interface profile

Interface selector

Port configuration

Access policies from 6.0(1)

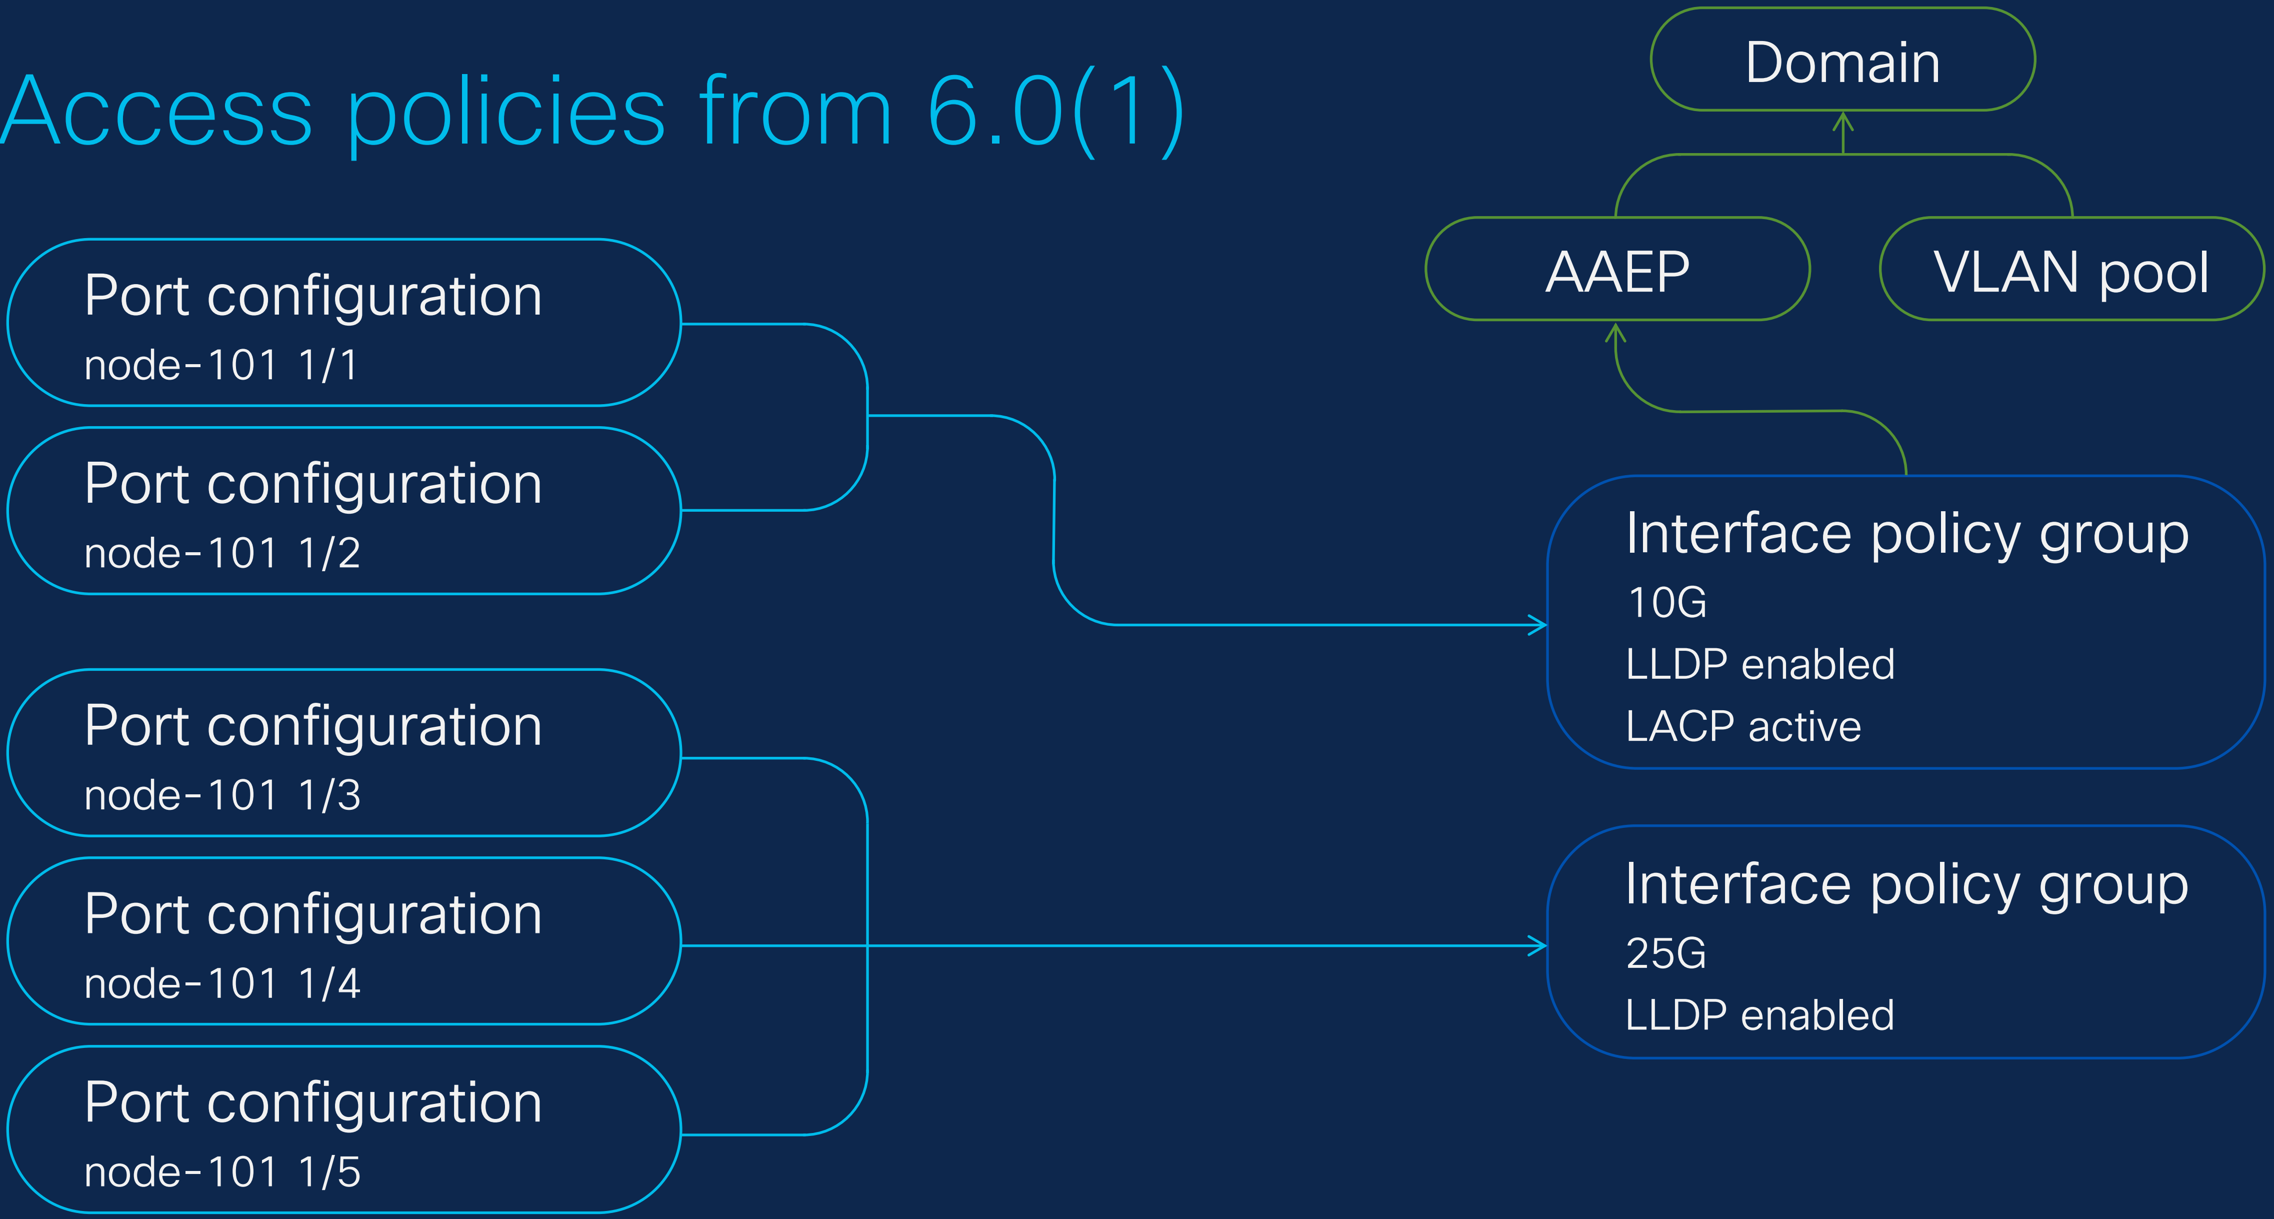

Port configuration
node-101 1/1

Port configuration
node-101 1/2

Port configuration
node-101 1/3

Port configuration
node-101 1/4

Port configuration
node-101 1/5

Domain

AAEP

VLAN pool

Interface policy group
10G
LLDP enabled
LACP active

Interface policy group
25G
LLDP enabled

Access policies from 6.0(1)

