

# What You Make Possible



# Cisco NX-OS Software Architecture

BRKARC-3471

*TOMORROW  
starts here.*



# Housekeeping

- We value your feedback- **don't forget to complete your online session evaluations after each session** & complete the Overall Conference Evaluation which will be available online from Thursday
- Visit the World of Solutions
- Please remember this is a 'non-smoking' venue!
- **Please switch off your mobile phones**
- Please make use of the recycling bins provided
- Please remember to wear your badge at all times

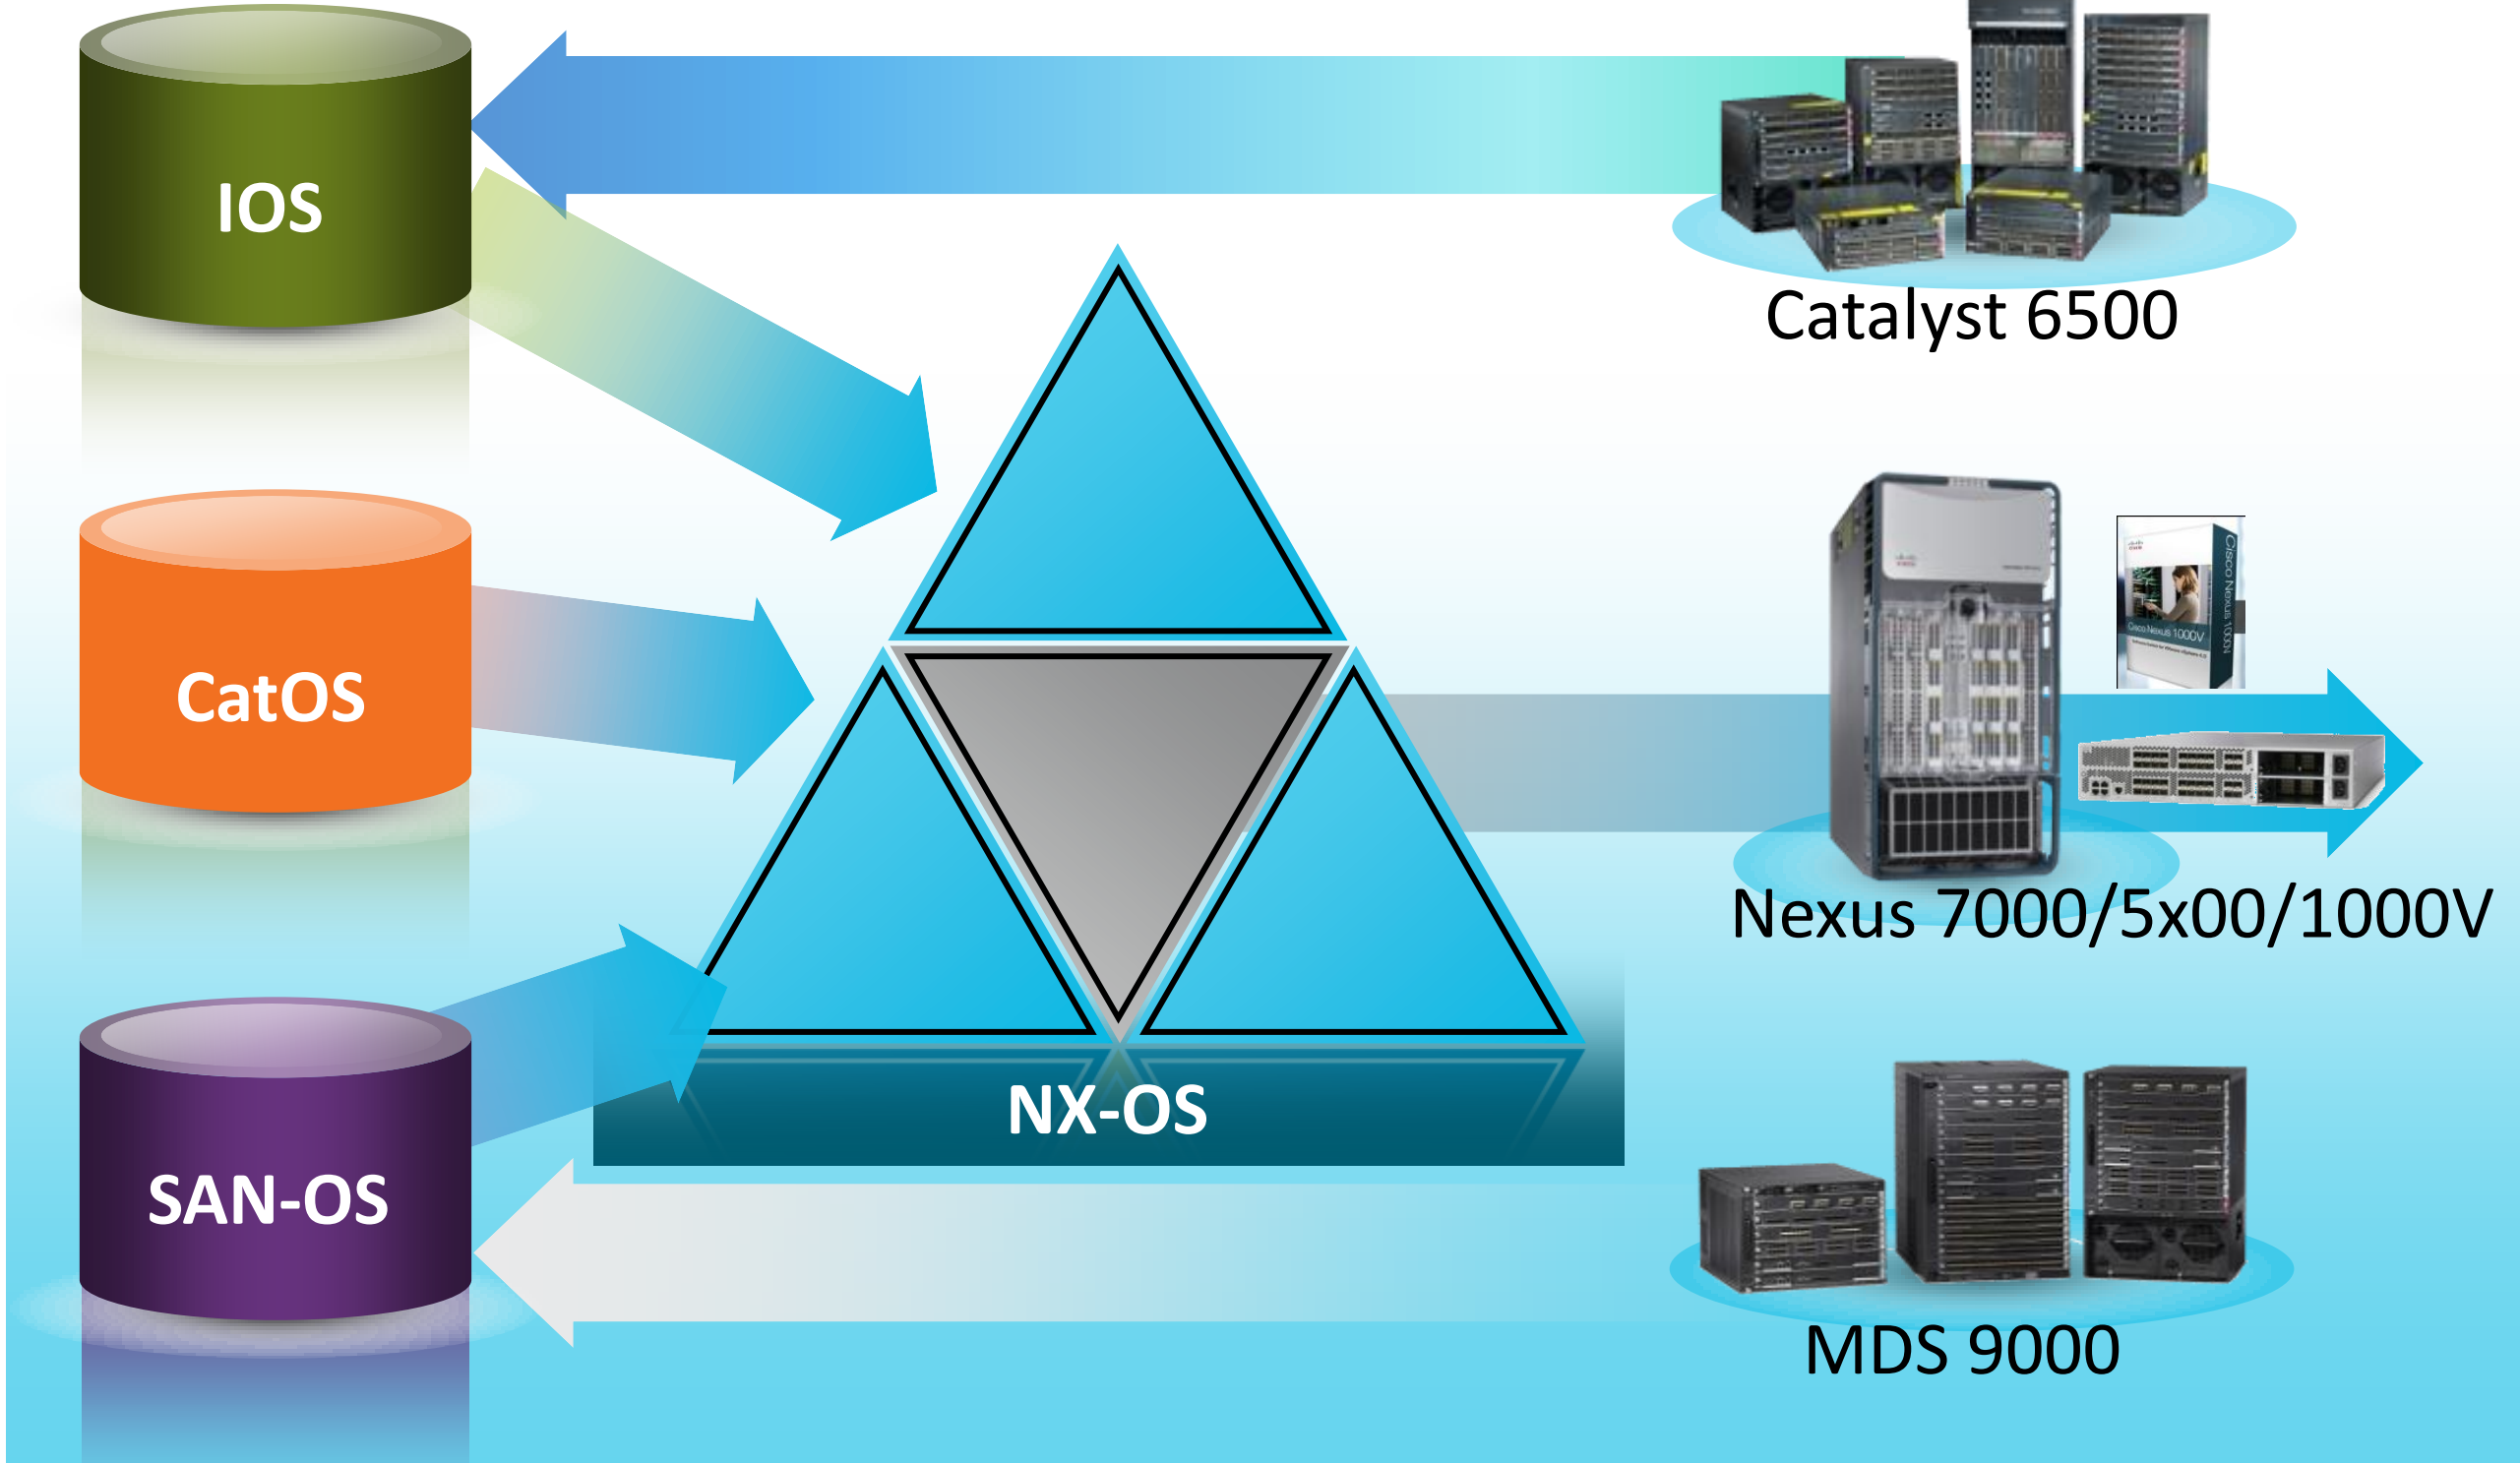
# Session Agenda

- **NX-OS Origins & Overview**
- NX-OS Modular Architecture
- High-Availability Infrastructure
- High-Availability Features & Capabilities
- Command Line Interface
- Operational & Management Features
- Licensing & Lifecycle
- Innovation
- Conclusion





# NX-OS: Designed for the Data Centre





# Cisco NX-OS Adoption

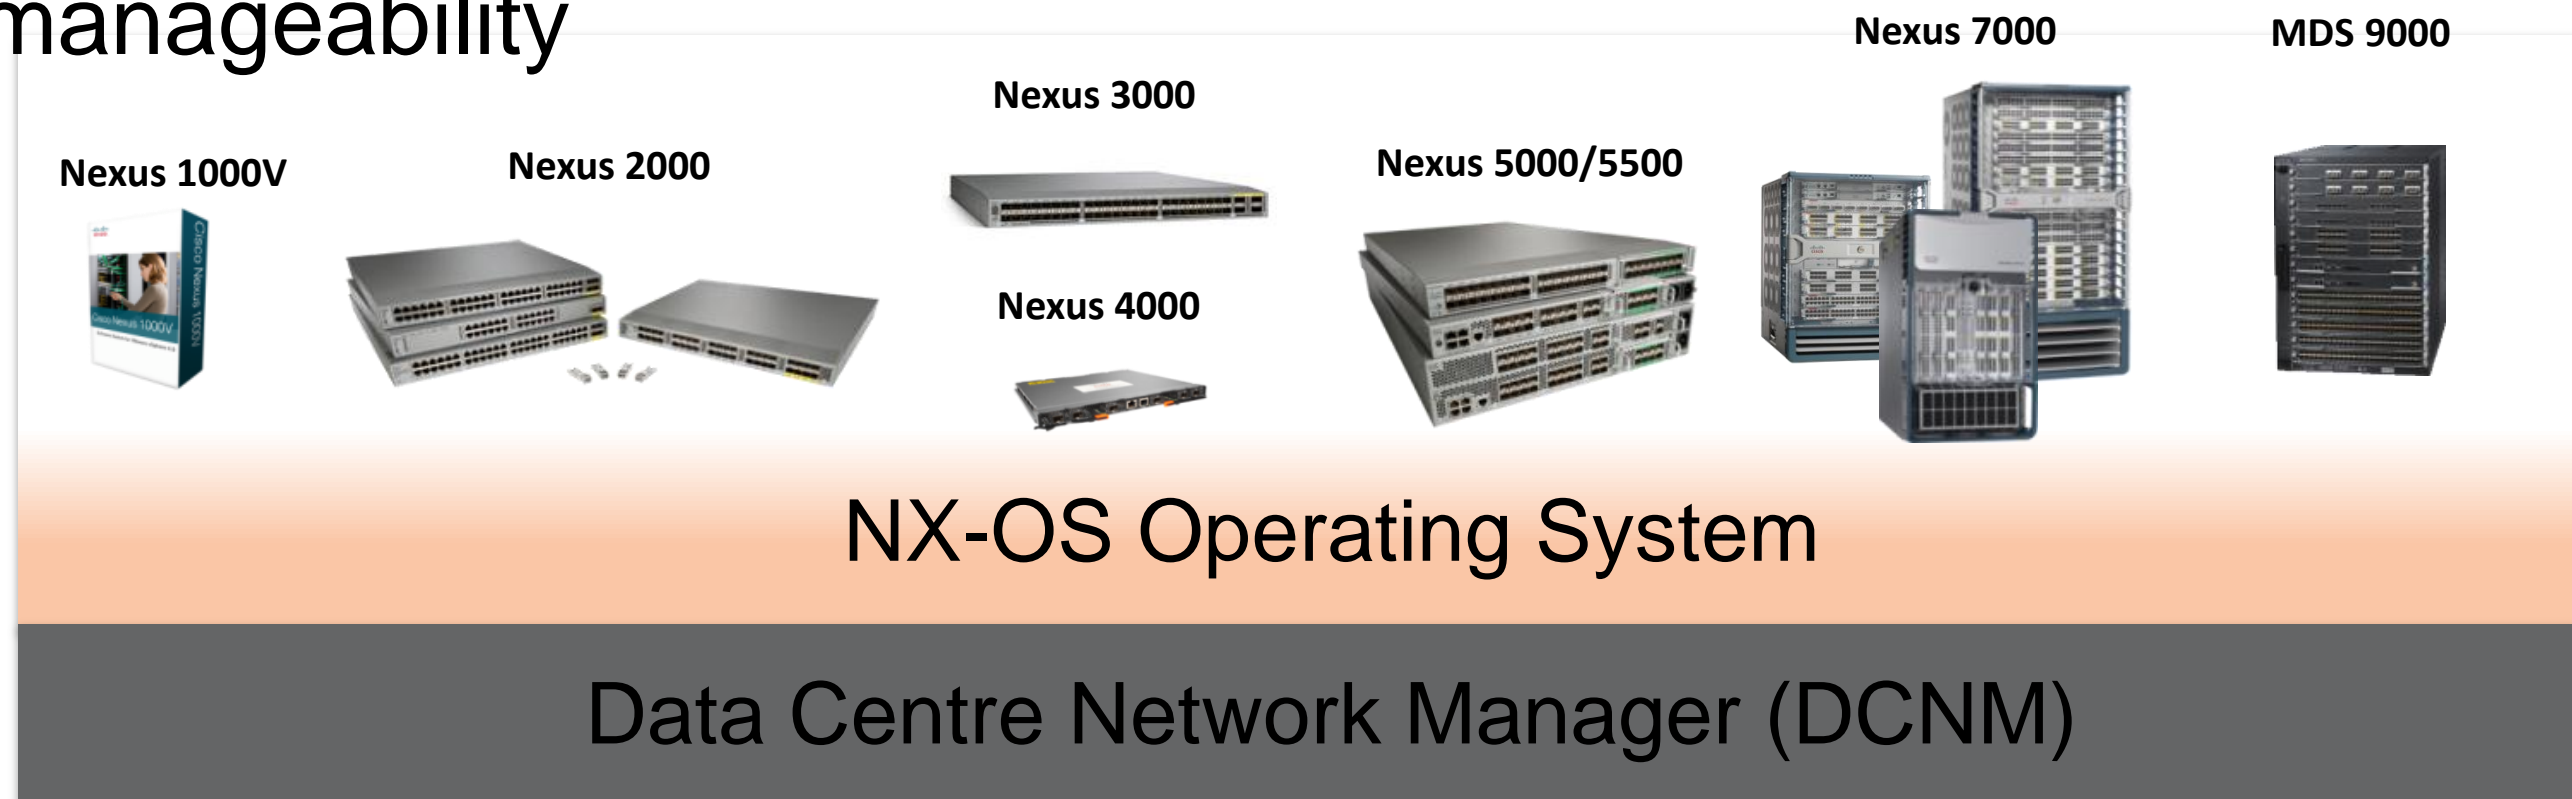
- Shipping for 5+ years
- 23,000+ customers
- 11,000,000 ports
- Validated Design Guides and Case Studies
- Differentiating features driving the adoption and being deployed
- Major Certification and Deployment milestones





# The Cisco Unified Fabric Family

- Complete data centre class switching portfolio
- Consistent operating system across all platforms
- Infrastructure scalability, transport flexibility and operational manageability



# Cisco NX-OS Highlights

Designed to Meet the Operational Needs of the Data Centre

- **Feature Rich Operating System**
  - Comprehensive L2 and L3 feature set
- **Modular, Multi-Threaded/Processor**
  - Highly scalable unprecedented uptime
- **Intelligent IOS-Like CLI**
  - Little or no retraining required
- **Zero Service Disruption**
  - Maintenance ≠ Downtime
- **Virtualisation Support**
  - Industry first virtualised network OS, VM-FEX
- **Layer 2 and Layer 3 Multipathing**
  - Resilient scalable Layer 2 and Layer 3 domains
- **Storage and Ethernet Convergence**
  - FCoE, iSCSI, HPC
- **Advanced Management Infrastructure**
  - XML and Web Services



# Comprehensive Data Centre Feature Set




## Layer 3

- Distributed IPv4/IPv6 Hardware Forwarding
- OSPF, EIGRP, IS-IS, BGP, RIP, PBR
- PIM-SM, SSM/Bidir, MSDP, MP-BGP, IGMP/MLD
- 16-way ECMP
- (HSRP, GLBP, VRRP) + Object Tracking
- MPLS (new)
- BFD



## Virtualisation

- VRF-lite
- Virtual Device Contexts (VDCs)
- LISP



## High Availability

- In-Service Software Upgrade (ISSU)
- Non-Disruptive Stateful Supervisor Switchover (SSO)
- Stateful process restarts
- Graceful Process Restart



## Operational Manageability

- GOLD, Smart Call Home, EEM w/ TCL
- NetFlow, NDE v5/v9, FNF CLI
- SPAN, ERSPAN, VACL Capture
- Wireshark
- SNMP/MIB
- NETCONF/XML
- Configuration checkpoint & rollback



## Layer 2

- Distributed Hardware Based Layer 2
- PVRST, MST
- STP Guards, Bridge Assurance, UDLD
- 802.1ad/LACP Portchannels
- Private VLANs
- Virtual Port Channel (vPC)
- Overlay Transport Virtualisation (OTV)
- Data Centre Bridging (DCB)
- Layer 2 Multipathing (FabricPath/TRILL)



## Security

- ACLs, VACLs, PACLs
- Cisco TrustSec & LinkSec (CTS/802.1AE)
- CoPP & Rate Limiters
- DHCP snooping, DAI, IP source guard
- 802.1x & Port Security
- Storm control
- Unicast RPF check

## Storage Area Networks

- FCoE
- FIP & FIP Snooping

## Quality of Service

- Ingress/Egress queuing with WRED
- Marking Policies & Mutation
- Ingress/Egress “1-rate 2-colour” & “2-rate 3-colour” policing
- Colour-aware policing
- MQC CLI model



# Comprehensive Data Centre Feature Set

## Innovation for the Data Centre




### Layer 3

- Distributed IPv4/IPv6 Hardware Forwarding
- OSPF, EIGRP, IS-IS, BGP, RIP, PBR
- PIM-SM, SSM/Bidir, MSDP, MP-BGP, IGMP/MLD
- 16-way ECMP
- (HSRP, GLBP, VRRP) + Object Tracking
- MPLS (new)
- BFD



### Virtualisation

- VRF-lite
- Virtual Device Contexts (VDCs)
- LISP



### High Availability

- In-Service Software Upgrade (ISSU)
- Non-Disruptive Stateful Supervisor Switchover (SSO)
- Stateful process restarts
- Graceful Process Restart



### Operational Manageability

- GOLD, Smart Call Home, EEM w/ TCL
- NetFlow, NDE v5/v9, FNF CLI
- SPAN, ERSPAN, VACL Capture
- Wireshark
- SNMP/MIB
- NETCONF/XML
- Configuration checkpoint & rollback



### Layer 2

- Distributed Hardware Based Layer 2
- PVRST, MST
- STP Guards, Bridge Assurance, UDLD
- 802.1ad/LACP Portchannels
- Private VLANs
- Virtual Port Channel (vPC)
- Overlay Transport Virtualisation (OTV)
- Data Centre Bridging (DCB)
- Layer 2 Multipathing (FabricPath/TRILL)



