



Připravte své datové centrum na příští dekádu

2. 2. 2022 | 9:00–12:30

Possibilities

Networking modul

2. 2. 2022 | Připravte své datové centrum na příští dekádu



Q&A

slido.com

#160 378

Passcode: 43r05p

Agenda

- EoS Nexus 5000, 2000
- Jednotný networking pro všechny workloady
- SAN zrychluje na 64G
- Nexus Dashboard – pomocník provozu DC sítě



Q&A

slido.com

#160 378

Passcode: 43r05p

EoS Nexus 5000, 2000



Q&A

slido.com

#160 378


Passcode: 43r05p

Nexus je mrtev! Ať žije Nexus!

- Od 5.5.2021 byl ukončen prodej posledních Nexus 5000 (5672UP, 5672UP-16G, 5696Q)
 - Konec SW maintenance 5.5.2022 (bezpečnostní patche do 4.5.2024)
 - Náhradníci jsou v řadě Nexus 9300 (93180YC-FX, 9336C-FX2-E)
- Podobný osud čeká i poslední zástupce Nexus 2000 (2248TP, 2348TQ, 2348UPQ) k datu 24.8.2022
 - Konec SW maintenance 24.8.2023 (bezpečností patche do 23.8.2025)
 - Náhrada opět v řadě Nexus 9300 (9348GC-FXP, 93108TC-FX3P, 93180YC-FX3)

Jak dál s FEX архитектурou?

- FEX architektura již nevyhovuje
 - Limitovaná propustnost
 - Vyšší latence
 - Škálovací limity, podporované topologie
- FEX řešil problémy, které již dnes řešíme jinak
 - Leaf – Spine topologie + kontroler (APIC, DCNM) = modulární přepínač
 - Obecná automatizace (Ansible, Terraform, atp.), programovatelnost (Python, NX-API)
- Výhody
 - Plný výkon, nižší latence, funkcionalita a škálovatelnost
 - Lepší flexibilita




Nexus Switches with Unified Ports

Industry unique

Any Storage Transport on the same Switch

- Block storage over FC
- Block storage over FCoE
- Ethernet
- Block Storage over iSCSI
- Block storage over NVMe/TCP
- Block storage over RoCE
- File storage over TCP/IP/Ethernet
- Object storage over TCP/IP/Ethernet



Phased migration to/from IP/Ethernet block storage from/to FC or FCoE

IP/Ethernet block storage can be iSCSI, RoCEv2 or NVMe/TCP

Nexus Switches with Unified Ports

Nexus 9336C-FX2-E
(LS 3600 FX2 Heavenly)



36 QSFP ports with breakout support
92 UP or 8/16/32G FC ports

Nexus 93360YC-FX2
(LS 3600 FX2 Heavenly)



96 x UP Ports – 1/10/25G Ethernet/FCoE or 4/8/16/32G FC
12 X QSFP Uplink Ports 40/100G Ethernet/FCoE

Nexus 93180YC-FX
(LS1800FX (Homewood))



48 x UP Ports – 1/10/25G Ethernet/FCoE or 4/8/16/32G FC
6 X QSFP Uplink Ports 40/100G Ethernet/FCoE

Nexus 9336C-FX2-E – FC/FCoE details

- FC support available in NX-OS 10.2(2)
- Densest FC switch with 92 line-rate 32G FC ports in 1RU
- 4G FC not available on N9336C-FX2-E
- Port 1 – 8 and 32 - 36 can't be converted to FC
 - 23 QSFP ports can be converted to FC resulting in connection to 92 devices
- QSFP ports work with Finisar 128G QSFP (DS-SFP-4X32G-SW). Two options are possible
 - MPO12-MPO12 cable for QSFP-to-QSFP port (Use-case: ISL connectivity at 4 x 32G FC)
 - MPO12-4LC Break-out cable for 4 lanes of 8/16/32 GFC breakout (Use-case: end-device connectivity)
- All the 4 lanes of the QSFP must be operating at the same speed (4 x 8G or 4 x 16G or 4 x 32G)
- F port at 8GFC not supported. 8GFC ISLs are supported.



36 QSFP ports with breakout support
92 UP or 8/16/32G FC ports

Nexus 9336C-FX2-E
(LS 3600 FX2 Heavenly)

QSFP specs: <https://ii-vi.com/product/128g-fibre-channel-parallel-100m-mmfc-gen2-qsf28-optical-transceiver/>

SAN designs using Nexus 9000 in NX-OS mode

FC Ports + NPIV parent

- Hosts – HBA & NIC
- N9K host ports – FC & Ethernet
- N9K uplink ports – FC for SAN, Ethernet for LAN
- Must have NPIV parent if N9K in NPV mode

FC/FCoE Ports + NPIV parent





- Hosts - CNA
- N9K host ports – FCoE
- N9K uplink ports – FC or FCoE depending on upstream SAN switch, Ethernet for LAN
- Must have NPIV parent if N9K in NPV mode

FC Ports + FC Switching

- Hosts - HBA & NIC
- N9K host ports - FC & Ethernet
- N9K uplink ports – FC for SAN, Ethernet for LAN
- NPIV parent not needed
- N9K must be in FC switching mode to directly connect to storage array


FC/FCoE Ports and Switching

- Hosts – CNA or HBA
- N9K host ports – FC/FCoE
- N9K uplink ports – FC/FCoE for SAN, Ethernet for LAN
- NPIV parent not needed
- N9K must be in FCoE switching mode to directly connect to storage array




SAN designs using Nexus 9000 in ACI mode

FC Ports + NPIV parent



- Hosts – HBA & NIC
- N9K host ports - FC & Ethernet
- N9K uplink ports – FC for SAN, Ethernet for LAN
- Must have NPIV parent

FC/FCoE Ports + NPIV parent



- Hosts - CNA
- N9K host ports – FCoE
- N9K uplink ports – FC or FCoE depending on upstream SAN switch, Ethernet for LAN
- Must have NPIV parent

- Follow the same design principles as N9K in NX-OS mode
- Must need a NPIV parent (MDS recommended) for FC services (zoning, device-aliases, FCNS, etc.)

For details of FC and FCoE on Nexus 9000

- Cisco Nexus 9000 Series NX-OS SAN Switching Configuration Guide
- Cisco Nexus 9000 Series NX-OS FC-NPV and FCoE-NPV Configuration Guide
- Cisco APIC Layer 2 Networking Configuration Guide
- Release notes
- Refer to [Cisco MDS 9000 Family Pluggable Transceivers Data Sheet](#) for FC SFP+ specifications, supported distance, cabling, etc.

FC/FCoE License requirements on Nexus 9000

Subscription

License	1G Fixed Platforms (GF)	10G/25G/40G/100G Fixed Platforms (XF)	Modular Platforms
Storage	Not Supported	C1-N9K-STRG-XF-3Y C1-N9K-STRG-XF-5Y	C1-N9K-STRG-M-3Y C1-N9K-STRG-M-5Y

NVMe/FC on Nexus 9000

- All Cisco switches with FC ports are certified for NVMe/FC
- All Cisco switches with FCoE ports are certified for NVMe/FCoE
- No config is required on Nexus and MDS switches for NVMe/FC and NVMe/FCoE
 - The capability of the end-devices is automatically detected

Jednotný networking v DC

CISCO *Engage*



Q&A

slido.com

#160 378

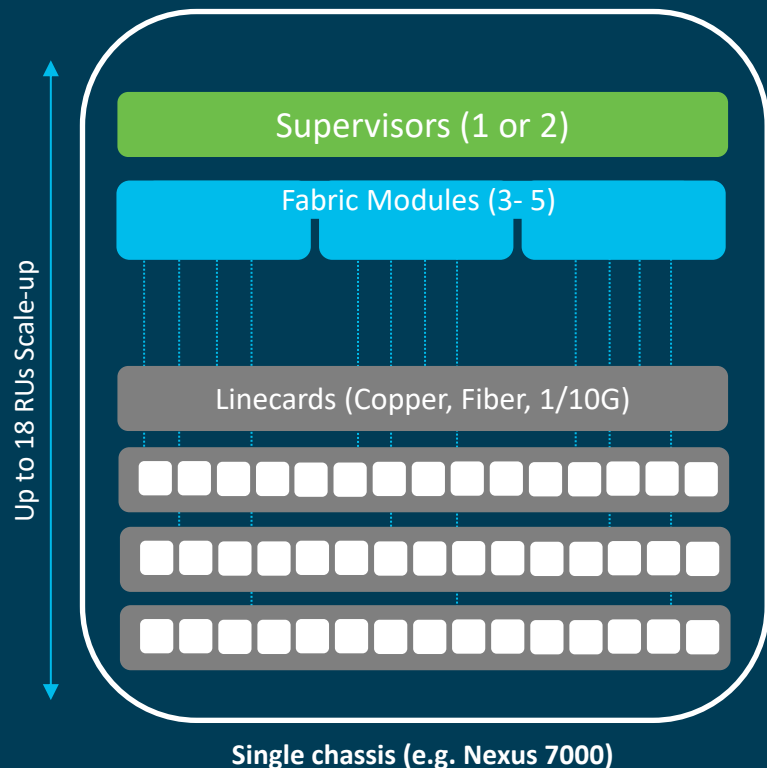
Passcode: 43r05p

ACI: One Network, any location



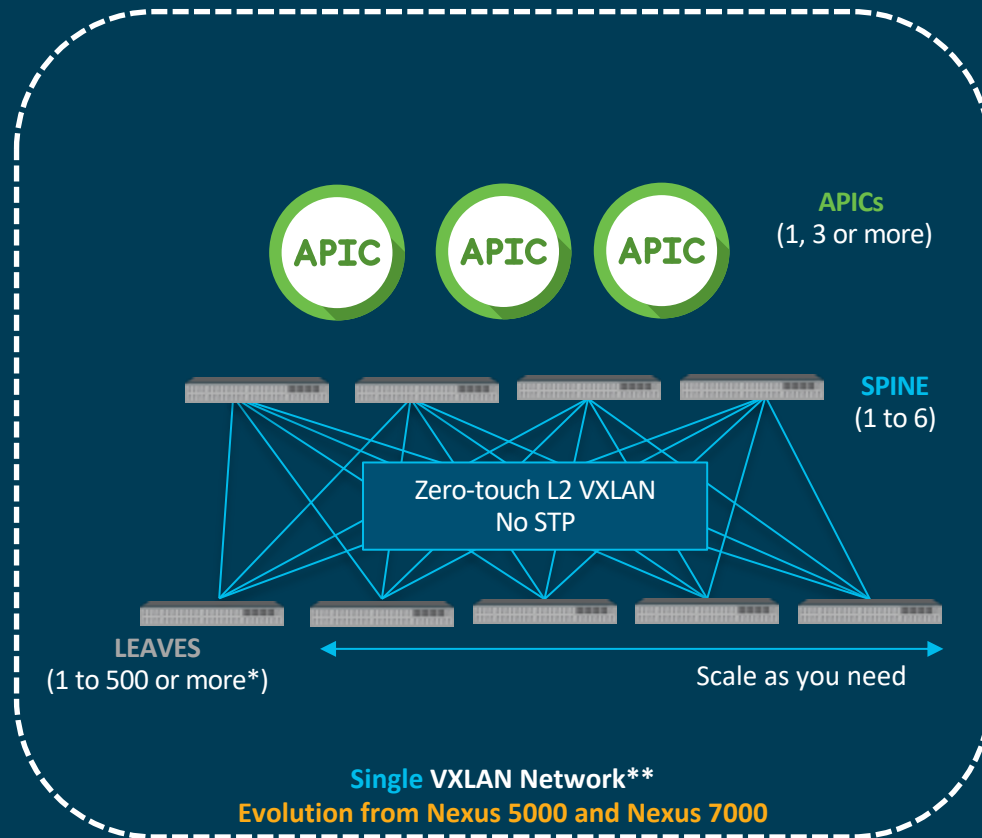
The DC network **before**

Classic modular switching



The DC network **NOW**

ACI



ACI: The on-prem elements

APICs

3 Recommended for **Production**
At least 1 **physical** APIC required

SPINES

LEAVES

Virtual/Container networking integration
included

LICENSING

Only applies to **Physical** Leaves
(no licensing on APICs nor Spines*)