Under the hood

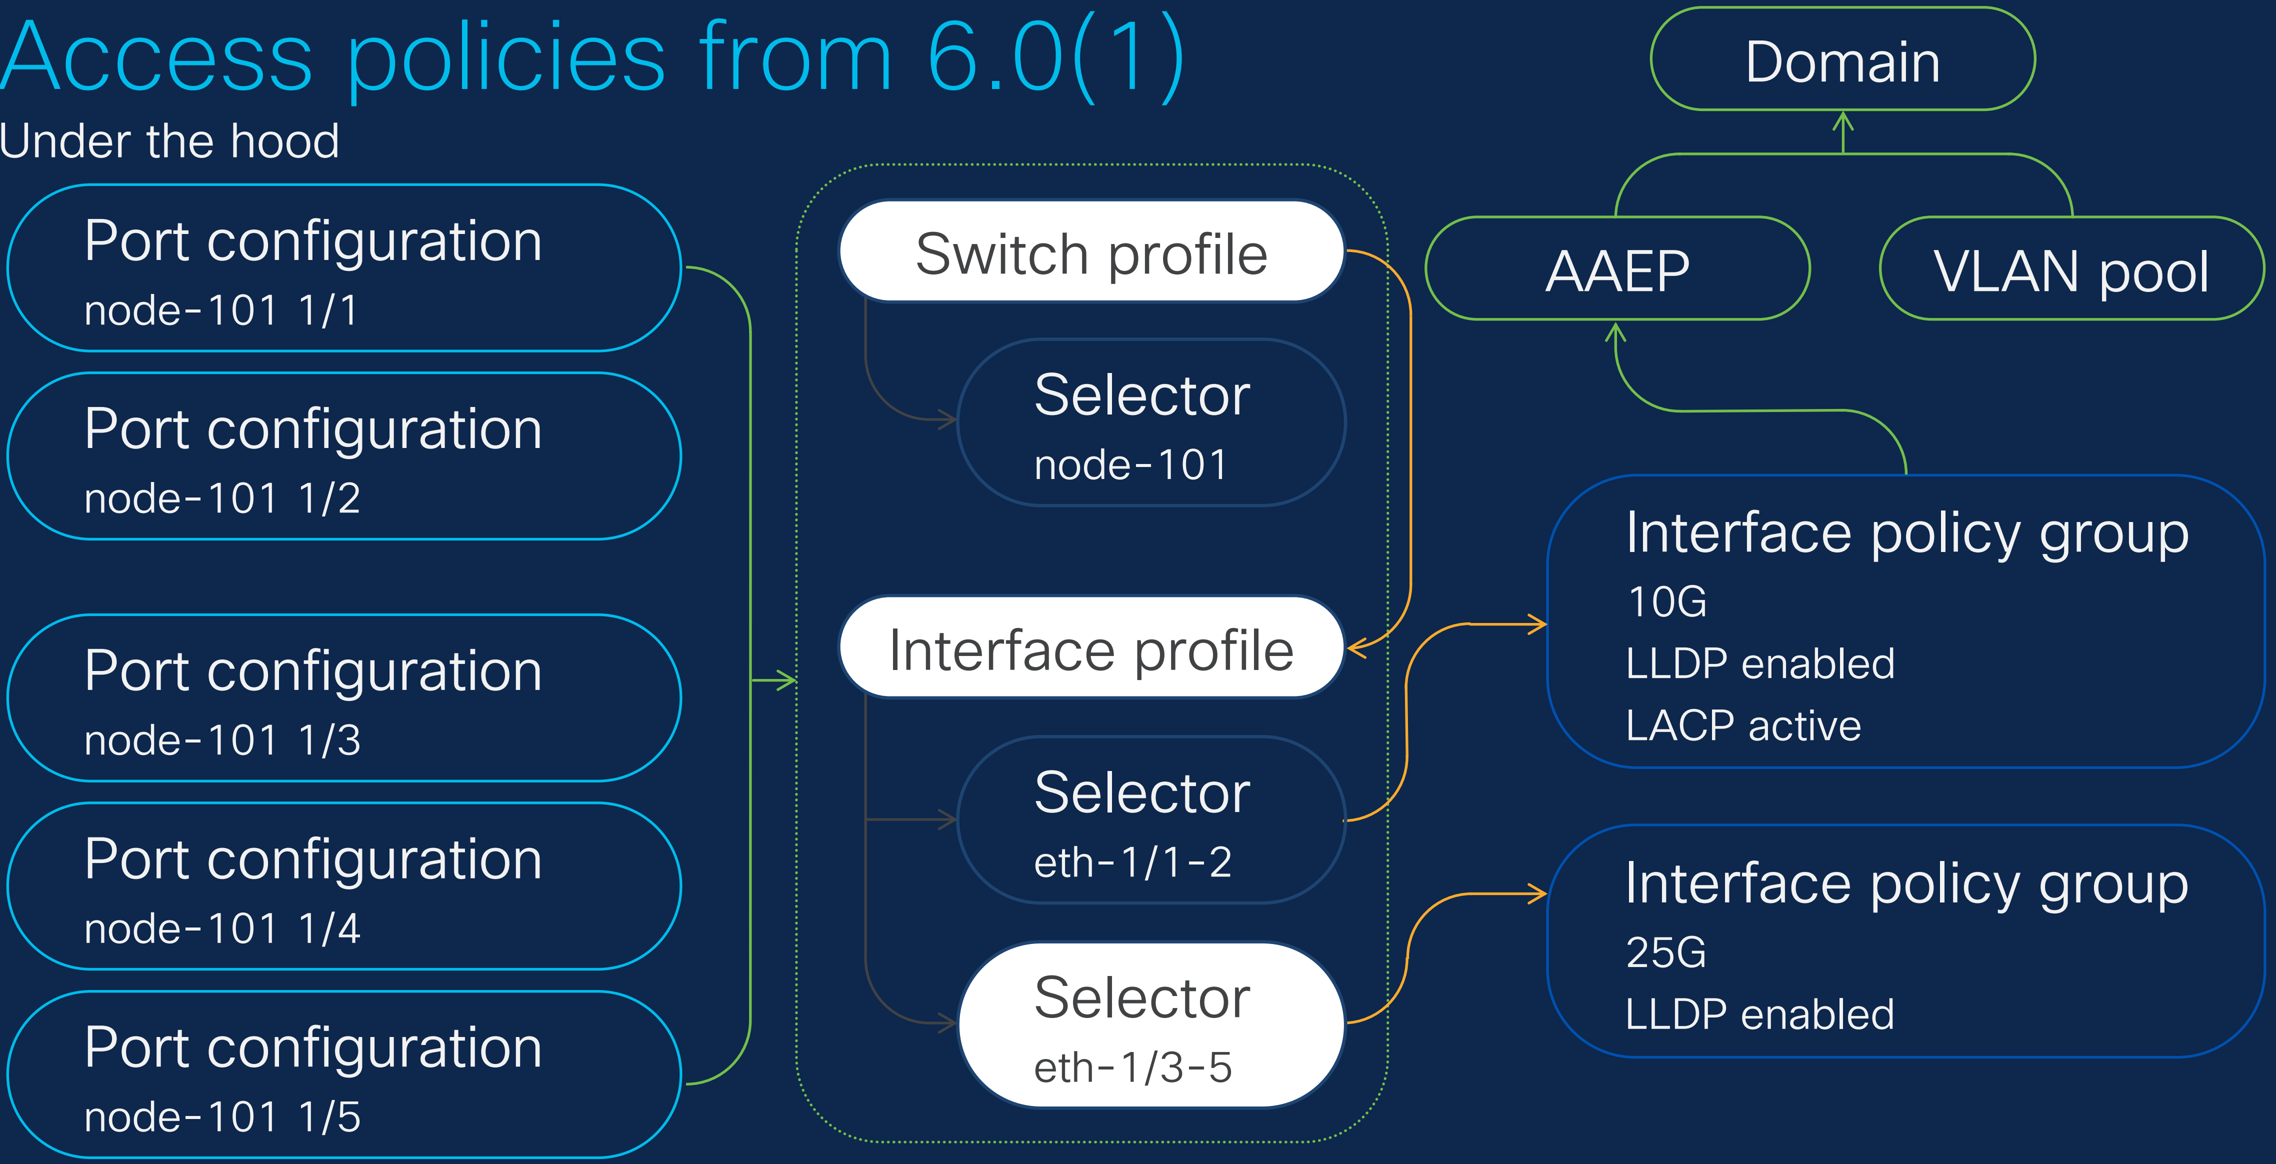

Port configuration
node-101 1/1

Port configuration
node-101 1/2

Port configuration
node-101 1/3

Port configuration
node-101 1/4

Port configuration
node-101 1/5

Switch profile

Selector
node-101

Interface profile

Selector
eth-1/1-2

Selector
eth-1/3-5

Domain

AAEP

VLAN pool

Interface policy group
10G
LLDP enabled
LACP active

Interface policy group
25G
LLDP enabled

APIC translates port configuration to read-only profiles that are optimized and match the best practices

Summary and Considerations

- The new interface configuration uses a new object infraPortConfig.
- In each infraPortConfig, users simply specify one node ID and one interface ID along with the interface policy group.
    - ➢ No need to learn profiles and selectors.
- infraPortConfig supports all configurations* that used to be done by the profiles/selectors directly.
    - ➢ No need to switch back and forth between the new and old way.
    - ➢ As users use the new way in the GUI, the configuration seamlessly adopt the new way even if the interfaces are currently configured by the old way
        - ➢ backward compatible
- Under the hood, infraPortConfig creates system-generated profiles/selectors that are read-only.
    - ➢ No need for users to manage, check or modify them directly.

* Access/Fabric policies, Eth, FC, PC, vPC, FEX, breakouts and up/downlink conversion

