Cisco Nexus Smart Switches with DPUs

A window to a transparent secure datacenter network architecture

Maurizio Portolani
Distinguished TME

BRKDCN-2643

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 13, 2025.

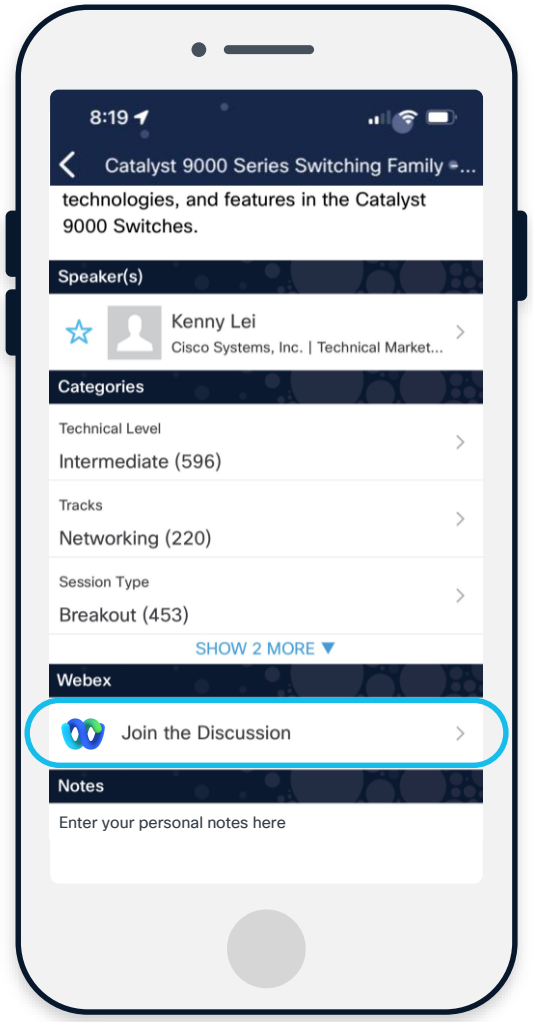

Agenda

01 Introduction to Smart Switches

02 Introduction to Hypershield

03 Configuration Example

04 Packet Walk

05 Design / Use Cases

06 Target Scale

DISCLAIMER

Any information provided in this document regarding future functionalities is for informational purposes only and is subject to change including ceasing any further development of such functionality. Many of these future functionalities remain in varying stages of development and will be offered on a when-and-if available basis, and Cisco makes no commitment as to the final delivery of any of such future functionalities. Cisco will have no liability for Cisco's failure to deliver any or all future functionalities and any such failure would not in any way imply the right to return any previously purchased Cisco products.

What is a Cisco Nexus N9300 Smart Switch

- A Nexus N9300 Smart Switch is a Silicon One based NXOS switch equipped with a Distributed Processing Unit (DPU) which is a High-performance, programmable packet processing engine
- DPUs provide a middle ground between software functions running in CPUs and hardware-based forwarding in ASIC
- A DPU consists of a programmable data plane, memory, specialized acceleration engines, and embedded CPUs for control plane/management

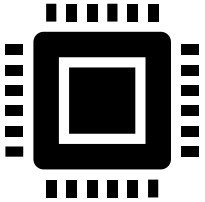

- DPU Functions:
  - Distributed stateful firewall
  - Encryption services (e.g., IPsec)
  - Telemetry/flow data collection offload
  - Stateful NAT/PAT

Cisco N9324C-SE1U Smart Switch

- 24 x 100G QSFP ports (2.4T)
- Silicon One E100 NPU
- 4 x AMD “Elba” DPUs (800G capacity)
- 16-core 2.5GHz Intel CPU
- 32GB DRAM per DPU (128GB total)

N9324C-SE1U

N9324C-SE1U Architecture

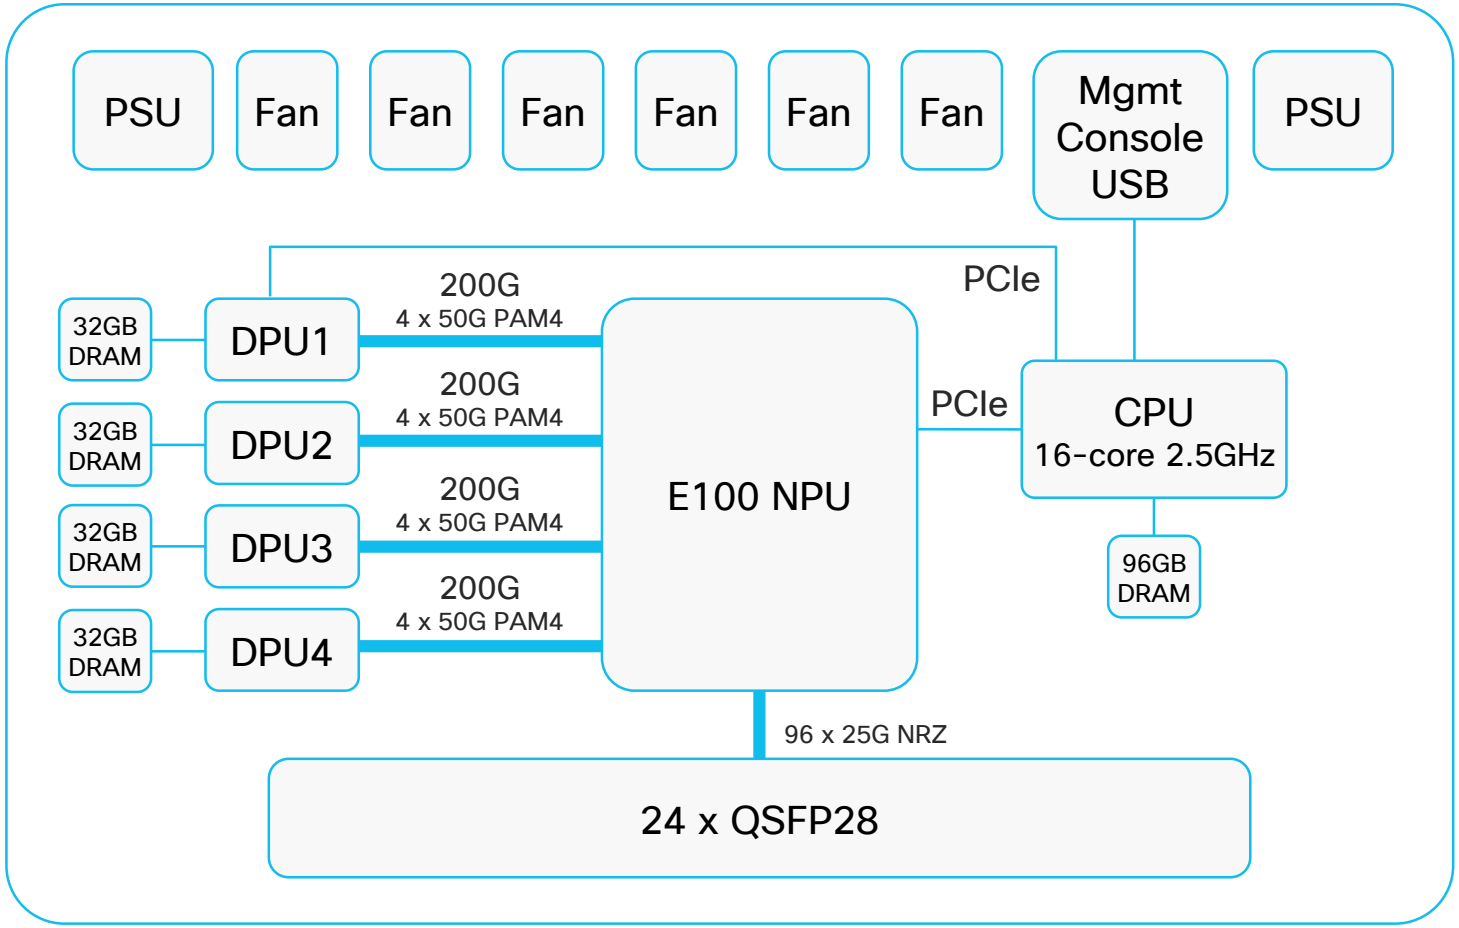

N9324C-SE1U

Cisco N9348Y2C6D-SE1U Smart Switch

- 48 x 25G SFP28 ports + 2 x 100G QSFP28 ports + 6 x 400G QSFP-DD ports (3.8T)
- Silicon One E100 NPU
- 2 x AMD “Giglio” DPUs (800G capacity)
- 16-core 2.5GHz Intel CPU
- 96GB system DRAM
- 64GB DRAM per DPU (128GB total)

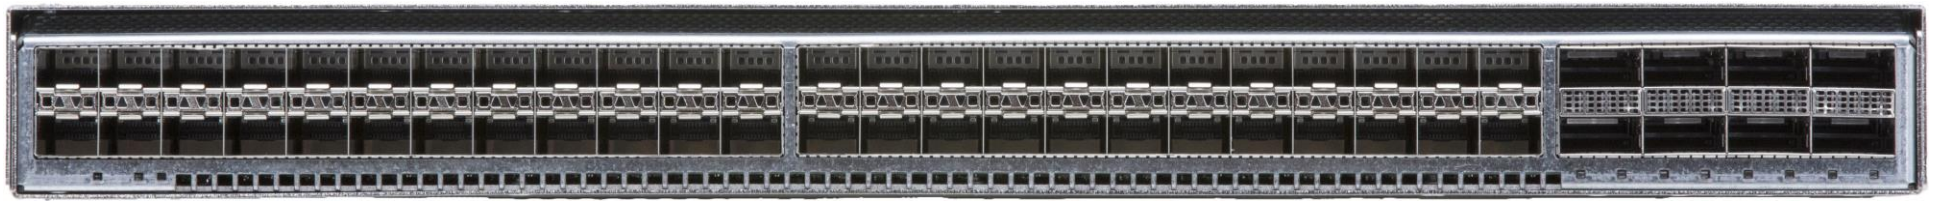

N9348Y2C6D-SE1U

N9348Y2C6D-SE1U Architecture

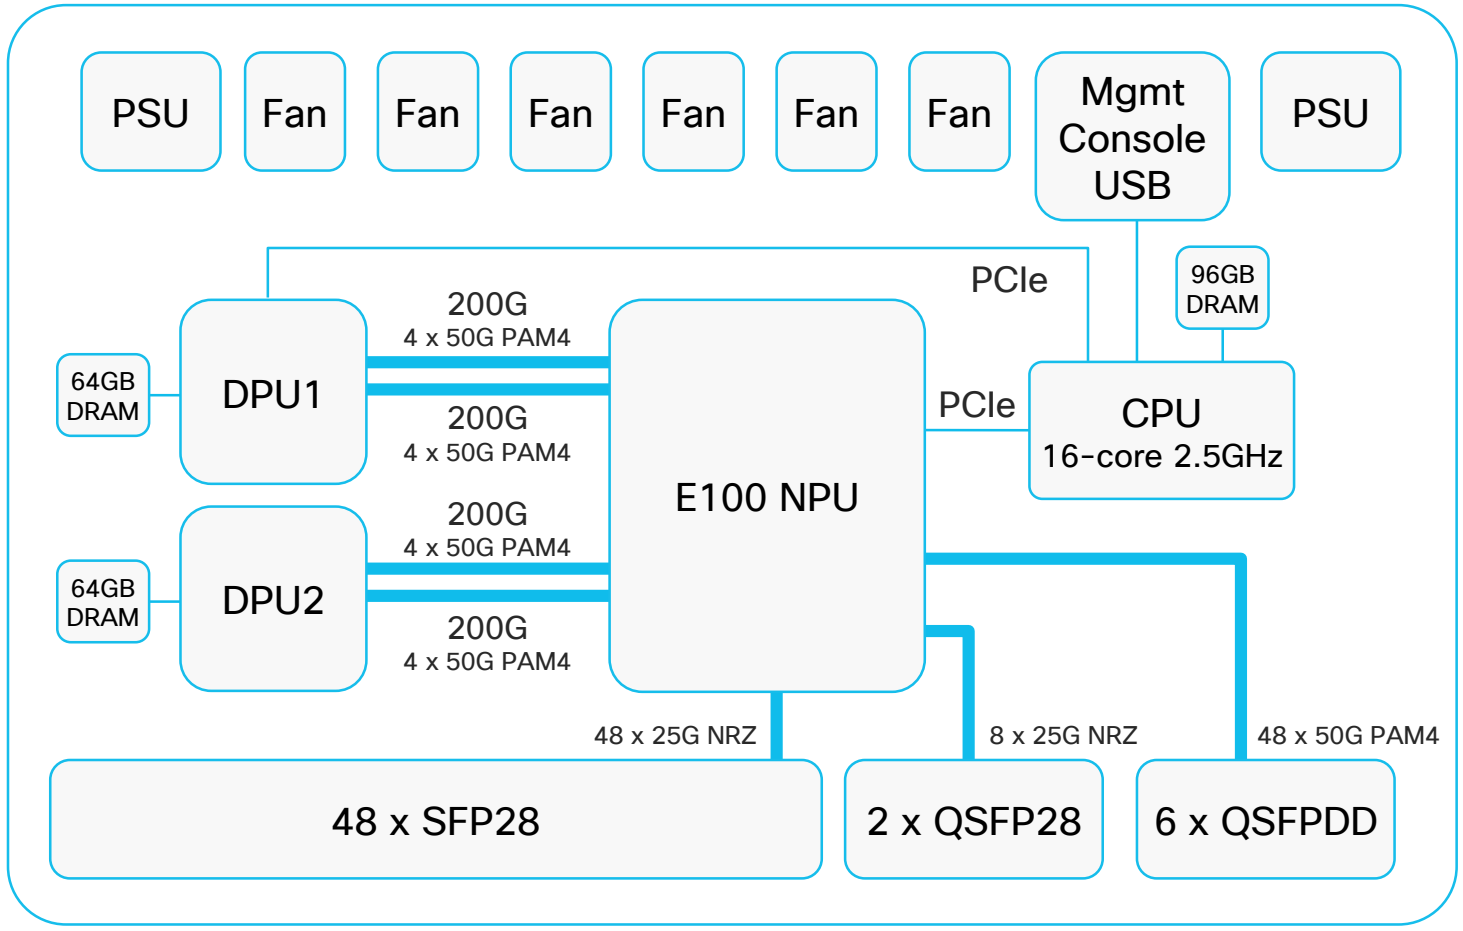

N9348Y2C6D-SE1U

AMD DPU

- Dual 200Gbps connectivity
- 144 P4-programmable pipeline
- 16 ARM cores
- Offloads:
  - IPSEC,
  - compression, decompression
  - deduplication
  - [...]