### Security

- ACLs, VACLs, PACLs
- Cisco TrustSec & LinkSec (CTS/802.1AE)
- CoPP & Rate Limiters
- DHCP snooping, DAI, IP source guard
- 802.1x & Port Security
- Storm control
- Unicast RPF check

### Storage Area Networks

- FCoE
- FIP & FIP Snooping

### Quality of Service

- Ingress/Egress queuing with WRED
- Marking Policies & Mutation
- Ingress/Egress “1-rate 2-colour” & “2-rate 3-colour” policing
- Colour-aware policing
- MQC CLI model



# Nexus Certifications for NX-OS 5.1

- IPv6 Ready Logo Phase I Certified



<https://www.ipv6ready.org/db/index.php/public/logo/01-000556/>

- FIPS 140-2 Certified

– Completed in April 2011 - Cert# 1533, 1534



– <http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/140val-all.htm>

- EAL4 Common Criteria Certified

– Completed in April 2011

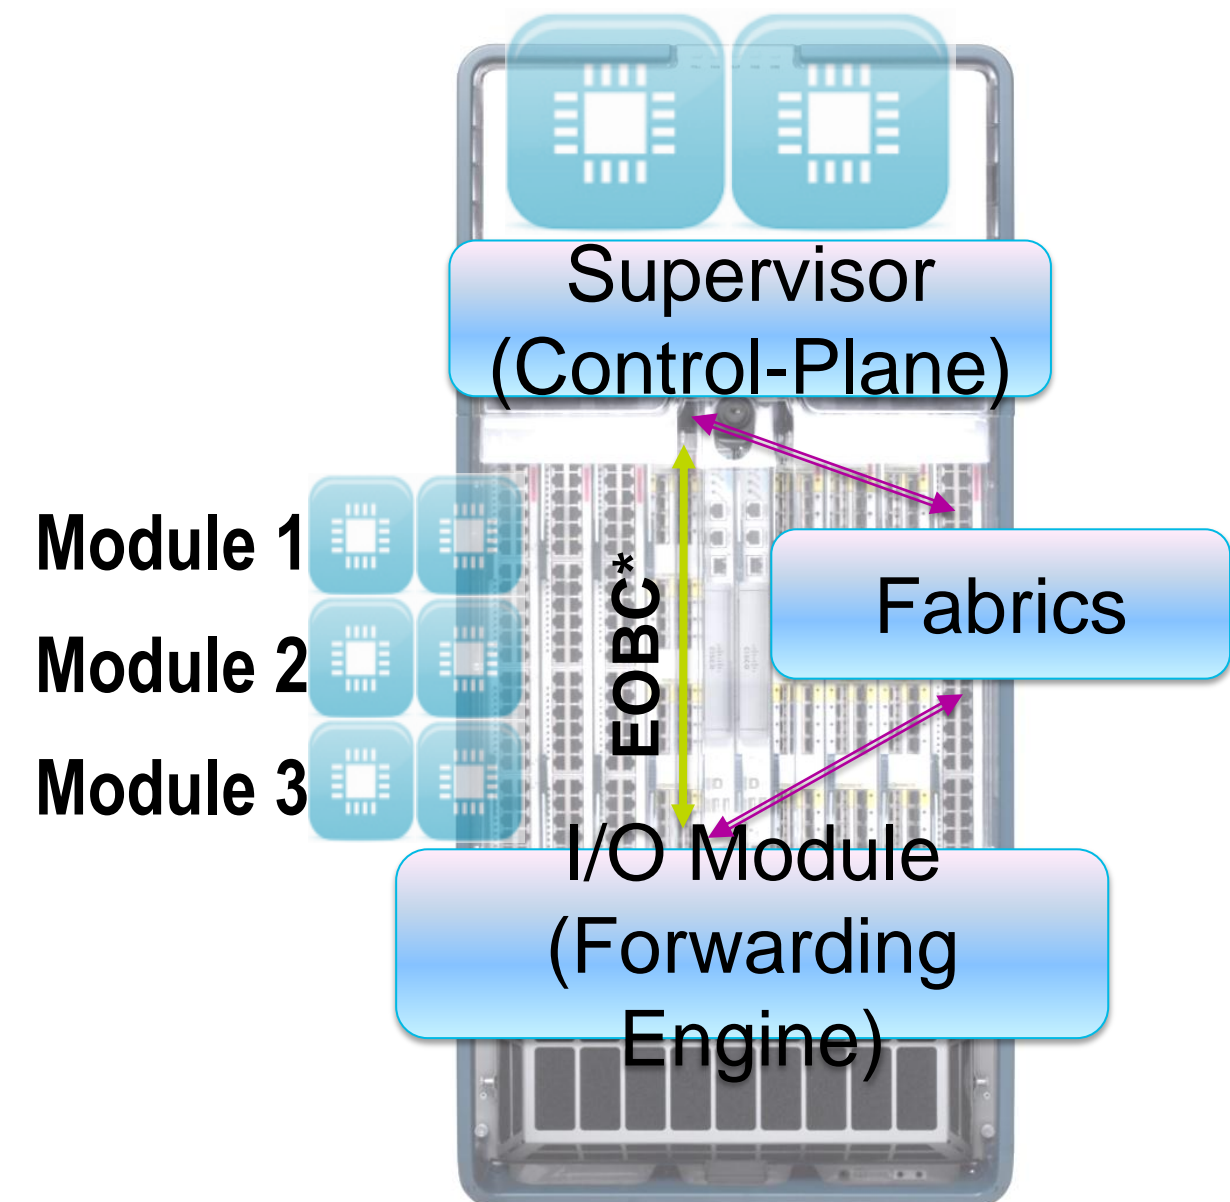
– <http://www.niap-ccevs.org/st/vid10349>



# NX-OS Distributed Architecture

## Distributed Forwarding and Control-plane

- OS designed to leverage distributed hardware architecture
- Fabric & forwarding engine removed from supervisor
- Each I/O module has independent control-plane and forwarding hardware
- Control-plane & data-plane separation (same on Nexus 5x00)
- Fully distributed system for non-disruptive SSO & ISSU (SSO only available on dual-sup Nexus 7000)



\*EOBC: Ethernet Out Of Band Channel

# Session Agenda

- NX-OS Origins & Overview
- **NX-OS Modular Architecture**
- High-Availability Infrastructure
- High-Availability Features & Capabilities
- Command Line Interface
- Operational & Management Features
- Licensing & Lifecycle
- Innovation
- Conclusion

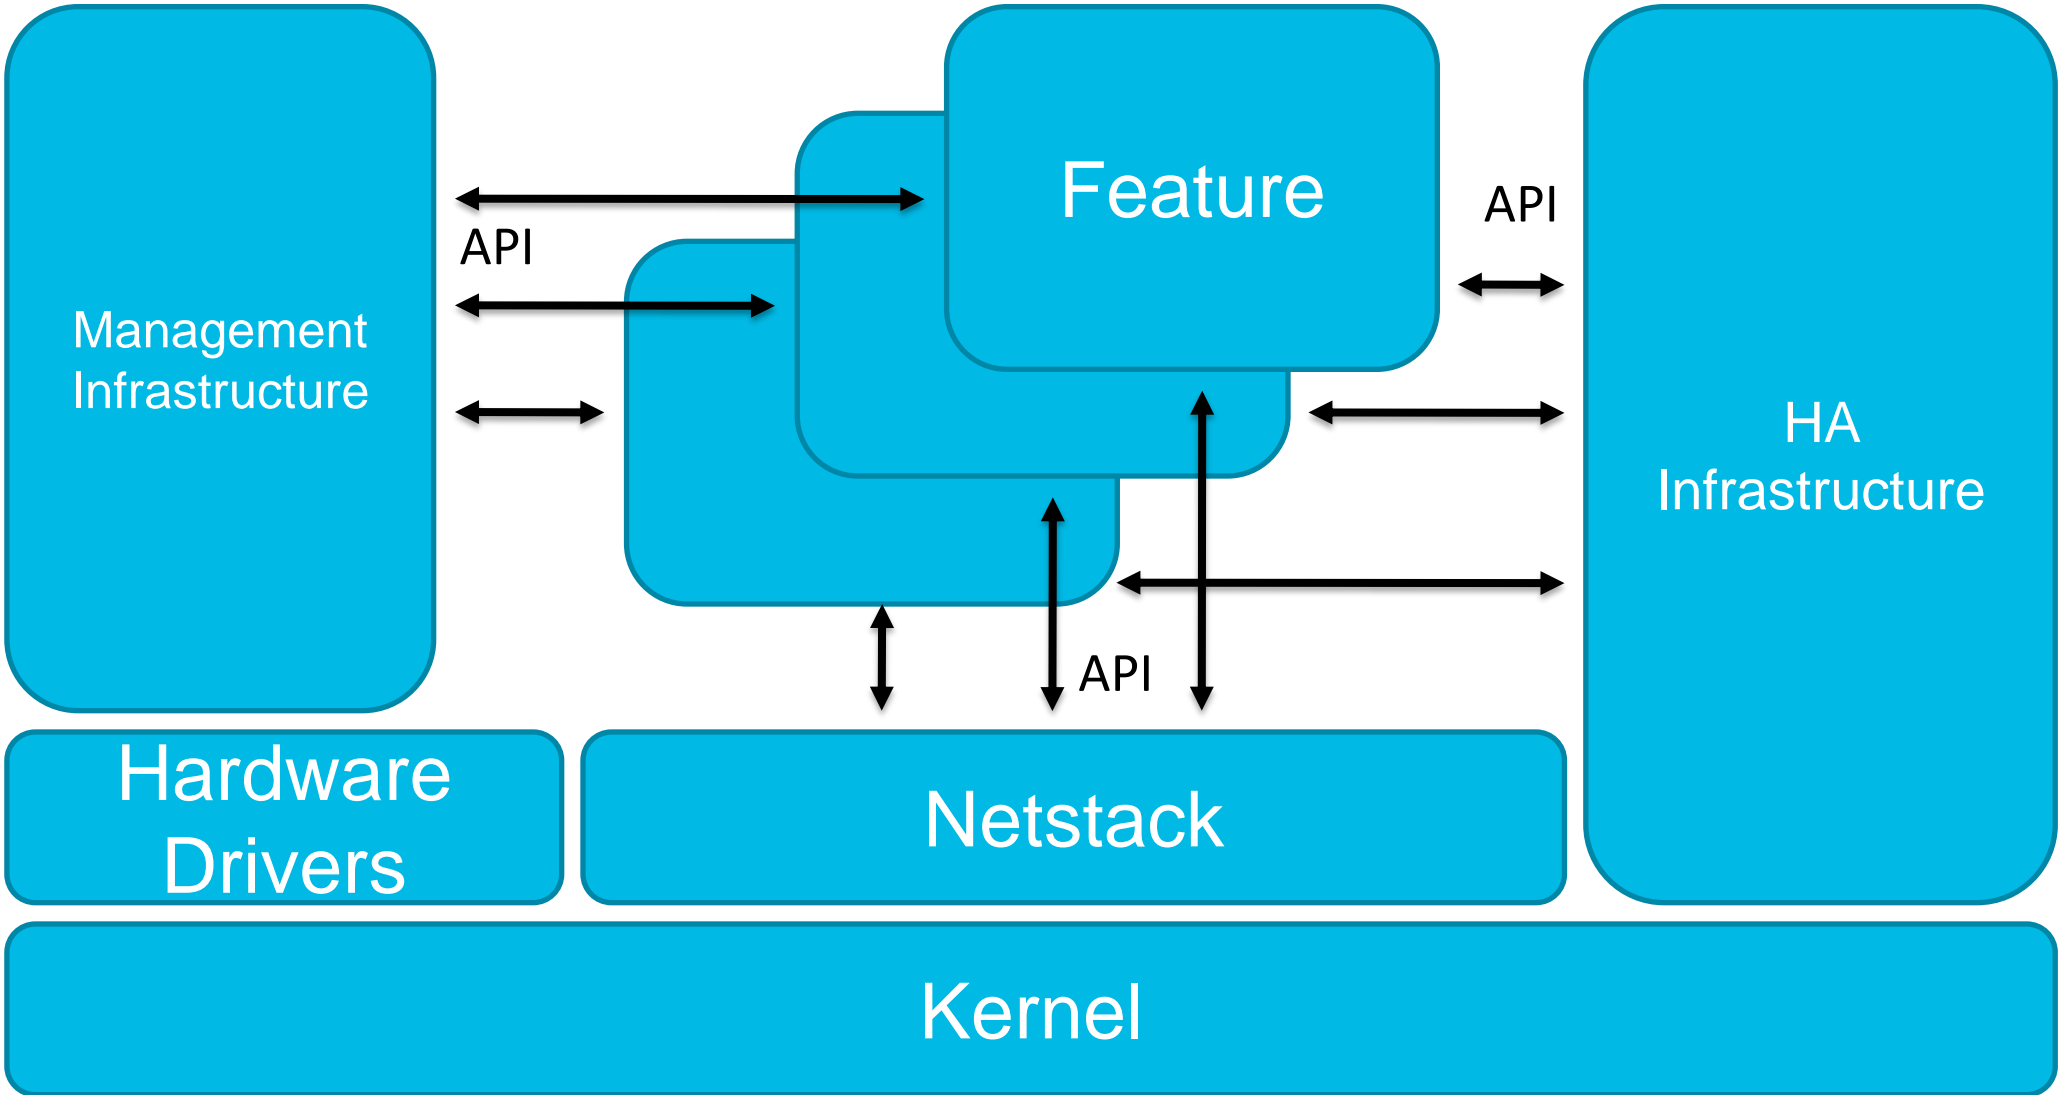


# NX-OS Modular Architecture

## Faster Defect Resolution

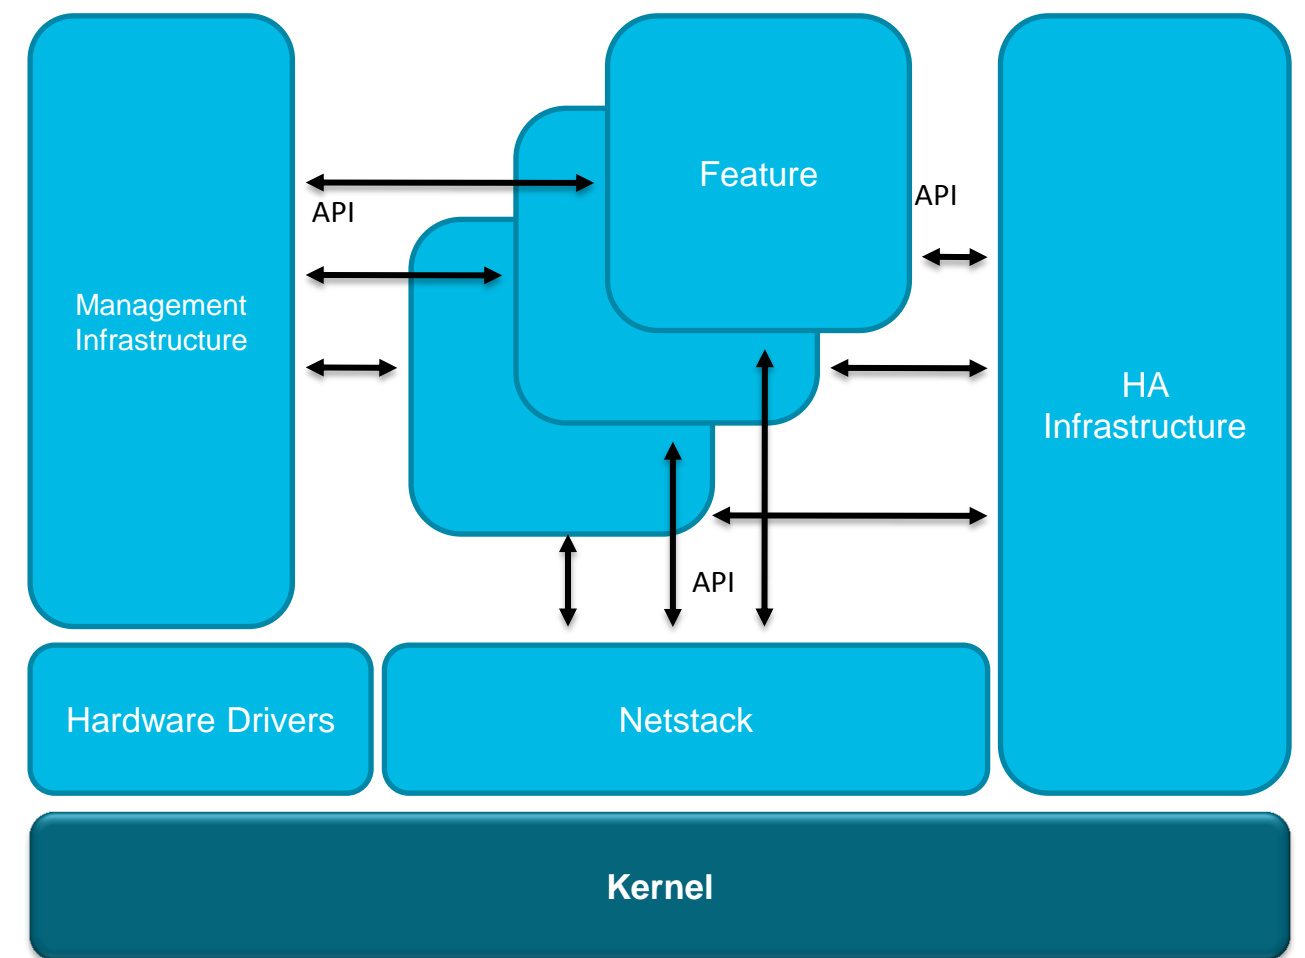
### Feature Velocity

### Consistency



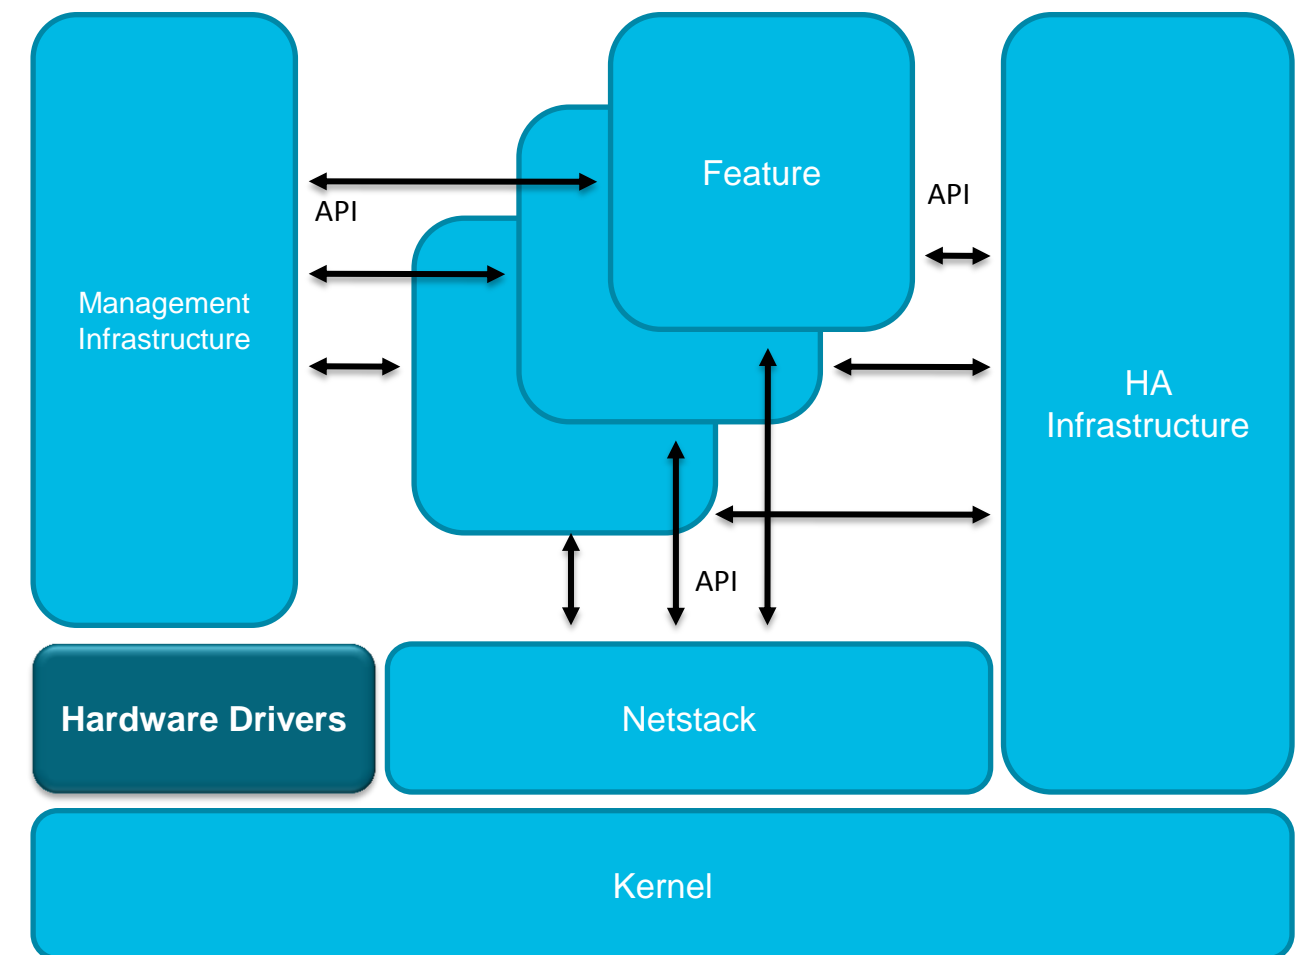
# NX-OS Kernel

- Linux 2.6 kernel
- Brings the benefits of Linux
  - Resilient Pre-emptive Multitasking
  - Multi-threaded
  - Scalable Multi-CPU/Core support
  - Constant development and enhancement