Možnosti nasazení APIC

Flexible controller deployments

Virtual APIC on ESXi in ACI 6.0.2F

Remote APIC cluster

Shipping

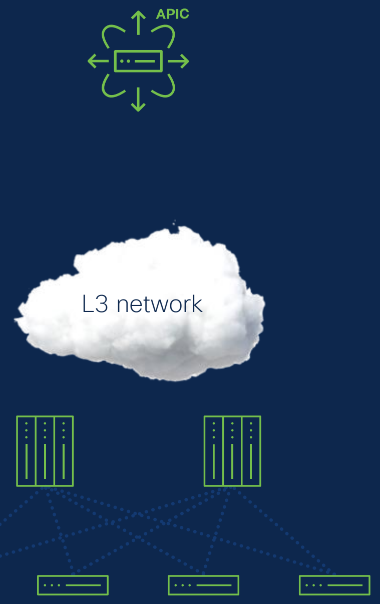

Deploy APIC cluster in a remote secure zone

Virtual APIC cluster

ACI 6.0.2F

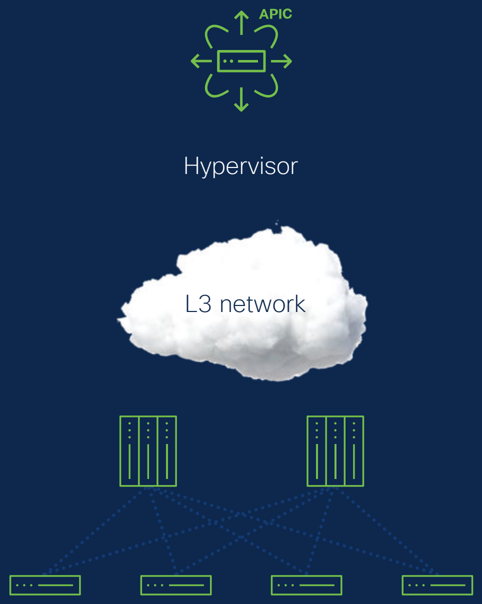

Deploy APIC cluster on ESXi hypervisor over L3 network

Cloud hosted APIC cluster

ACI 6.0.3F

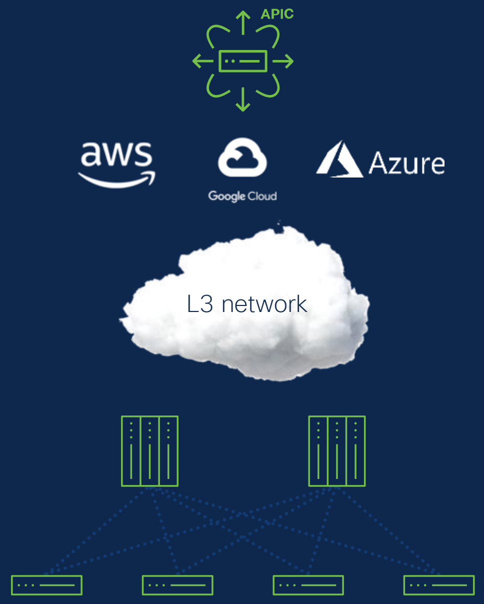

Deploy APIC cluster in the cloud to manage on-premises fabric

Flexible controller deployments

vAPIC in ESXi

Medium form factor (same scaling as physical APIC*)

APIC on customer managed ESXi

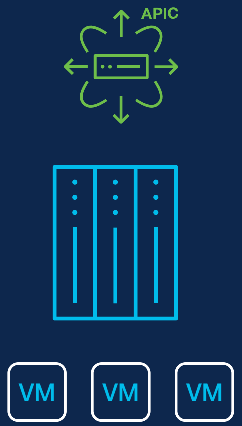

- 16 vCPU of 3 GHz or Higher
- 96 GB of RAM
- Disk 1: SSD or NVMe – 120GB (root disk)
- Disk 2: SSD or NVMe – 360GB and
- 2 Interfaces.
  - OOB 1Gbps or higher
  - Infra / Inband 10Gbps or higher.

APIC M4/L4

Next generation APIC controller appliance

Based on UCS Gen 6

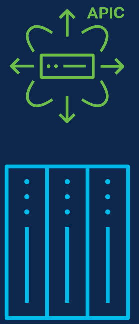

Old APIC Bootstrap (via console on each APIC)

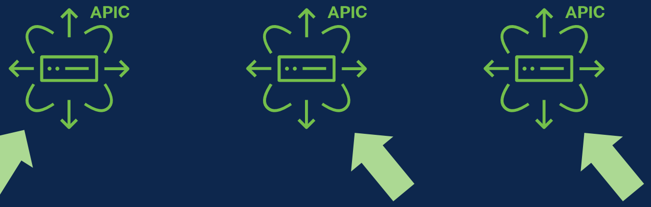

On APIC1

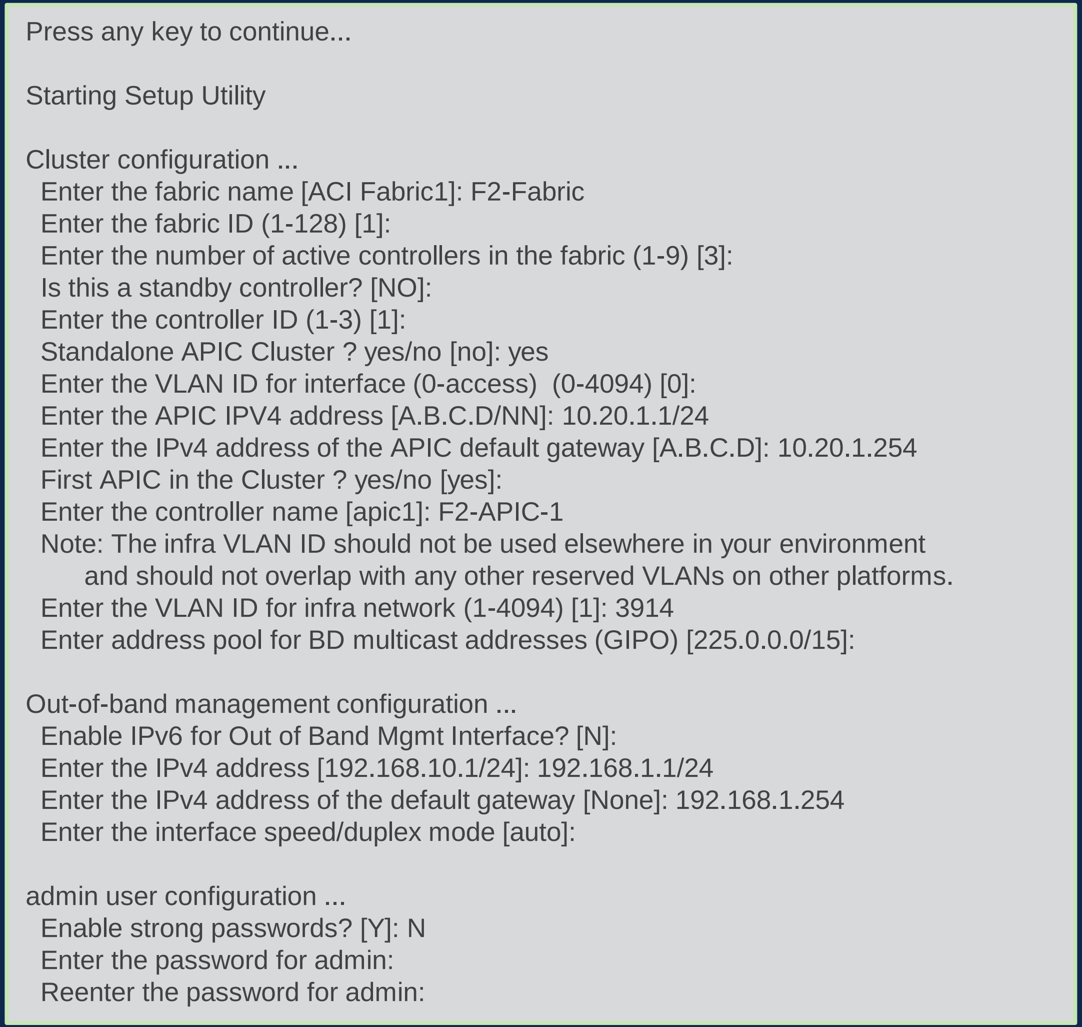

```
Press any key to continue...

Starting Setup Utility

Cluster configuration ...