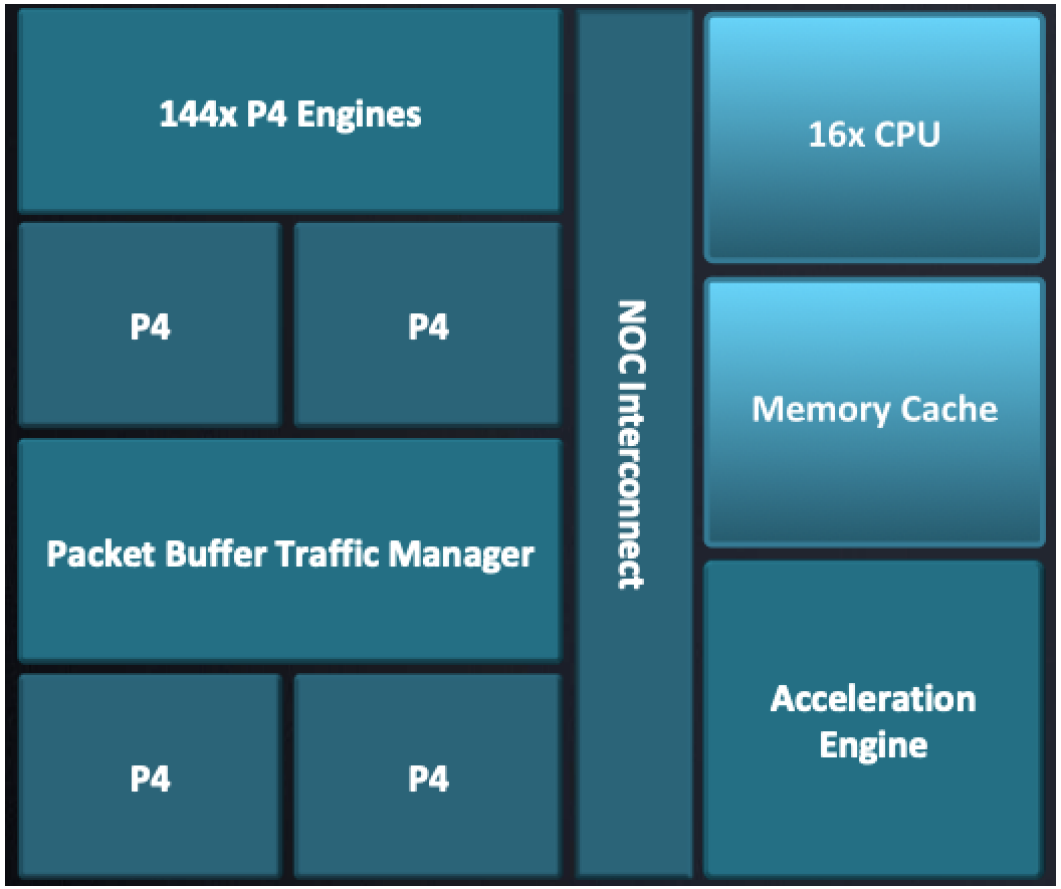

Port Speed and Optics Support

- Native speeds: 40G and 100G
- Breakout support:
  - 4x25G
  - 4x10G
  - 2x50G (*)
- QSA:
  - 10G QSA – Yes
  - 25G QSA – No
  - 1G QSA – No

- Optics:
- 100G Optics
  - QSFP28-100G-SR4
  - QSFP28-100G-PSM4
  - QSFP28-100G-CWDM4
  - QSFP28-100G-LR4
  - QSFP28 AOC, 1, 3, 5, 7, 10, 15, 30m
  - QSFP28 – 100G DR, 100G FR
- 40G and Optics
  - QSFP-40G-LR4
  - QSFP AOC, 1, 3, 5, 7, 10, 15, 30m
  - QSFP-40G-LR4-S
  - QSFP-40G-SR-BD
  - QSFP-40G-SR4-S
  - QSFP-40G-SR4

100G and 40G Copper Cable support (1,3,5)

(*) 2x50G was tested with third party cables: Mellanox QSFP-100G-PCC.

Software Components

- NXOS handles images for both Cisco NXOS and the DPU
- NXOS provides the L3 connectivity for the Hypershield Agent
- The Hypershield agent receives Firewall configurations and policies from Hypershield and pushes them to the DPU complex that in-turn configures firewall rules on DPUs.
- The Hypershield Agent also collects firewall flow information from the DPU complex, construct firewall flow logs and exports them

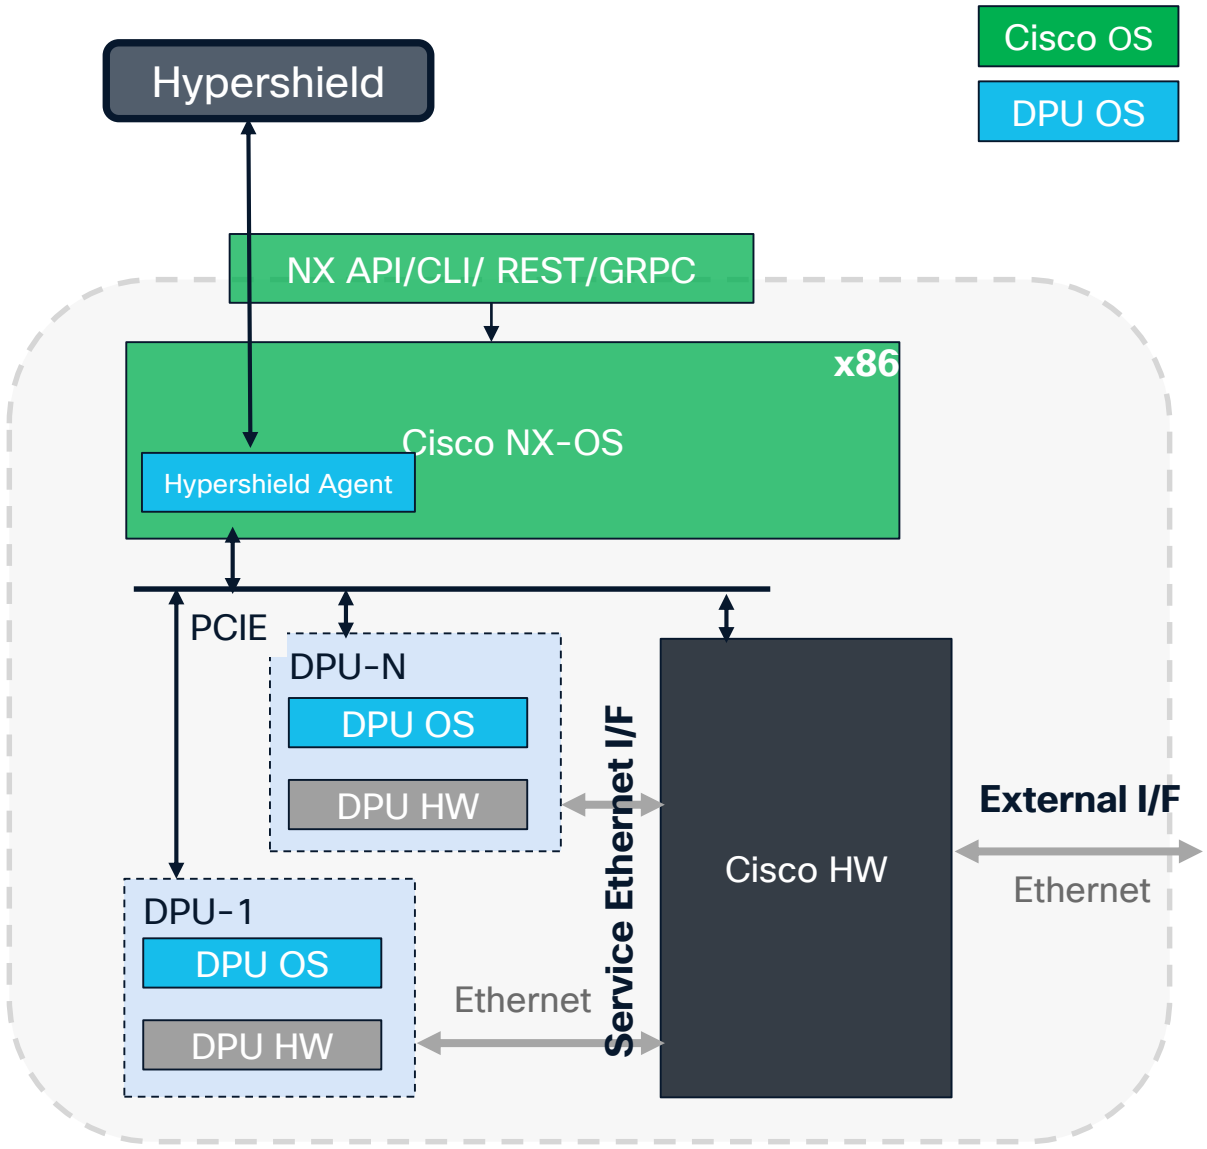

The Cisco Nexus 9300 Smart Switch versus a Firewall

	Nexus 9300 Smart Switch	Firewall
Coordination of policies	Autonomous segmentation, compensating controls, etc...	Central Console, L7 functions
Performance/Cost	Higher Performance for less	Less performance
Latency	Lower latency (*)	Higher
Power Consumption	Lower than Firewall+Switch combined	Higher
Stateful Inspection	TCP termination and TCP state tracking	TCP termination and TCP stater tracking
MACSEC	MACSEC on Switch ports	NA
Deep Packet Inspection	ALG	More features

(*) Expected Latency = ~2.5us (Silicon One) + 3us (DPU) + 2.5us (Silicon One) << FW Latency