# NX-OS Platform Specific Portion

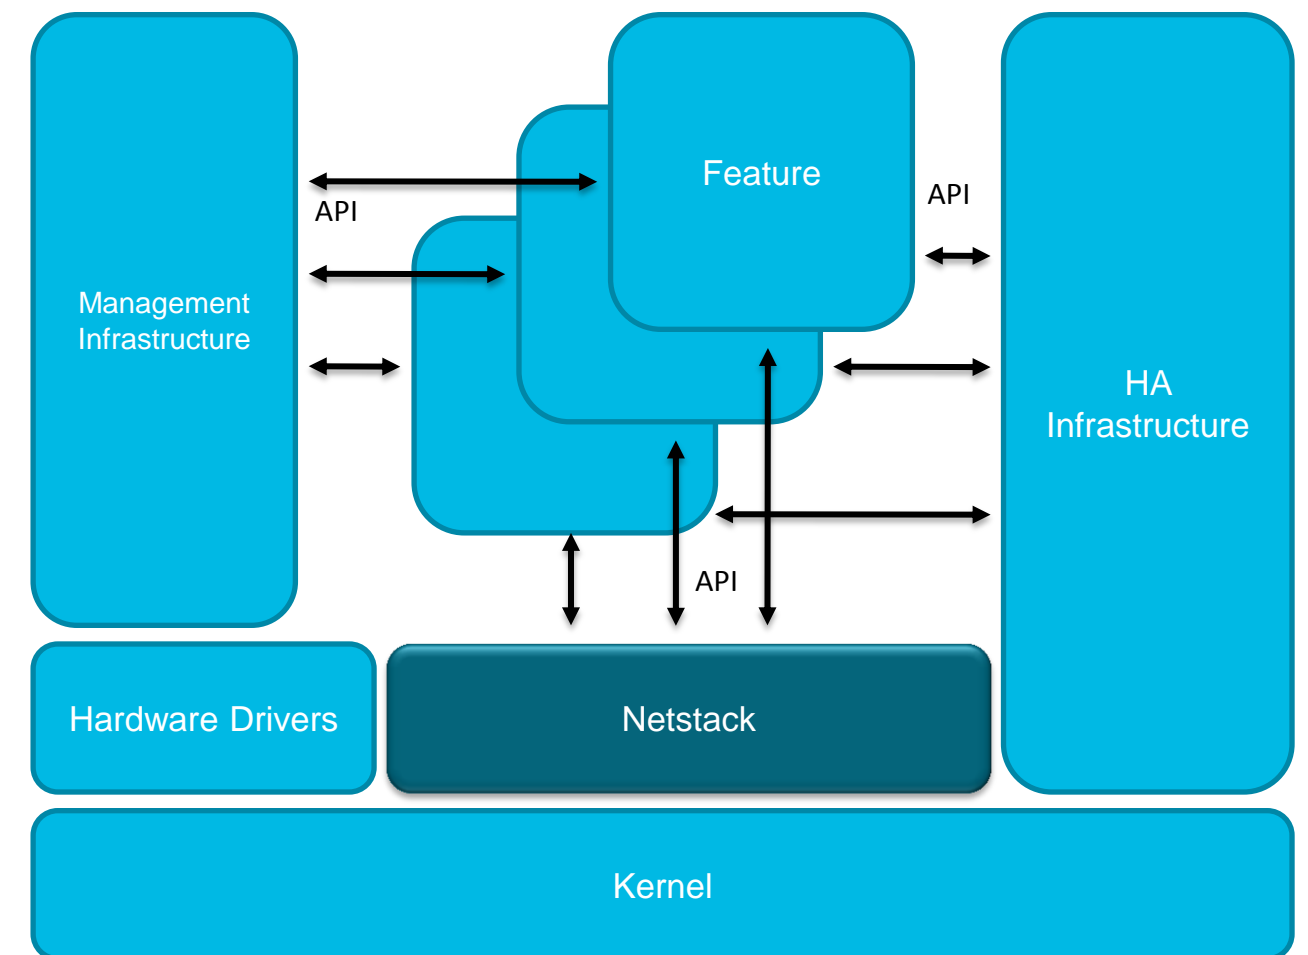
- Chipset specific code
- Provide Hardware Abstraction Layer (HAL)
- Ported per platform





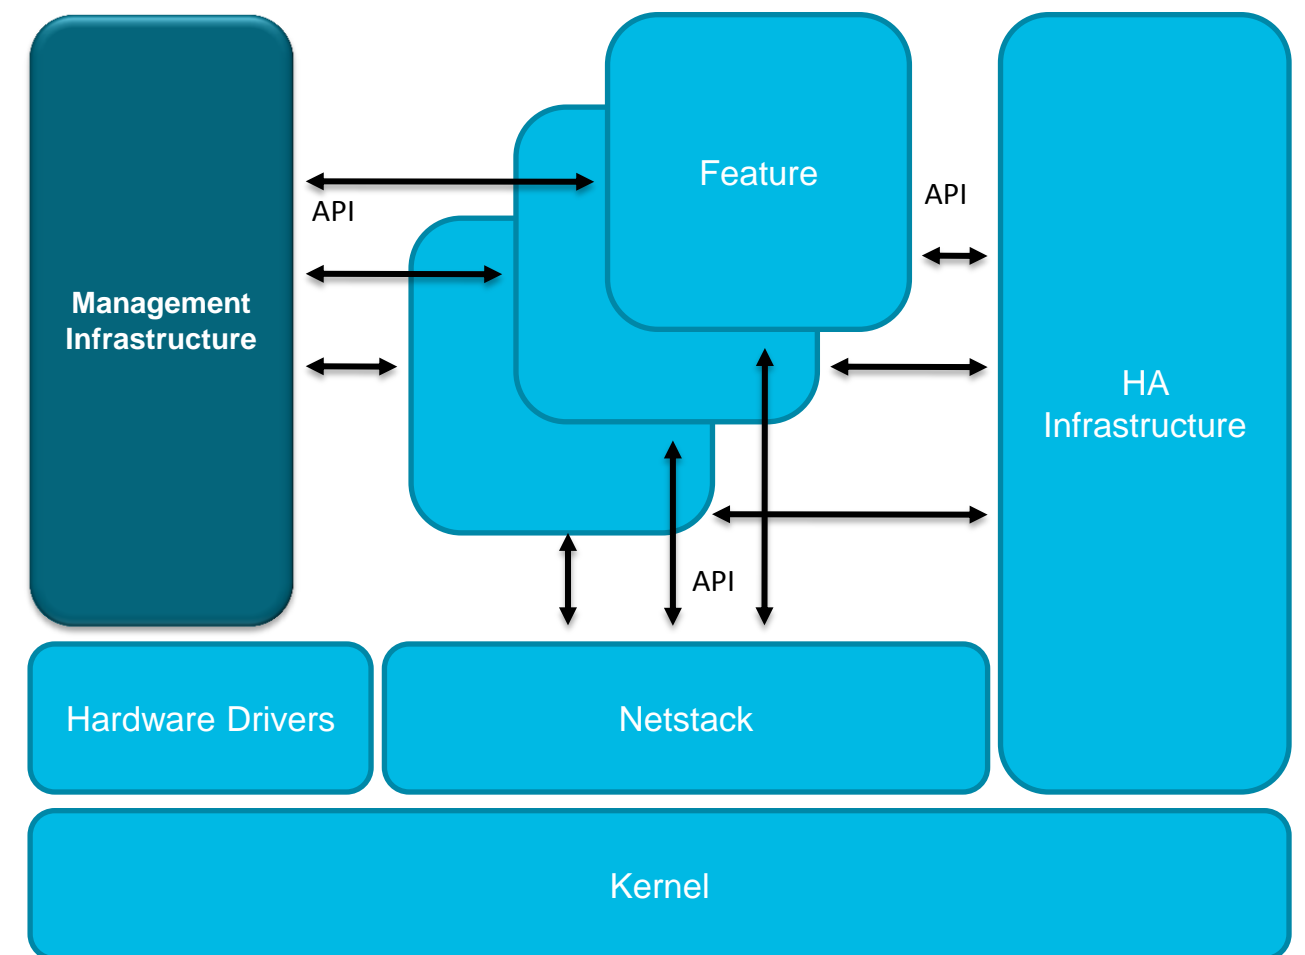
# NX-OS Netstack

- Complete network stack implemented in user space
  - L2 Packet Management/ARP
  - IPv4/IPv6
  - ICMPv4/ICMPv6
  - TCP/UDP & Socket Library
- Added Functionality
  - Virtualisation (VDCs/VRFs)
  - High-Availability (SSO)
- Added system stability



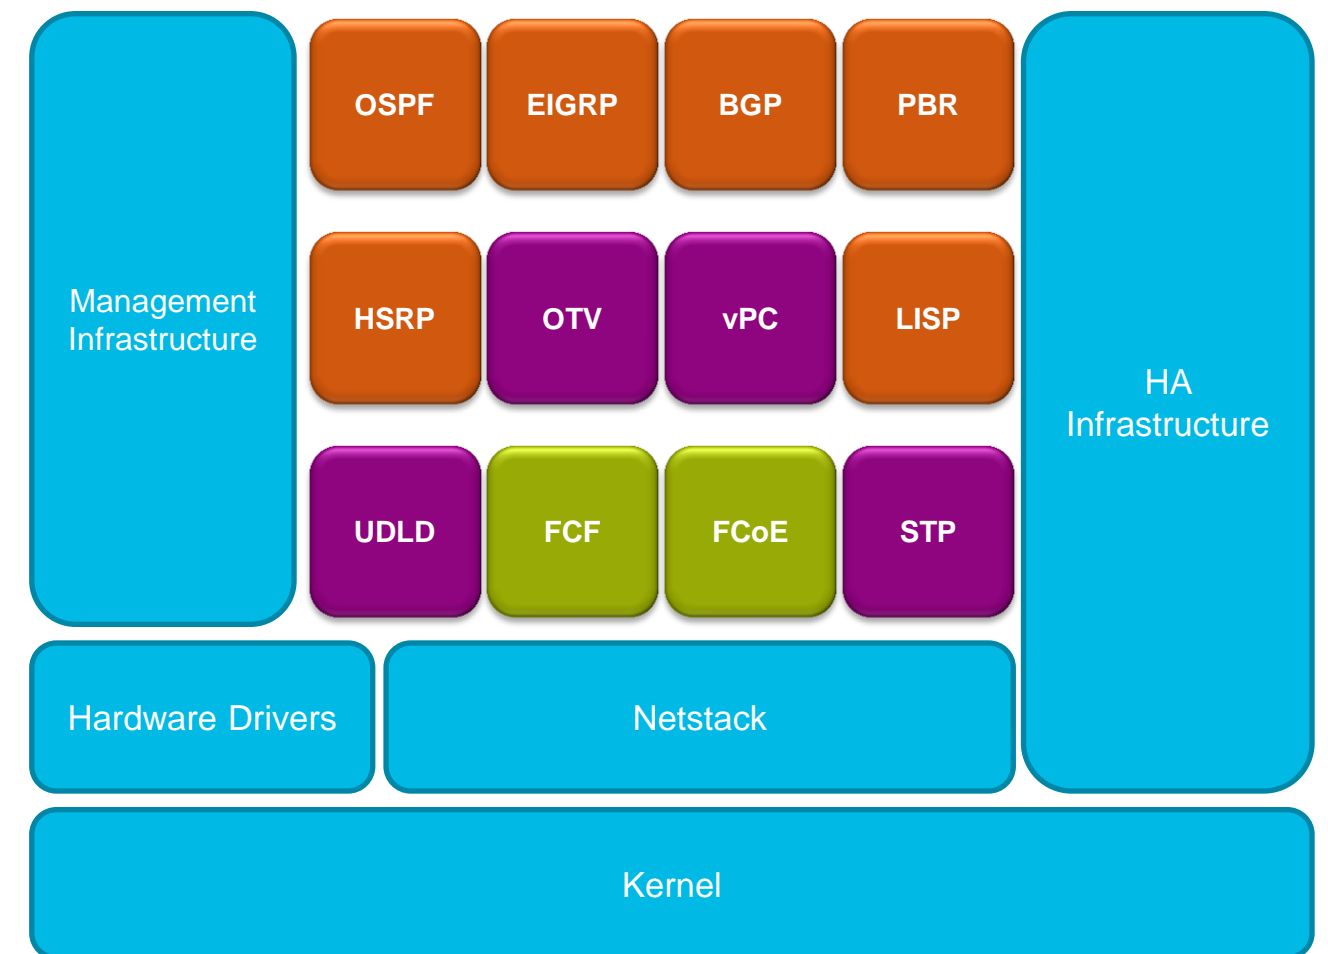
# NX-OS Management Infrastructure

- Provides CLI and configuration interfaces
- Provides SNMP agent
- Provides NETCONF/XML interface



# NX-OS Feature/Service Granularity

- Highly granular implementations
- Each service is an individual memory protected process
  - Including multiple instances of particular service
- Effective fault isolation between services
- Individually Monitored & Managed