  Enter the fabric name [ACI Fabric1]: F2-Fabric
  Enter the fabric ID (1-128) [1]:
  Enter the number of active controllers in the fabric (1-9) [3]:
  Is this a standby controller? [NO]:
  Enter the controller ID (1-3) [1]:
  Standalone APIC Cluster ? yes/no [no]: yes
  Enter the VLAN ID for interface (0-access)  (0-4094) [0]:
  Enter the APIC IPV4 address [A.B.C.D/NN]: 10.20.1.1/24
  Enter the IPv4 address of the APIC default gateway [A.B.C.D]: 10.20.1.254
  First APIC in the Cluster ? yes/no [yes]:
  Enter the controller name [apic1]: F2-APIC-1
  Note: The infra VLAN ID should not be used elsewhere in your environment
        and should not overlap with any other reserved VLANs on other platforms.
  Enter the VLAN ID for infra network (1-4094) [1]: 3914
  Enter address pool for BD multicast addresses (GIPO) [225.0.0.0/15]:

Out-of-band management configuration ...
  Enable IPv6 for Out of Band Mgmt Interface? [N]:
  Enter the IPv4 address [192.168.10.1/24]: 192.168.1.1/24
  Enter the IPv4 address of the default gateway [None]: 192.168.1.254
  Enter the interface speed/duplex mode [auto]:

admin user configuration ...
  Enable strong passwords? [Y]: N
  Enter the password for admin:
  Reenter the password for admin:
```