Hypershield: Comprehensive Security

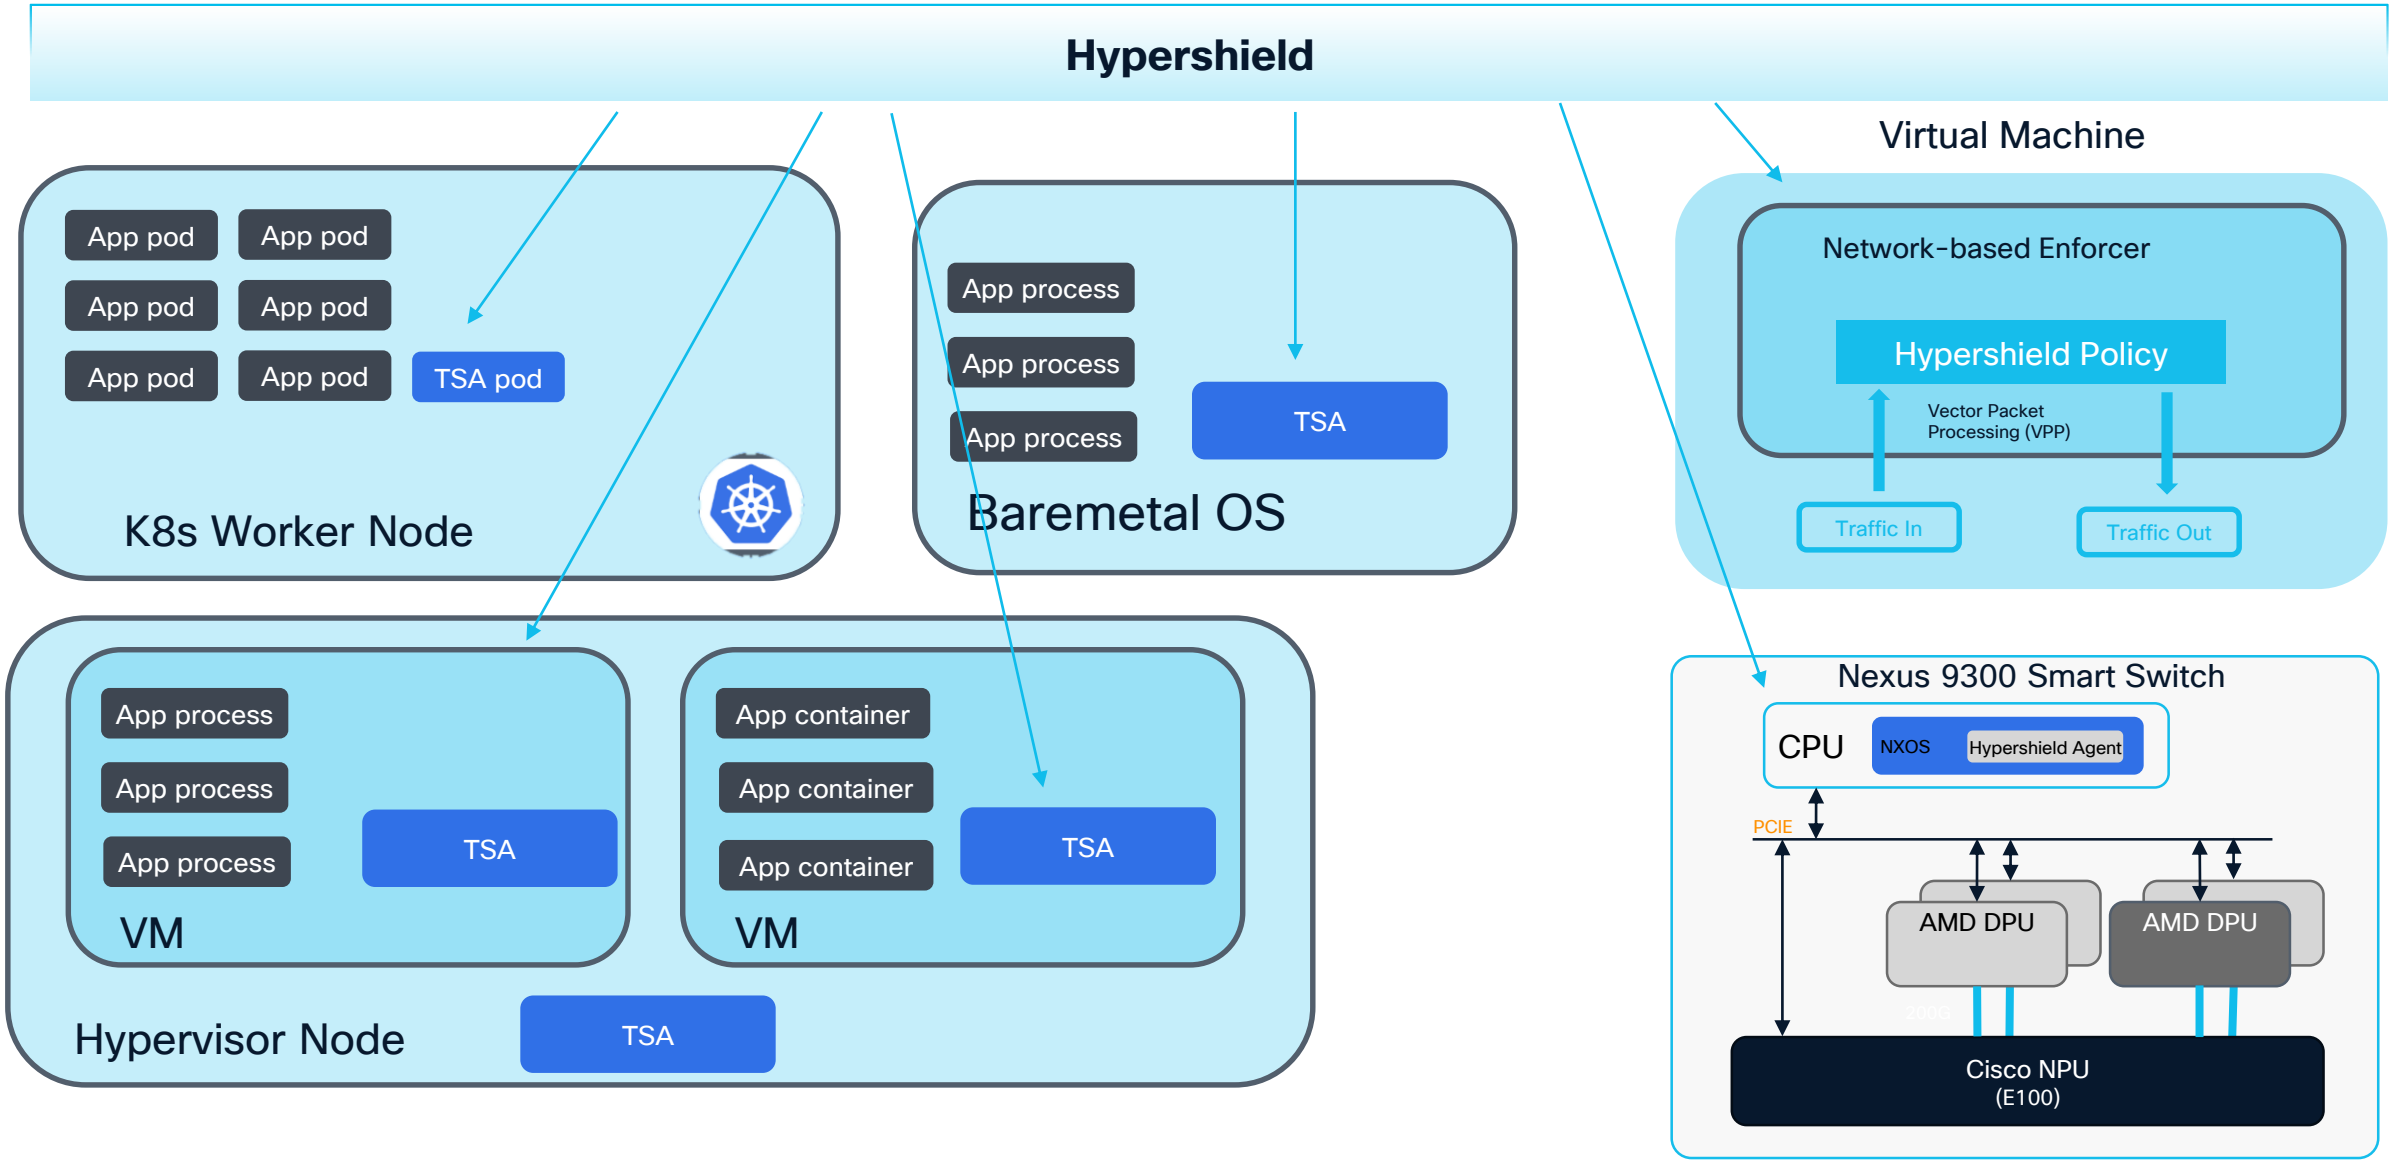

Hypershield defined policies

Policies cover both network and workloads

PARC = Principal Action Resource Condition

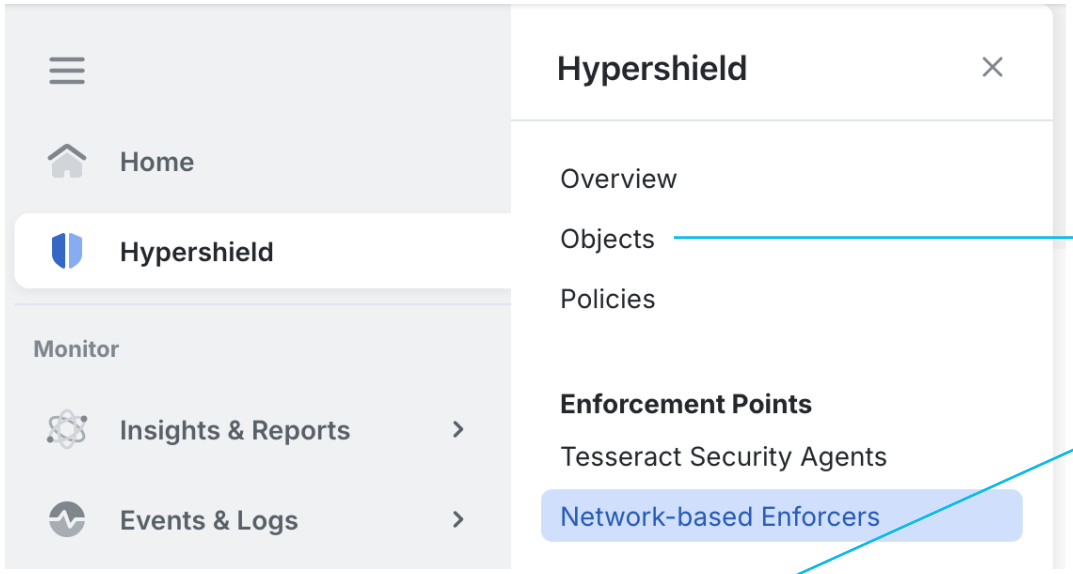
Hypershield
Overview
Objects
Policies
Enforcement Points
Tesseract Security Agents
Network-based Enforcers
Home
Hypershield
Monitor
Insights & Reports
Events & Logs

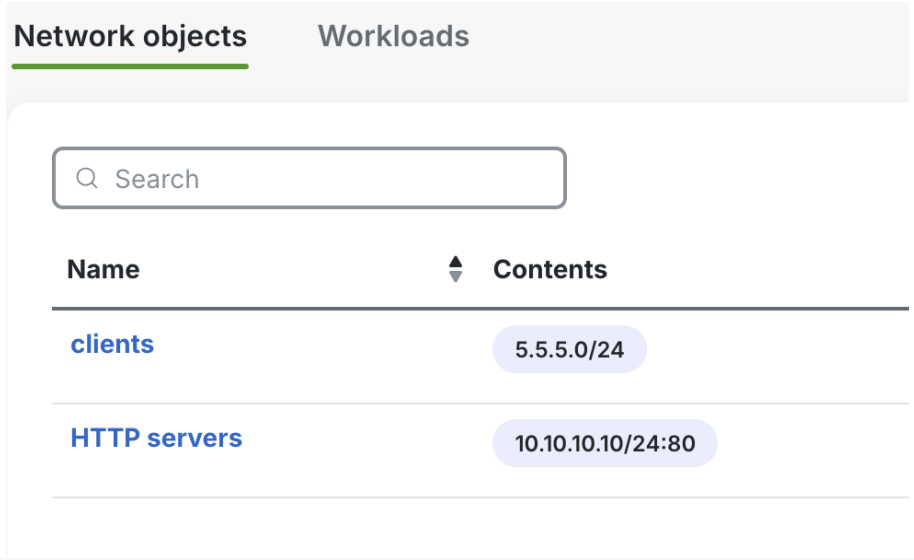
Network objects
Workloads
Search
Name
Contents
clients
5.5.5.0/24
HTTP servers
10.10.10.10/24:80

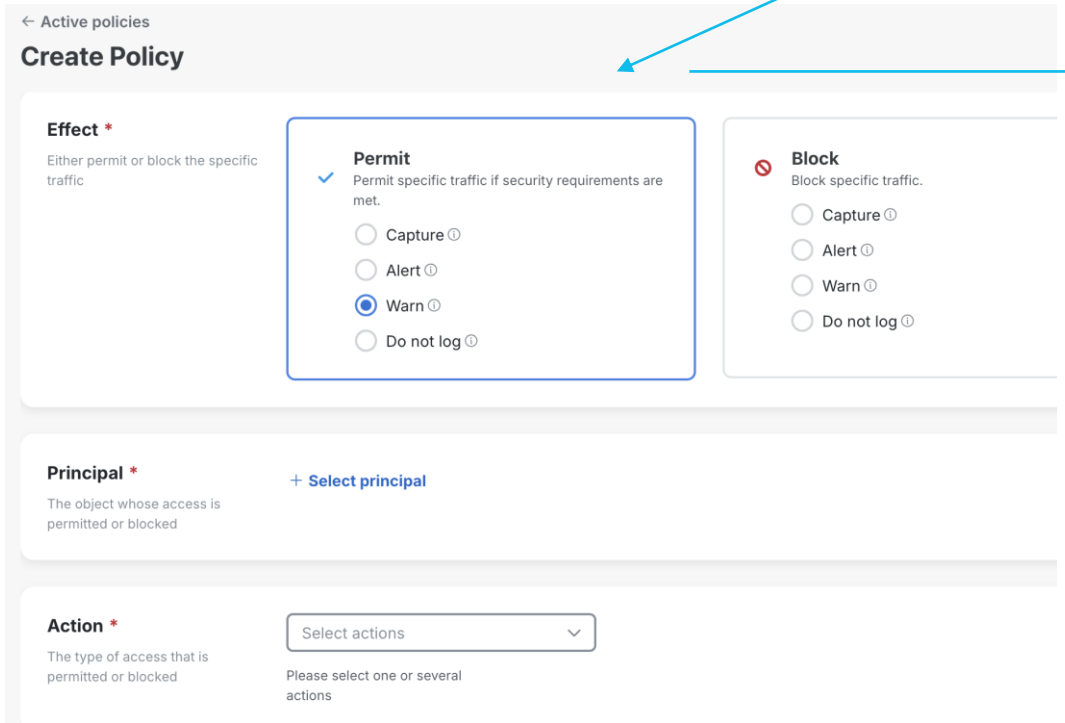
Active policies
Create Policy
Effect *
Either permit or block the specific traffic
Permit
Permit specific traffic if security requirements are met.
Capture
Alert
Warn
Do not log
Block
Block specific traffic.
Capture
Alert
Warn
Do not log
Principal *
The object whose access is permitted or blocked
+ Select principal
Action *
The type of access that is permitted or blocked
Select actions
Please select one or several actions

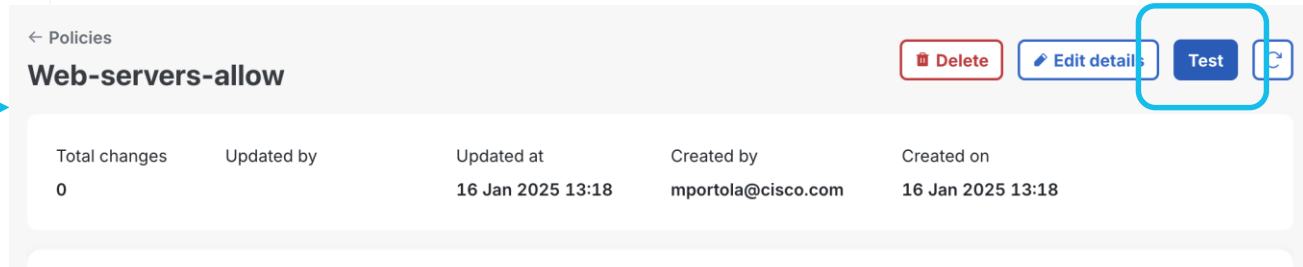
Policies
Web-servers-allow
Delete
Edit details
Test
Total changes
0
Updated by
Updated at
16 Jan 2025 13:18
Created by
mportola@cisco.com
Created on
16 Jan 2025 13:18

Testing & Deploying Policy

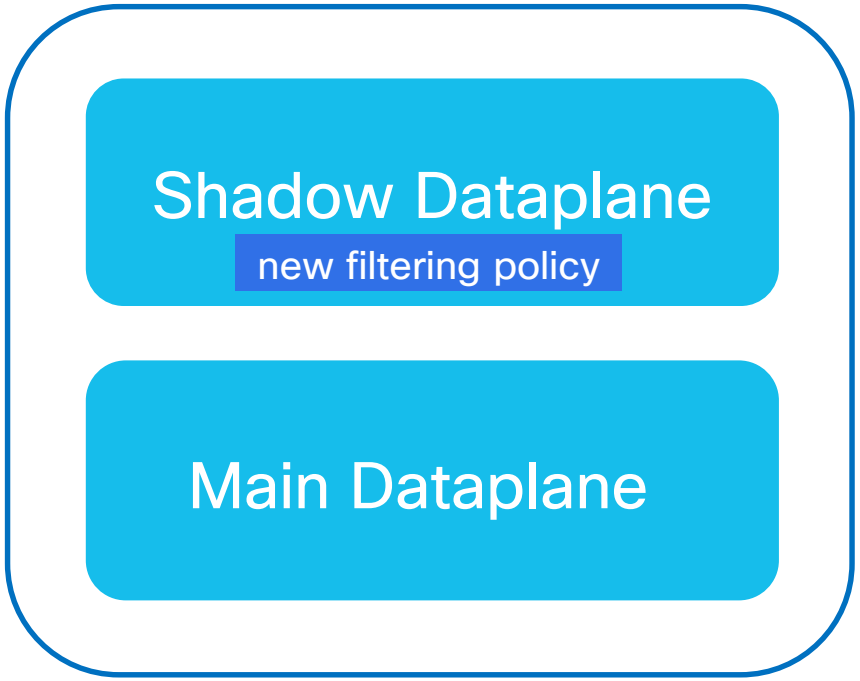

Latency min	0%	within 5%	Passed	*No difference*
Latency max	0%	within 5%	Passed	*No difference*
Latency avg	0%	within 5%	Passed	*No difference*

Flow

Metric	Difference	Threshold for success	Result	Reason
Flow Age max	0.01%	within 5%	Passed	Difference as expected
Flow Age avg	0.01%	within 5%	Passed	Difference as expected
ipv4 flows	0%	within 0%	Review	*No difference*
icmp flows	0%	within 0%	Review	*No difference*
total allowed flows	0%	within 0%	Review	*No difference*

Hit Counts Review

Hit counts are tracked to show any discrepancies between the number of times a policy was exercised between the primary dataplane and this update currently on the shadow dataplane.

Policy	Number of Hits (Primary dataplane)	Number of Hits (Shadow dataplane)	Difference	Review
Allow ping from App to DB	n/a	12	n/a	Needs review

Policy Rule LifeCycle

Test Reports of deployed policy kept in History tab.

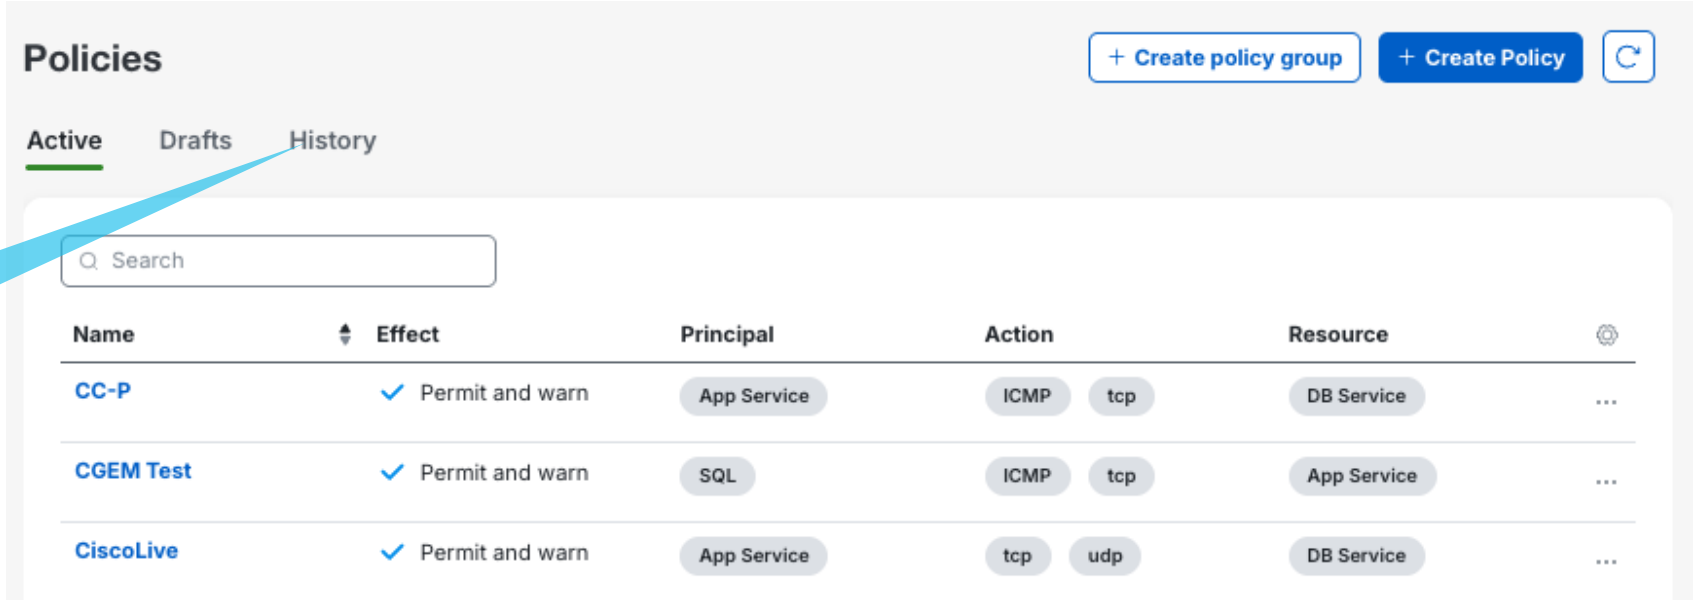

How does Policy end up in an Enforcer?