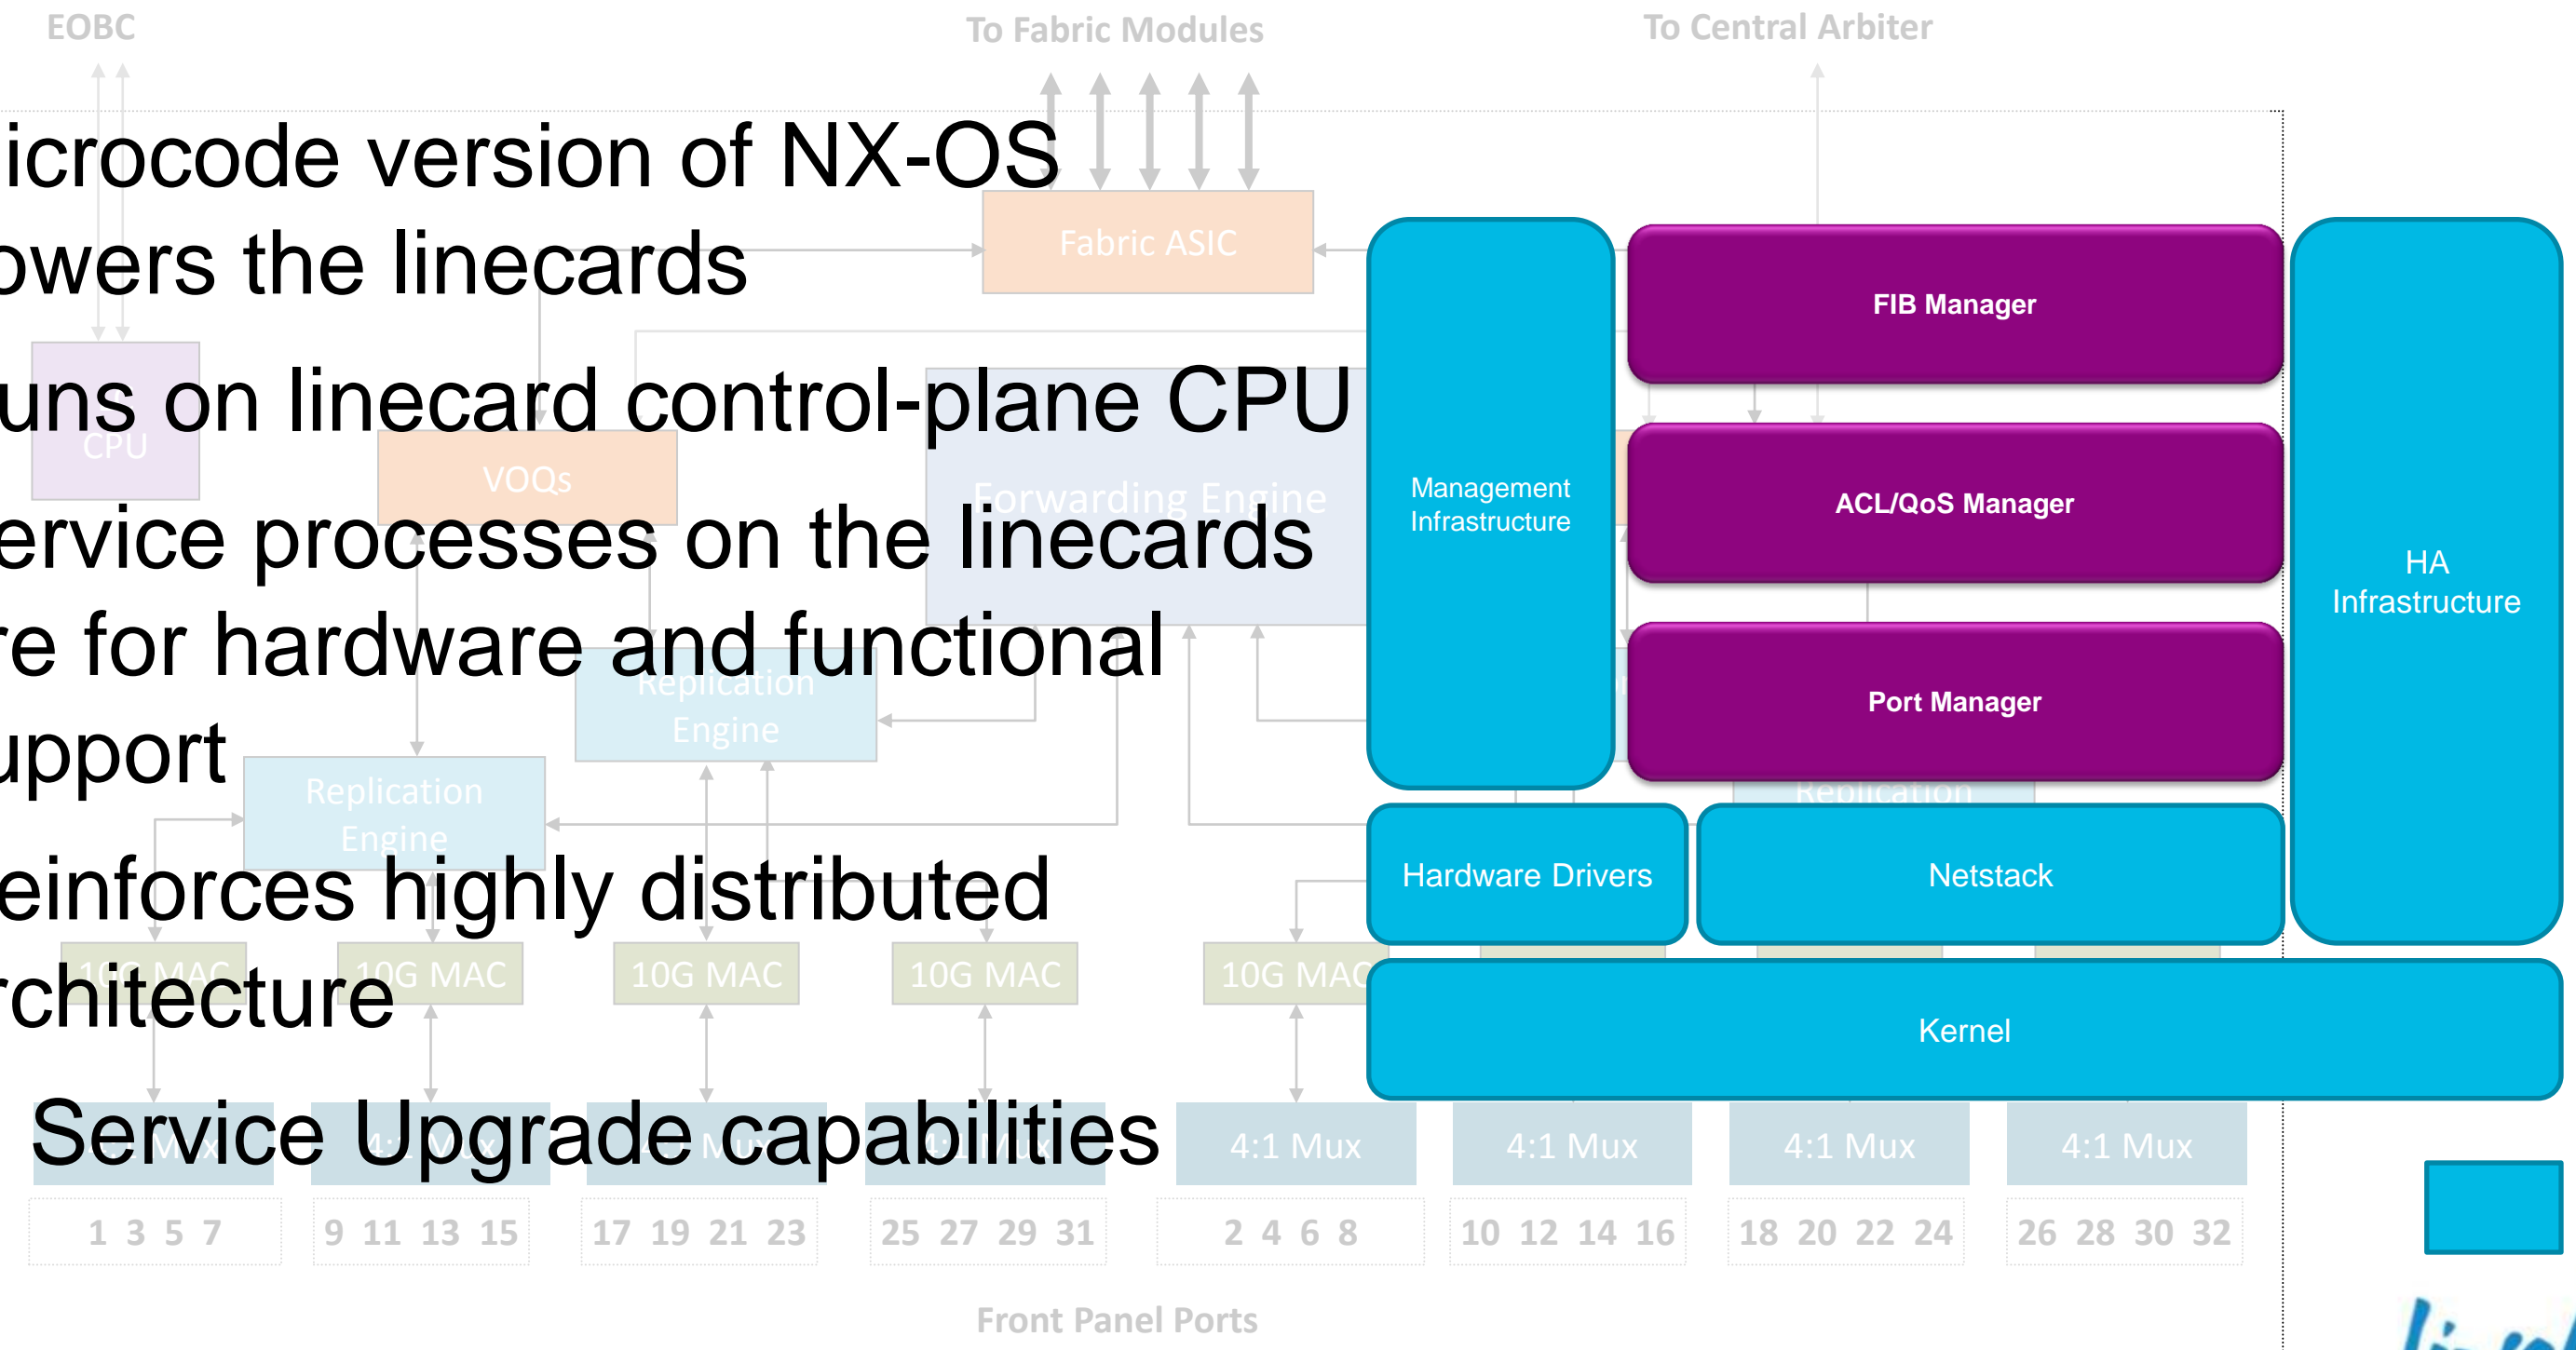


# NX-OS Conditional Features

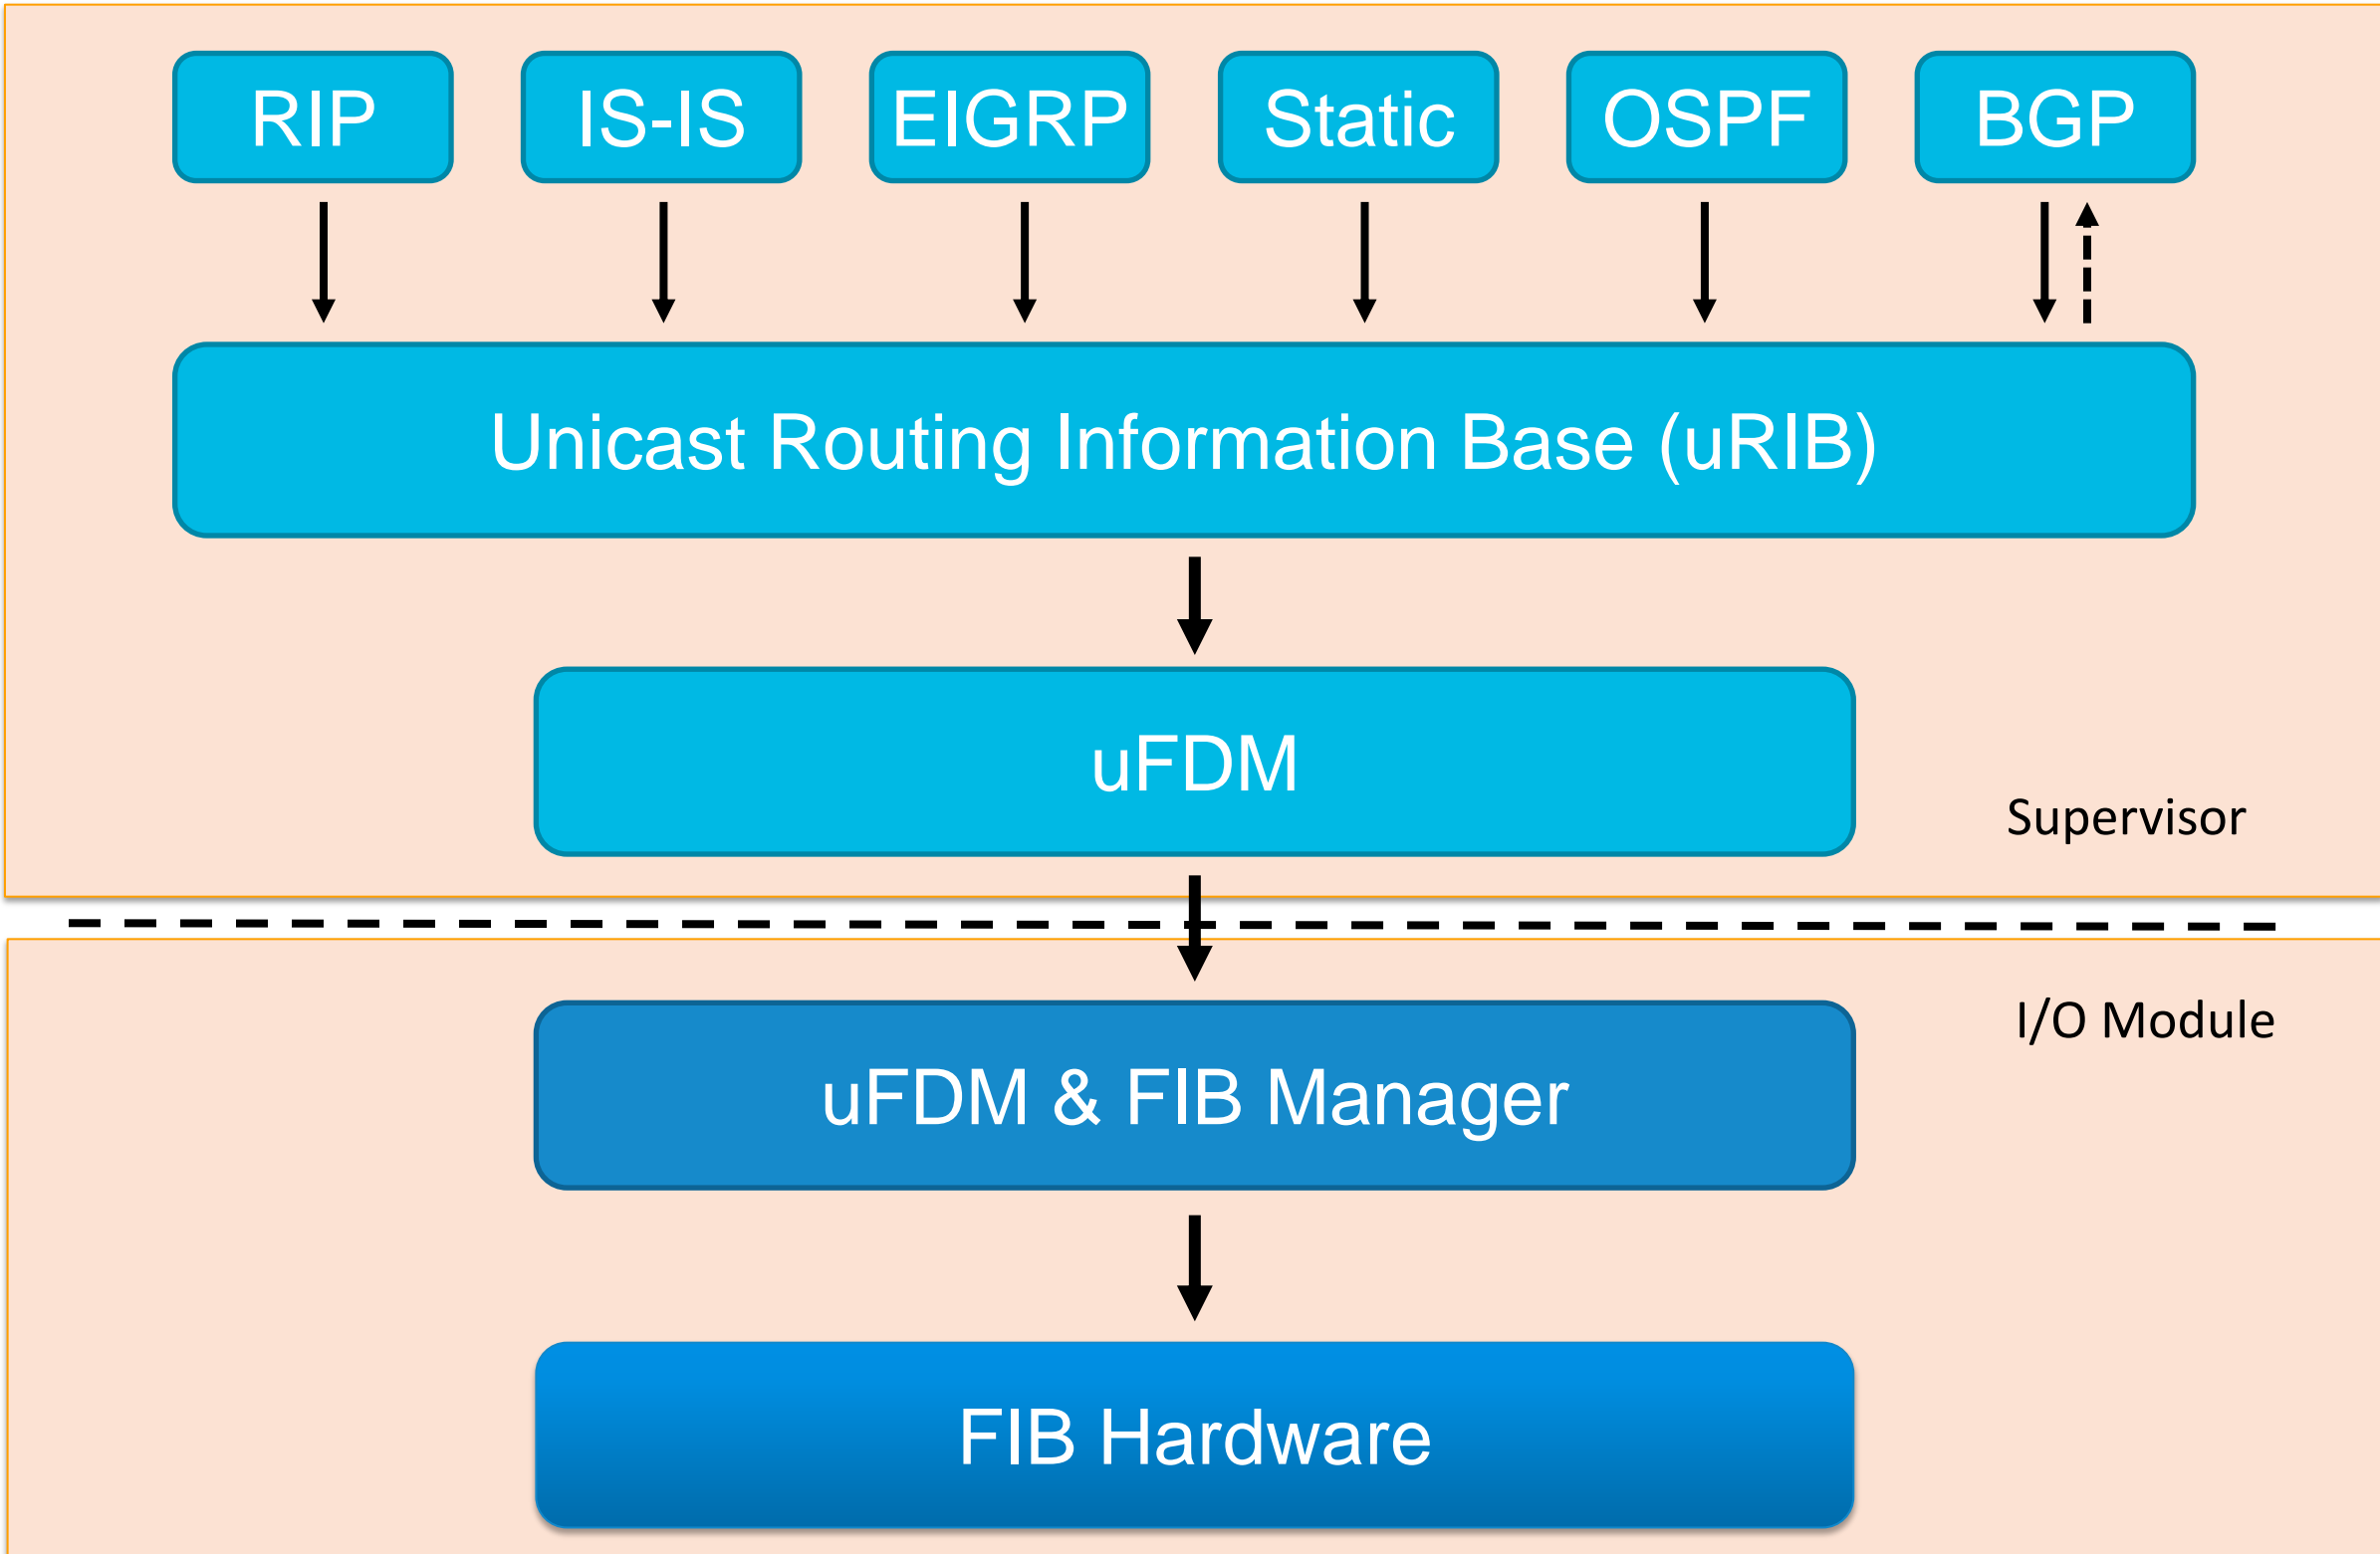
- Services (Protocols/Features) can be explicitly enabled/disabled
  - N7K-1(config)# feature ?
    - bgp Enable/Disable Border Gateway Protocol (BGP)
    - dot1x Enable/Disable dot1x
    - eigrp Enable/Disable Enhanced Interior Gateway Routing Protocol
    - hsrp Enable/Disable hsrp
    - igmp Enable/Disable Internet Group Management Protocol (IGMP)
- Disabling a service:
  - Releases associated resources
  - Removes associated CLI
  - Removes associated configuration

# NX-OS Runs on the Linecard

- Microcode version of NX-OS powers the linecards
- Runs on linecard control-plane CPU
- Service processes on the linecards are for hardware and functional support
- Reinforces highly distributed architecture
- In Service Upgrade capabilities



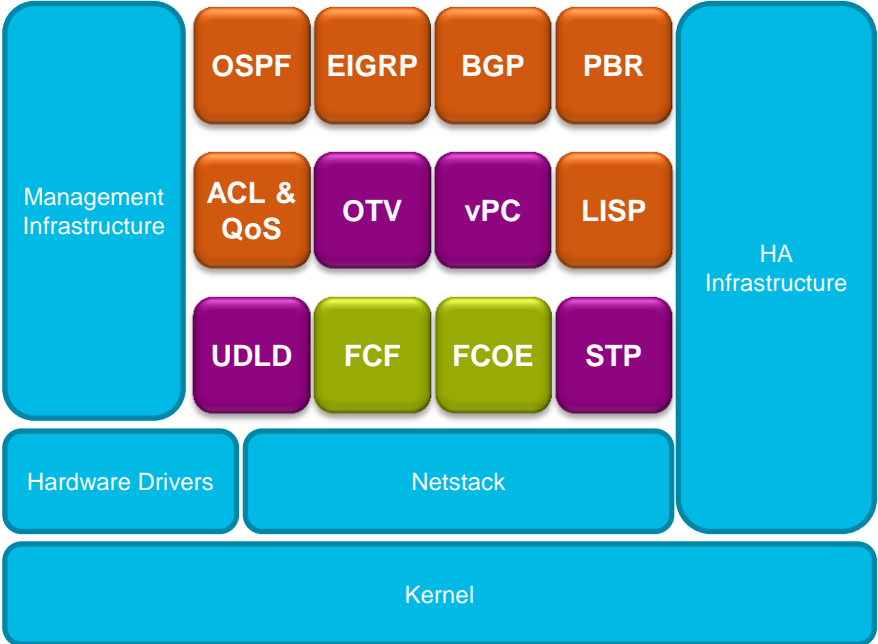
# Control-Plane/Data-Plane Separation



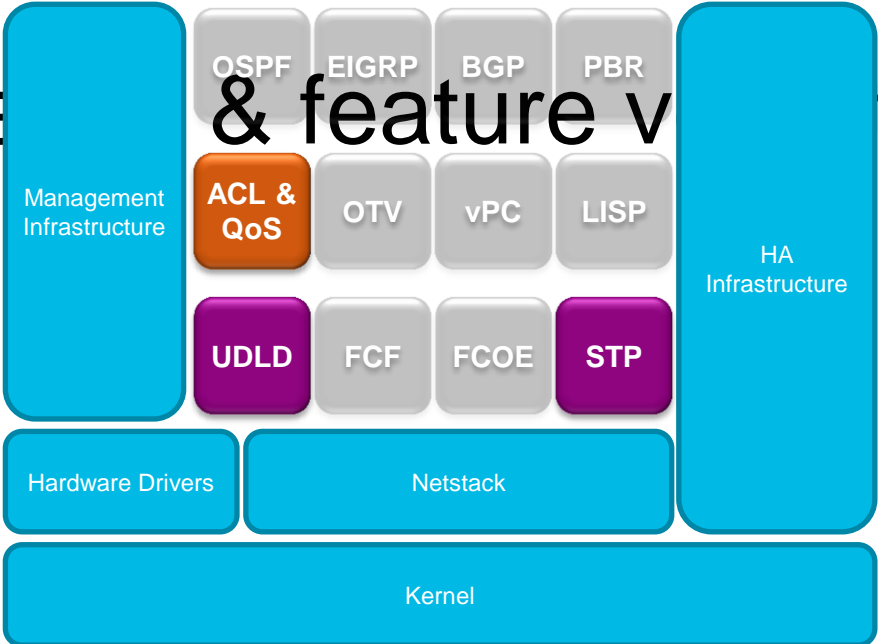


# NX-OS Platform Packaging and Delivery

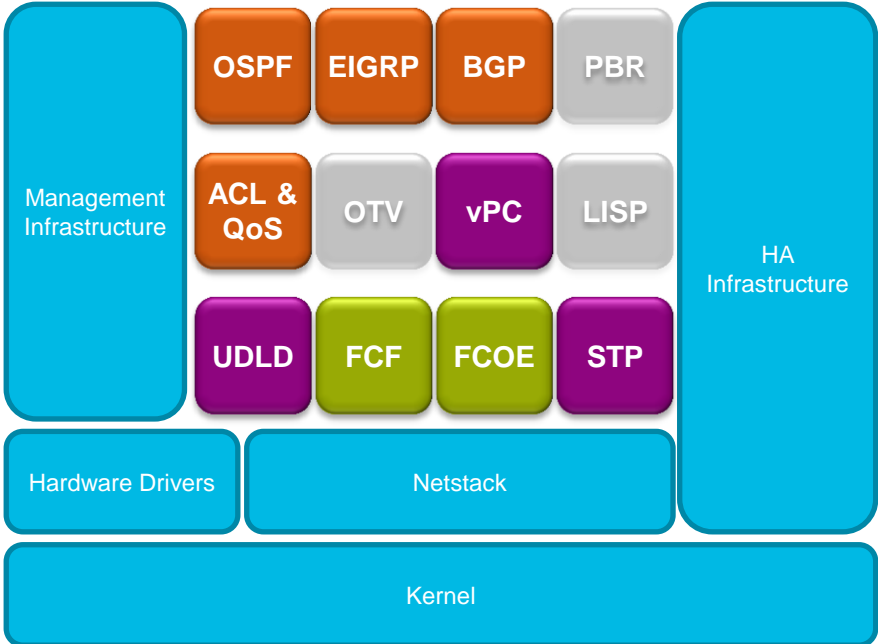
- Modular nature of NX-OS allows delivery of “permutations” based on hardware capabilities
- Kernel, core infrastructure code, and APIs remain consistent
- Minimises development
- Maximises code re-use & feature variety



Nexus 7000



Nexus 1000v



Nexus 5500

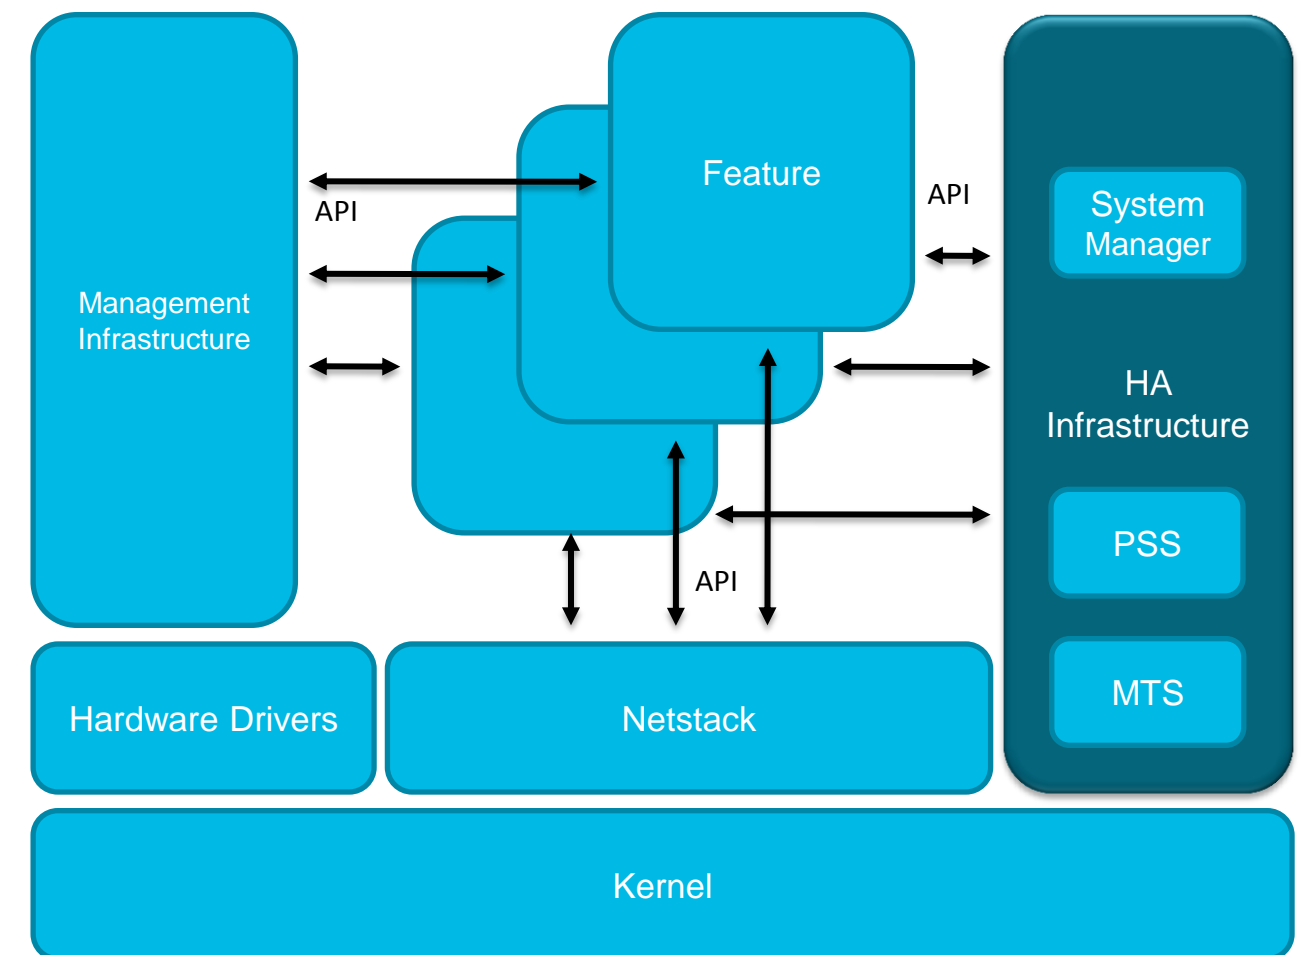
# Session Agenda

- NX-OS Origins & Overview
- NX-OS Modular Architecture
- **High-Availability Infrastructure**
- High-Availability Features & Capabilities
- Command Line Interface
- Operational & Management Features
- Licensing & Lifecycle
- Innovation
- Conclusion



# NX-OS High-Availability Infrastructure

- Actually composed of 3 sub-services
  - System Manager
  - Message & Transaction Service (MTS)
  - Persistent Storage Service (PSS)