On APIC2, 3

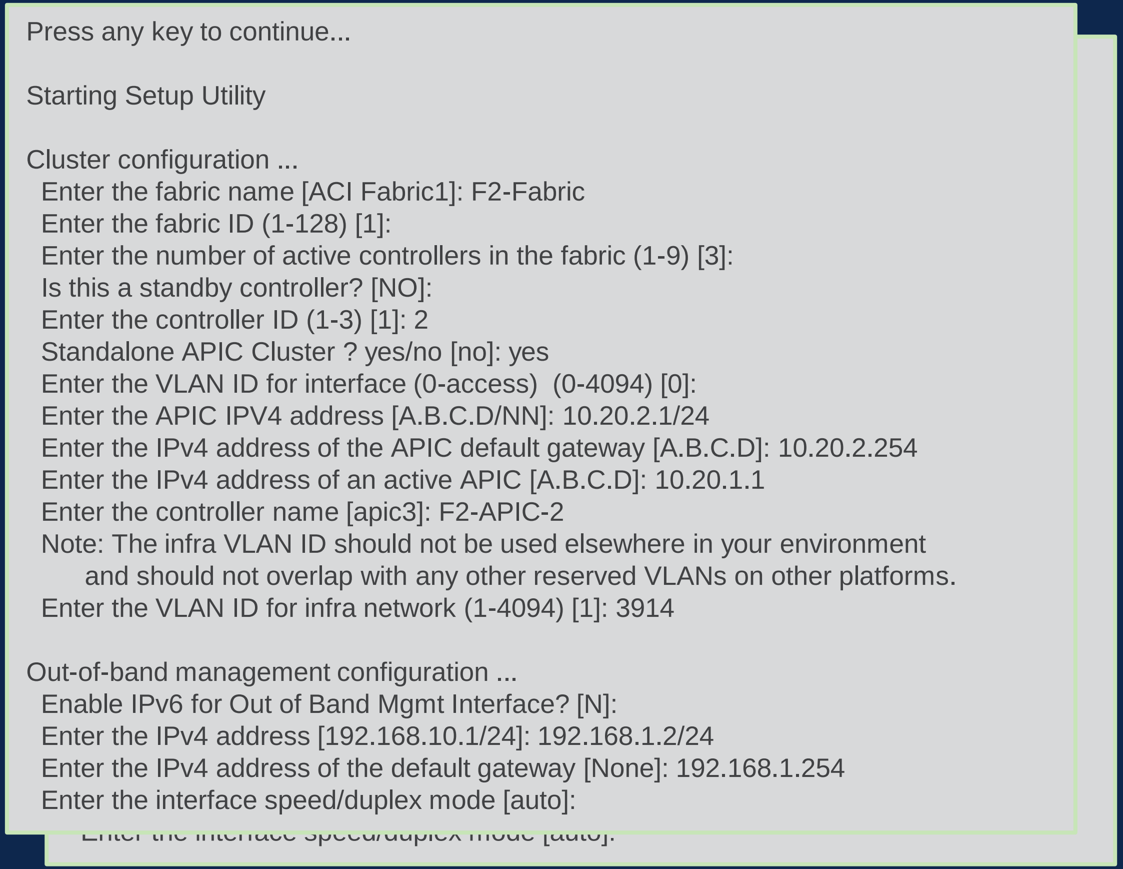

```
Press any key to continue...

Starting Setup Utility

Cluster configuration ...
  Enter the fabric name [ACI Fabric1]: F2-Fabric
  Enter the fabric ID (1-128) [1]:
  Enter the number of active controllers in the fabric (1-9) [3]:
  Is this a standby controller? [NO]:
  Enter the controller ID (1-3) [1]: 2
  Standalone APIC Cluster ? yes/no [no]: yes
  Enter the VLAN ID for interface (0-access)  (0-4094) [0]:
  Enter the APIC IPV4 address [A.B.C.D/NN]: 10.20.2.1/24
  Enter the IPv4 address of the APIC default gateway [A.B.C.D]: 10.20.2.254
  Enter the IPv4 address of an active APIC [A.B.C.D]: 10.20.1.1
  Enter the controller name [apic3]: F2-APIC-2
  Note: The infra VLAN ID should not be used elsewhere in your environment
        and should not overlap with any other reserved VLANs on other platforms.
  Enter the VLAN ID for infra network (1-4094) [1]: 3914

Out-of-band management configuration ...
  Enable IPv6 for Out of Band Mgmt Interface? [N]:
  Enter the IPv4 address [192.168.10.1/24]: 192.168.1.2/24
  Enter the IPv4 address of the default gateway [None]: 192.168.1.254
  Enter the interface speed/duplex mode [auto]:
```

New APIC Bootstrap (All through APIC1)

NEW

6.0(2)