Physical

L-Size (Recommended for
1200+ physical leaf ports)
M-size (Recommended for
<1200 physical leaf ports)

Virtual

VMWare VMs
(Recommended for 2-4 leaves, 2
VMs + 1 Physical APIC)

Modular



Nexus 9500
(w/9700 LCs)

Fixed
(NX-OS
Capable)

Nexus 9300
(9332C, 9364C
9316D-GX)

vmware Hyper-V



Fixed
(NX-OS
Capable)

Nexus 9300
(100M/1/10/25/40/50/100/400G)

Premier

Assurance & Insights

Advantage

Multisite
Remote-Leaf

Essentials

Single/Multi-pod

Add-ons

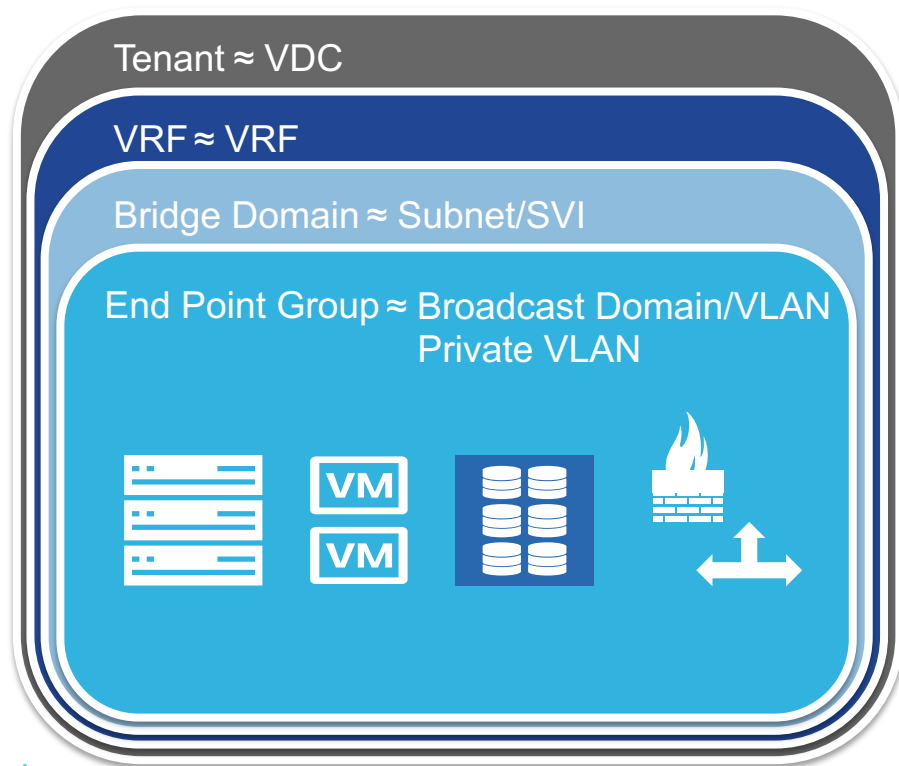


Storage



Encryption

ACI Policy Model



Contracts ≈ Access Lists



L2 External EPG ≈ 802.1q Trunk

L3 External EPG ≈ L3 Routed Link

Three approaches to using EPGs in ACI

EPG/BD = VLAN

Create a BD and one EPG for each existing VLAN.

Common strategy to lift-and-shift traditional configurations.

Simpler for migration, complex for Micro Segmentation.

EPG = App Tier

Create one EPG for each application Tier.

Flat-network design, many apps can share a single BD and IP subnet.

Fantastic for Green Field and automated deployments.

Hybrid (Combination)

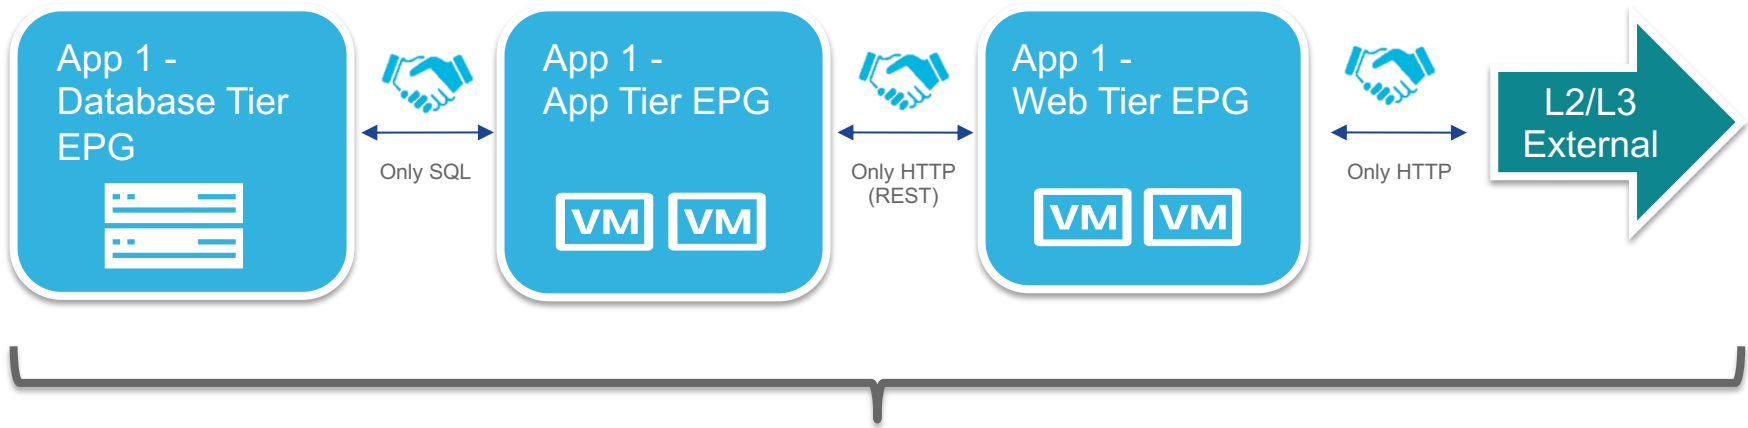
New Apps and Legacy Apps share the same Fabric.

Tenant and VRF sharing.

or

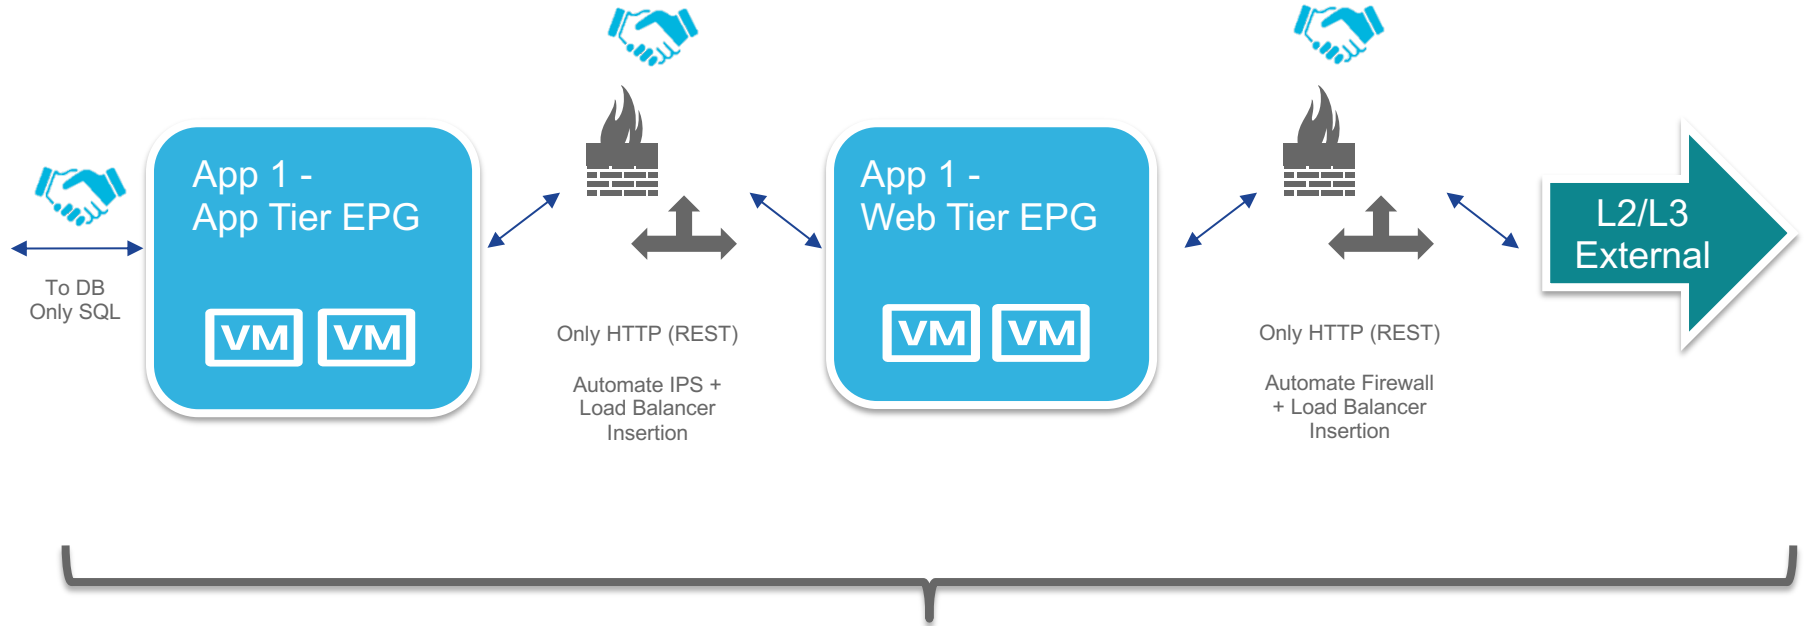
Dedicated Tenant/VRF and leaking.