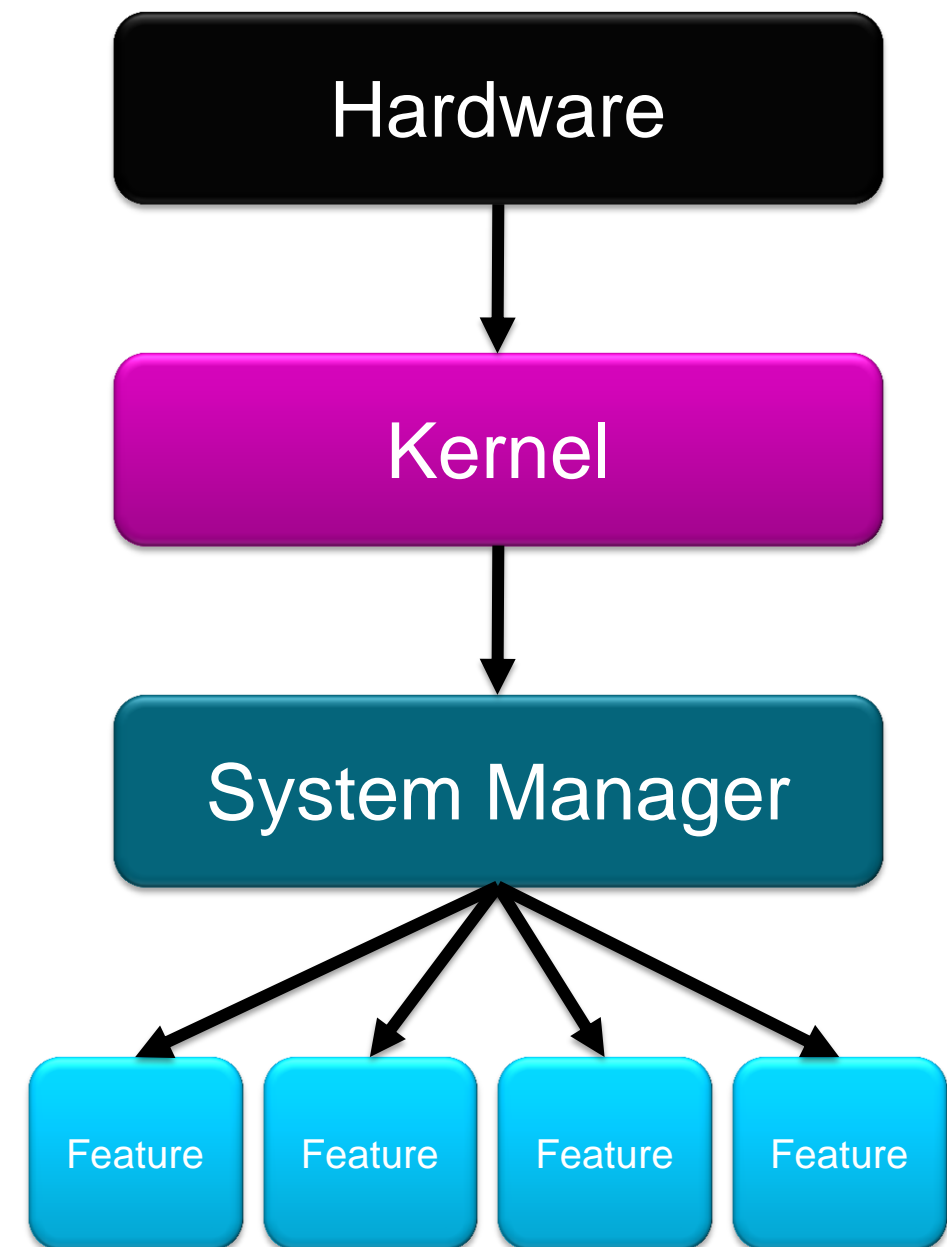
# NX-OS System Manager

- Centre of service management and fault recovery
- Acts like Unix-like 'init' process.
- Starts up configured features/services
- Heartbeats received from services



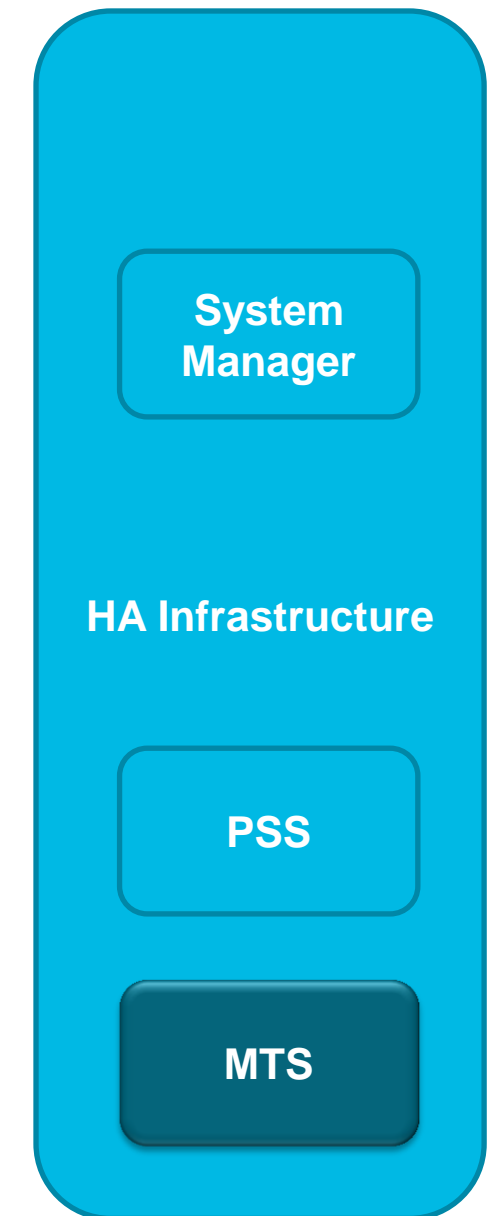
# Hierarchical Fault Detection & Recovery

- System Manager monitors services
  - Exit codes for crashes
  - Heartbeat for freeze-ups
  - Can kill/restart child processes
- Kernel monitors System Manager
- Hardware monitors Kernel



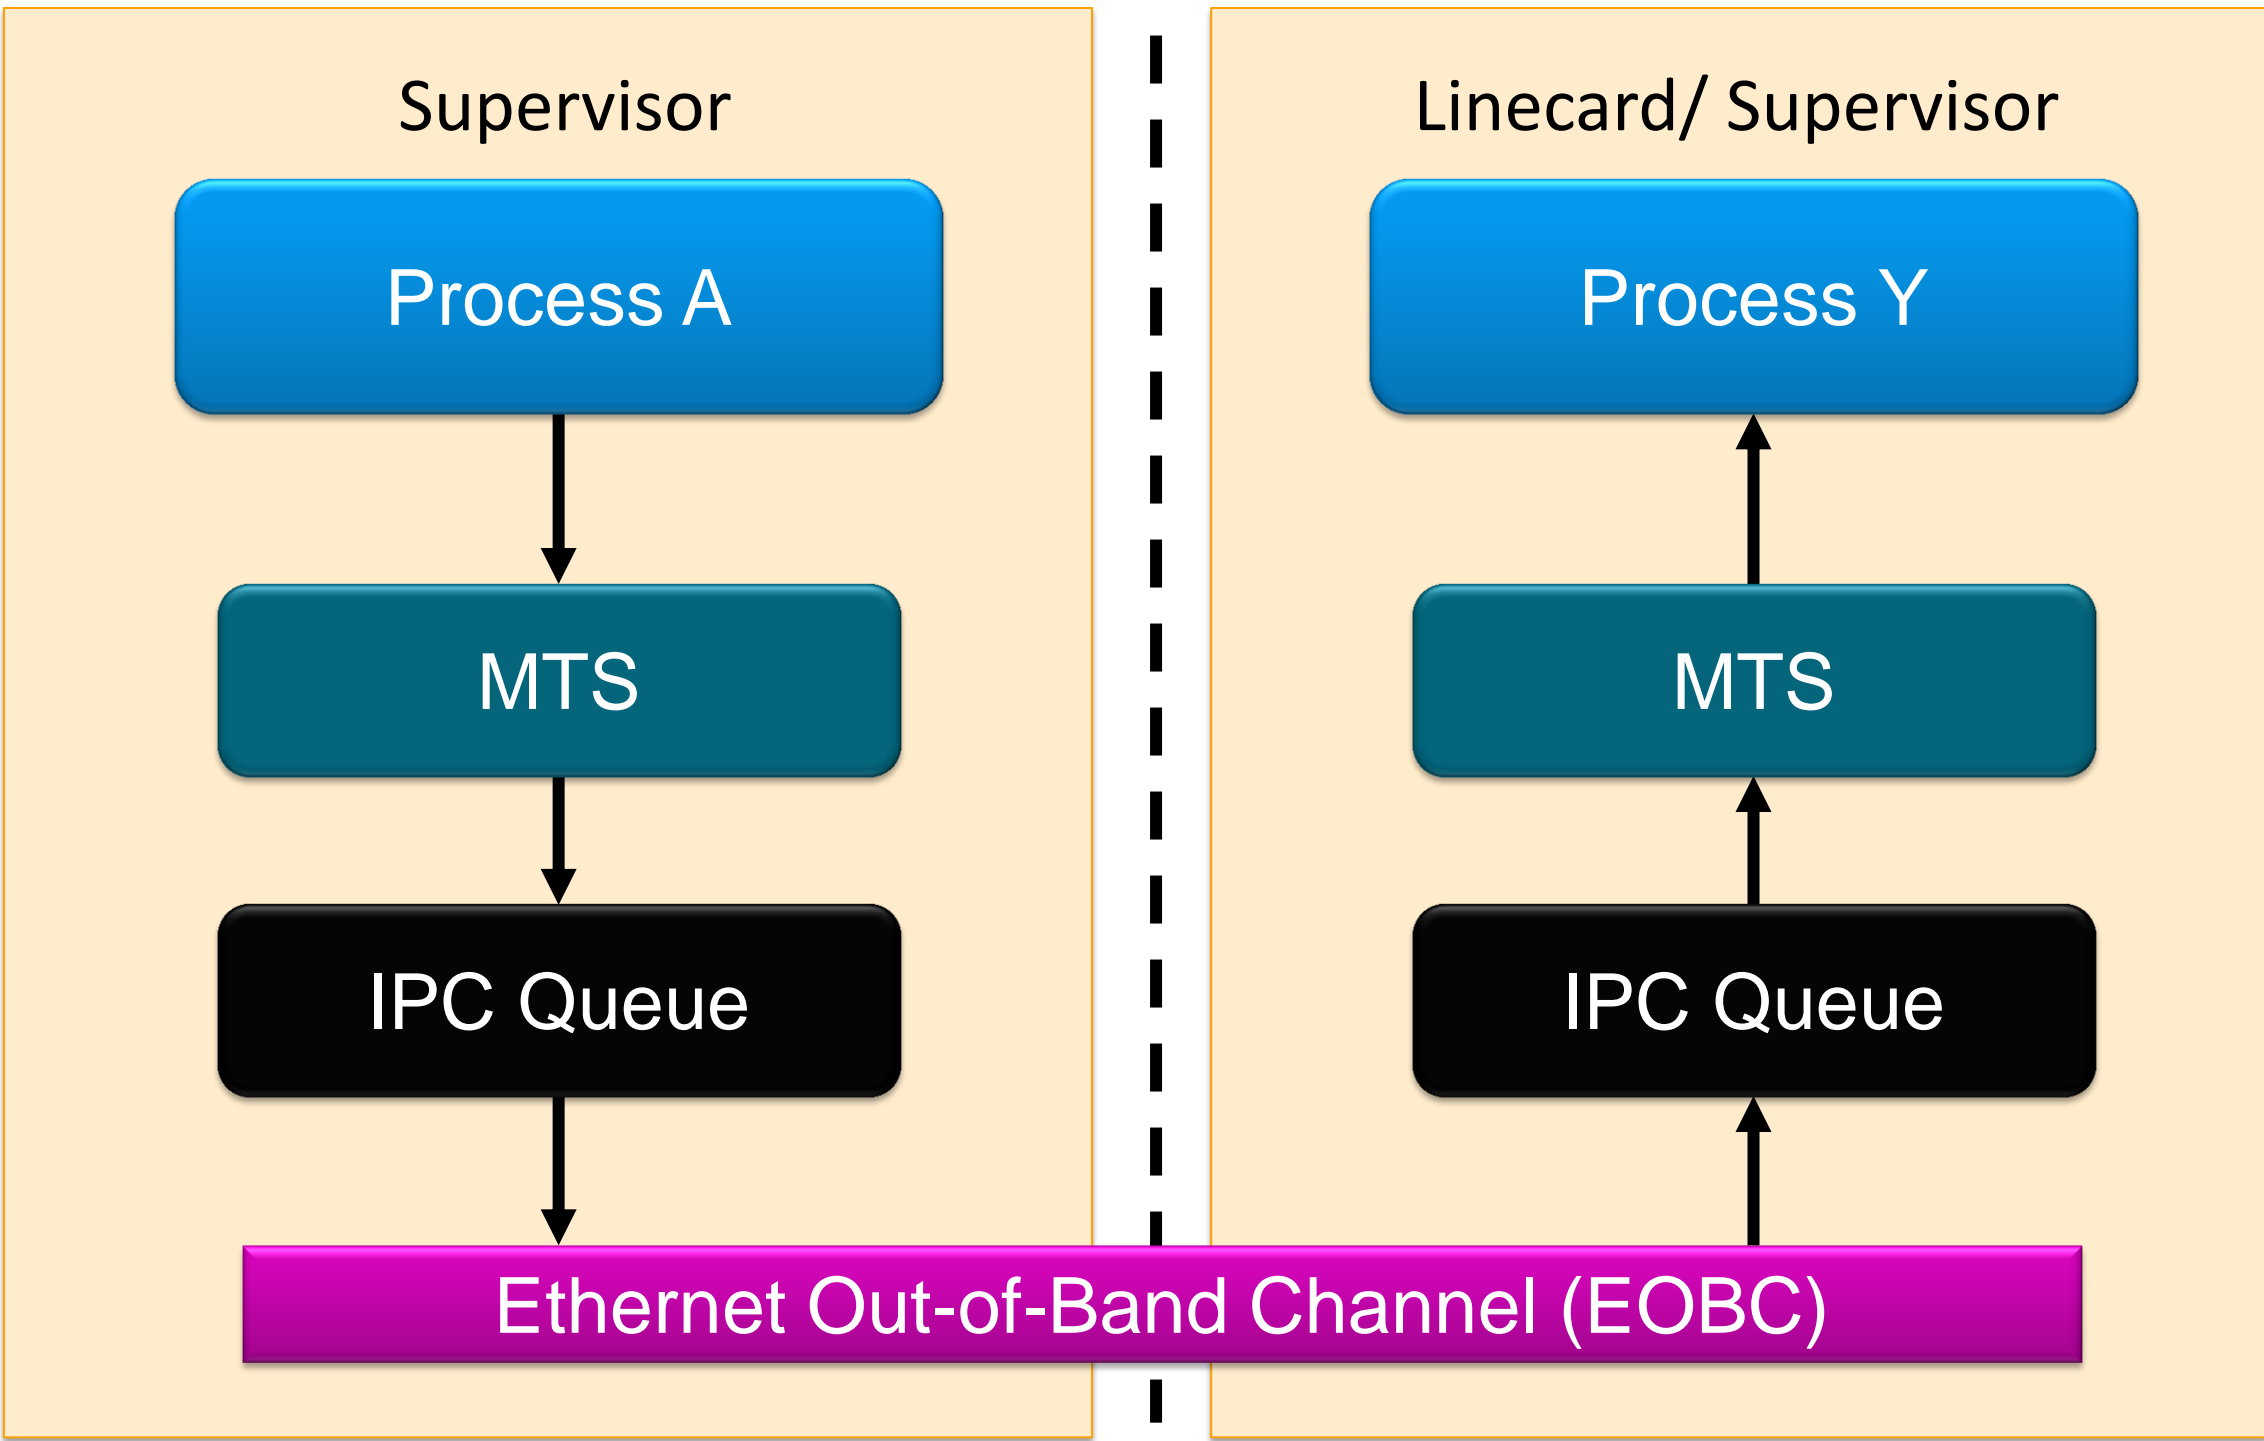
# NX-OS Message & Transaction Service

- Message relay system for IPC communications
- Provides reliable unicast & multicast delivery
- Used for service-to-service and module-to-module messaging



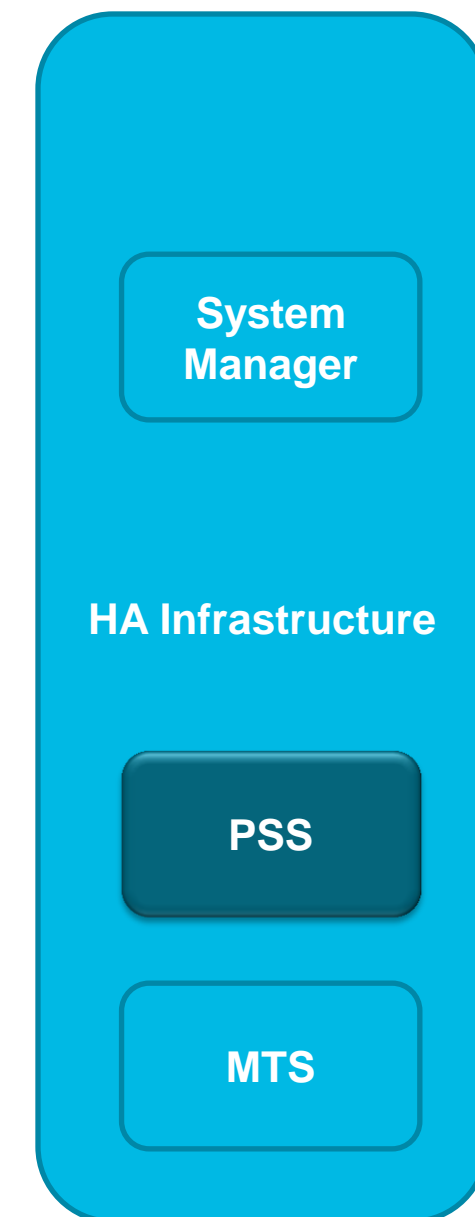


# Process Messaging Across Slots



# NX-OS Persistent Storage Service

- Lightweight key/value database
- Provides store options for DRAM or NVRAM
- API for services to store data
- Used to maintain runtime data/state