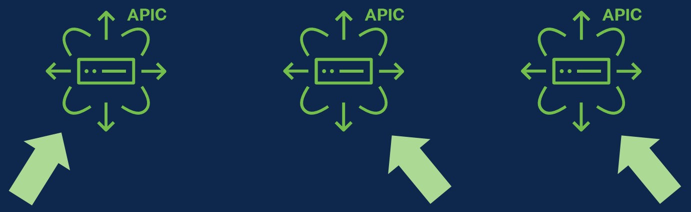

On APIC1 (only password and OOB)

```
Press any key to continue...

Starting Setup Utility
APIC Version: 6.0(2h)
Welcome to APIC Setup Utility
Press Enter Or Input JSON string to bootstrap your APIC node.

admin user configuration ...
  Enter the password for admin [None]:
  Reenter the password for admin [None]:
Out-of-band management configuration ...
 Enter the IP Address [192.168.10.1/24]: 192.168.1.1/24
  Enter the IP Address of default gateway [192.168.10.254]: 192.168.1.254
Would you like to edit the configuration? (y/n) [n]:

System pre-configured successfully.
Use: https://192.168.1.1 to complete the bootstrapping
Press F9 to enable debug shell
```

On APIC2, 3

No Console Input

The rest is via the GUI or API on APIC1 through CIMC IP of other APICs

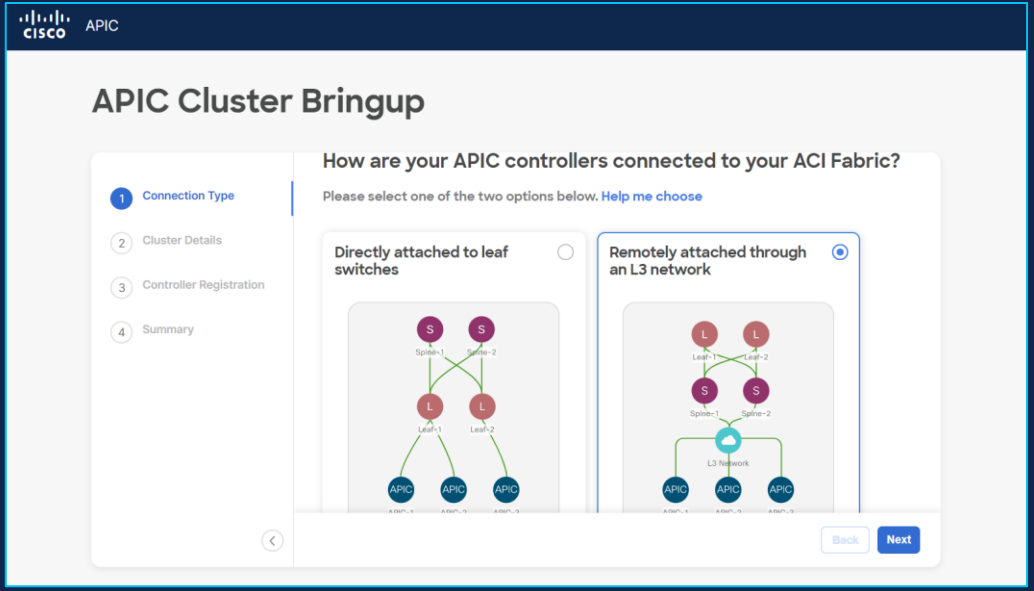

APIC

APIC Cluster Bringup

How are your APIC controllers connected to your ACI Fabric?

Please select one of the two options below. Help me choose

1 Connection Type
2 Cluster Details
3 Controller Registration
4 Summary

Directly attached to leaf switches ○

Remotely attached through an L3 network ●

Auto Firmware Update on APIC discovery

Automatically upgrade a new APIC joining the cluster to the same version as other APICs

Use Case 1: APIC Replacement

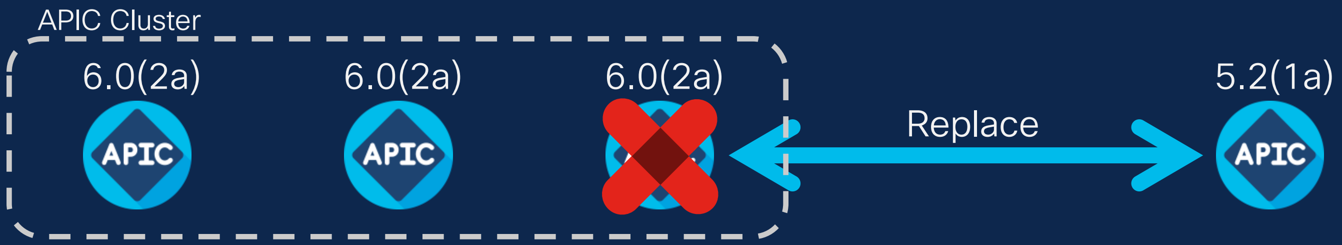

Use Case 2: Cluster Expansion

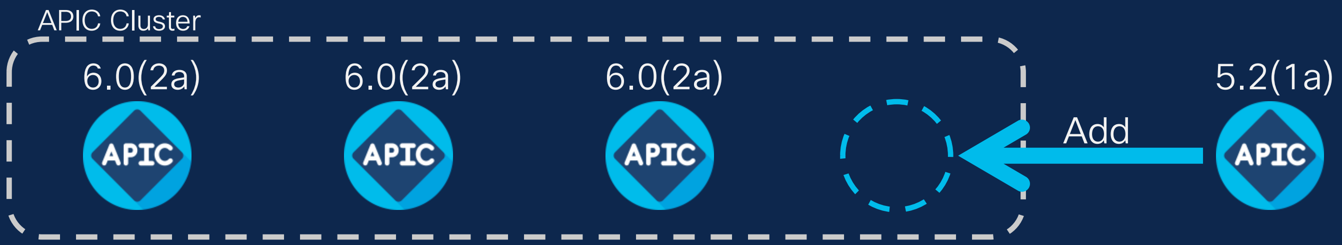

Automatizace ACI

Cisco ACI: A platform built for automation

Simplicity and speed: A single API call to deliver a datacenter-wide construct (like a VRF)