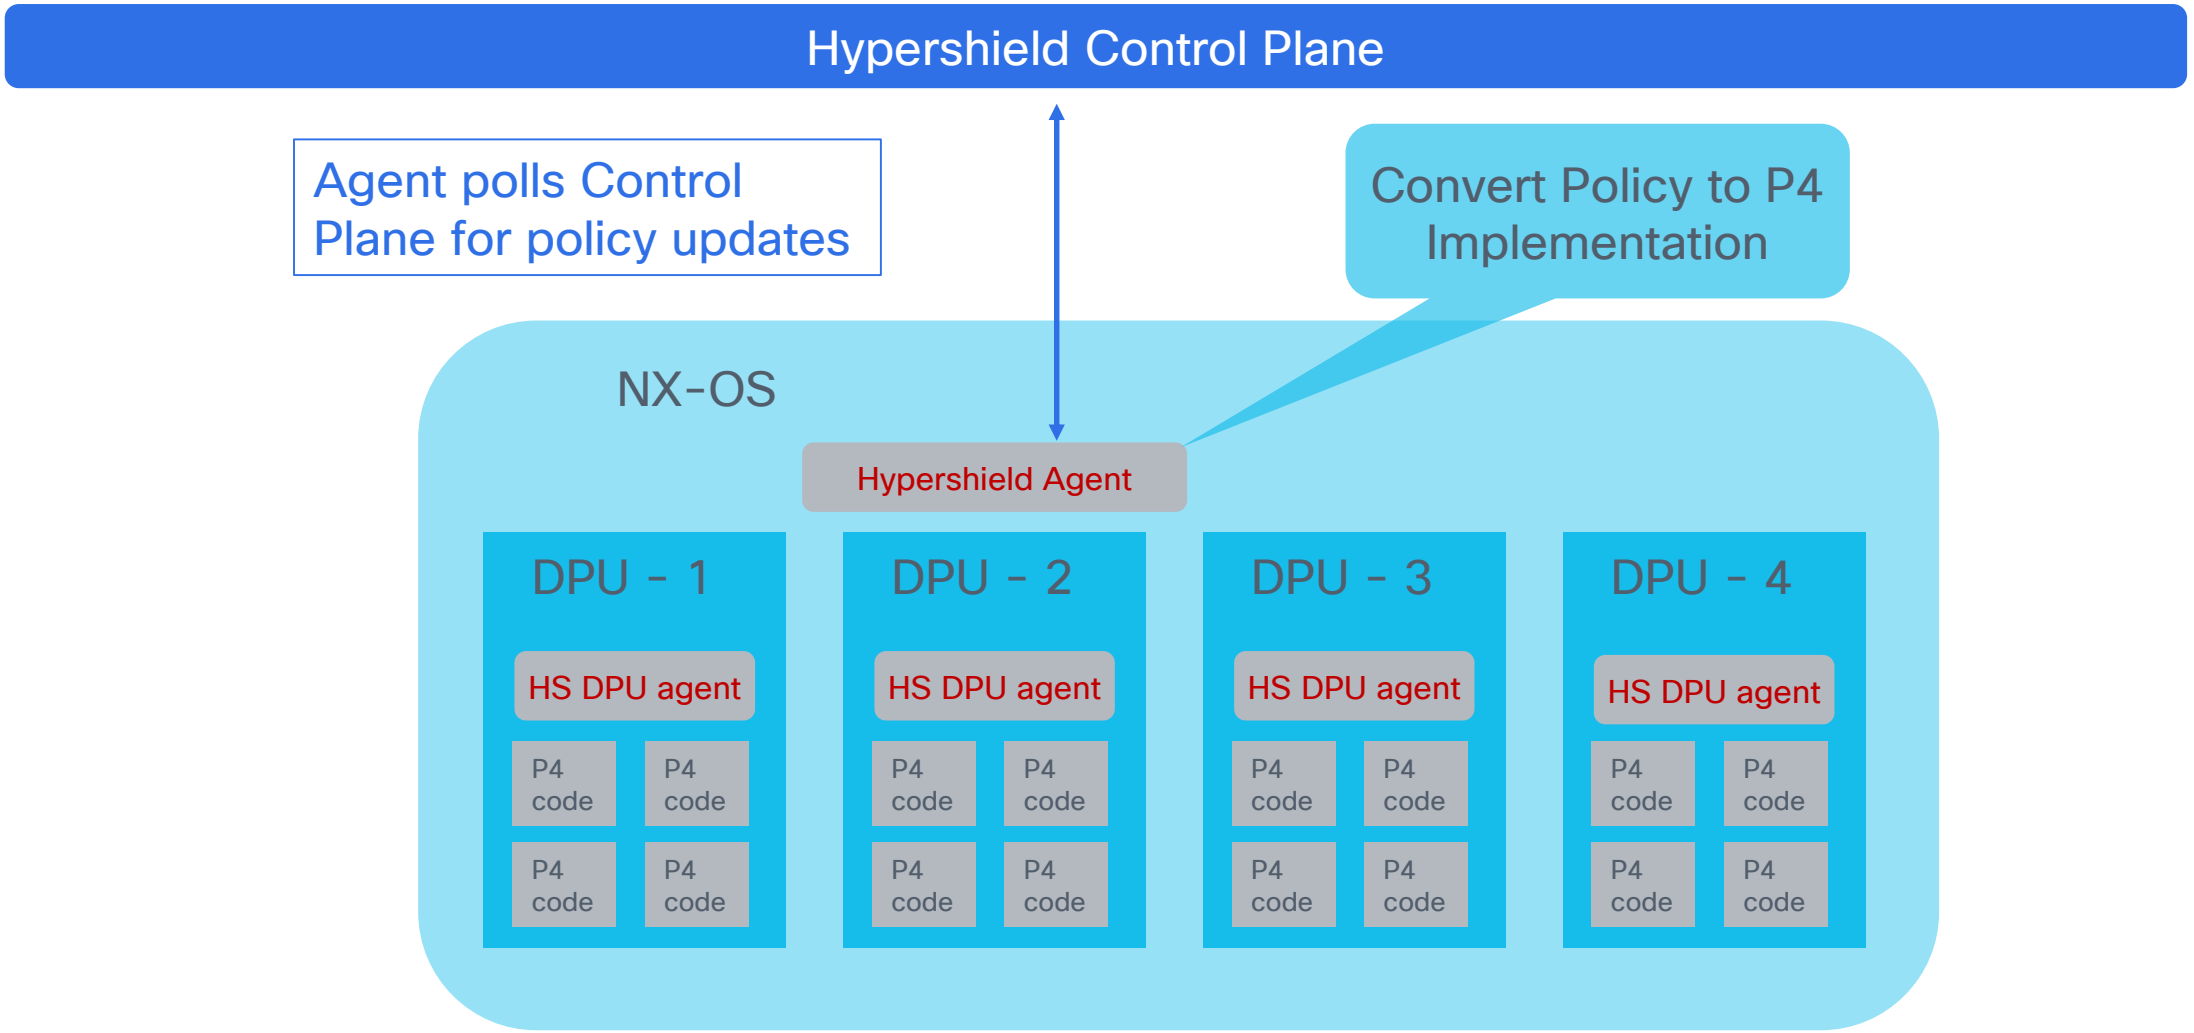

Policy Decision Process

- No Rule Ordering -> All policy rules are evaluated
- Default Deny
- Deny Wins, except for conditionals
- Apply Effect(s)

Implementation

- Depends on enforcer
- Obviously not evaluating every rule for each decision

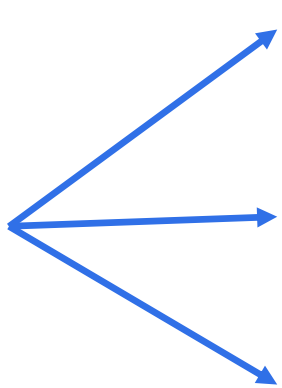

DPU Switch: DPU plus VPP

VM Enforcer: VPP

End Point: ebpf

- Both deny and allow policies are sorted separately
- Deny will override any Permit policies.
- The priority is:
  - Deny (Most specific to least specific ordered)
  - Permit (Most specific to least specific ordered)
  - Default Deny (Default, no policy needed)