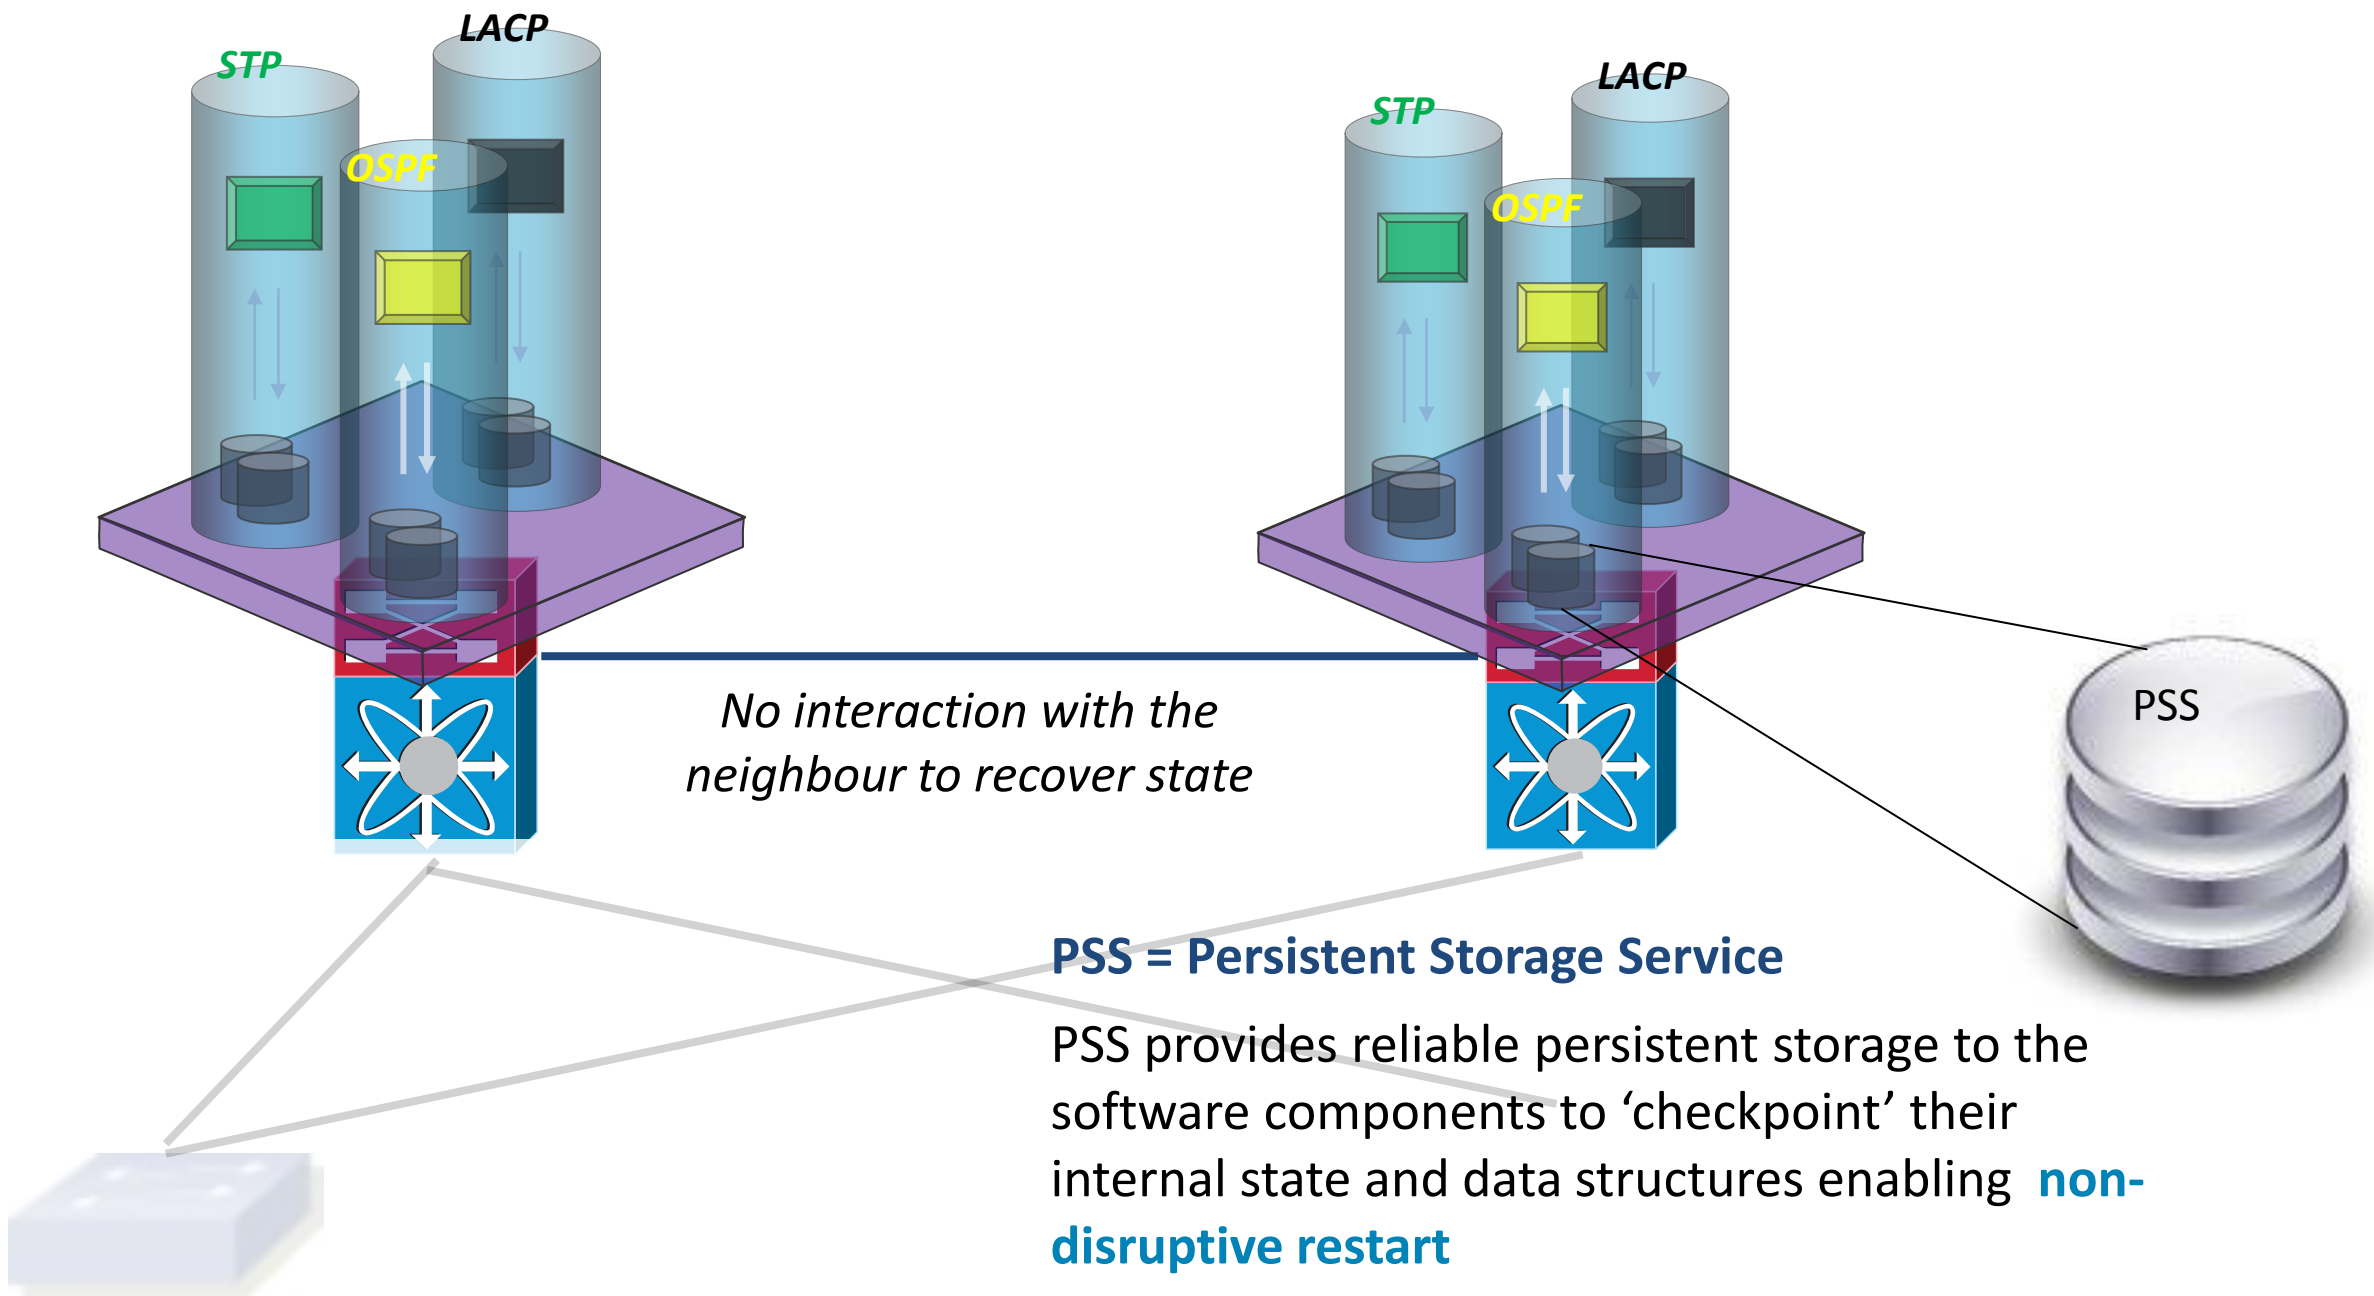
# Session Agenda

- NX-OS Origins & Overview
- NX-OS Modular Architecture
- High-Availability Infrastructure
- **High-Availability Features & Capabilities**
  - Conclusion
- Command Line Interface
- Operational & Management Features
- Licensing & Lifecycle
- Innovation



# NX-OS High Availability

## Stateful Process Restart\*



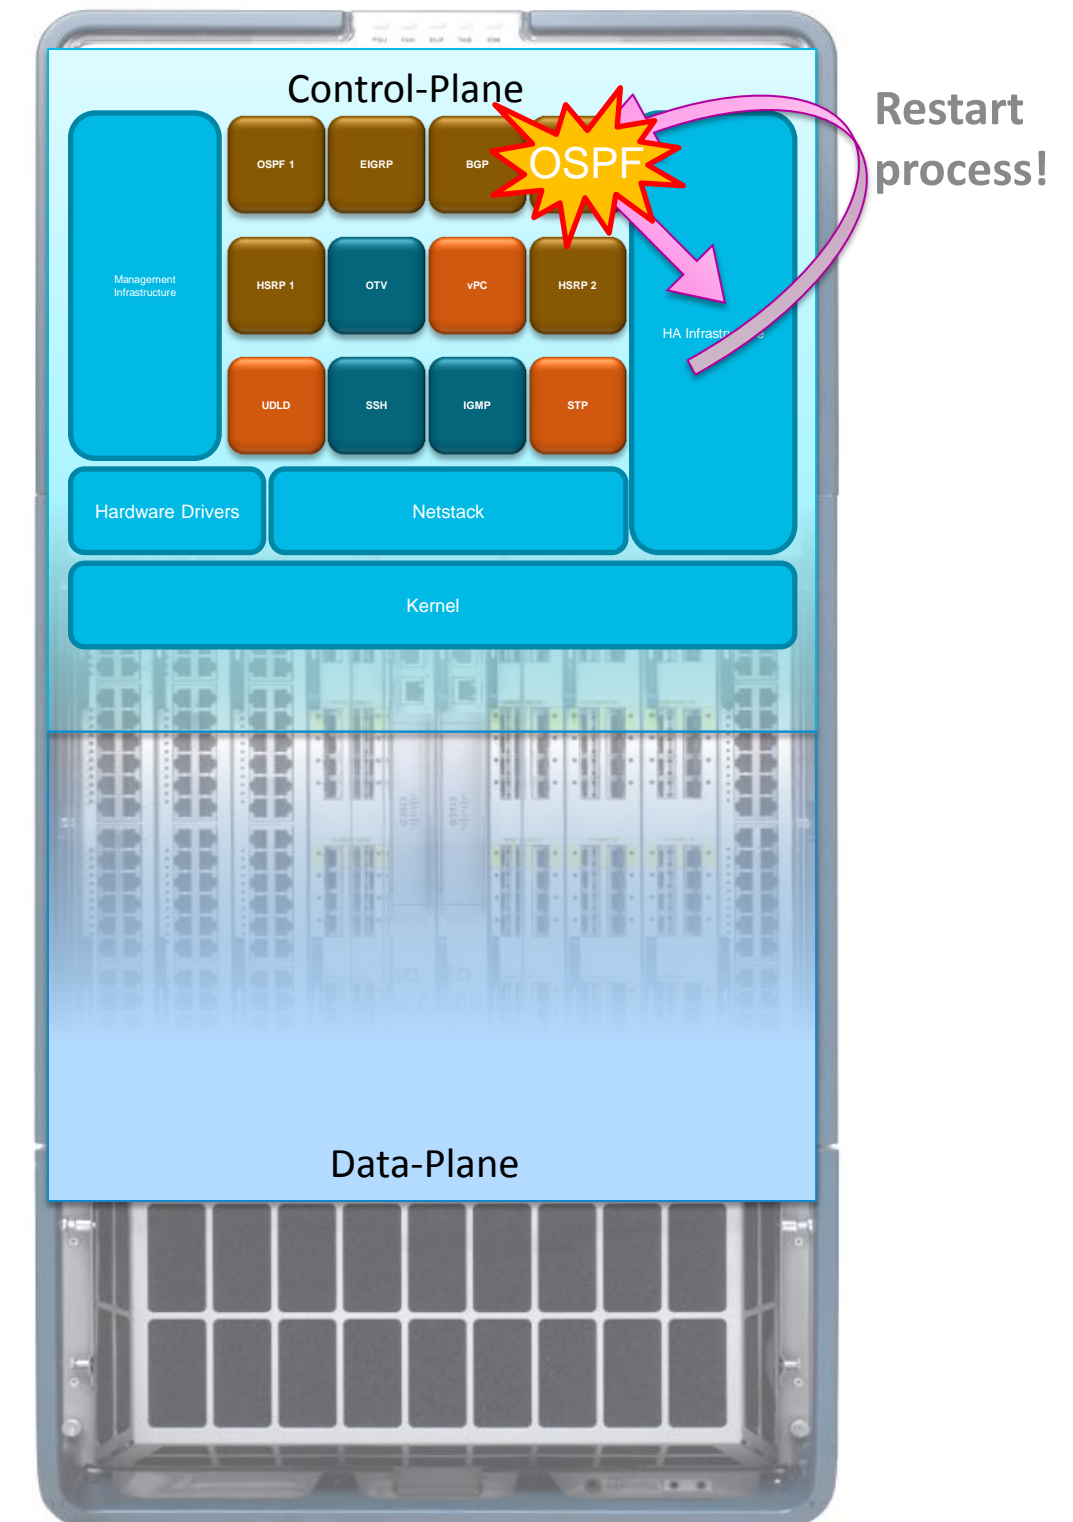
\*NSR with stateful restart is supported for OSPF v2 and v3 as well as ISIS

# NX-OS Stateful Process Restart

- NX-OS services checkpoint their runtime state to the PSS for recovery in the event of a failure

If a fault occurs in a process...


- HA manager determines best recovery action (restart process, switchover to redundant supervisor)
- Process restarts with no impact on data plane
- State is recovered, operation resumes
- Total recovery time: ~10s ms