Advanced ACI Policy Model – Micro Segmentation



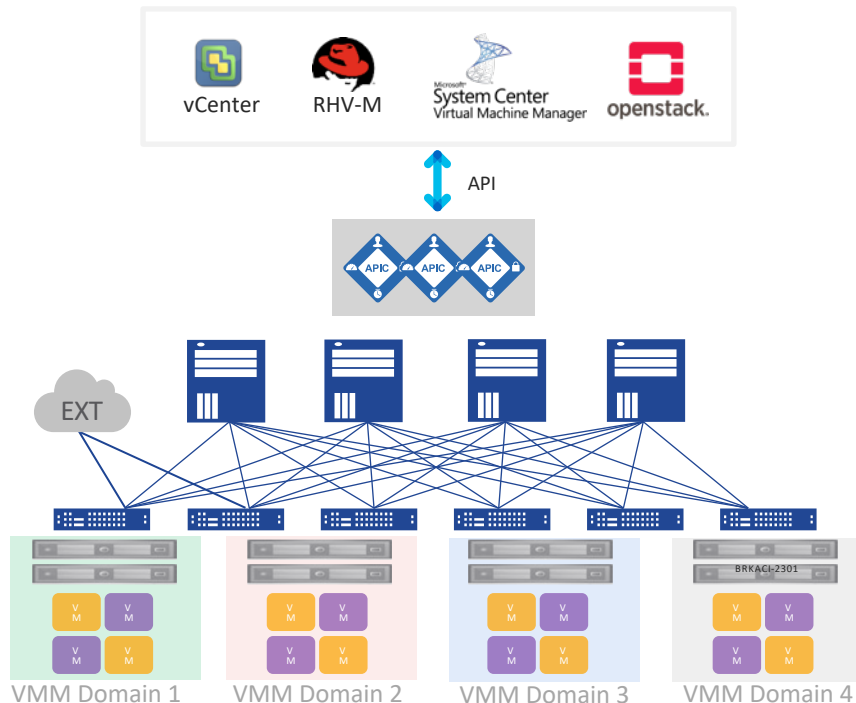
Application Profile

Advanced ACI Policy Model – Service Insertion



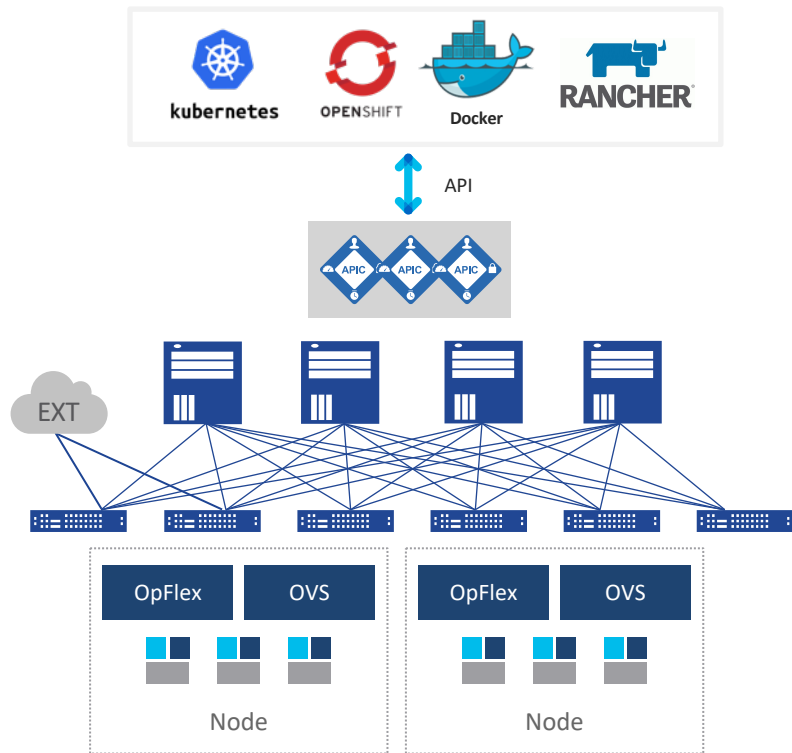
Application Profile with Service Graphs

Virtual Machine Manager (VMM) Domains for Hypervisor Integration



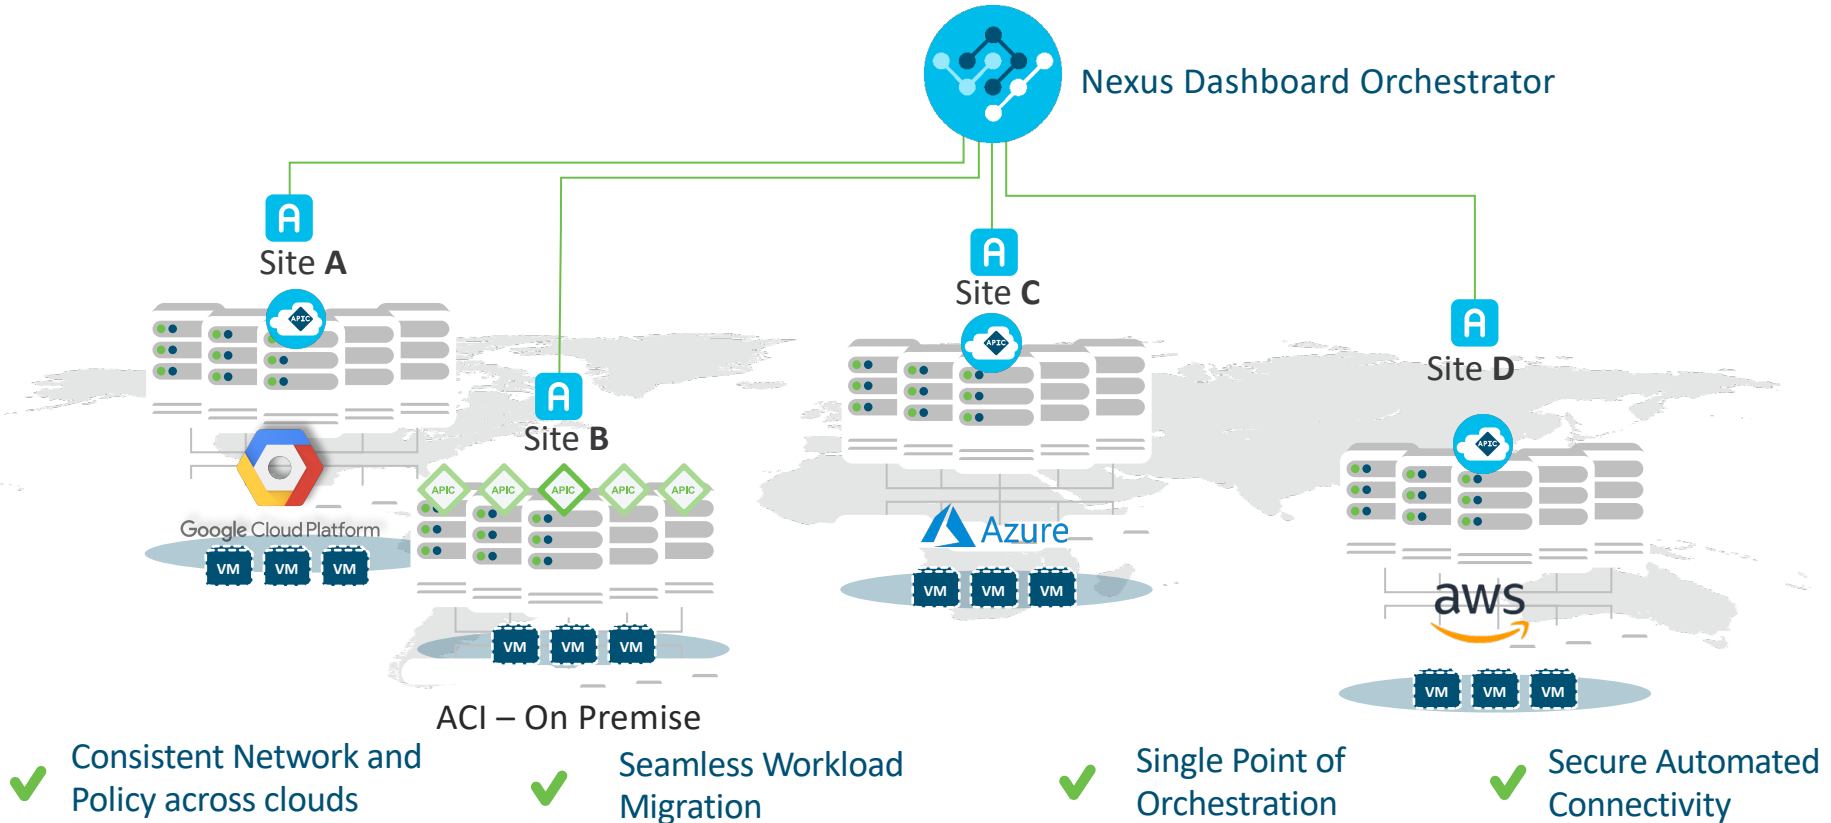
- API Relationship is formed between APIC and Virtual Machine Manager (VMM)
- APIC obtains virtualization inventory, performs virtual and physical correlation
 - APIC has visibility of hypervisor hosts, Virtual Machines, vNICs and more.
- APIC manages the hypervisor virtual switch (VDS, OVS, etc ...) via either API or Opflex
- Multiple VMMs supported on a single ACI Fabric simultaneously

Containers / PaaS Integration with ACI CNI Plugin



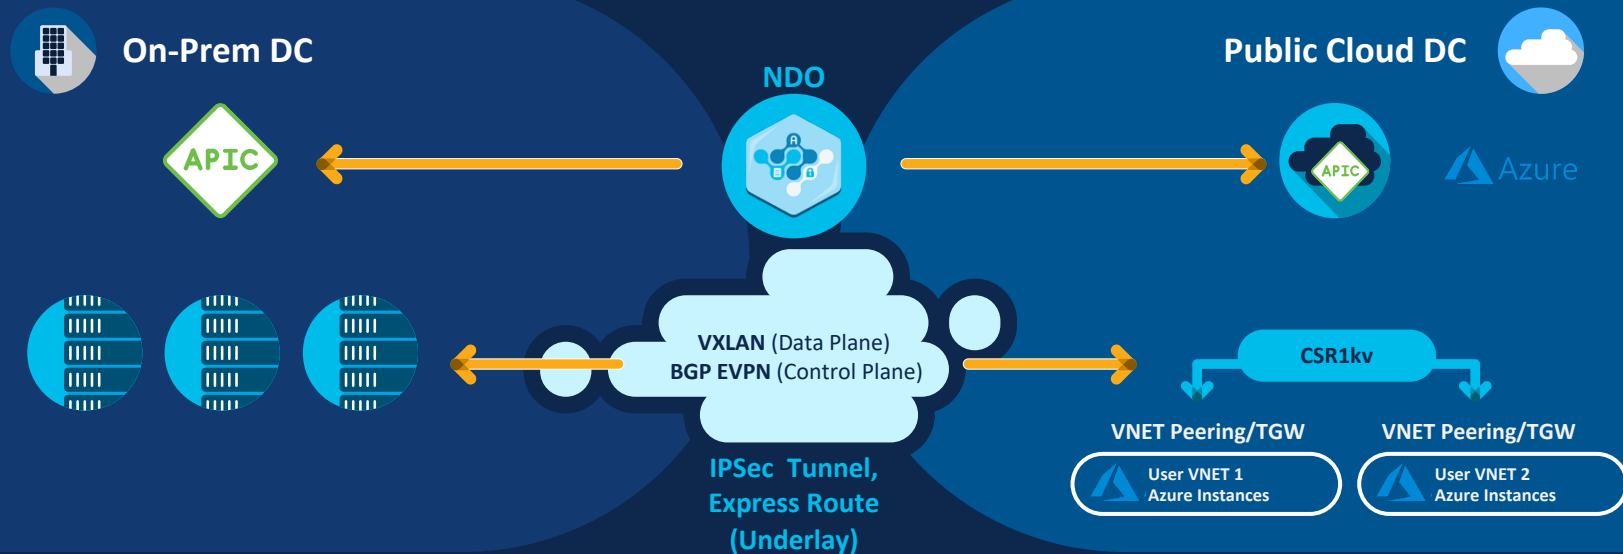
- Integration with kubeapi server, APIC obtains inventory of nodes and Kubernetes objects.
- ACI CNI Plugin Implements:
 - Distributed OVS Load Balancer for *ClusterIP* services.
 - Distributed HW Load Balancer for *LoadBalancer* services.
 - Distributed Kubernetes Network Policies
- Secure multi-tenancy and ACI policies
- Visibility: Live statistics in APIC per container, health metrics