Configuration on the Switch

Service Acceleration Configuration for L3 Traffic

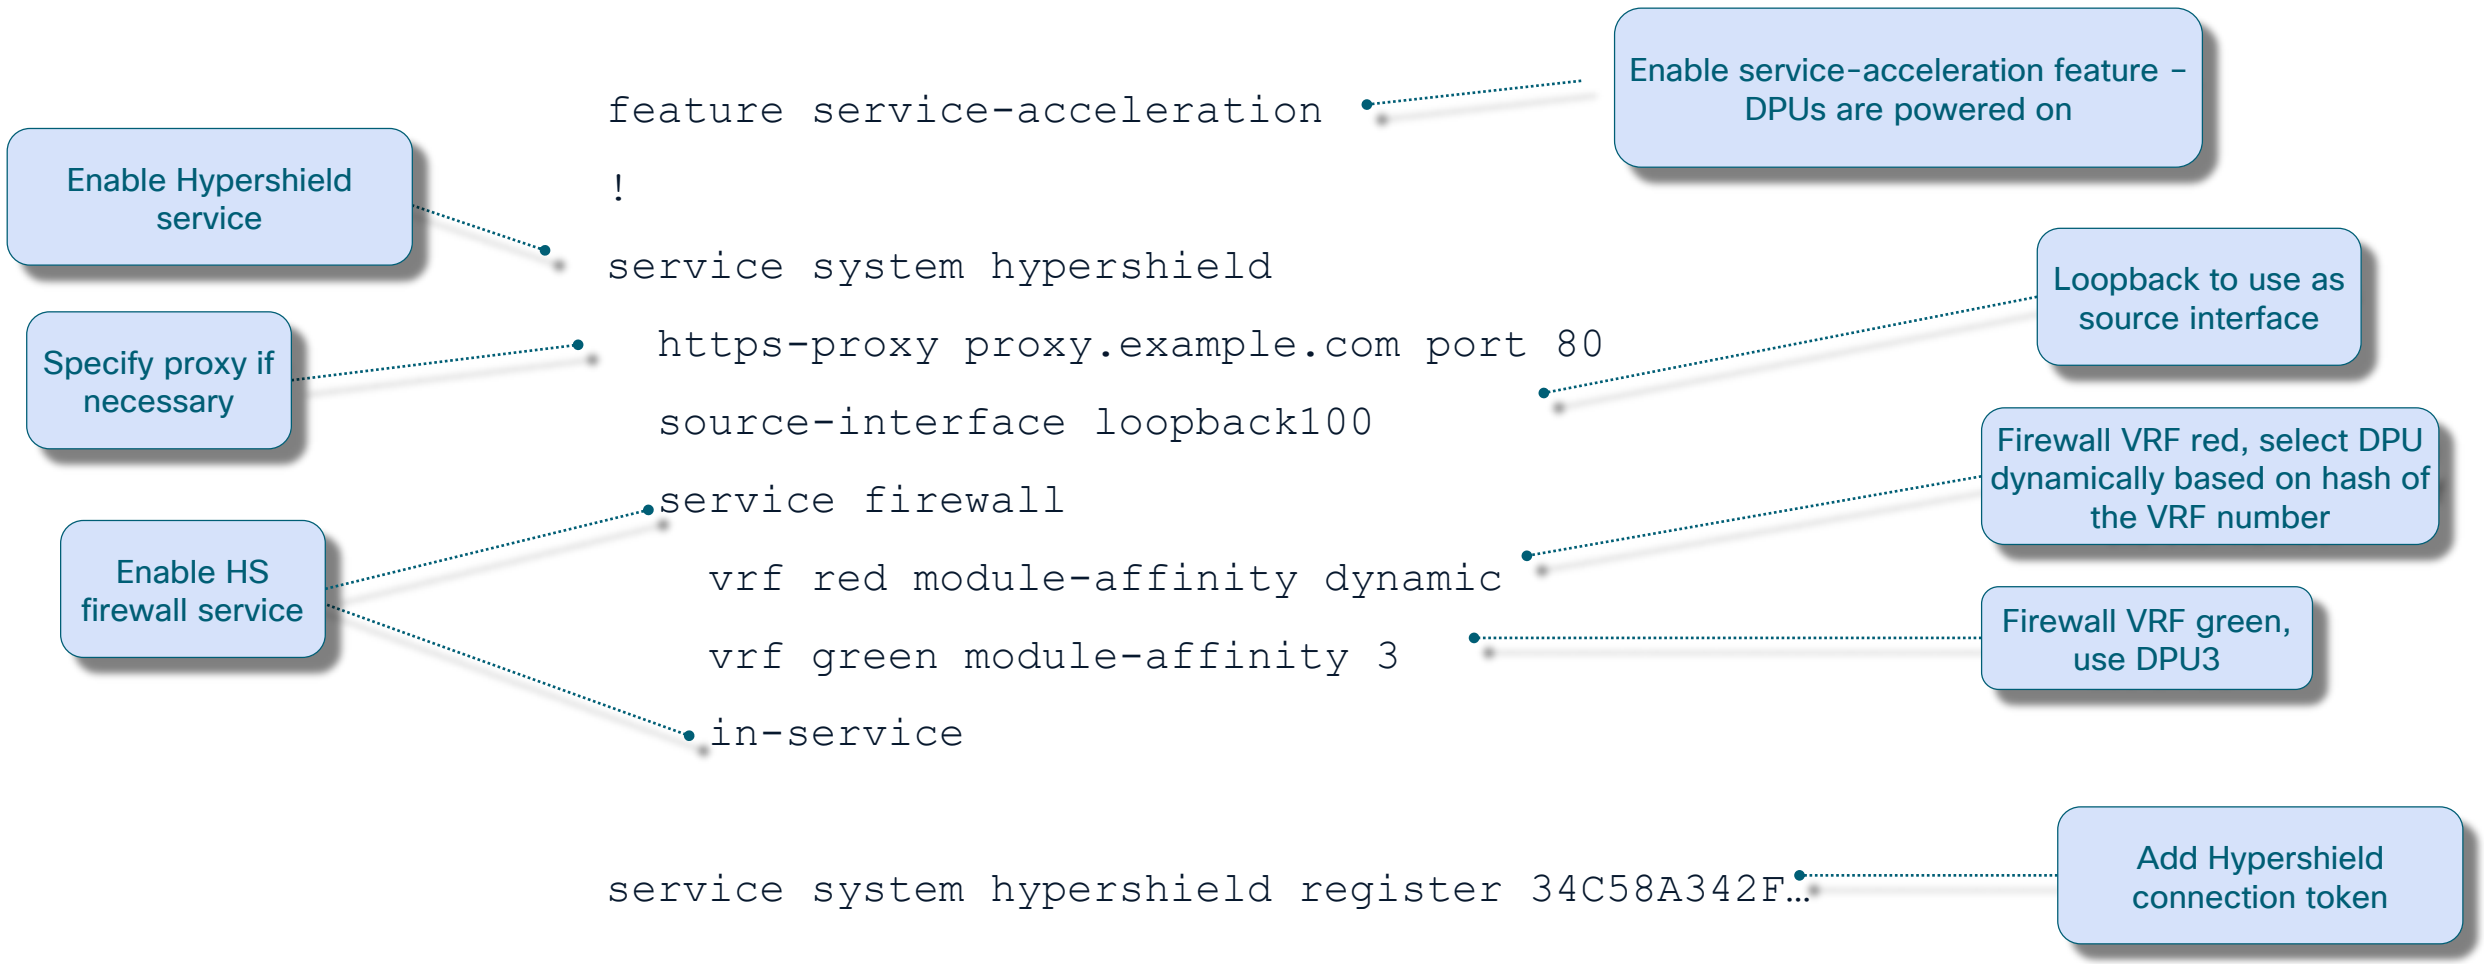

Service Acceleration Configuration for L2 Traffic

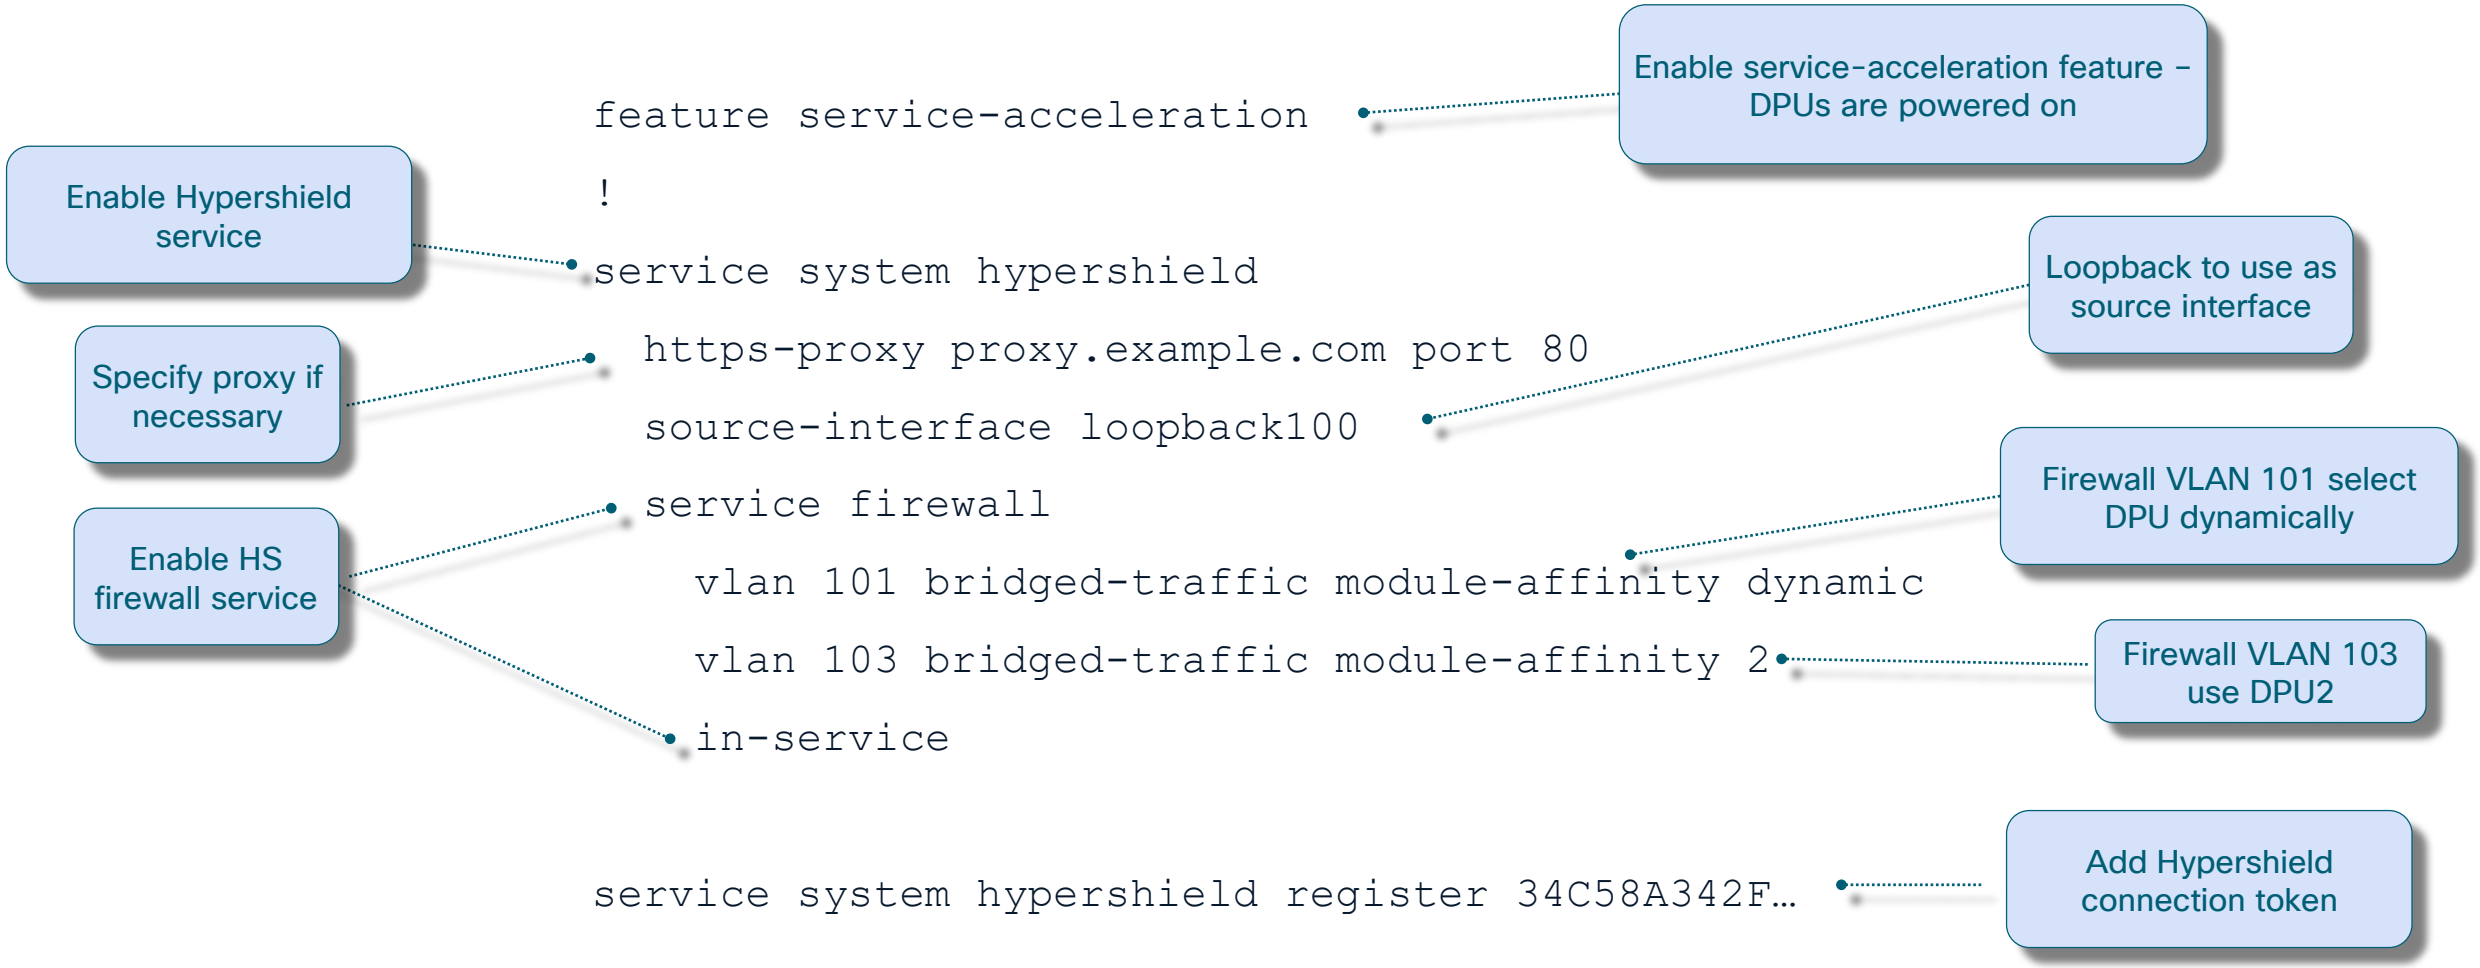

Enable Service Acceleration

- If "feature service-acceleration" is not configured, the DPUs are powered off. The switch functions as a NXOS switch.

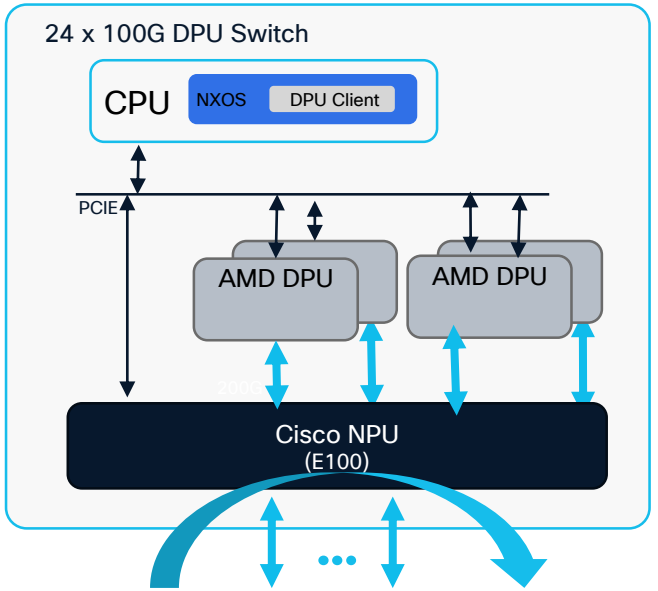

- "feature service-acceleration" enablement powers up the DPUs but to complete the configuration you need to define which VRFs traffic should be redirected to the DPU