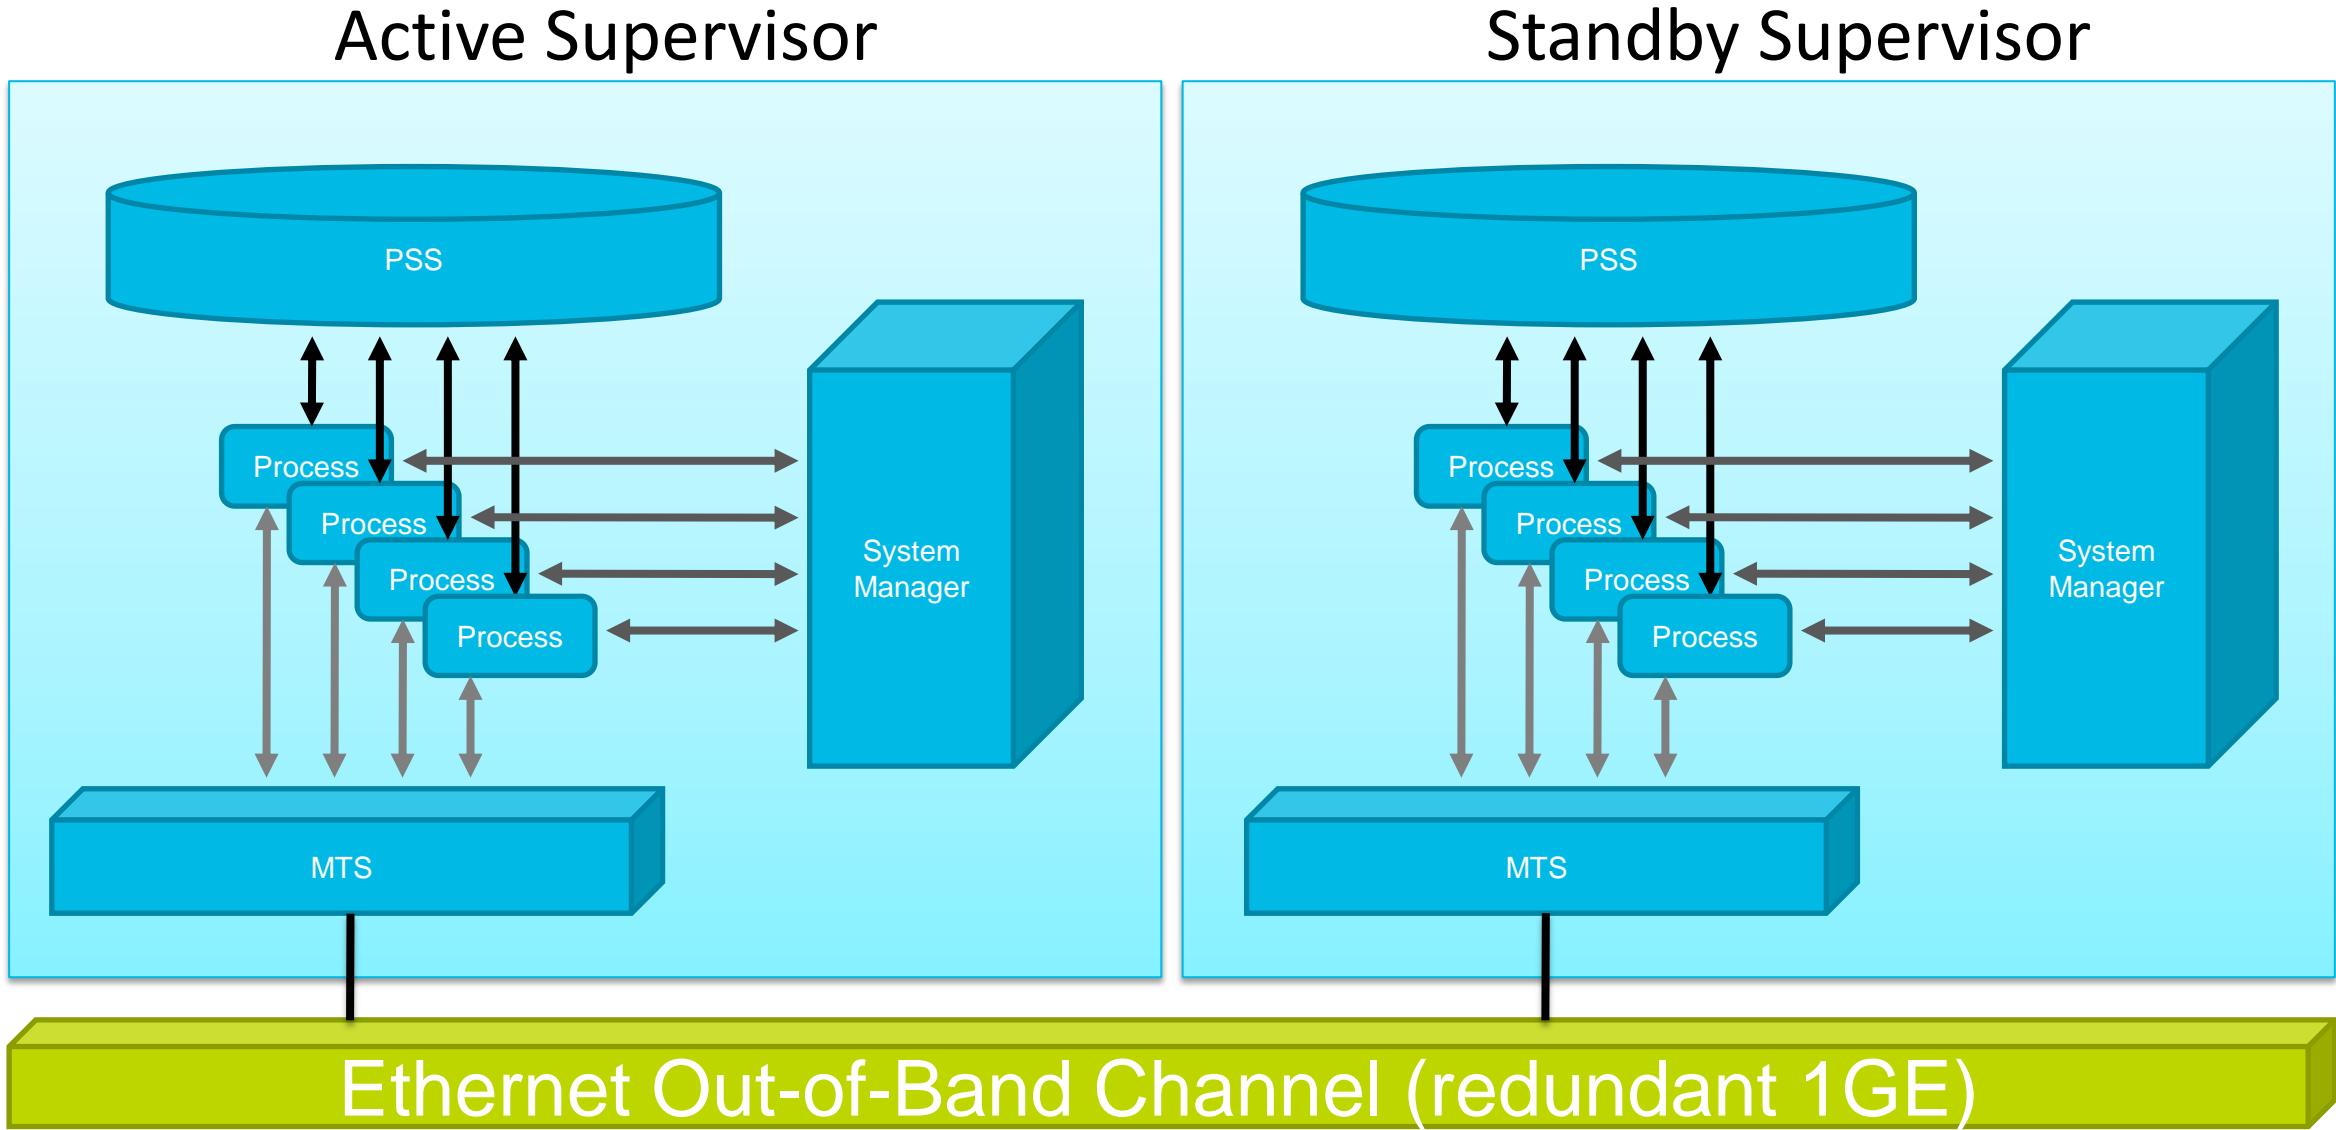
# NX-OS High Availability

## Graceful Restart & NSF



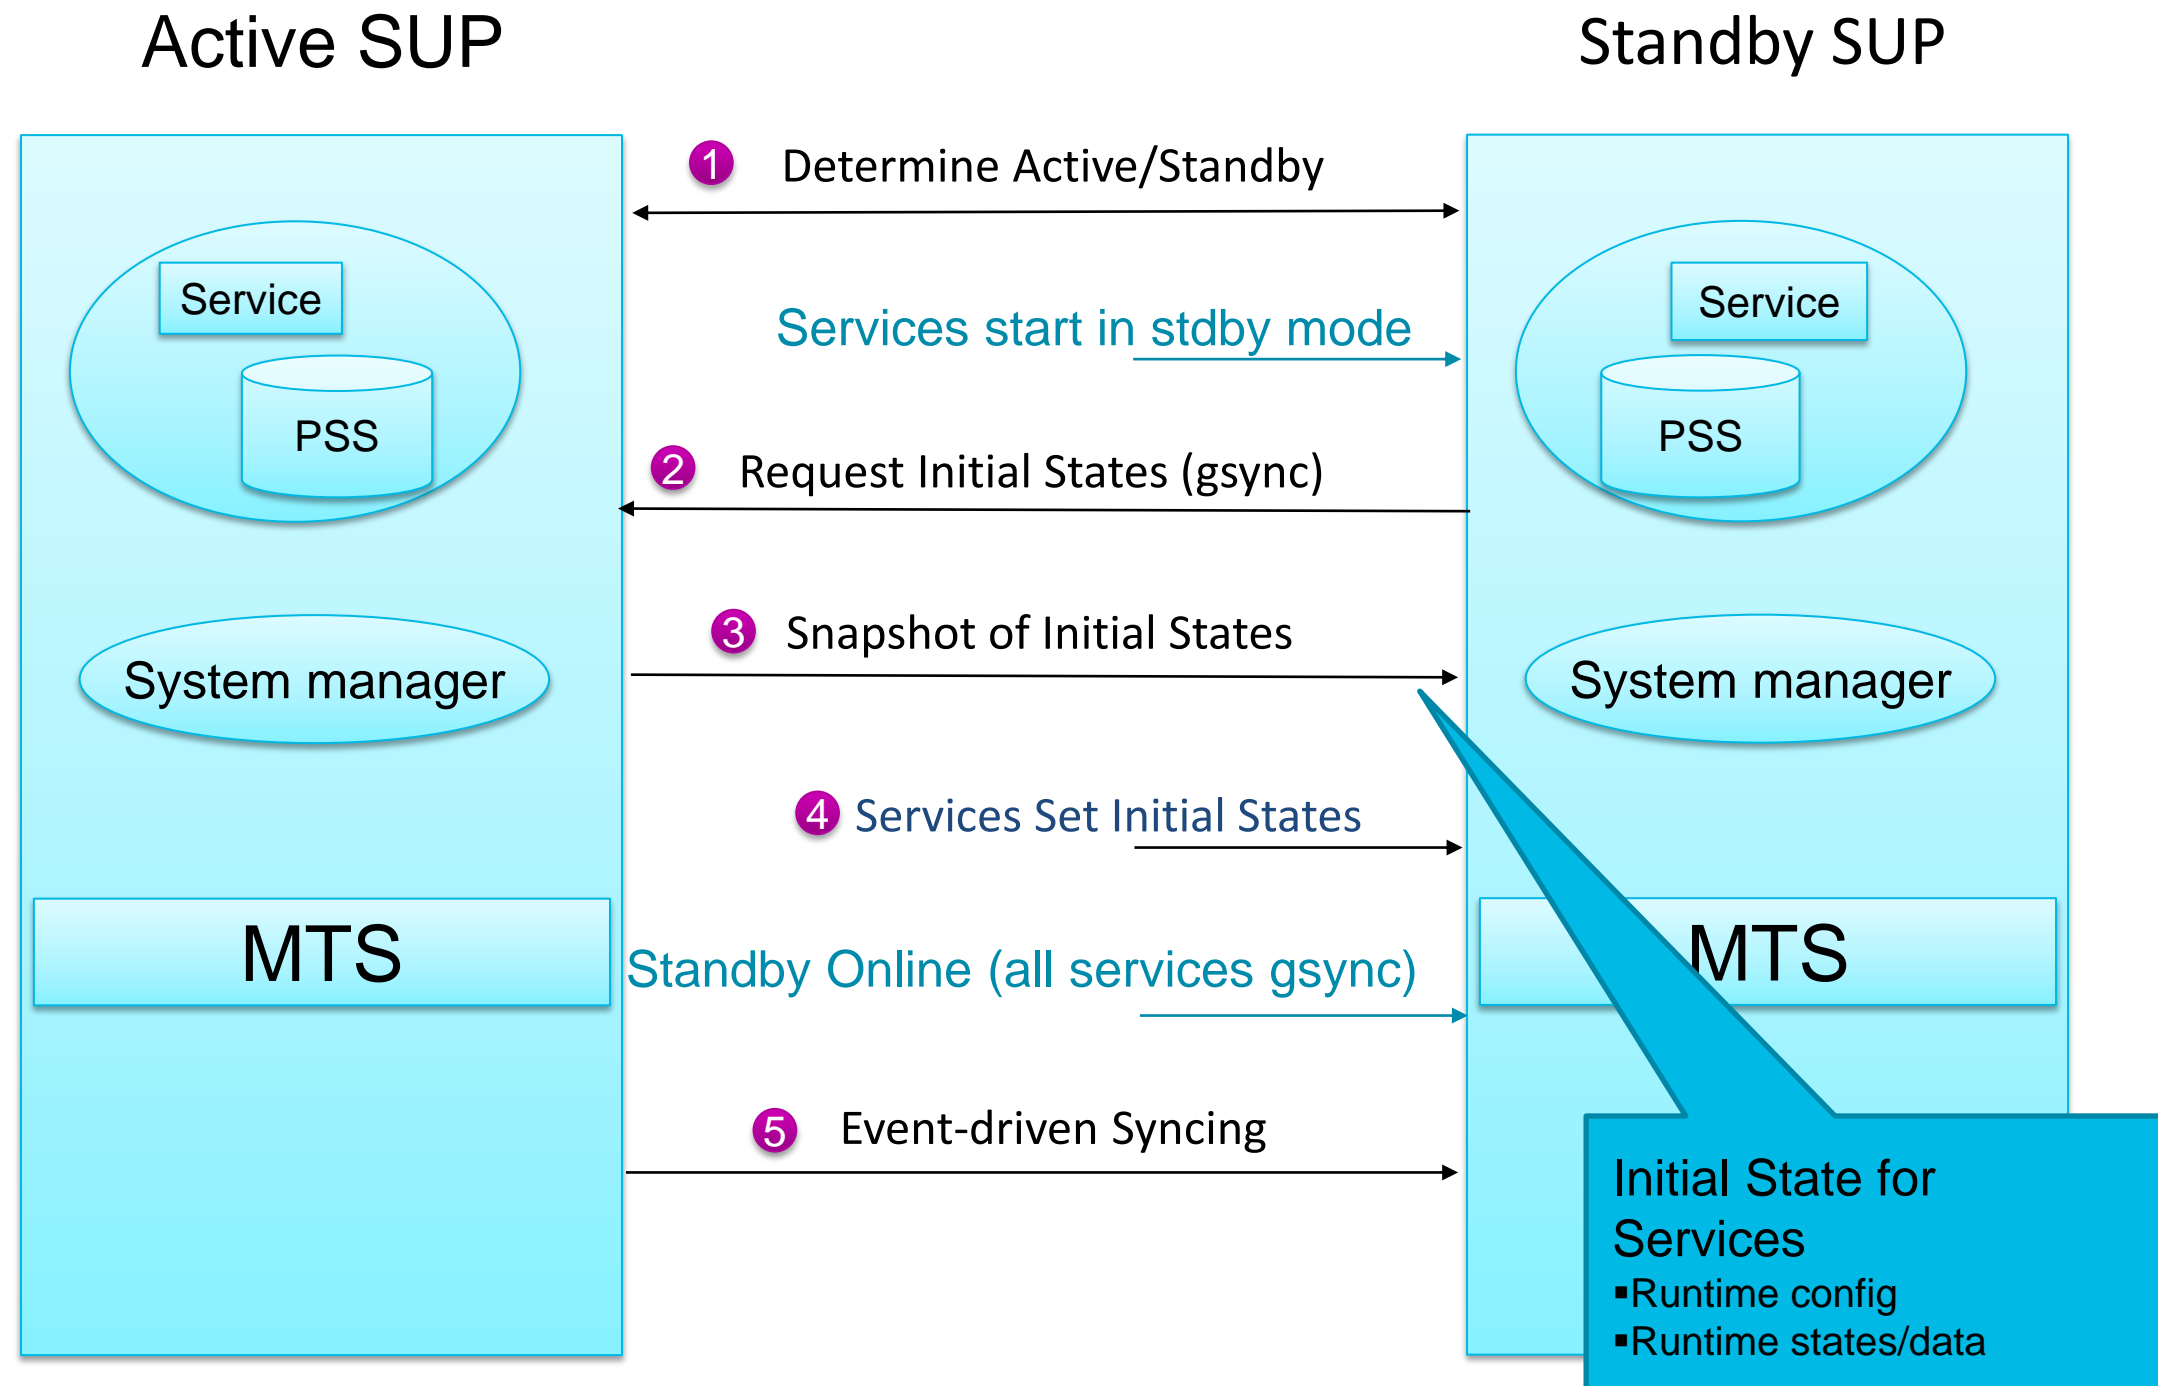
# NX-OS High Availability

## Redundant Supervisor Model\*



\*Applies to those system with dual supervisors

# Active and Standby Supervisor Syncing



# Stateful Supervisor Switchover

- Active/Standby

```
Nexus-Dual-Sup# show system redundancy status
Redundancy mode
-----
      administrative:  HA
      operational:    HA

This supervisor (sup-1)
-----
      Redundancy state:  Active
      Supervisor state:  Active
      Internal state:    Active with HA standby


Other supervisor (sup-2)
-----
      Redundancy state:  Standby
      Supervisor state:  HA standby
      Internal state:    HA standby
```

- Fast switchover time – State is already in place
- Switchover initiated if:
  - ☹ repeated critical process restart failures
  - ☹ kernel failures
  - ☹ supervisor hardware failure detected by Diagnostics

# In-Service Software Upgrade

```
N7K# install all kickstart bootdisk:5.0-kickstart system bootdisk:5.0-system  
N7K#
```

Release 5.0



- 3 Perform SSO
- 4 Upgrade standby supervisor
- 5 Reload standby supervisor

- 1 Upgrade standby supervisor
- 2 Reload standby supervisor

Release 5.0


- 6 Upgrade LCs in series\*

\* Parallel upgrade of the I/O modules supported on the Nexus 7000 from 5.2





# Hitless ISSU on the Nexus 5x00



- Difference in the detailed operation from Nexus 7k
  - Single supervisor/control-plane vs. dual supervisor
  - L3 ISSU not Supported Today on the 5k
  - ISSU not possible if non-edge STP designated port (only works in the access)
- During ISSU, control plane functions are temporarily suspended.
- Control Plane restored within 80 seconds
- Hitless ISSU of the Nexus 2000s (Nexus 5x00/Nexus 7000)
- Supported from NX-OS 4.2(1)N1

# Session Agenda

- NX-OS Origins & Overview
- NX-OS Modular Architecture
- High-Availability Infrastructure
- High-Availability Features & Capabilities
- Command Line Interface
- Operational & Management Features
- Licensing & Lifecycle
- Innovation
- Conclusion



# NX-OS CLI Highlights

## Improved Over IOS Model

- Hierarchy Independent CLI allows 'show' commands to be executed from exec-mode or config-mode

```
N7K(config)# show interface e1/1
Ethernet1/1 is up
  Hardware: 10/100/1000 Ethernet, address: 001b.54c1.5d44 (bia 001b.54c1.5d44)
  MTU 1500 bytes, BW 100000 Kbit, DLY 10 usec,
    reliability 255/255, txload 1/255, rxload 1/255
<snip>
```

- Support for CIDR 'slash' notation for IPv4/IPv6 masks

```
N7K(config)# int e1/1
N7K(config-if)# ip address 192.168.0.1/23
```

- Mode-aware CLI history

```
N7K# show cli history ?
<CR>
  config-mode      Display history of config commands only
  exec-mode        Display history of exec commands only
  this-mode-only   Display history from current mode only
  unformatted      Display just the commands
N7K# show cli history config-mode
12 05:20:34      int e1/1
13 05:20:42      where detail
```

# Review Configuration with Flexibility

- Display feature-specific configuration

```
N7K# show running-config ntp

ntp server 171.68.10.80 use-vrf management
ntp server 171.68.10.150 use-vrf management
ntp source 172.26.244.101
clock format 12-hours
clock format show-timezone
```

- Exclude features with lengthy configuration (e.g. ACL, QoS, etc.)

```
N7K# show running-config exclude
aaa                cert-enroll        diagnostic        ntp                track
acllog             cfs                    eem              radius             vshd
aclmgr             cmp                    ipqos            rpm
callhome           copp                   license           security
cdp                dhcp                   monitor           spanning-tree
```

- Compare between startup- and running-configuration

```
N7K# copy running-config startup-config
[#####] 100%
N7K# config terminal
Enter configuration commands, one per line.  End with CNTL/Z.
N7K(config)# feature telnet
N7K(config)# sh running-config diff
*** Startup-config
--- Running-config
*** 1,11 ****
    feature lacp
--- 1,11 ----
+ feature telnet
  feature lacp
```

Identify the line number and difference between startup-config and running-config

# NX-OS Running-config Permutations

```
N7K# show running-config ?
<CR>
>          Redirect it to a file
aaa        Display aaa configuration
all        Current operating configuration with defaults
am         Display am information
arp        Display arp information
bgp        Display bgp information
callhome   Display callhome configuration
cdp        Display cdp configuration
cmp        Display CMP information
copp       show running config for copp
dhcp       Display dhcp snoop configurations
diagnostic Display diagnostic information
diff       Show the difference between running and startup
configuration
dot1x      Display dot1x configuration
eem        Show the event manager running configuration
eigrp      Display eigrp information
icmpv6     Display icmpv6 information
igmp       Display igmp information
interface  Interface configuration
ip         Display ip information
ipqos      show running config for ipqosmgr
```

- “show running-config” provides many enhancements



# IOS to NX-OS Conversion Tool

- Available today on cisco.com  
<http://tools.cisco.com/nxmt>
- Migrate Catalyst 6500/4500 configuration to Nexus 7000/5x00
- Visit booth 2940 in the World of Solutions for a demo

The screenshot shows the 'IOS-NXOS Migration Tool' web interface. The browser address bar displays the URL: [http://tools.cisco.com/nxmt/design.htm?\\_flowId=conversion-flow&\\_flowExecutionKey=e1662](http://tools.cisco.com/nxmt/design.htm?_flowId=conversion-flow&_flowExecutionKey=e1662). The page features a navigation menu with options like 'Solutions', 'Products & Services', 'Ordering', 'Support', 'Training & Events', and 'Partner Central'. The main content area is titled 'IOS-NXOS Migration Tool' and includes a 'Conversion Template' section. This section contains a form with the following fields: 'Template Name' (text input with value 'Template A'), 'Comments' (text area with value 'some comments'), 'Source IOS Version' (dropdown menu with value 'Cat 6K'), and 'Target NX-OS Version' (dropdown menu with value 'Nexus 7K'). Below the form are 'Prev', 'Next', and 'Cancel' buttons. The footer of the page includes copyright information: '© 1992-2009 Cisco Systems, Inc. All rights reserved.' and links for 'Contacts', 'Feedback', 'Help', and 'Site Map'.

# Session Agenda

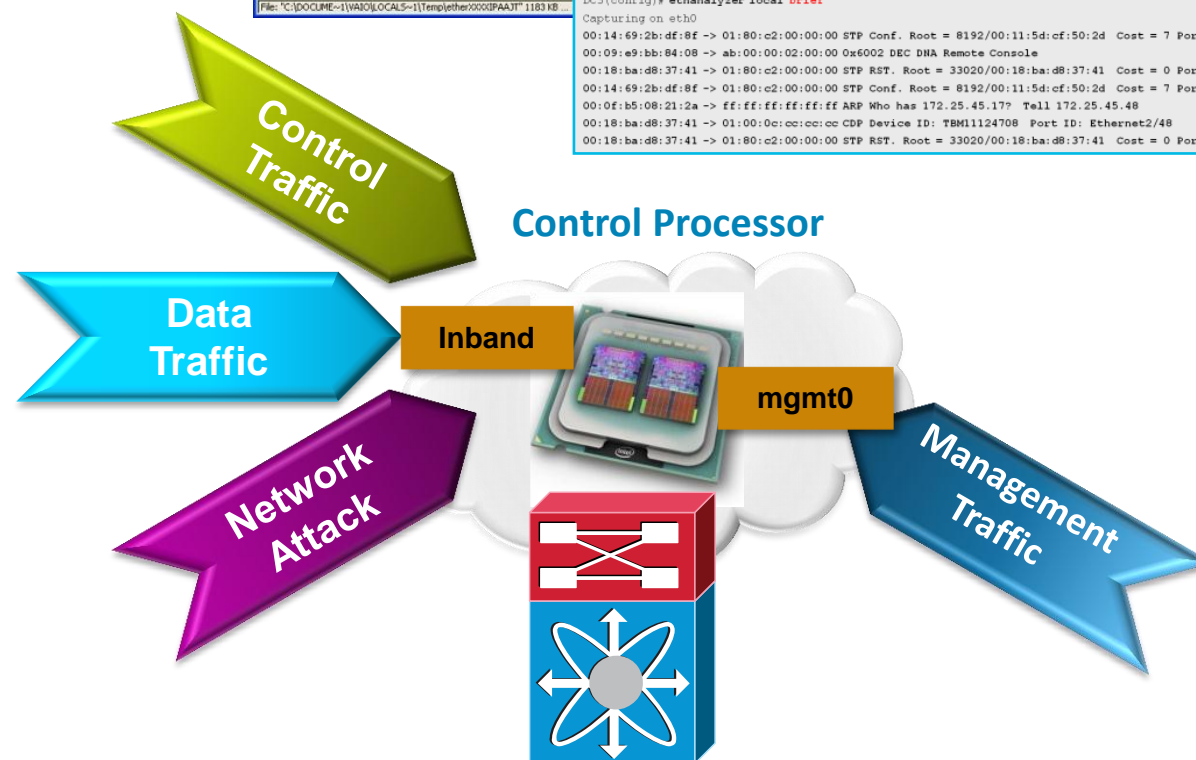
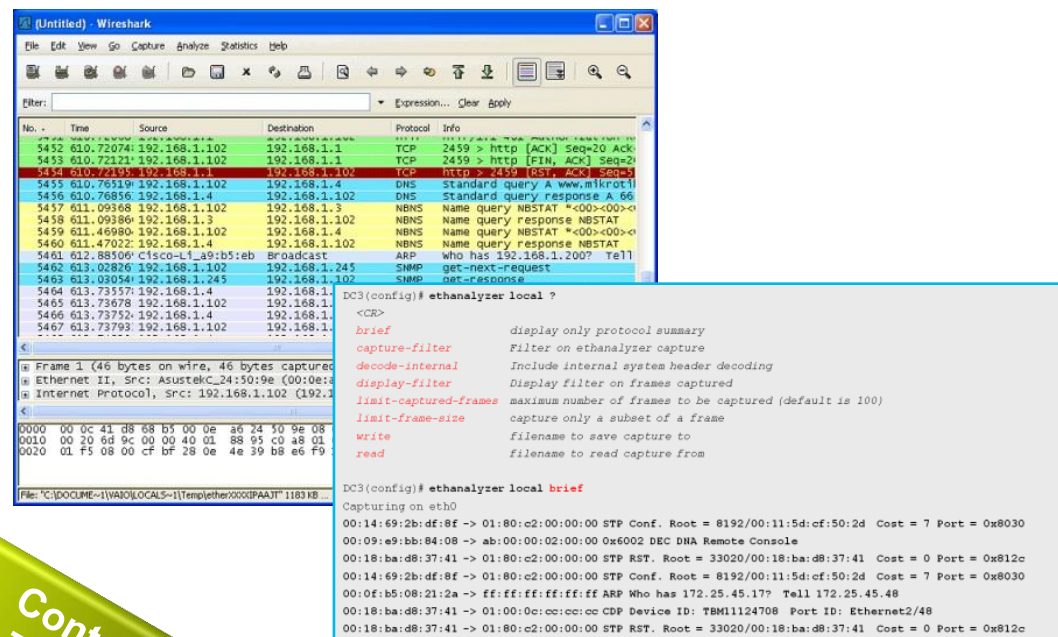
- NX-OS Origins & Overview
- NX-OS Modular Architecture
- High-Availability Infrastructure
- High-Availability Features & Capabilities
- Command Line Interface
- Operational & Management Features
- Licensing & Lifecycle
- Innovation
- Conclusion



# Embedded WireShark Analyser

## EthAnalyser

Real-time, on-the-device protocol analyser provide ultimate visibility into various traffic hitting CPU from remote locations



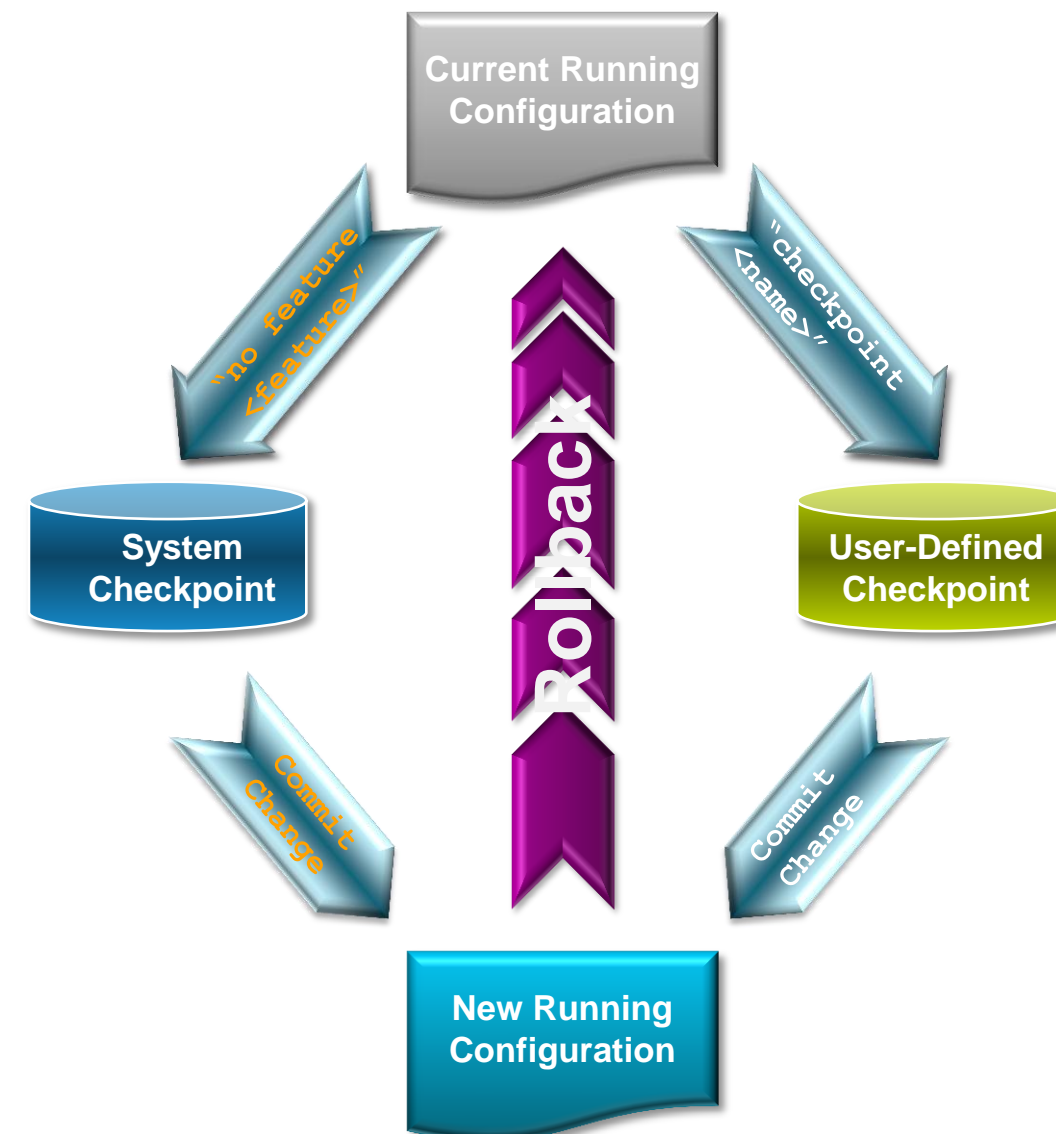
- Monitor traffic from inband and mgmt0 interfaces to the Control Processor
- Extensive capture and display options, including to file (.pcap)
- Capture rules/filters

# NX-OS Configuration

## Checkpoint & Rollback

- Facilitate change-management with configuration snapshots

- System and user generated checkpoints
- System checkpoint automatically created when any conditional features are disabled
- User-defined checkpoint can be initiated from CLI
- Rollback to any checkpoint allows easy recovery





# Checkpoint & Configuration Rollback

## Examples

- System-checkpoint created automatically upon feature removal

```
N7K(config)# no feature vpc
N7K(config)# sh checkpoint summary
System Checkpoint Summary
-----
1) system-fm-vpc:
Created by admin
Created at Fri, 16:51:40 06 Nov 2009
Size is 24,567 bytes
Description: None
```

← Default name for system-checkpoint, 'system-fm-xxx'

- User-defined checkpoint with description simplifies configuration management

```
N7K# checkpoint 2009-11-06 description SQL DC ACL Update
N7K# sh checkpoint summary
User Checkpoint Summary
-----
1) 2009-11-06:
Created by admin
Created at Fri, 18:33:41 06 Nov 2009
Size is 25,773 bytes
Description: SQL DC ACL Update
```