ACI Anywhere



Cloud ACI - Extensions to Azure

Simplified Multi-cloud Connectivity



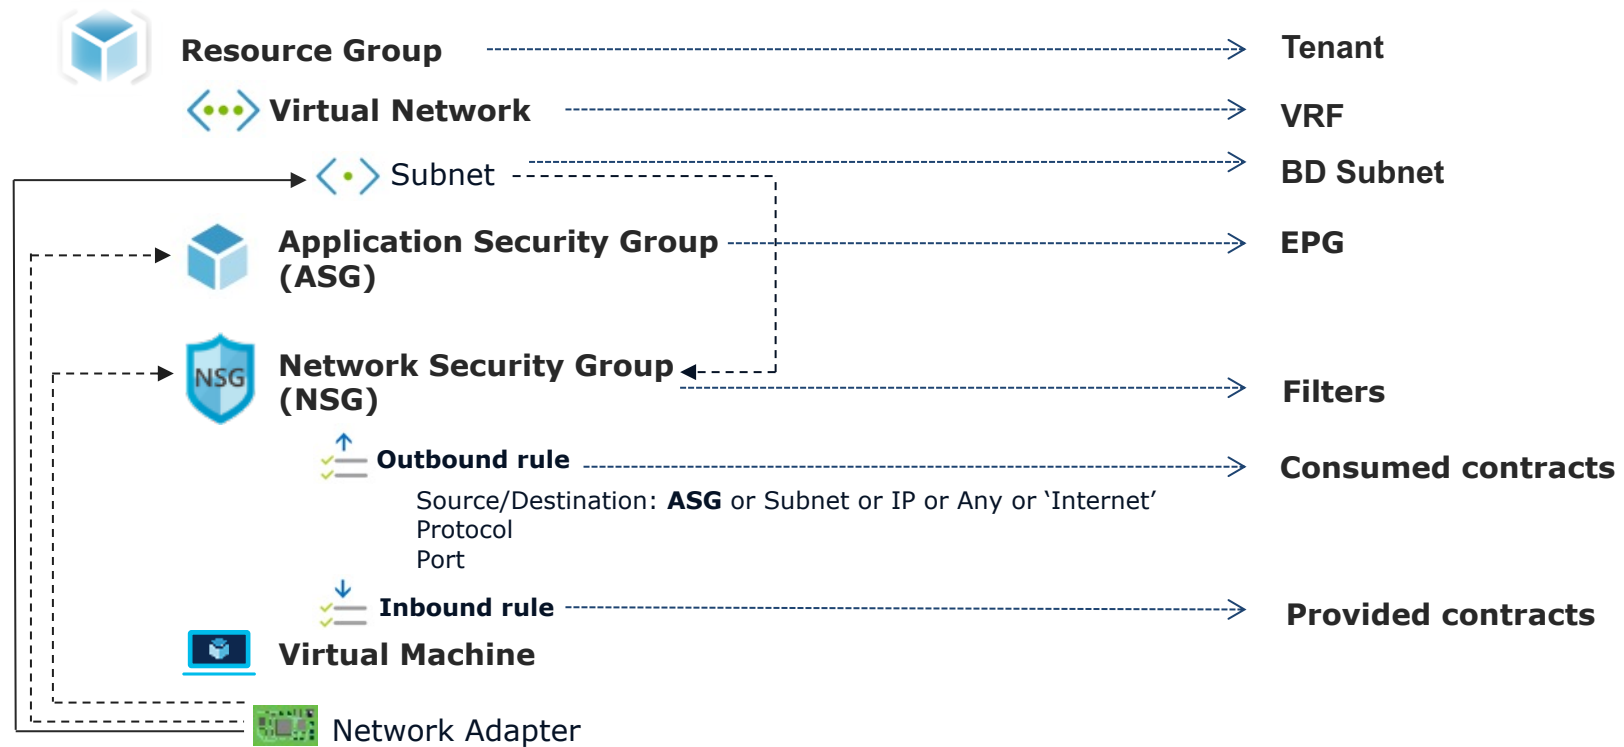
Automated network
extension to Azure

Encrypted
tunnel, Express
Route options

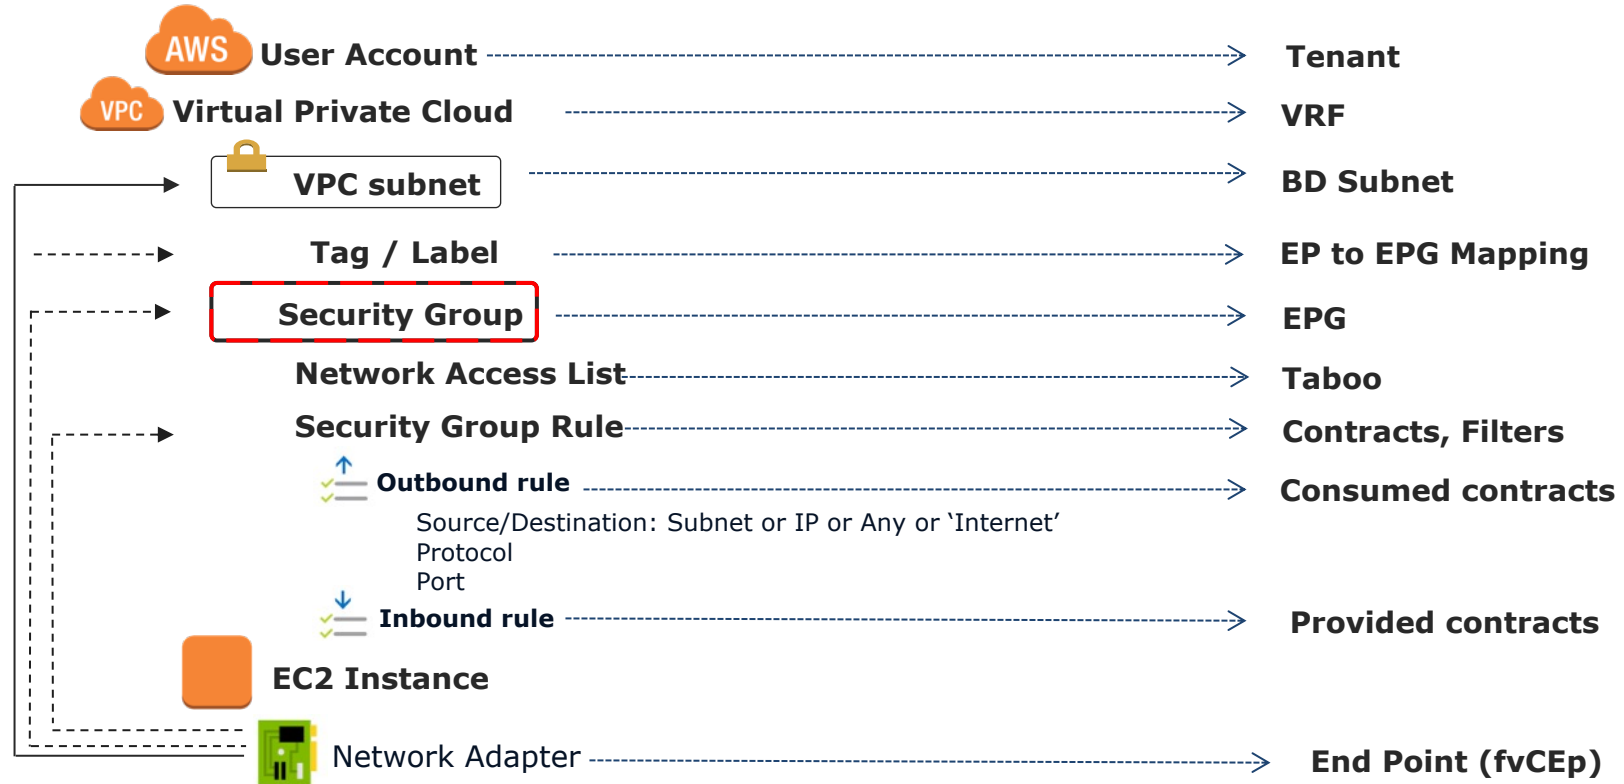
Azure Multi-region,
Multi-account
Support