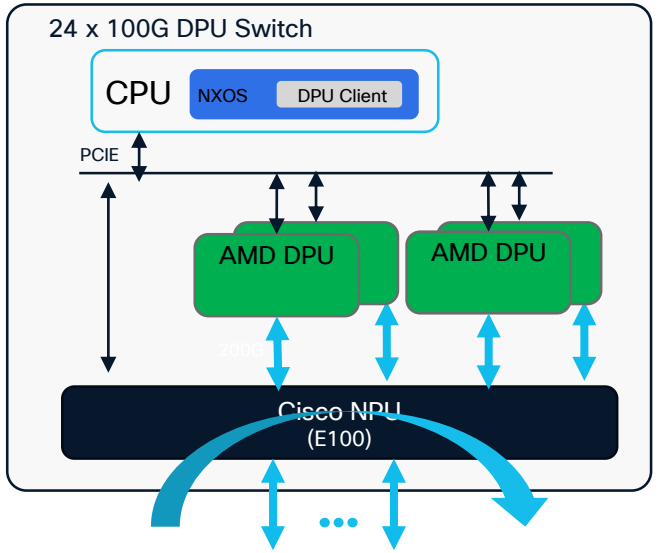

Assign VRF traffic to the DPU Configuration for L3 Traffic

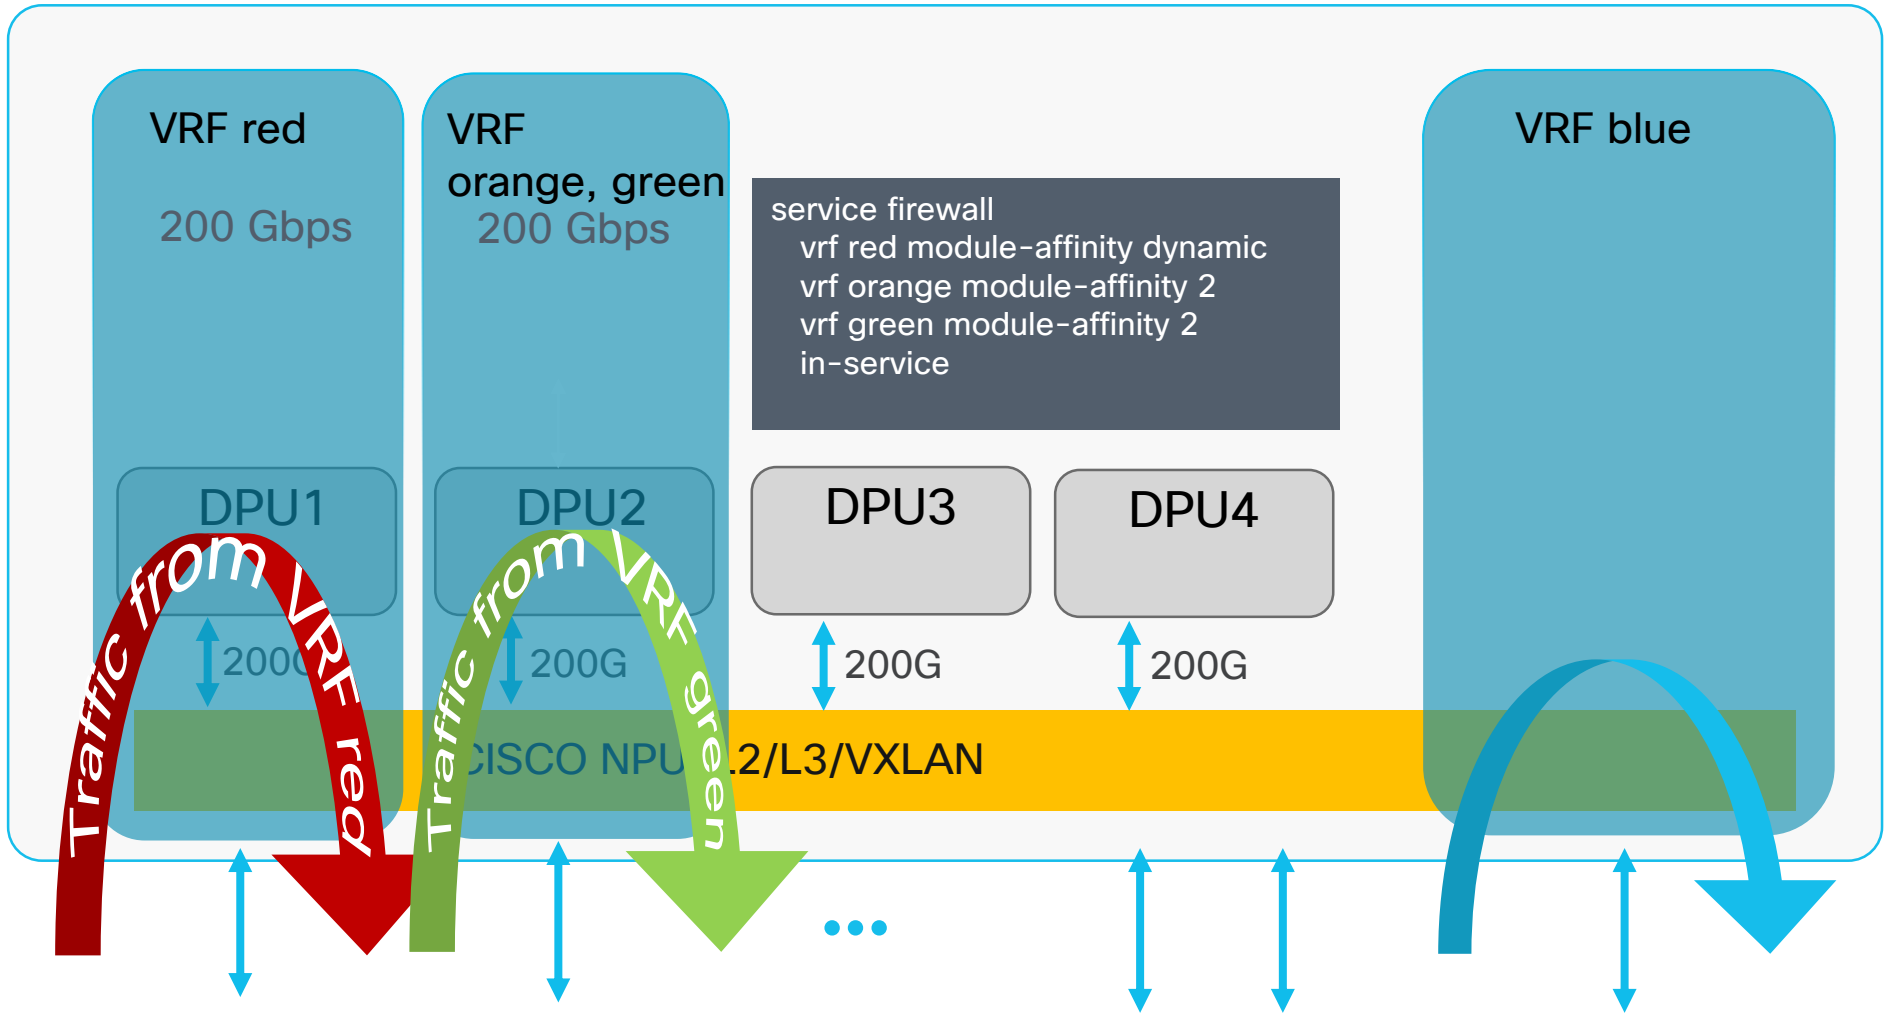

Assign VLAN traffic to the DPU Configuration for L2 Traffic

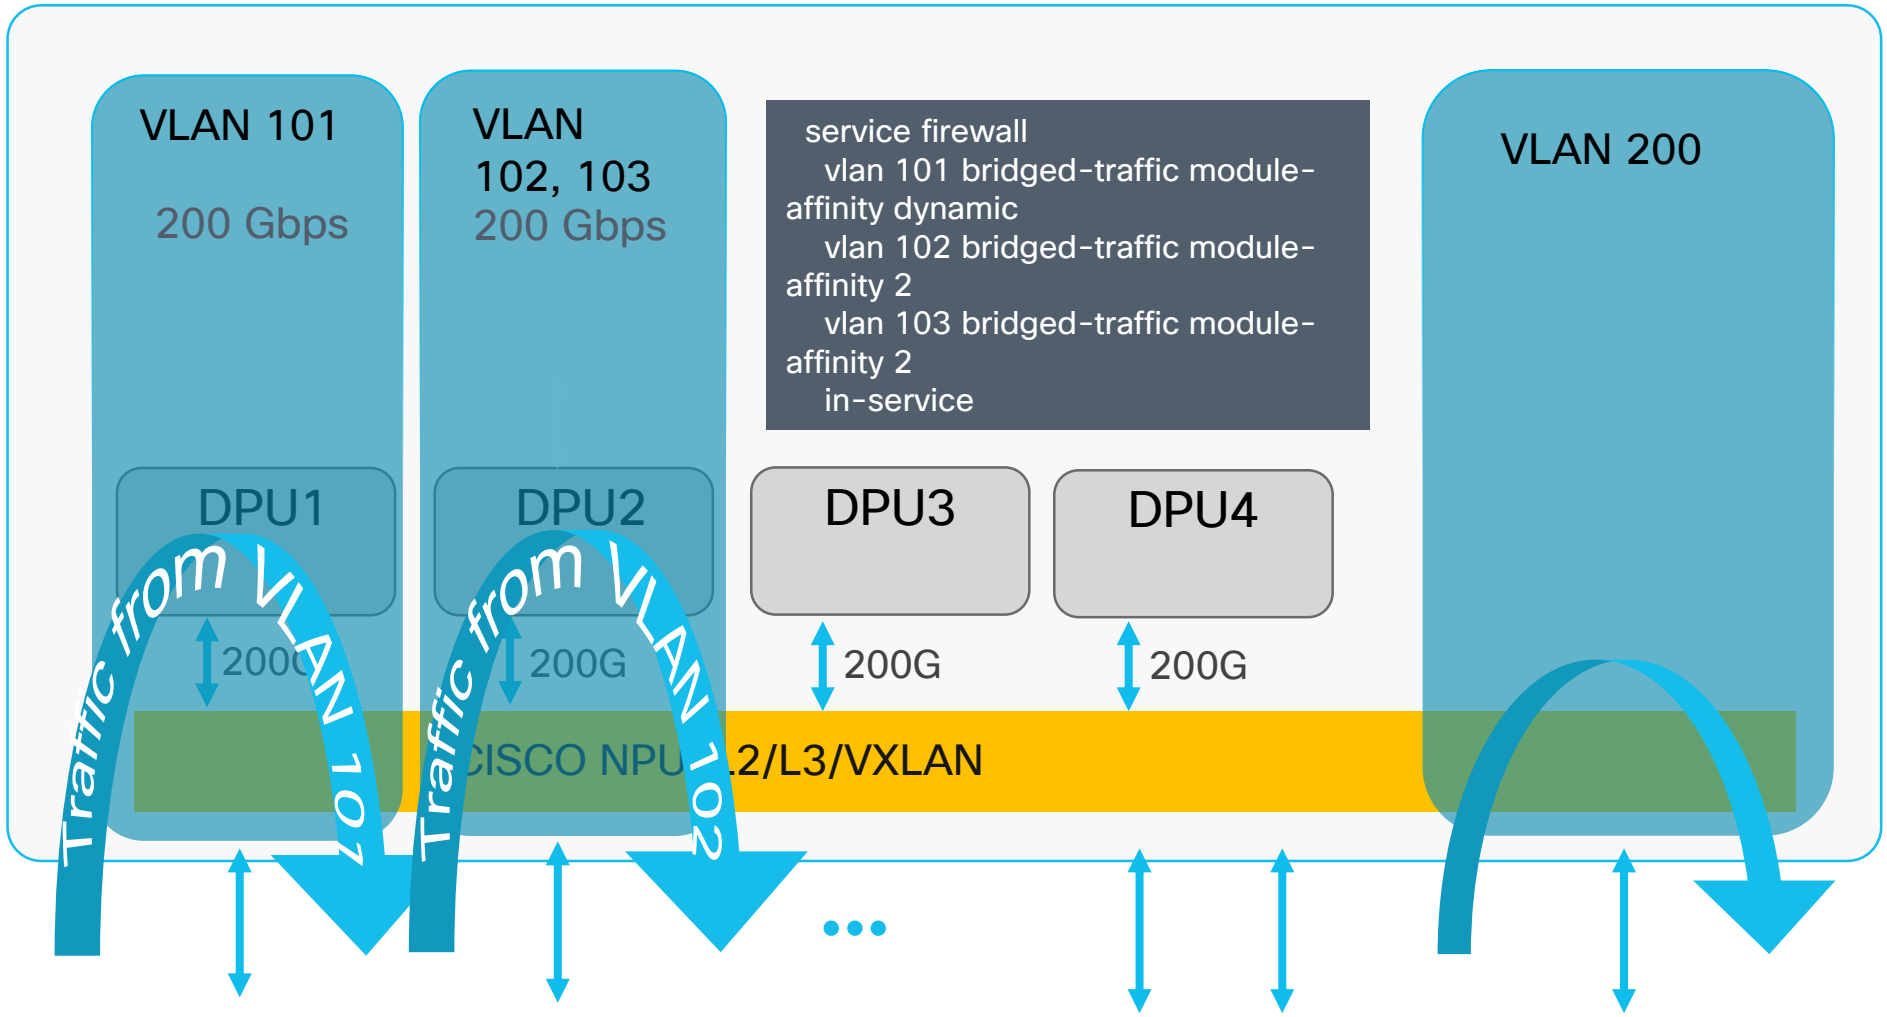

Packet Walk

Smart Switch Represented as a Switch + a Firewall VRF Redirect

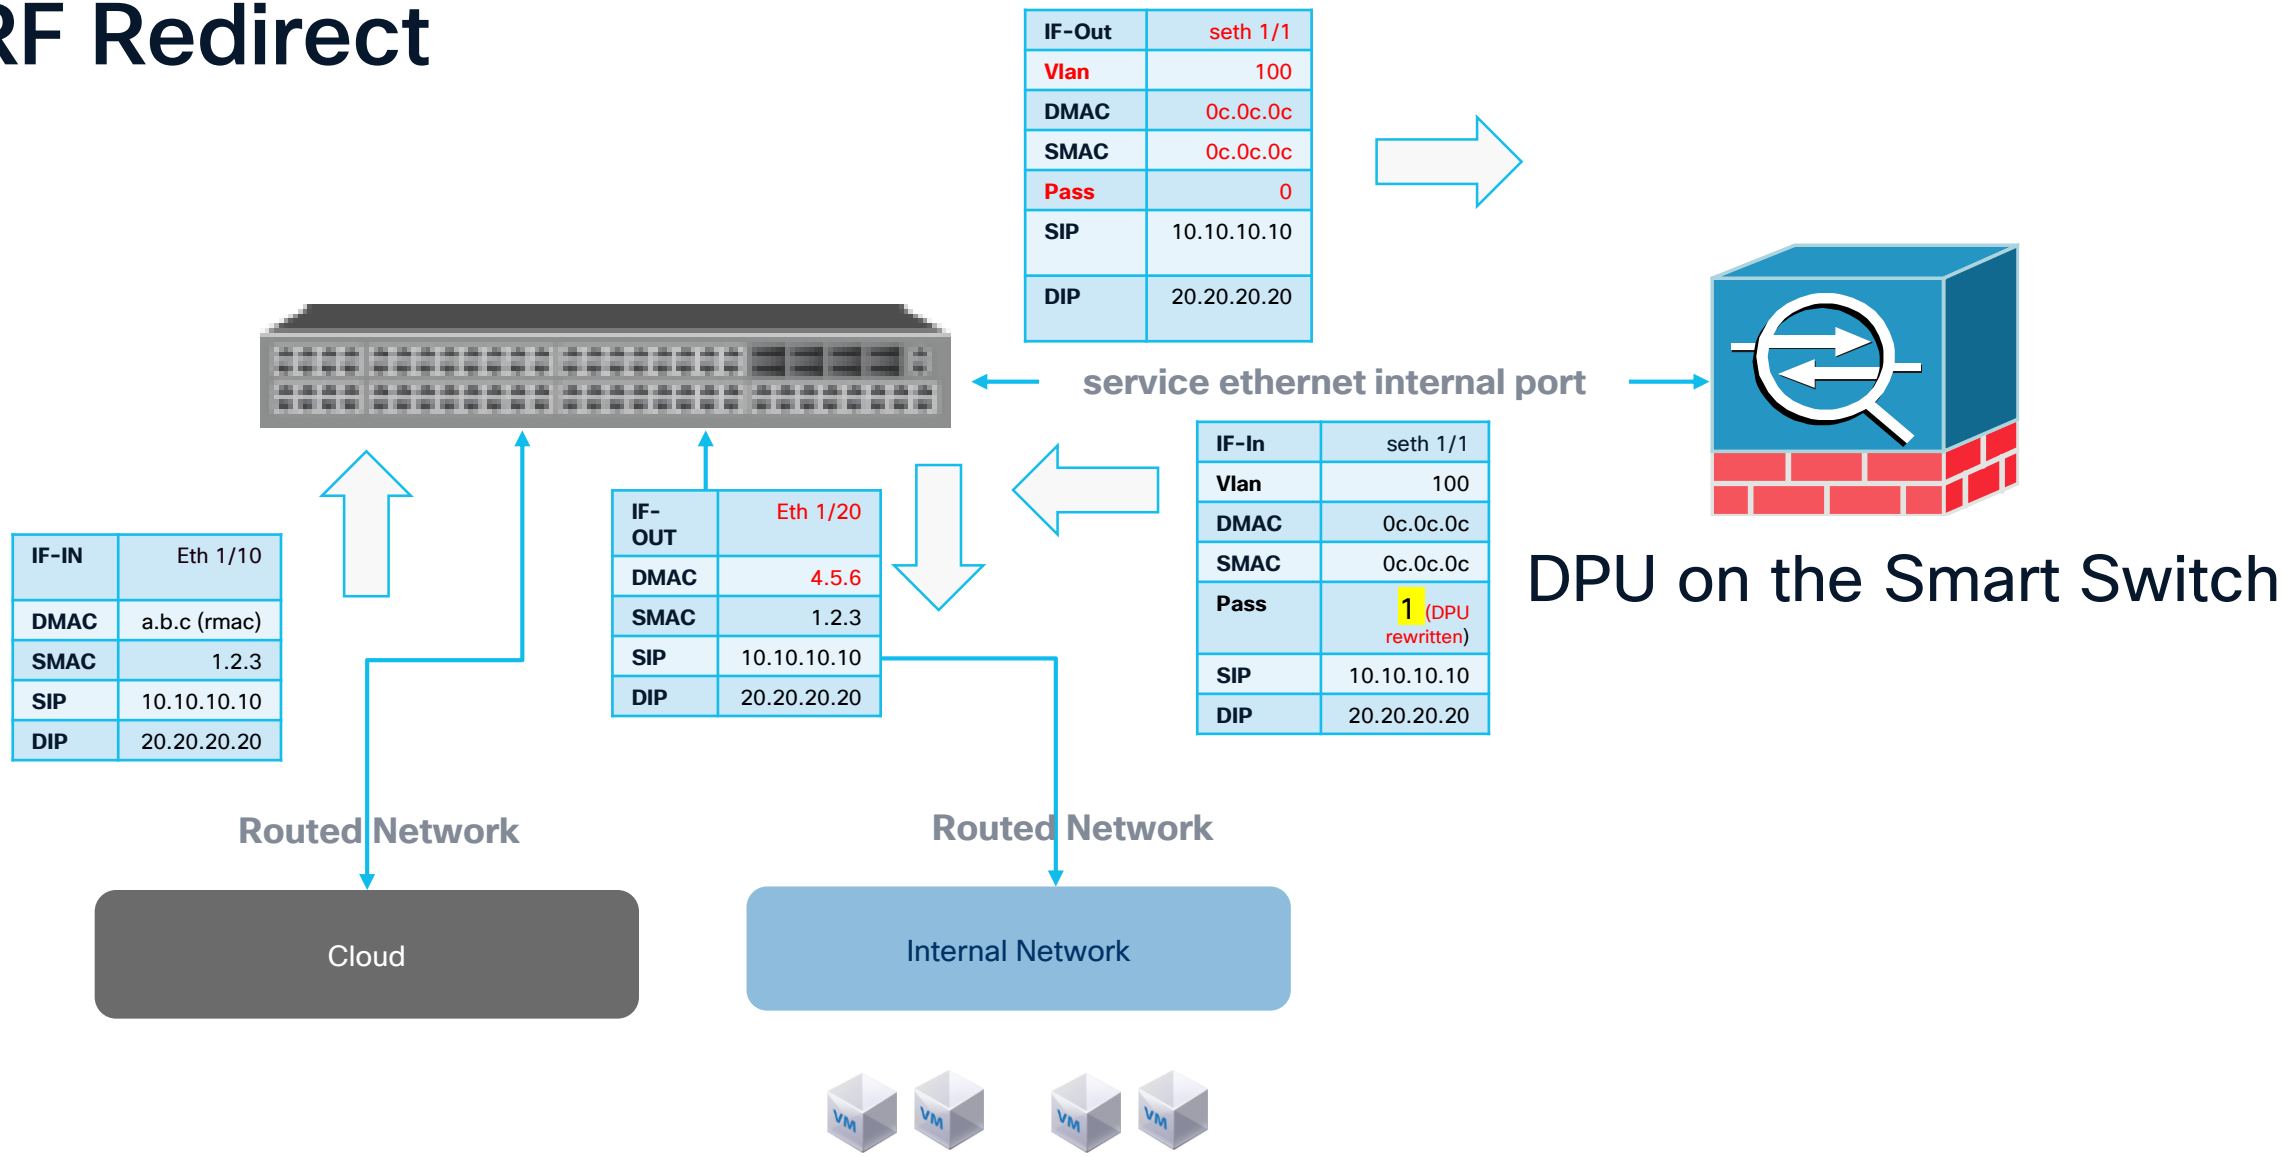

DPU on the Smart Switch

Smart Switch Represented as a Switch + a Firewall
L2 Redirect

IF-Out	seth 1/1
Vlan	10
DMAC	0c.0c.0c
SMAC	0c.0c.0c
Pass	0
SIP	10.10.10.10
DIP	10.10.10.20