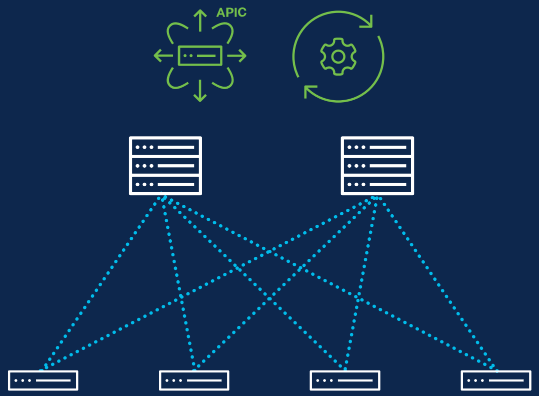

Day 0

Out of the box automation

- Hardware
- Fabric
- Underlay
- Parts of overlay

ACI fabric is automatically deployed based on industry best practices

Avoid 100s of design decisions required to deploy a traditional fabric

Day 1

Operations ticketing queue

ServiceNow ticket – create VRF

Ansible | Terraform | Python

APIC API

Hundreds of switches and 1000s of ports

IT operations lifecycle: Provision, automate and operate

RedHat Ansible and HashiCorp Terraform with Cisco ACI

An infrastructure-as-code practice

140+ Ansible modules for ACI

180+ Terraform resources for ACI

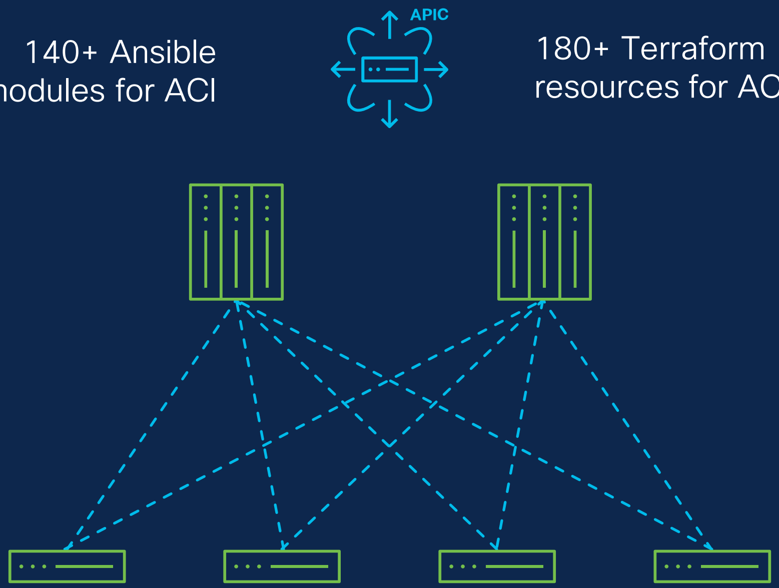

Automate fabrics end-to-end with Ansible & Terraform

Capabilities

- Comprehensive coverage of the ACI REST API
- Common configuration management language
- Simple CI/CD
  - Ansible Automation Platform, GitHub, and GitHub Actions (CI)
  - Terraform Cloud for Business

Benefits

- Leverage ACI's object model and single configuration API for 100s of Nexus 9000s
- ACI as on-prem anchor for hybrid cloud
- Certified by both RedHat, HashiCorp, and Cisco Cloud Networking BU

Cisco ACI: A purpose-built Infrastructure as Code platform

Nexus-as-Code

Complex L3out simplified to 23 lines of code; dependencies managed

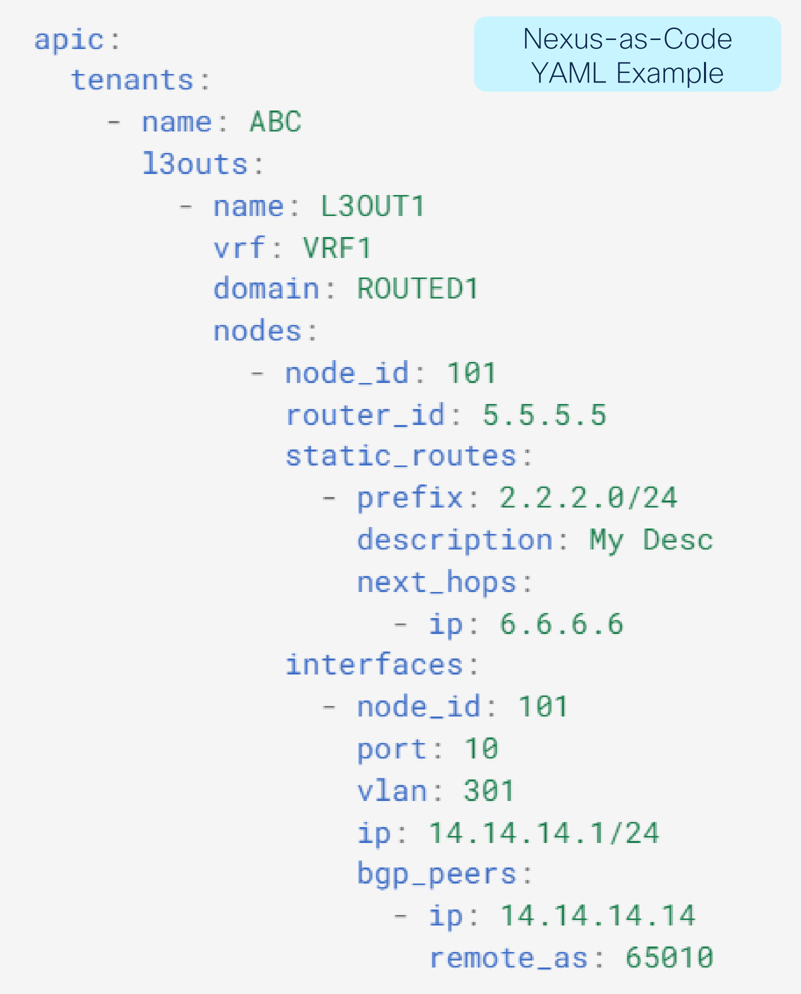

Nexus-as-Code YAML Example

```yaml
apic:
  tenants:
    - name: ABC
      l3outs:
        - name: L3OUT1
          vrf: VRF1
          domain: ROUTED1
          nodes:
            - node_id: 101
              router_id: 5.5.5.5
              static_routes:
                - prefix: 2.2.2.0/24
                  description: My Desc
                  next_hops:
                    - ip: 6.6.6.6
              interfaces:
                - node_id: 101
                  port: 10
                  vlan: 301
                  ip: 14.14.14.1/24
                  bgp_peers:
                    - ip: 14.14.14.14
                      remote_as: 65010
```