End-to-end control plane for route
and policy exchange

Policy Mapping - Azure



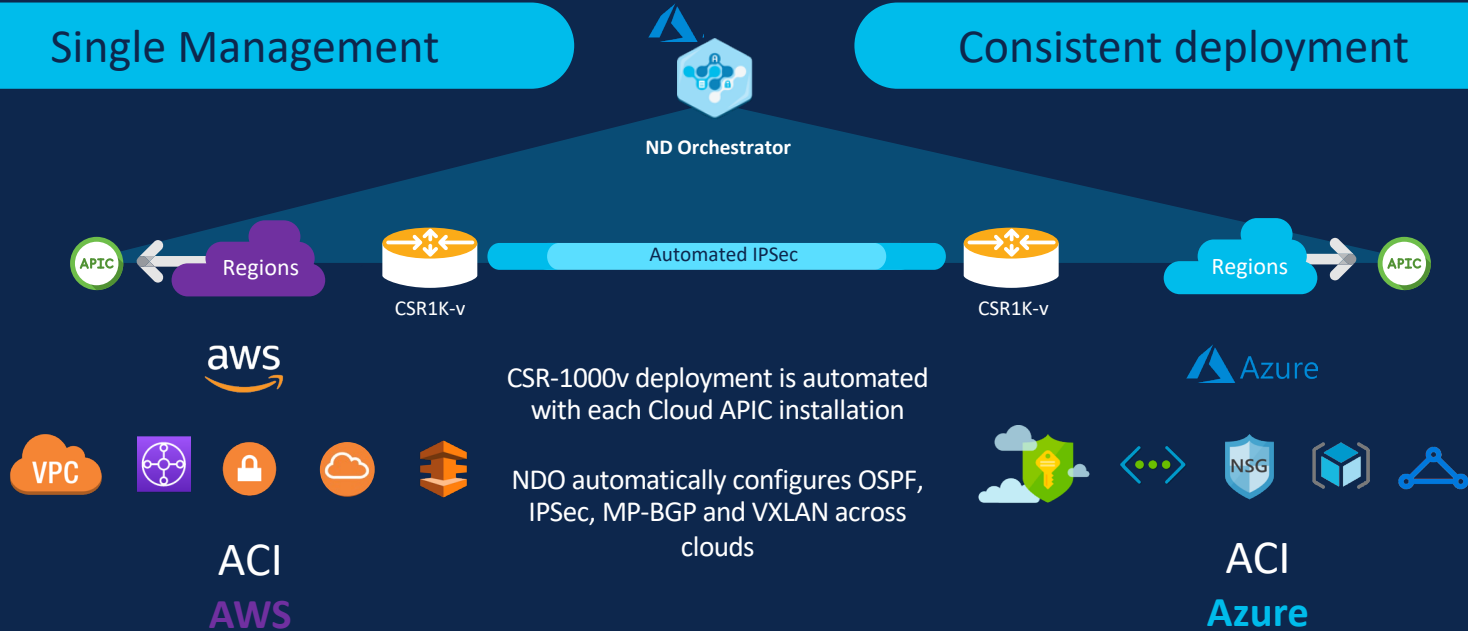
Policy and Resource Mapping - AWS



Cloud Only Solution

Single Management

Consistent deployment



Enable Multicloud Faster: Learn One Network, Manage and Interconnect multiple

SAN zrychluje na 64G

CISCO *Engage*



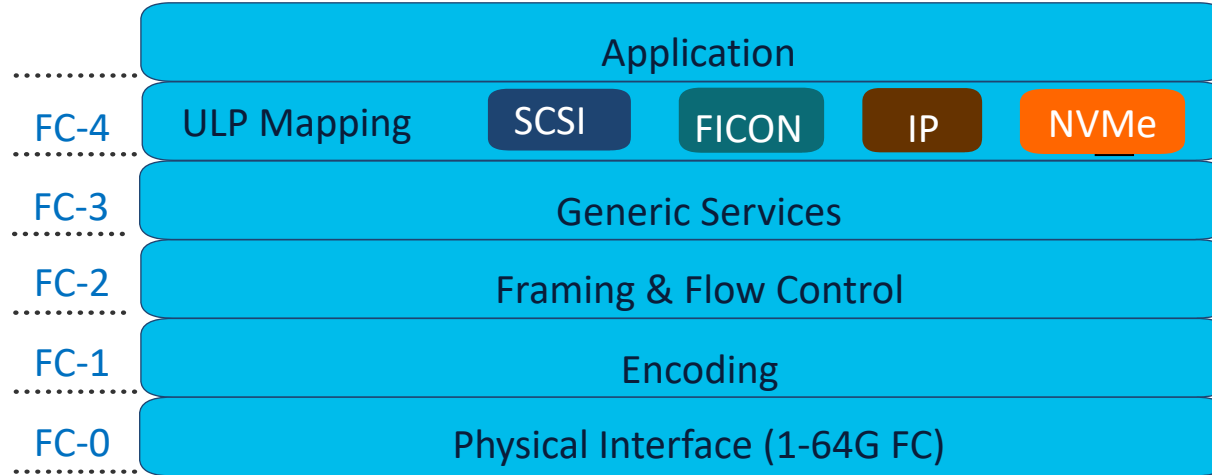
Q&A

slido.com

#160 378

Passcode: 43r05p

Fibre Channel Architecture & NVMe

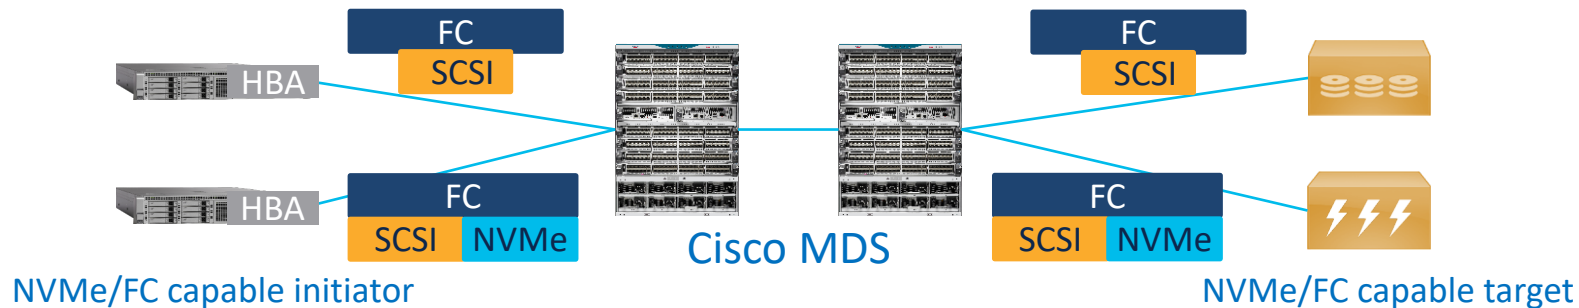


- NVMe/FC extends benefits of NVMe over Fibre Channel fabric
- Utilize FC benefits like plug-n-play, fabric server, name server, zone server, etc.
- NVMe/FC, SCSI-FCP & FICON can be transported concurrently in the same fabric

NVMe/FC – Phased & Seamless transition

Traditional FC-SCSI capable initiator

Traditional FC-SCSI capable target

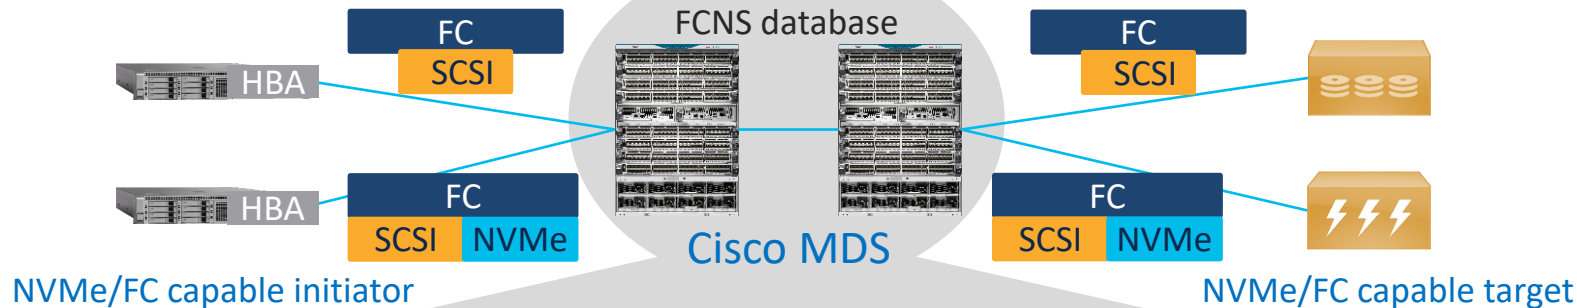


- Dual-stack end-devices and Multiprotocol switching in Cisco MDS and Nexus
 - Concurrent support of NVMe & SCSI transport
- SCSI or NVMe capability of end-devices is auto-detected and advertised
 - Similar to the existing plug-and-play architecture of Fibre Channel
 - No config on Cisco MDS and Nexus to enable or disable NVMe/FC. It's always on.
- NVMe/FC switching does not require additional license on Cisco MDS
- NVMe/FC is independent of FC speed. Higher speeds recommended.

NVMe/FC – Phased & Seamless transition

Traditional FC-SCSI capable initiator

Traditional FC-SCSI capable target



```
MDS9700# show fcns database vsan 160
VSAN 160:
```


FCID	TYPE	PWWN	(VENDOR)	FC4-TYPE:FEATURE
0x590020	N	10:00:00:90:fa:e0:08:5d	(Emulex)	scsi-fcp:init NVMe:init
0x590140	N	21:00:00:24:ff:7f:06:39	(Qlogic)	scsi-fcp:init NVMe:init

(showing entries only for dual-stack NVMe capable initiators. Other devices will look similar)


- End-devices register Upper Layer Protocol (ULP) with FCNS database, to be advertised to other end-devices in the same zone

Cisco MDS 9000 Switch Family

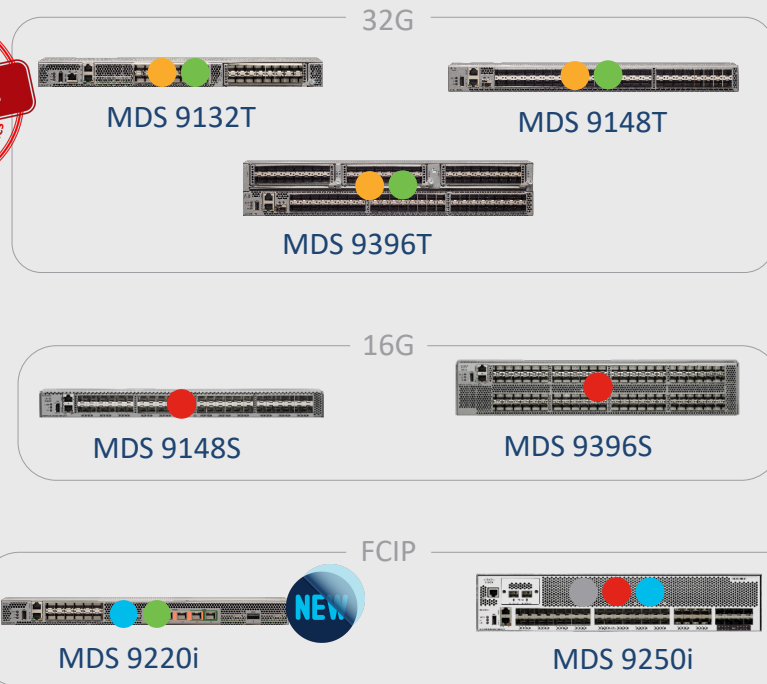
SAN Directors



Modules



Fabric Switches



8/16/32/64G FC

SAN Analytics


FCIP

4/8/16/32G FC

2/4/8/10/16G FC

FICON

64G module for MDS 9700 directors



- Quad-rate 48 ports – 8/16/32/64G FC*
- All ports have dedicated bandwidth (line-rate, non-oversubscribed, non-blocking)
- Requires Fab-3 fabric modules and Supervisor-4 (upgrade is non-disruptive)
- Next-gen SAN Analytics optimized for high-parallelism of NVMe
- 8 ports with FC TrustSec 256-bit AES encryption
- 48,000 B2B credits per module to absorb micro-bursts from all-flash NVMe arrays

* 64G FC SW optics available Q2 CY22

Port type	B2B credits			Extended B2B credits		
	Min	Max	Default	Min	Max	Default
F Port	1	500	100	N/A	N/A	N/A
E Port	2	1000	1000	2001	16000	3500

Accelerate the performance of All-flash NVMe storage



Limit the culprit devices using

Dynamic Ingress Rate Limiting

to prevent congestion and slow drain in storage area networks (SAN)



Cisco Patent



End-device independent
Upgrading of end-devices is not needed



Adaptive
DIRL dynamically adjusts as per the traffic profile of the host



No side effects
Rate limits congested hosts only. Other non-congested hosts and storage ports are not impacted



Easy adoption
DIRL is available on MDS switches after a software-only upgrade.




Affordable
No additional license needed



Topology independent
DIRL works in edge-core, edge-core-edge, or collapsed core (single switch fabric) topologies

[Slow-Drain Device Detection, Troubleshooting, and Automatic Recovery Whitepaper](#)

Cisco SAN Analytics Advantages with 64 GFC



- Support high-parallelism of NVMe
- Visibility into
 - Every IO flow - Virtual Machine-Initiator-Target-LUN (VM-ITL flow) for SCSI operations
 - Every IO flow - Virtual Machine-Initiator-Target-Namespace (VM-ITN flow) for SCSI operations
 - Every port
 - Every speed
 - Every transaction and Every IO
- All the capabilities of 32GFC SAN Analytics. Additionally,
 - ASIC integrated traffic inspection and metric calculation
 - Programmable on-board NPU for future extensive with software-only upgrade
 - NVMe/FC parsing & key lookup support (namespace ID and connection ID)
 - ASIC timestamping for increased accuracy of latency metrics
 - 1GbE port on module for low granularity metric directly from the module (SW support is TBD)