VLAN information is preserved

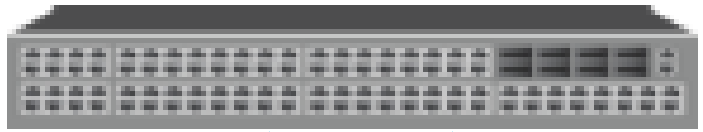

service ethernet internal port

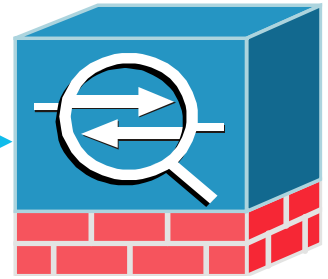

IF-In	seth 1/1
Vlan	10
DMAC	0c.0c.0c
SMAC	0c.0c.0c
Pass	1 (DPU rewritten)
SIP	10.10.10.10
DIP	10.10.10.20

DPU on the Smart Switch

IF-IN	Eth 1/10
VLAN	10
DMAC	4.5.6
SMAC	1.2.3
SIP	10.10.10.10
DIP	10.10.10.20

IF-OUT	Eth 1/20
VLAN	10
DMAC	4.5.6
SMAC	1.2.3
SIP	10.10.10.10
DIP	10.10.10.20

App (Vlan 10)
10.10.10.0/24
(GW: 10.10.10.1)

Redirect from NPU to DPU

VRF-Lite Vlan carried as 802.1q tag, to serve as identifier for VRF Context

NXOS Dynamically Generated Configuration:

```
interface seth1/3.100
  vrf member red
  encapsulation dot1q 100

epbr service __red_dpu_redir type dpu
  vrf red
  service-end-point module 1 vlan 100
```

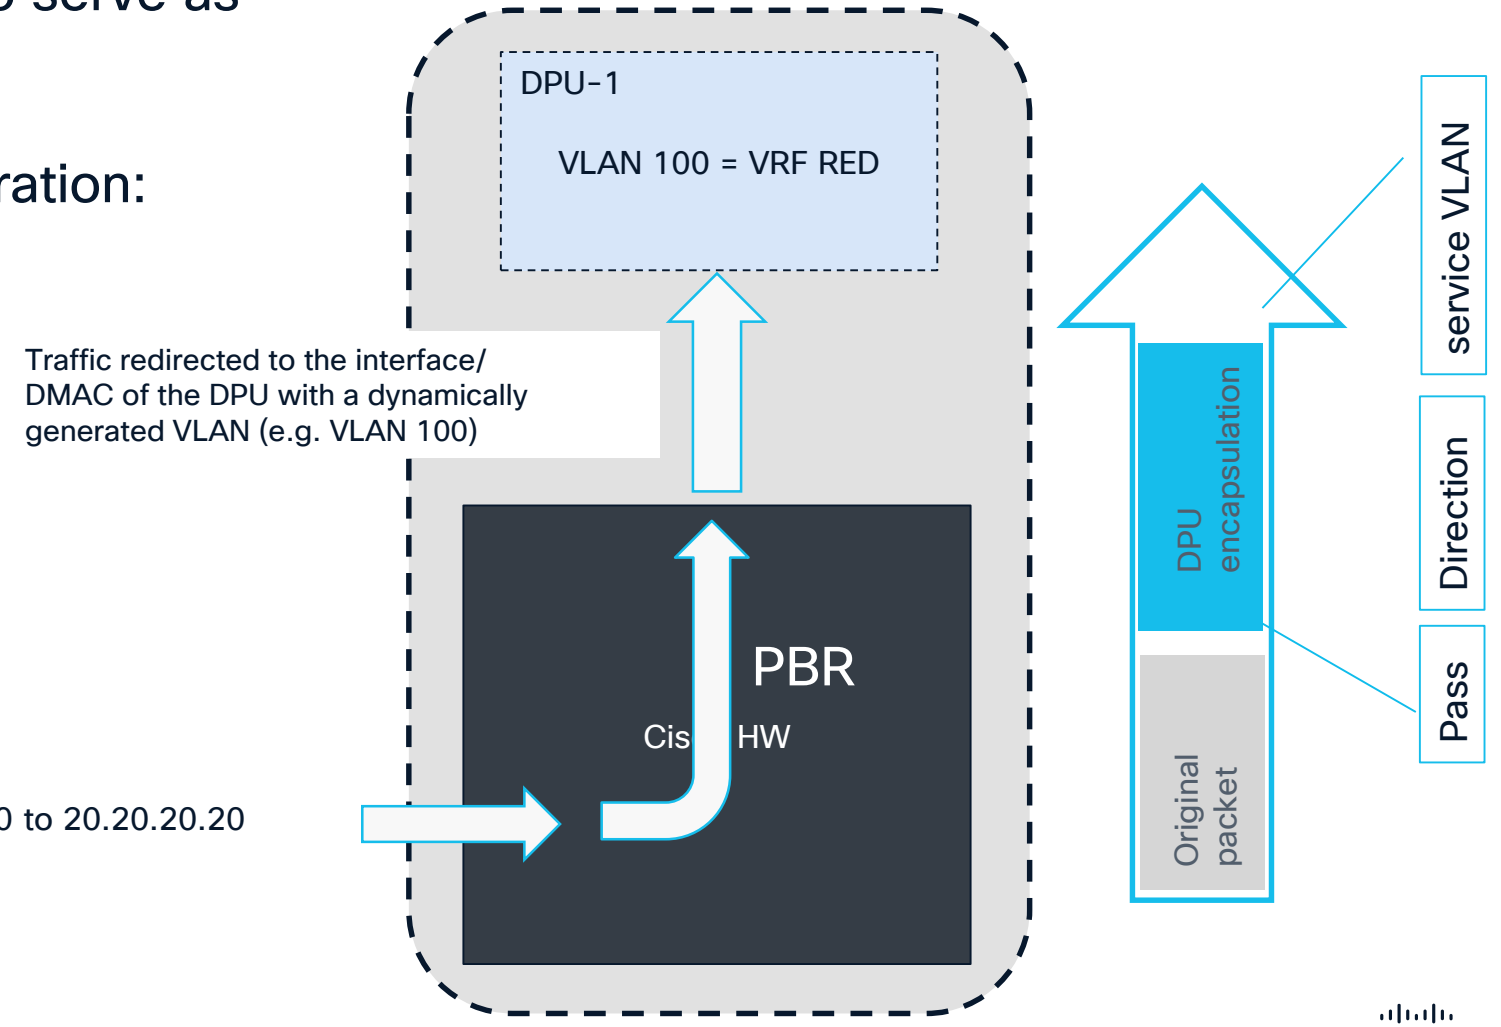

DPU Day in the Life of a Packet

- Packet comes to the DPU P4 pipeline
- If the flow is not found in the flow-table, packets are punted to VPP to setup the flows
- Once the flow is installed, all subsequent packets are forwarded in P4 except:
  - SYN, SYN-ACK, ACK and FIN, FIN-ACK, ACK => VPP
  - These are also synchronized to the peer in case of HA
  - ALG (TFTP/FTP/RPC/RTSP/DNS) – P4 does not support parsing of ALGs, so all the packets are punted to VPP to be parsed.

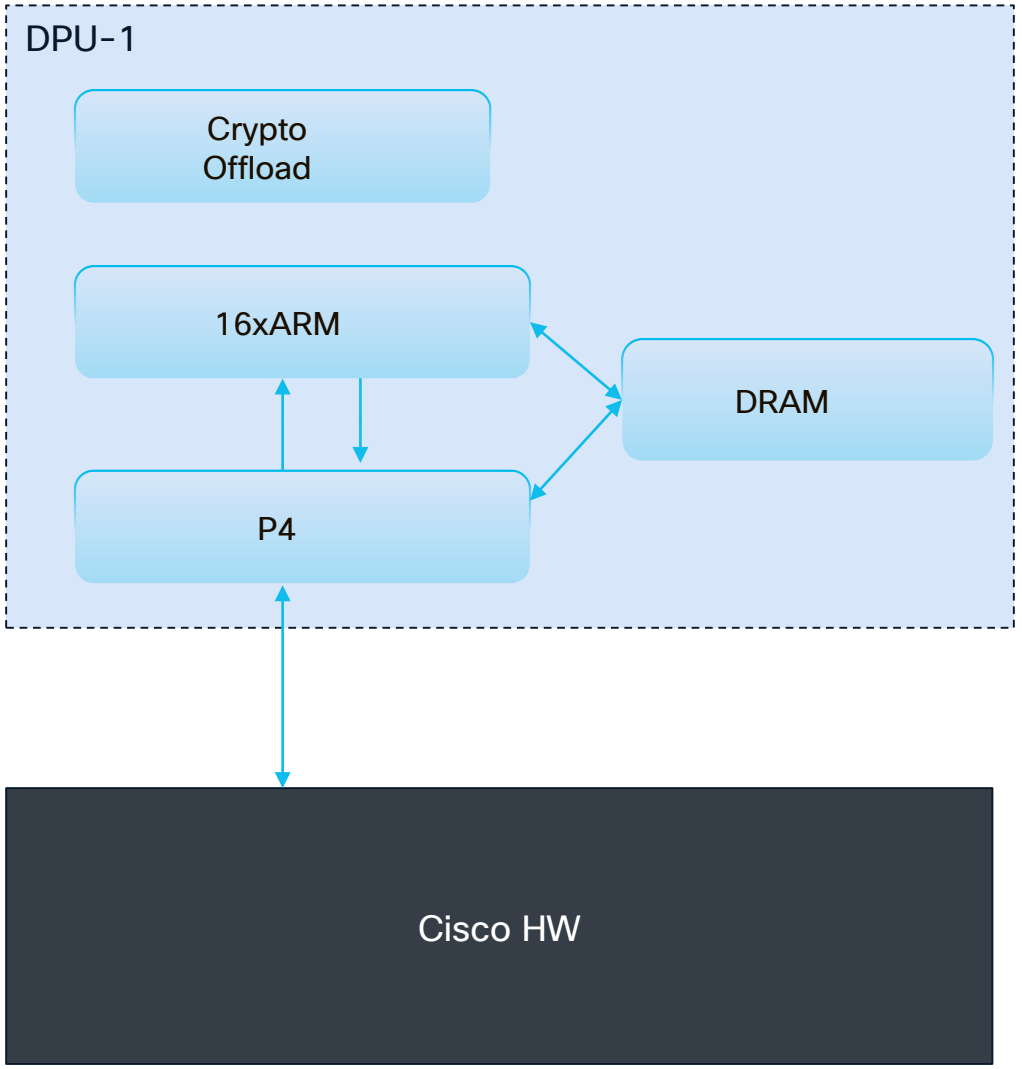

Features by release

Switching Software Features – NXOS 10.5(3s)

N9324C switch support

Feature set with DPUs enabled:

- Layer 3 IPv4/IPv6 forwarding with multi-VRF support
- VRF-based redirection to DPUs
  - No inter-VRF flow (same VRF in and out)
- Supported interface types – routed port, routed subinterface, routed port-channel
- Supported routing protocols – BGP, IS-IS, OSPF, EIGRP, static
- DPU lifecycle management (software upgrade, DPU health, etc.)

Key Features in the Next Release

- Stateful HA
- L2 support
- SVIs
- First Hop Redundancy Protocols
- vPC
- Integration with Nexus Dashboard

Use cases

What are the Smart Switches positioned for?

- Cloud Edge
- Zone-based Firewall
- Datacenter Interconnect
- and
- Top of the Rack deployments
- All the designs that we present here are assuming the use of the future NXOS 10.6(1s) release

Cloud Edge

- 10.5(3s) deployments would need to ensure traffic symmetricity via ECMP
- 10.6(1s) will introduce Active/Active stateful HA

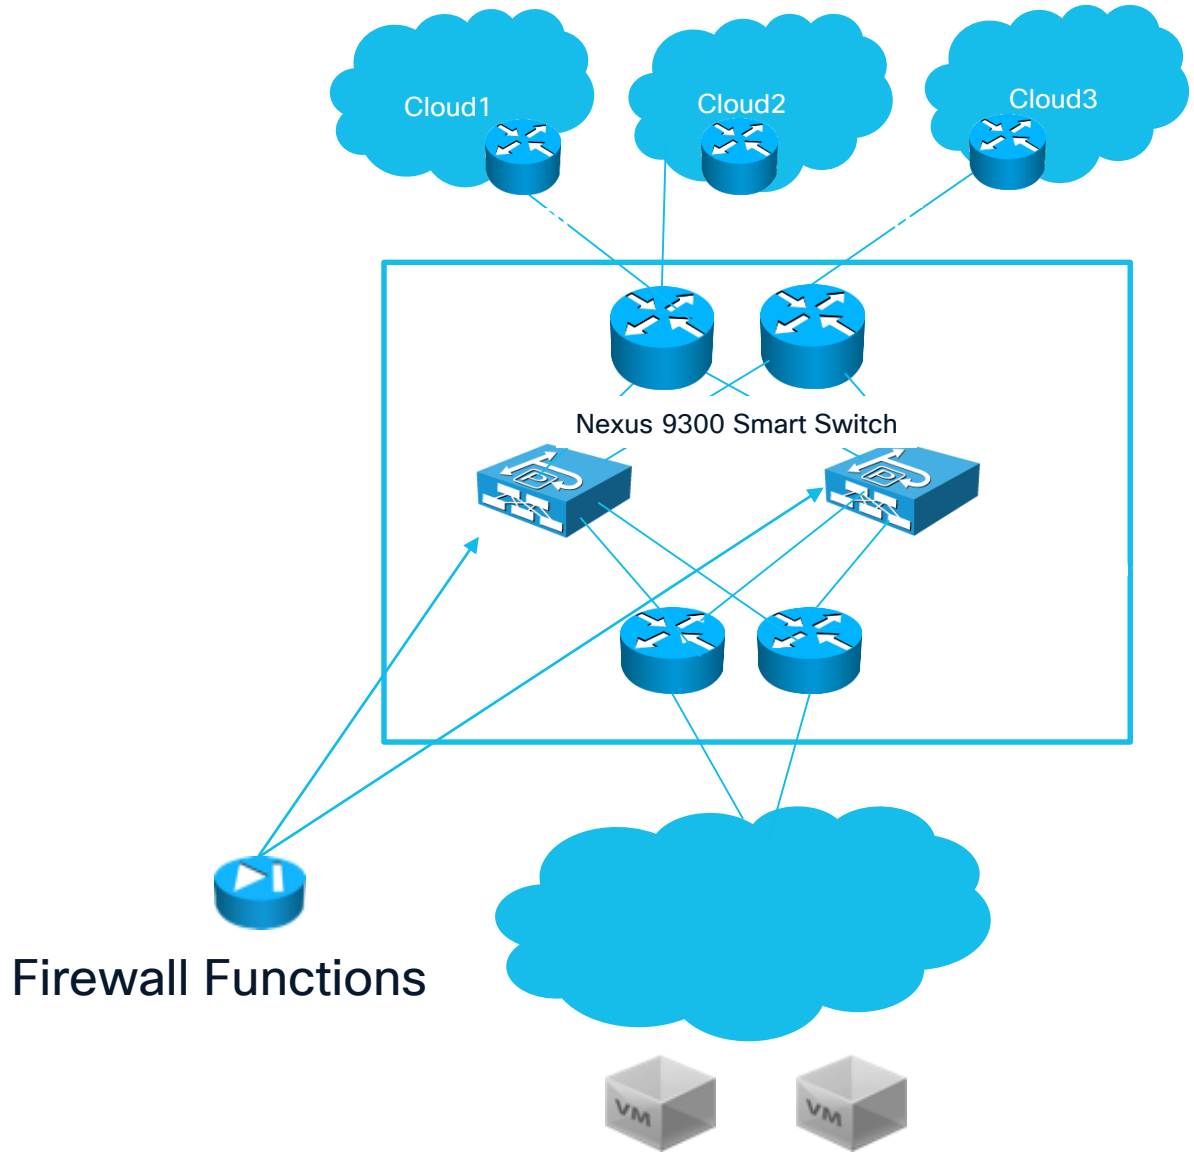

Zone Based Firewall Deployments

- Using the Smart Switch as a Firewall
- Inserting it into a Fabric with VXLAN EVPN or ACI
- For instance sending the traffic to the Smart Switch with Policy Based Redirect
- This requires 10.6(1s) for HA
- 10.6(1s) is not yet available