Cisco supports entire stack

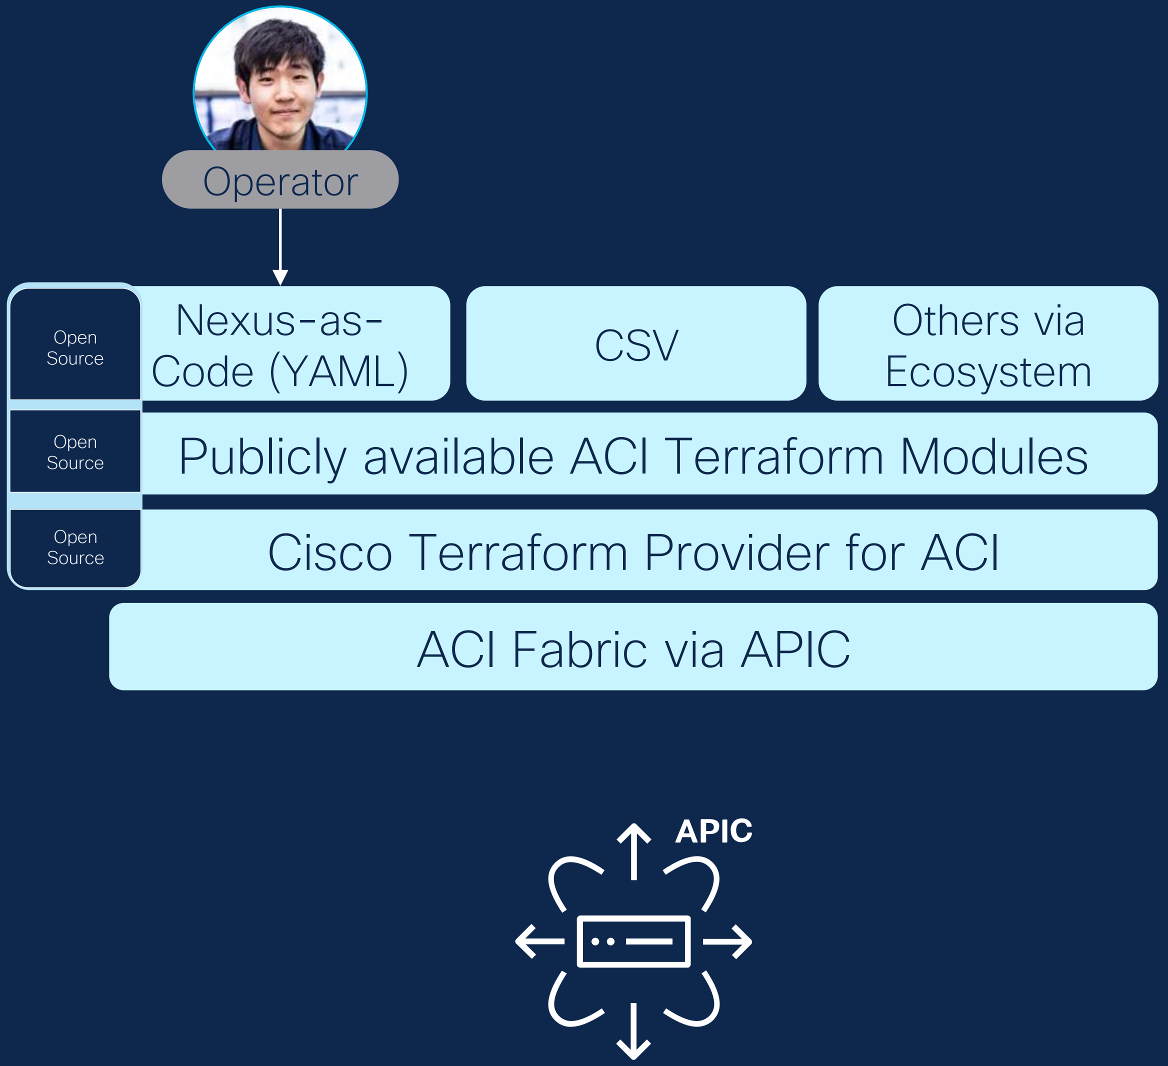

Nexus-as-Code is an ACI abstraction to simplify consumption using a prescriptive Data Model

Symetrický PBR

Weighted based Symmetric PBR

- Prior to APIC release 6.0, there is no option to specify weight for each PBR destination. Thus, the assumption is that PBR destinations (service devices) in the same PBR policy have same/similar capacity to handle traffic.

- This is the ask to support weight for each PBR destination, so that Symmetric PBR can distribute traffic based on the weights. It can cover the situation where a PBR policy has the mix of service devices that have different capacities.

- For existing PBR related features, please refer PBR white paper.
  - https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739971.html

With weight per PBR destination

- Each PBR destination has weight configuration.
- The weight range is 1-10.
  - The total number of weights per PBR policy is up to 128 if PBR destinations are in a BD
  - The total number of weights per PBR policy is up to 64 if PBR destinations are in an L3Out

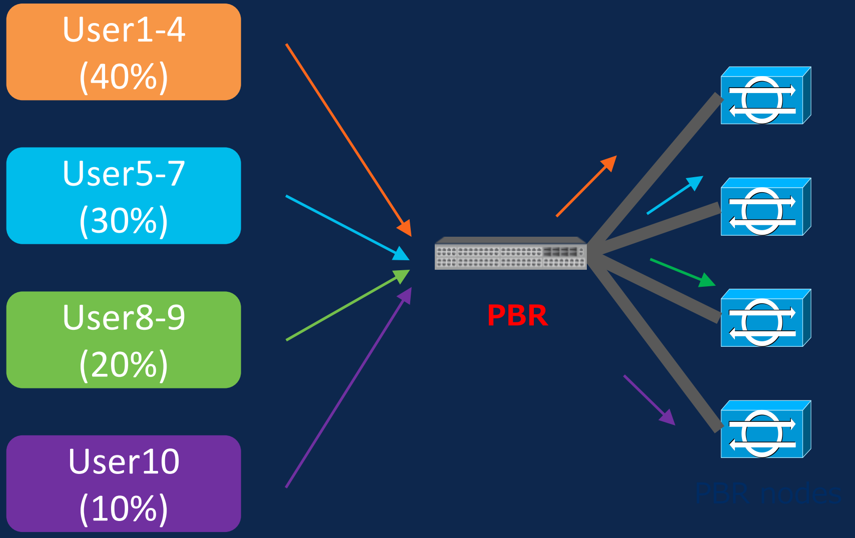

Destination	IP	Weight (optional)	%
Destination 1	10.1.1.1	4	40%
Destination 2	10.1.1.2	3	30%
Destination 3	10.1.1.3	2	20%
Destination 4	10.1.1.4	1	10%

Total # of buckets: 10

Užitečné odkazy

References

- Getting Started with Cisco ACI 6.0
  - https://dcloud2-lon.cisco.com/content/demo/571506

- ACI Recommended releases
  - https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/recommended-release/b_Recommended_Cisco_ACI_Releases.html

- ACI Upgrade tools
  - https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/all/apic-installation-aci-upgrade-downgrade/Cisco-APIC-Installation-ACI-Upgrade-Downgrade-Guide/m-pre-upgrade-downgrade-checklists.html
  - https://github.com/datacenter/ACI-Pre-Upgrade-Validation-Script
  - https://dcappcenter.cisco.com/pre-upgrade-validator.html

References

- Cisco ACI Endpoint Security Group (ESG) Design Guide
  - https://www.cisco.com/c/en/us/td/docs/dcn/whitepapers/cisco-aci-esg-design-guide.html

- New APIC Bootstrap (Getting Started Guide 6.0)
  - https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/6x/getting-started/cisco-apic-getting-started-guide-60x/initial-setup-60x.html#Cisco_Task_in_List_GUI.dita_b290f6c9-c72e-4152-9490-6d27cb412d45

- New APIC Bootstrap API
  - https://developer.cisco.com/docs/apic-rest-api-configuration-guide/#!bringing-up-a-cluster

- Cisco Nexus-as-Code
  - https://developer.cisco.com/docs/nexus-as-code

CISCO

The bridge to possible