Nexus Dashboard

CISCO *Engage*

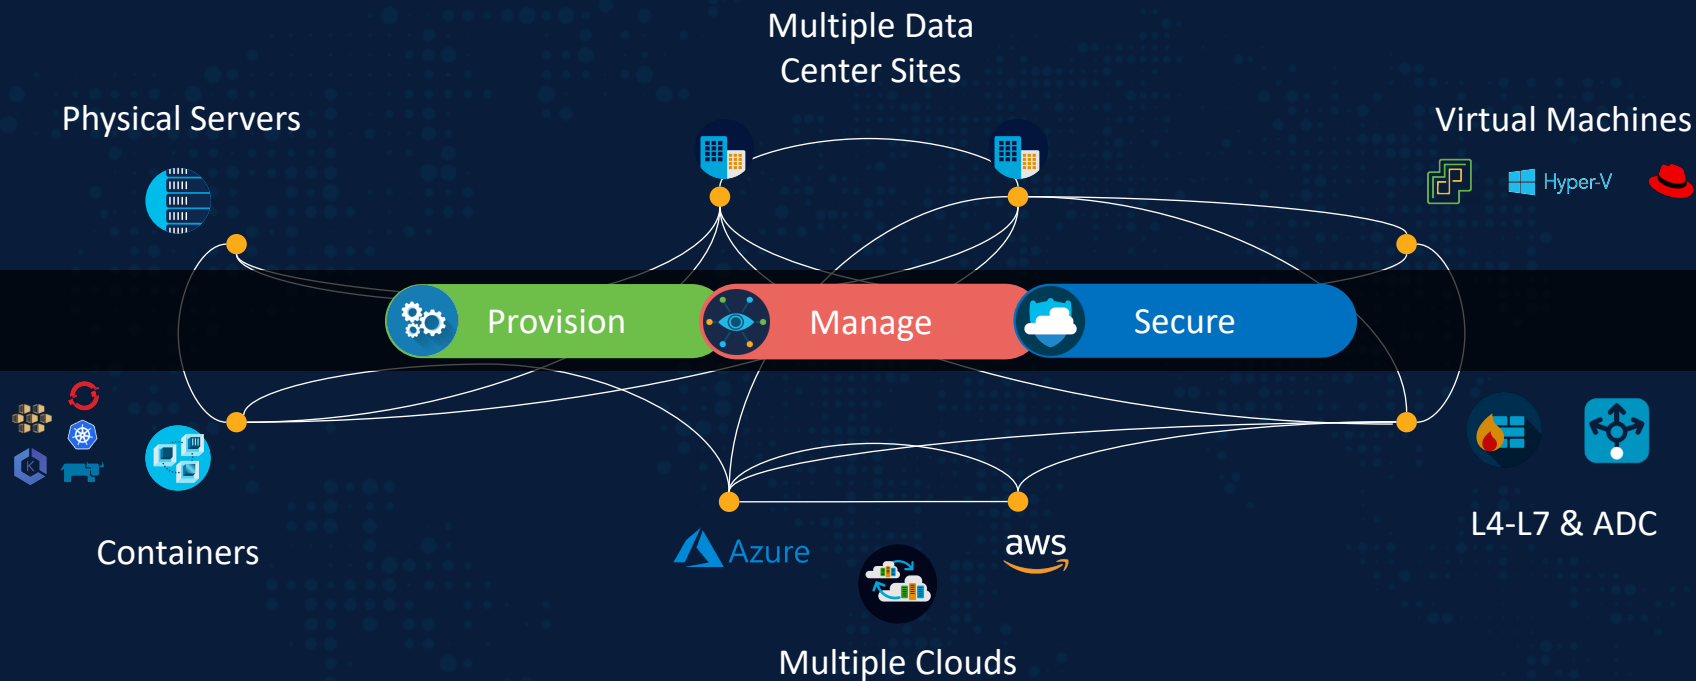


Q&A

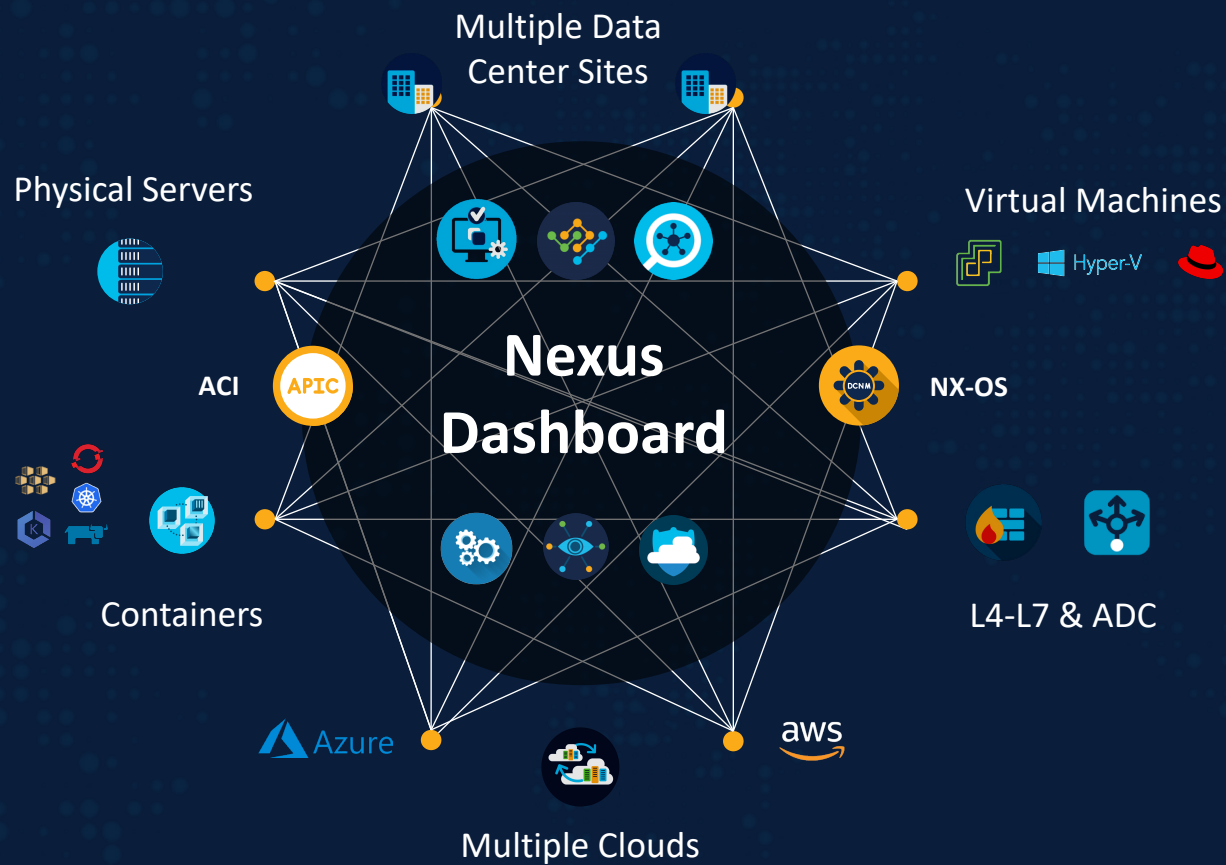
slido.com

#160 378

Passcode: 43r05p



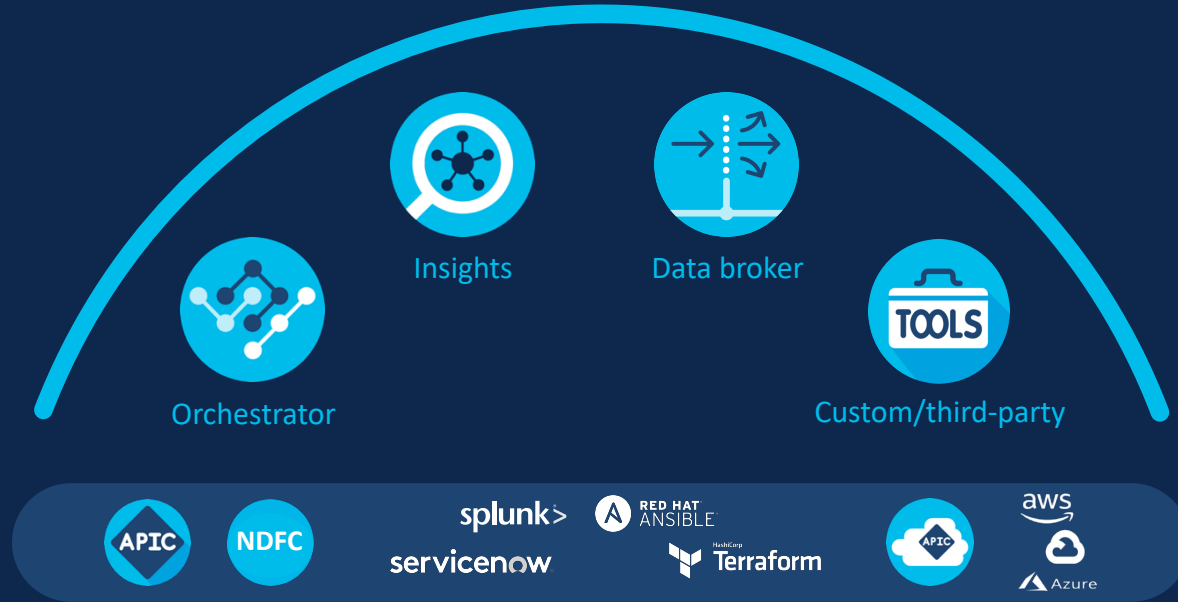
The network plays a **critical** role to connect **the business** anywhere



Configure the network **once**, deploy, monitor and secure **anywhere**

Nexus Dashboard

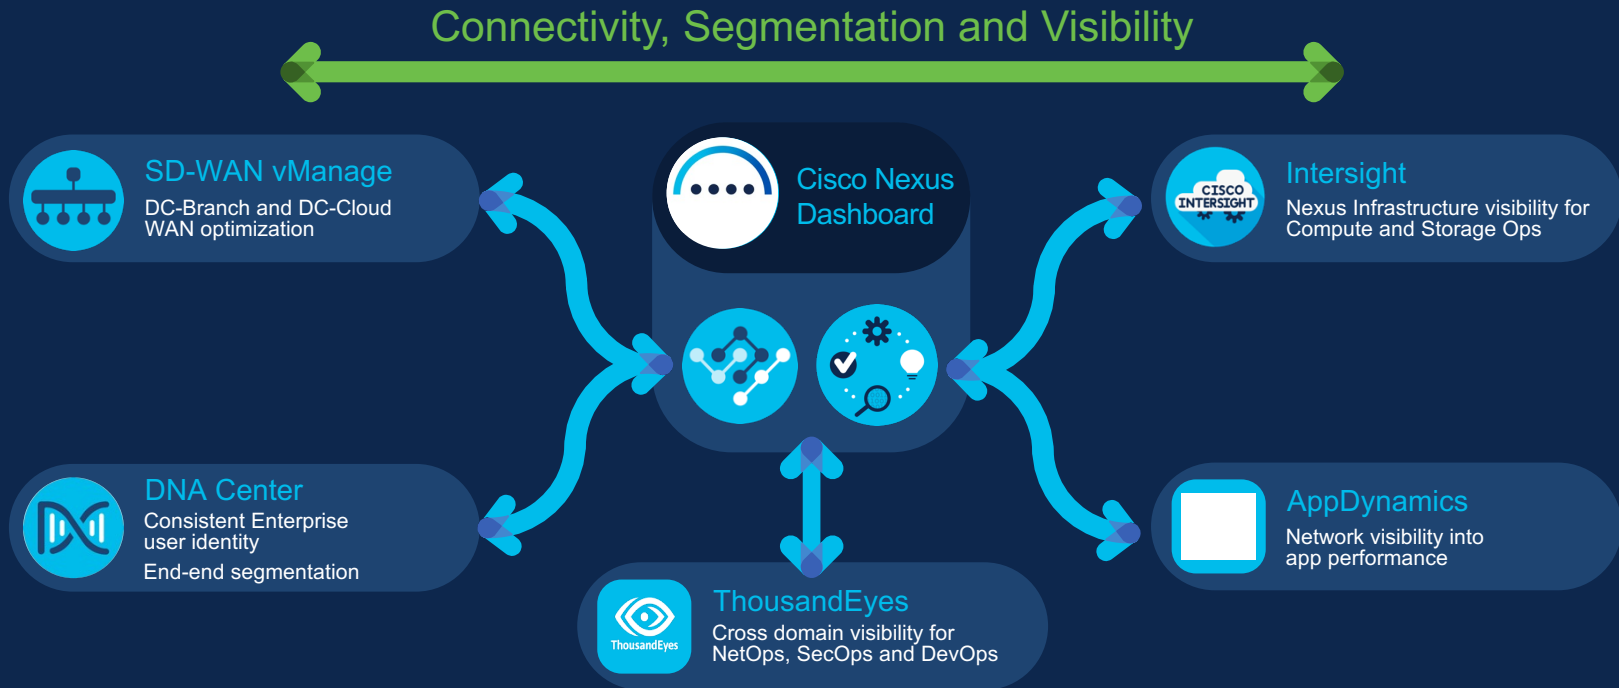
Simple to automate, simple to consume



Consume all services in one place

Nexus Dashboard

Service integrations



Nexus Dashboard

3rd Party Service integrations

Analytics/
monitoring

splunk[>]