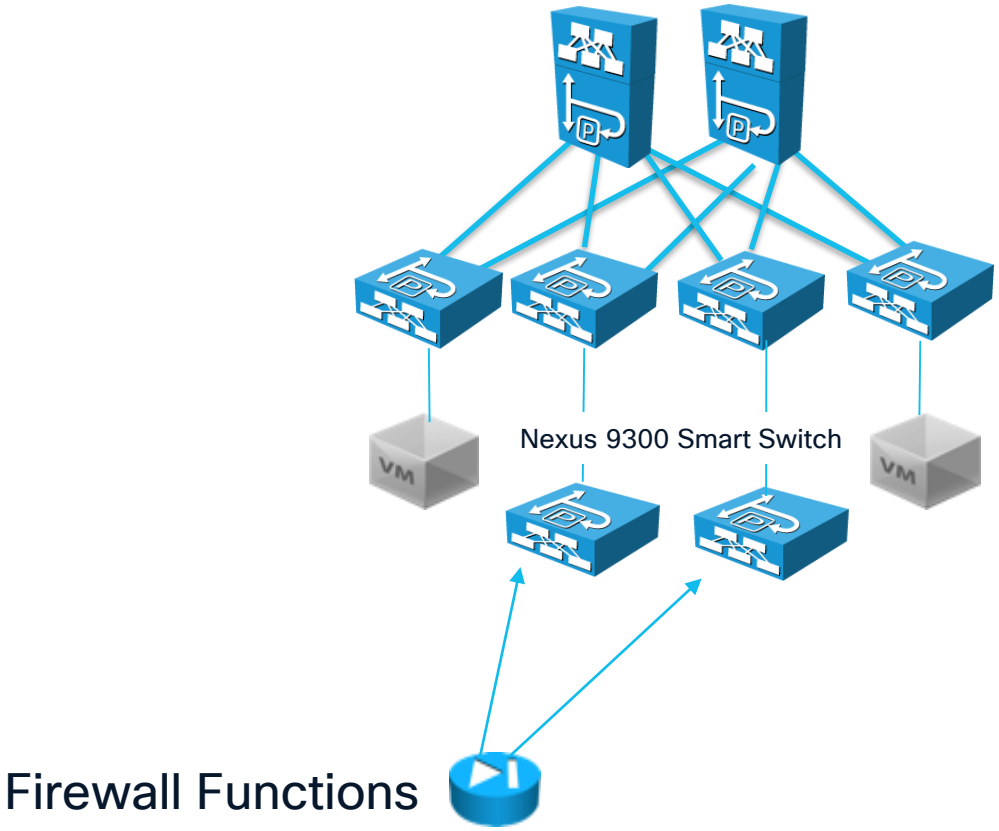

Nexus 9300 Smart Switch as a 2-arms firewall

- Starting with 10.6(1s) it will be possible to deploy the Nexus 9300 Smart Switch as a redundant pair with Active/Active stateful HA and with support for HSRP.
- This requires 10.6(1s) for HA
- 10.6(1s) is not yet available

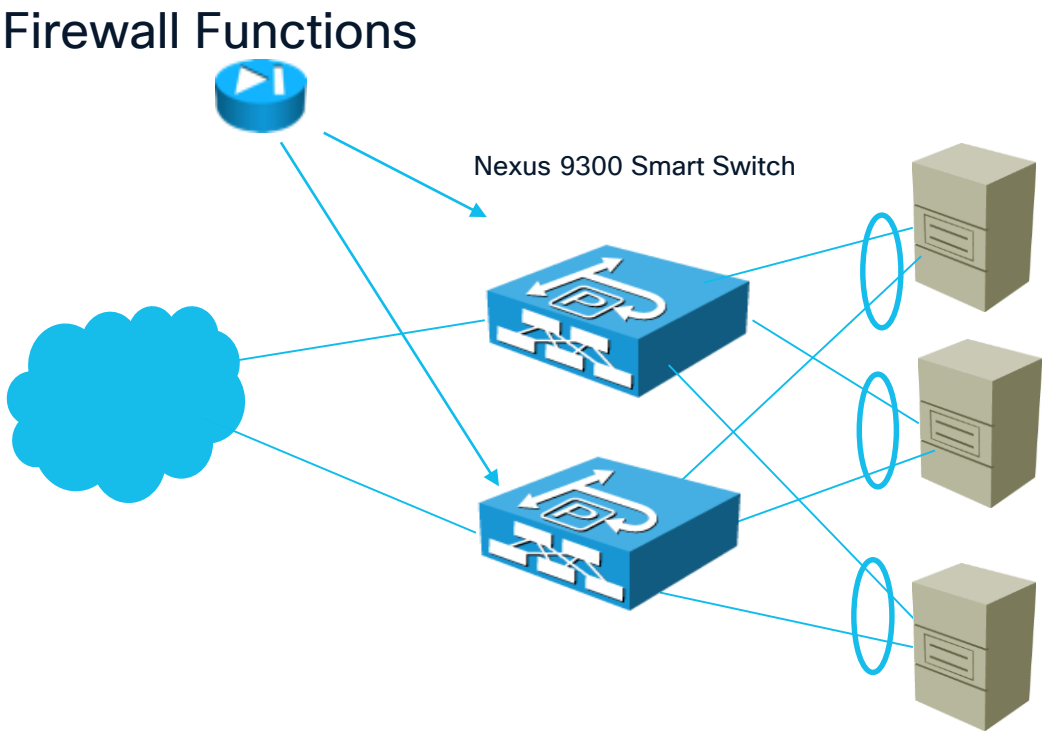

Top of the Rack Deployment in a Layer 3 Fabric Mobility scope: vPC Pair

This drawing is not a VXLAN EVPN fabric

This design requires 10.6(1s), not yet available

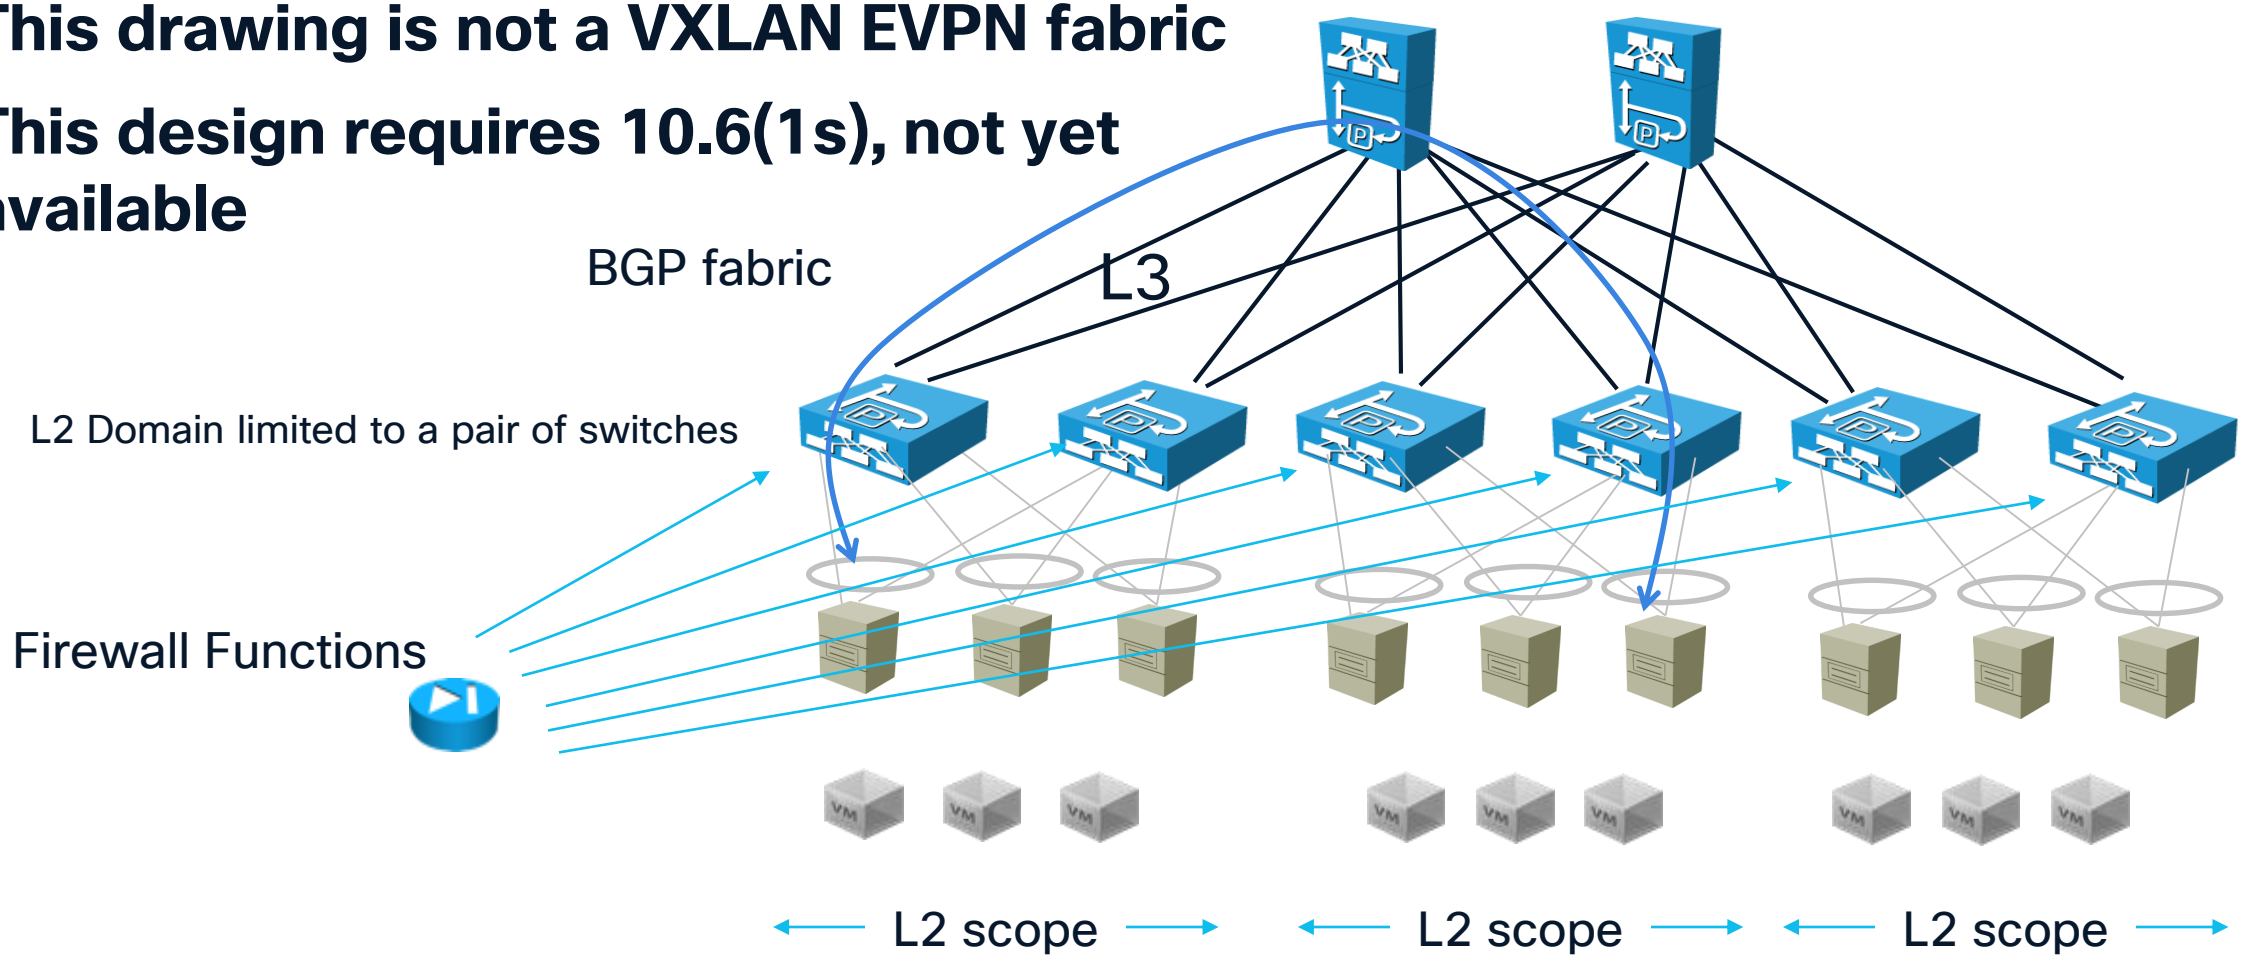

Target Scale

DPU Firewall Policy Scale Limits
Disclaimer: Pending QA Validation

N9324C-SE1U

Scale Parameter	Value
Throughput	~800Gbps
Number of FW sessions	32M per box (*)
Hypershield PARC policies	600K (pending QA validation)
Number of VRFs (**)	20 in the initial release, 100 at GA (256 per DPU in theory)
Connections Per second	80k per DPU

(*) flows are distributed among DPUs:
- each DPU can handle 8M flows (4M per direction) with the 9324C * 4 DPUs = ~32M per box
- Exact numbers are subject to QA validations but
- Dual Dataplane requirements are already factored in these calculations

(**) Hypershield doesn't support VRFs in the initial release

Related Sessions

Related Sessions

- Cisco Hypershield: Mastering Next-Generation Security - BRKSEC-2265
- Hypershield in a Nutshell: What it is and how it works - CISCOU-2059
- Cisco N9300 Smart Switch - CTF-1052
- Fusing Security Into the Network with Cisco Smart Switches - WOSSEC-2002
- Secure Data Center Networking | Cisco N9300 Smart Switch, Hypershield, and Nexus Dashboard - DEMAIDC-14
- Fusing Security and Networking | Hypershield, Smart Switch, and Hybrid Mesh Firewall - DEMAIDC-15
- Secure Data Center with N9300 Smart Switches - WOSDCN-1004
- Data Center Security In The Age Of AI for partners - Dive in Cisco Smart Switches and Hypershield - WOSGPE-1002
- A Look Inside the AI-Ready Secure Network: Powering the Next Decade of Innovation - KDDENT-1000

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 2026.

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

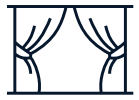

Visit the Cisco Showcase for related demos

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

Thank you

cisco Live !


CISCO