Cisco Nexus
Dashboard
Data Broker

ITSM/
workflows

servicenow[®]

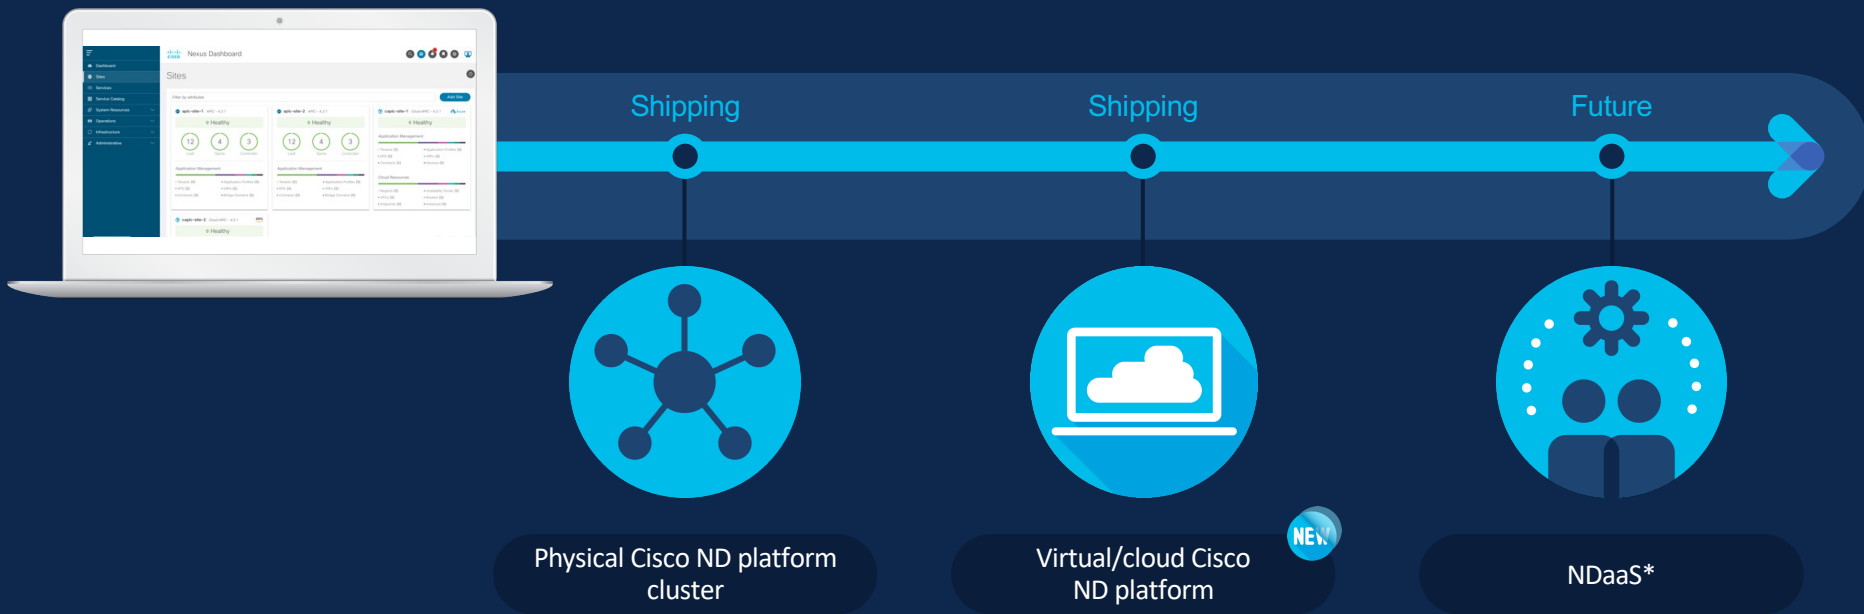
Infrastructure
as code

 RED HAT[®]
ANSIBLE[®]

 HashiCorp
Terraform

Nexus Dashboard

Deployment options



Nexus Dashboard Insights

Overview

Assurance

Am I doing correct configuration?
Are interdependencies known?
Does the change impact something
am not aware of?

Proactive advisories

Was the issue preventable?
Is the network exposed to known
vulnerabilities?
Can I get proactive advice?



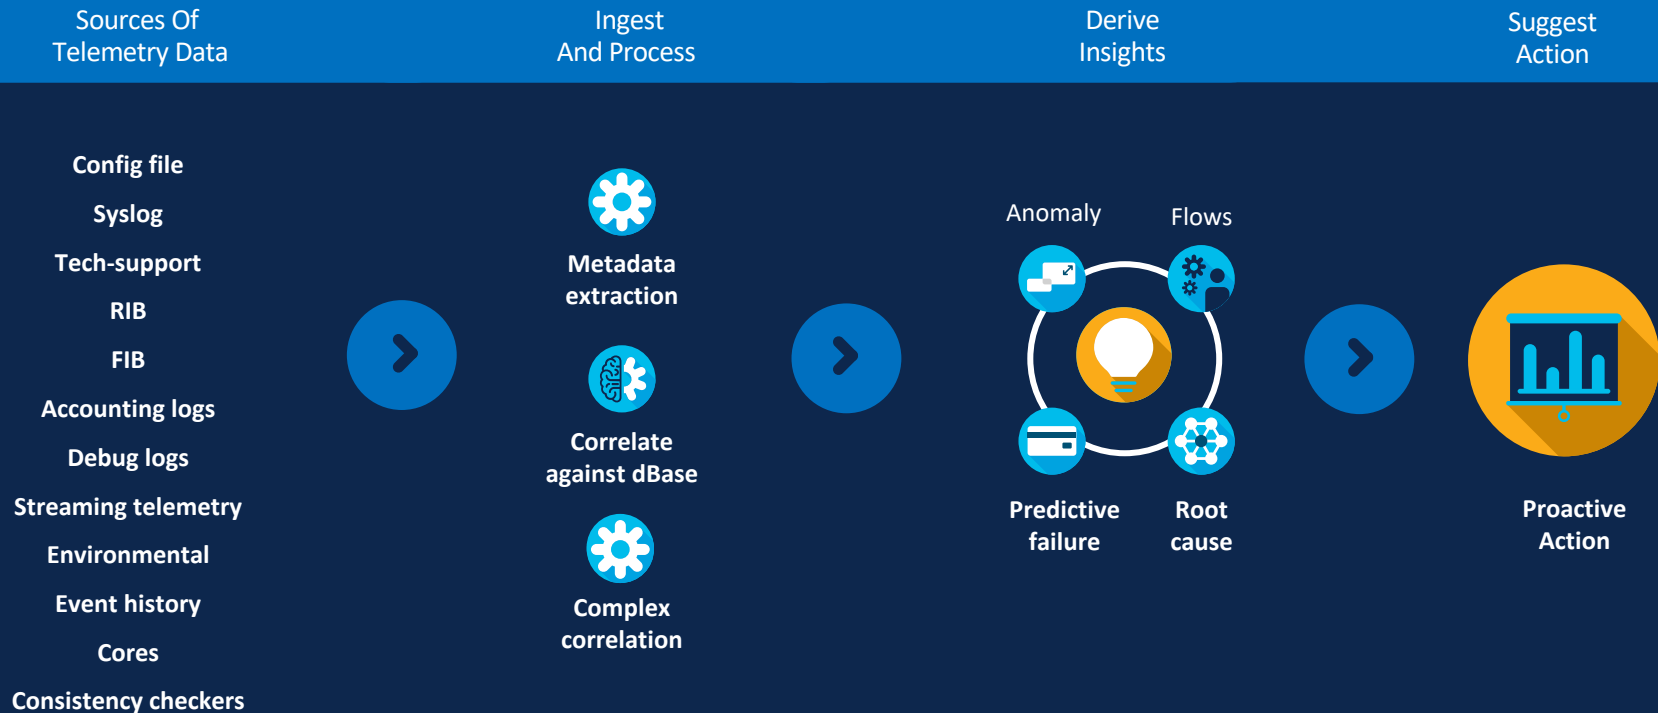
Troubleshooting

Where is the problem and what's the
blast radius?
How do I reduce MTTR?
How do I prove network is healthy?

Single-pane-of-glass

Can I get visibility across
datacenters?
Single point for monitor and
control?

Nexus Dashboard Insights



Proactive Operations : Increase Availability, Performance, And Simplify Operations

Nexus Dashboard Insights

Feature set

Assurance and compliance



Configuration assurance



Pre-change analysis



Delta analysis



Explorer



Change management



Communication compliance

Visibility and troubleshooting



One View



AppDynamics integration



Endpoint analytics



Connectivity analysis



Flow analytics



Capacity planning



Topology view



Microbursts detection

Advisory and maintenance



Upgrade assist



PSIRT notification



TAC assist



Advisories

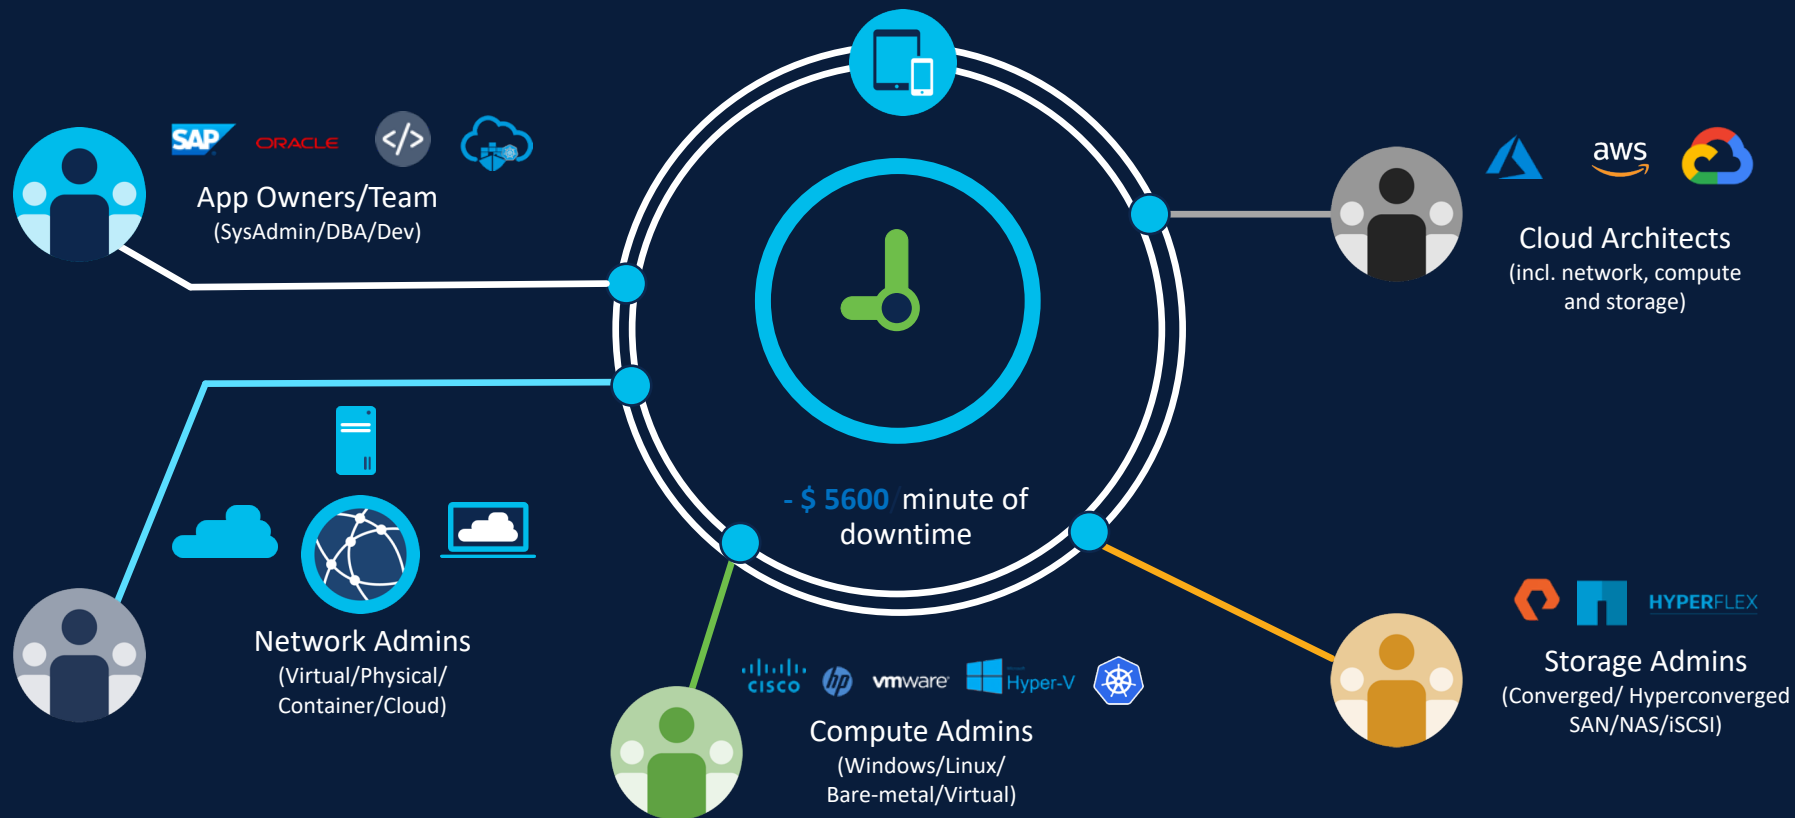


Field notices

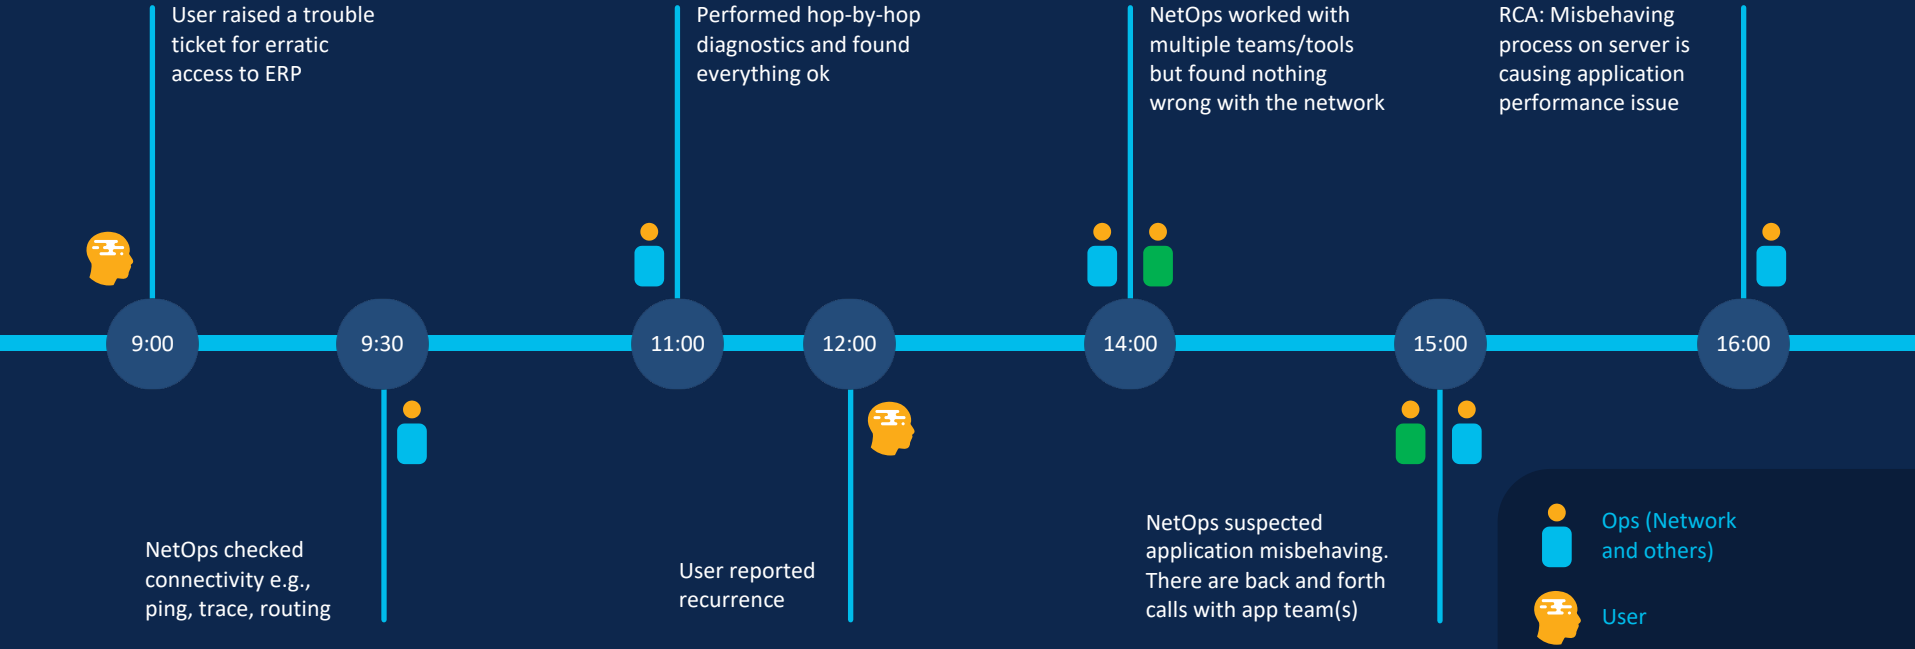


Email notifications

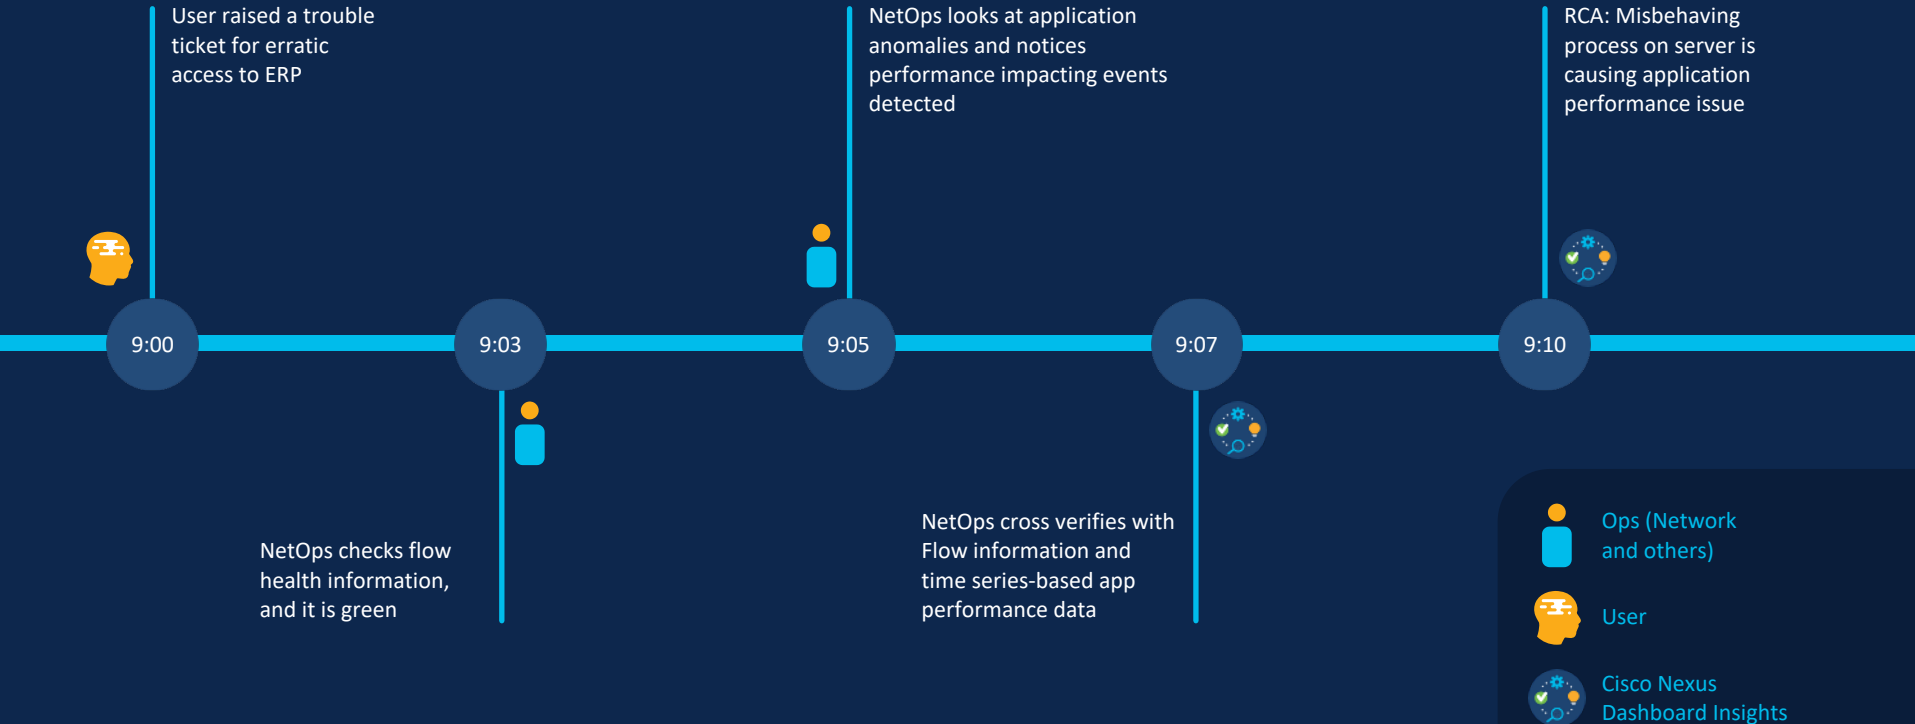
The problem: Finding where the problem is



Traditional troubleshooting workflow



Now with Cisco Nexus Dashboard Insights



How to buy Nexus Dashboard Insights

Software included with Switch subscription licenses

DCN Day2Ops

Cisco Nexus Dashboard Insights,
plus



Multi DC
(DCN Advantage)

or



Single DC
(DCN Essential)

DCN Advantage

Cisco Nexus Dashboard
Orchestrator


DCN Premier

Cisco Nexus Dashboard Insights
and Nexus Dashboard Orchestrator

Install Base (ACI/NX-OS)

New Purchase (ACI/NX-OS)

Solve Multicloud Networking Operations Challenges



Assurance



Troubleshooting



Proactive advisories



Need for **homogenous experience** across heterogenous cloud environments

Závěr