← Timestamp of checkpoint help configuration management

- Flexible option for configuration rollback

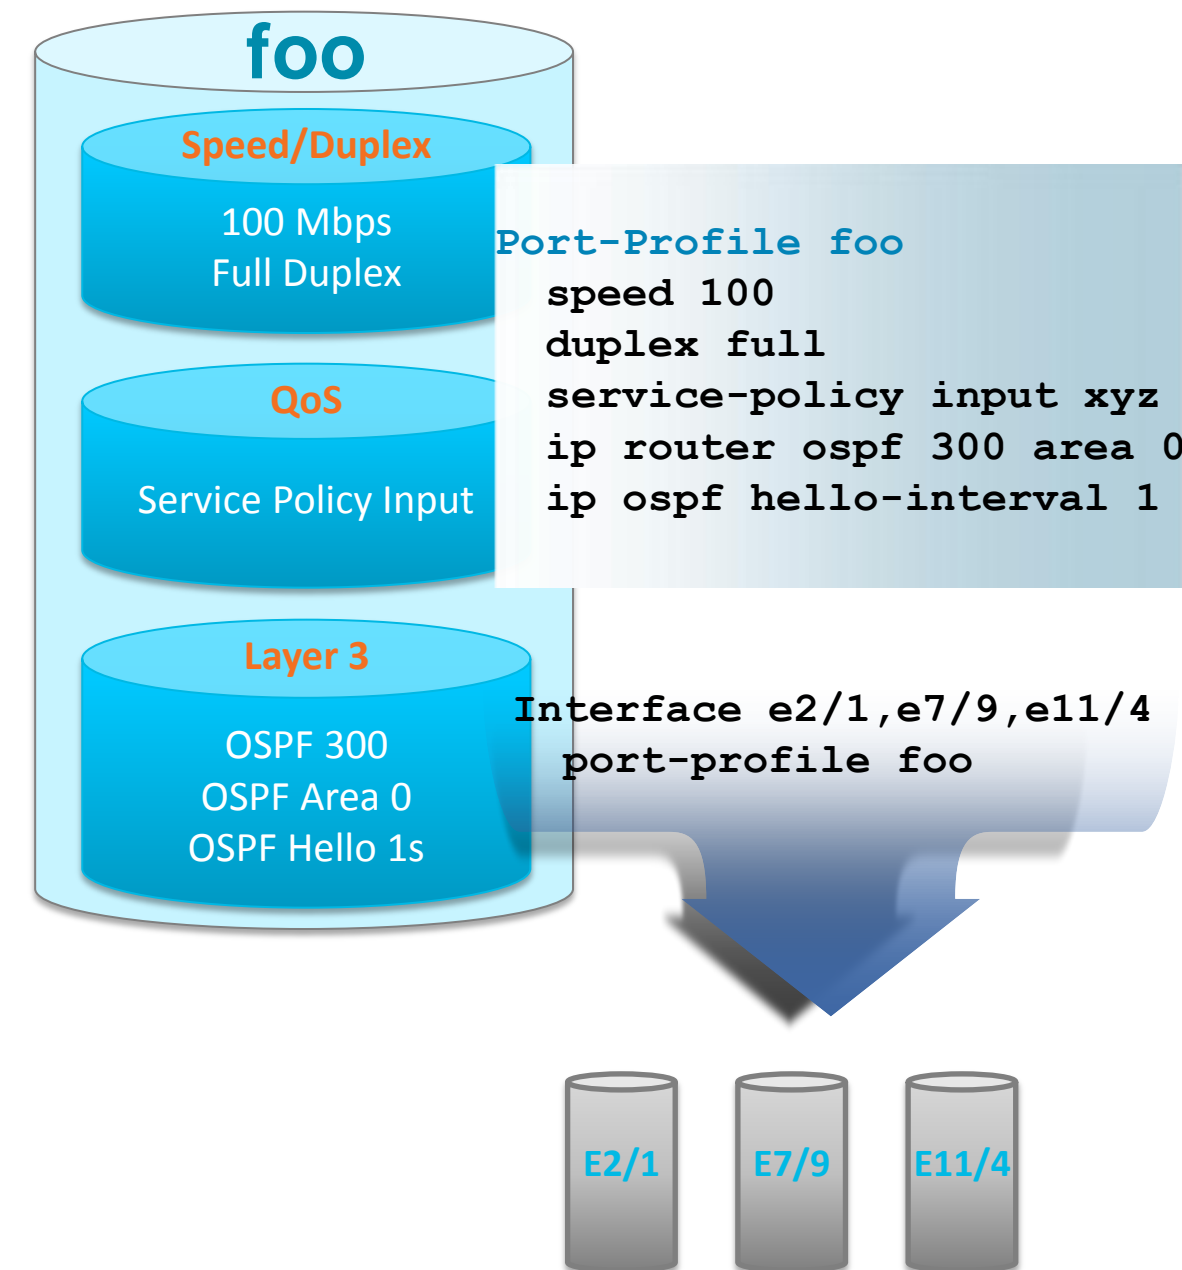
```
N7K# rollback running-config checkpoint 2009-11-11 ?
<CR>
atomic          Stop rollback and revert to original configuration
                 (default)
best-effort     Skip errors and proceed with rollback
stop-at-first-failure Stop rollback at the first error
verbose        Show the execution log
```



# NX-OS Port-Profiles

## Simplify the Configuration of Groups of Interfaces

- Enables the application of **common configuration** across groups of ports
- A port-profile can **inherit attributes** from other port-profiles (nested profiles)
- A change to a port-profile **automatically updates configuration** of all member ports
- Any interface command available on a Nexus interface can be a part of a port-profile e.g. ACL, L3, VLAN, etc.
- Configuration precedence/order:
  - **Default config. < Port-profile < Manual config.**



# NX-OS XML Integration

- Remote management via NETCONF/XML

```
<?xml version="1.0"?>
<rpc message-id="101"
  xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target>
      <running/>
    </target>
    <config>
      <xs:interface xmlns:xs="http://www.cisco.com/SANOS/1.0/interface">
        <xs:mgmt>
          <xs:Naming>
            <xs:intf>0</intf>
          </xs:Naming>
          <xs:ip>
            <xs:address>
              <xs:host>1.1.1.1</xs:host>
              <xs:netmask>255.255.255.0</xs:netmask>
            </xs:address>
          </xs:ip>
        </xs:mgmt>
      </xs:interface>
    </config>
  </edit-config>
</rpc>
```

- Pipe CLI command output to XML formatting

```
N7K# show int e1/1 | xml
<?xml version="1.0" encoding="ISO-8859-1"?>
<nf:rpc-reply xmlns:nf="urn:ietf:params:xml:ns:netconf:base:1.0"
  xmlns="if_manager">
<nf:data>
  <show>
    <interface>
      <TABLE_interface>
        <ROW_interface>
          <interface>Ethernet1/1</interface>
          <state>up</state>
          <eth_hw_desc>10/100/1000 Ethernet</eth_hw_desc>
          <eth_hw_addr>001b.54c1.5d44</eth_hw_addr>
          <eth_bia_addr>001b.54c1.5d44</eth_bia_addr>
          <eth_mtu>1500</eth_mtu>
          ...
        </ROW_interface>
      </TABLE_interface>
    </show>
  </interface>
</nf:data>
</nf:rpc-reply>
```

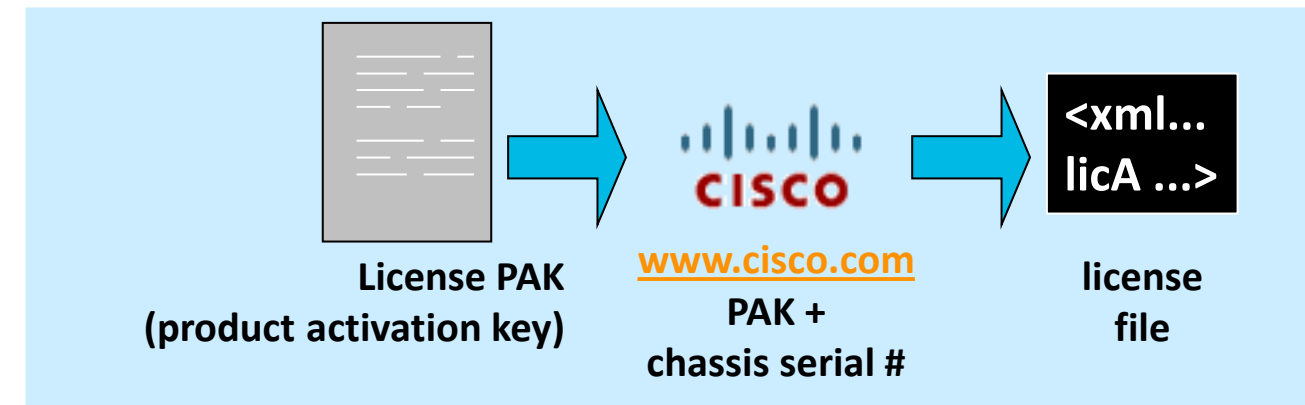
- Eliminate/Simplify screen scraping for output data
- Human-readable format – easier parsing
- Future-proofing through open and flexible standard protocol
- NetConf over SSH for security
- XML API allow easy integration with 3rd-party NMS applications

# Session Agenda

- NX-OS Origins & Overview
- NX-OS Modular Architecture
- High-Availability Infrastructure
- High-Availability Features & Capabilities
- Command Line Interface
- Operational & Management Features
- **Licensing & Lifecycle**
- Innovation
- Conclusion



# NX-OS Licensing



- Licenses are enforced on the switch

# show license host-id

License tied chassis serial # stored in dual redundant NVRAM modules on backplane

- Licenses are issued in the form of a digitally signed text file

# install license bootflash:N7K-1234.lic

## Grace Period

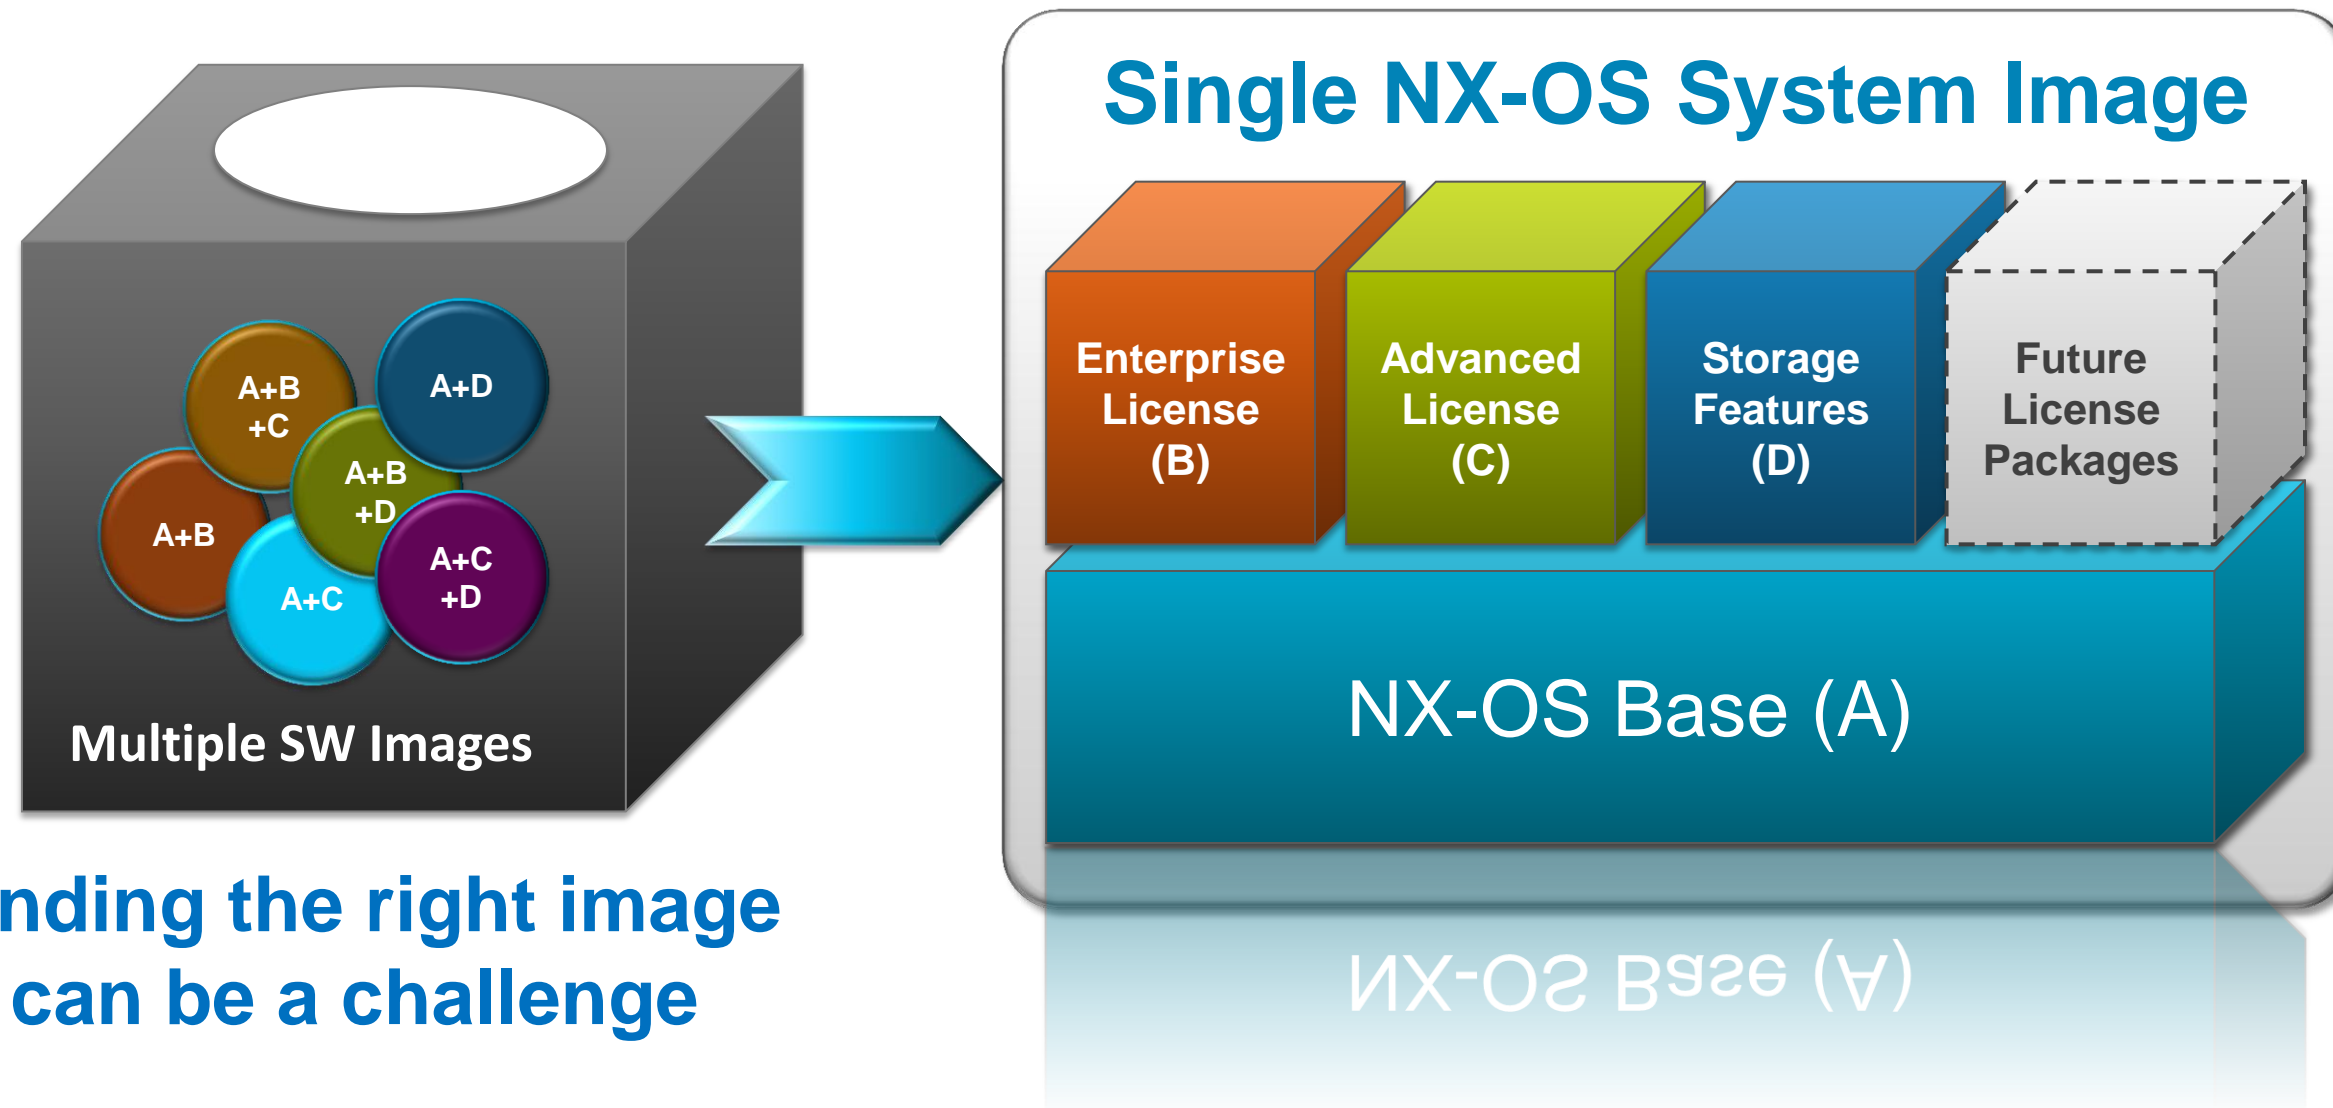
- Allows trying functionality without a licence for 120 days
- Periodic syslog, callhome and SNMP traps warning when grace period nears expiry

## Time-Bound Licenses

- License with expiry date
- Used for demo or as an emergency
- Periodic syslog, callhome and SNMP traps warning when time bound license nears expiry
- After expiry date feature will continue to run if grace period has not been exhausted

# NX-OS Platform Universal System Image

## License-based Feature Management

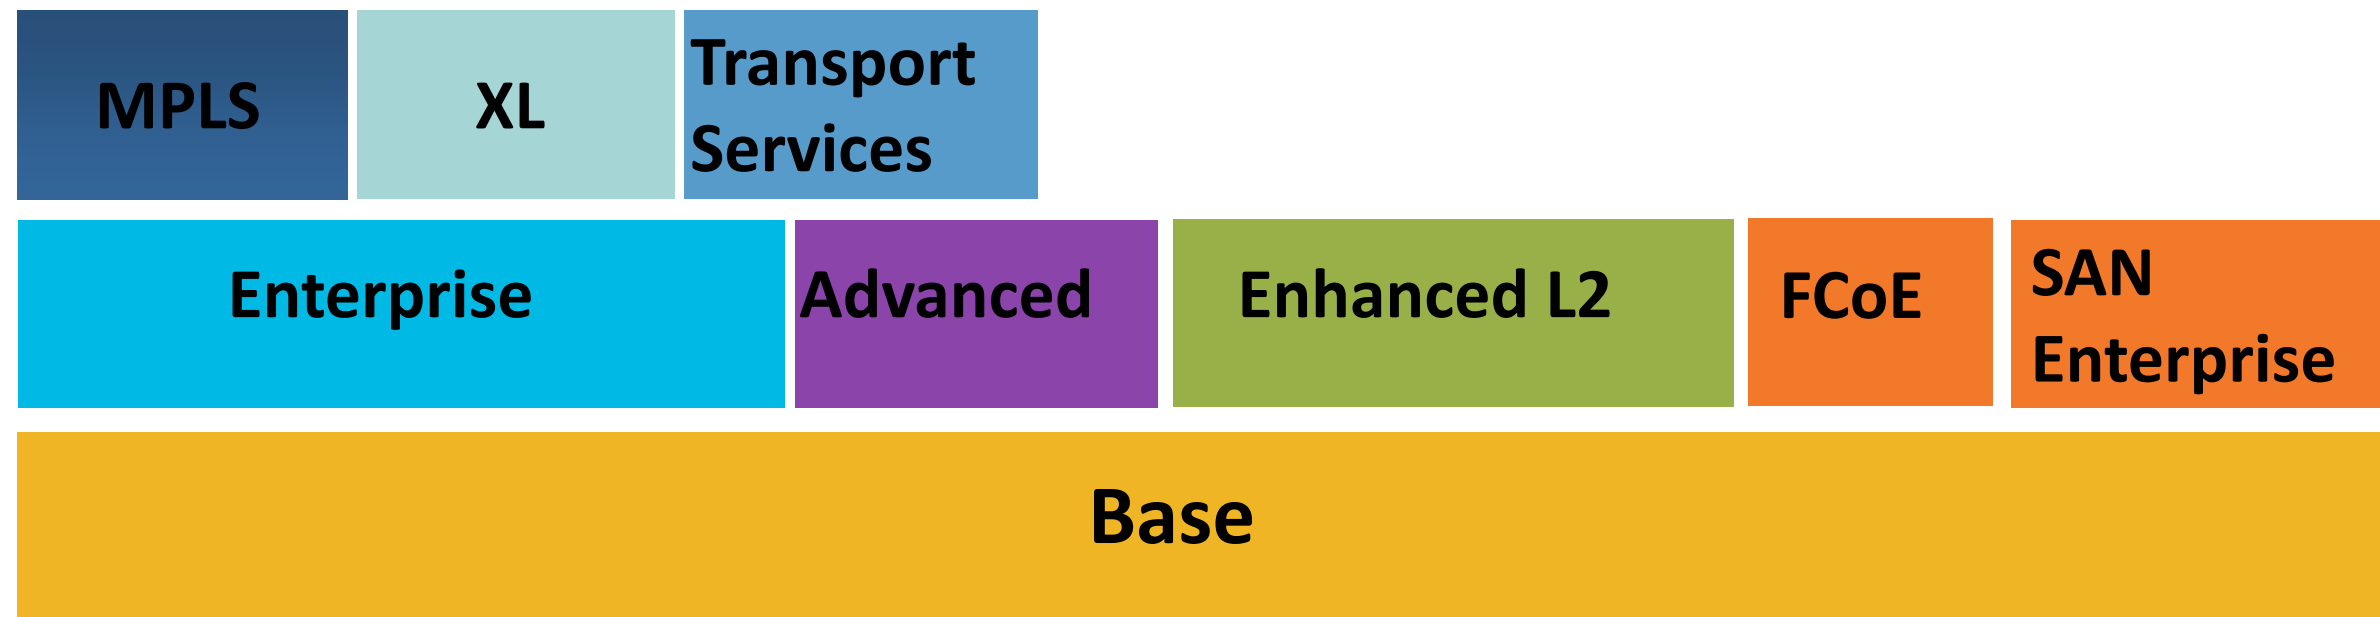


Finding the right image can be a challenge



# NX-OS Licensing

## Nexus 7000 Licenses Overview



- **Simplified Software Management**

8 NX-OS enforceable licenses enable full suite of functionalities for any switching deployment

- **Less Costly Software Upgrades**

NX-OS feature upgrades can be done by enabling a new license key, reducing truck-rolls

- **Enable Development of New Software Based Business Models**

Services on Demand - purchase upgrades as needed (pay-as-you grow)

- **Transport Services evolution:** *LISP* is added to the license starting from NX-OS 5.2

- **Enhanced L2 evolution:** *PONG* is added to the license starting from NX-OS 5.2



# NX-OS Software Packaging

## Nexus 7000 Overview

Scalable Services				
<ul style="list-style-type: none"> <li>IP routing</li> <li>OSPFv2</li> <li>OSPFv3</li> <li>IS-IS</li> <li>BGP for IPv4</li> <li>BGP for IPv6</li> <li>EIGRP for IPv4</li> <li>EIGRP for IPv6</li> <li>BFD</li> <li>IP Multicast</li> <li>PIM: Sparse, Bidir, ASM and SSM for IPv4 &amp; IPv6</li> <li>Multicast Source Discovery Protocol (MSDP) for IPv4</li> <li>PBR for IPv4 and IPv6</li> <li>GRE Tunnels</li> </ul>	<ul style="list-style-type: none"> <li>OTV</li> <li>LISP</li> </ul>	<ul style="list-style-type: none"> <li>FabricPath</li> <li>PONG</li> </ul>	<ul style="list-style-type: none"> <li>MPLS VPN</li> <li>LDP</li> <li>MPLS QoS</li> <li>TE/FRR</li> <li>mVPN</li> <li>MPLS OAM</li> <li>6PE/6VPE</li> </ul>	<ul style="list-style-type: none"> <li>Inter VSAN routing</li> <li>VSAN based access control</li> </ul>
	Transport Services			Storage Enterprise
	<ul style="list-style-type: none"> <li>VDCs</li> <li>TrustSec</li> </ul>			<ul style="list-style-type: none"> <li>Multi-Hop FCoE</li> <li>FCF</li> <li>FIP</li> </ul>
Enterprise LAN	Advanced LAN	Enhanced L2	MPLS	FCoE*

Base	vPC Port Profile WCCP Port Security GOLD EEM TACACS LACP ACL QoS STP STP Guards UDLD CDP CoPP uRPF IP Source Guard DHCP Snooping CMP ISSU SSO Dynamic ARP Inspection Smart Call Home SNMP 802.1x SPAN Netflow v5 and v9 IEEE1588
	* Per Module-based license



# NX-OS Software Packaging

## Nexus 5000 Overview

- SVI routed interfaces
- L3 routed ports on non-FEX interfaces
- Static Routing
- RIPv2
- EIGRP for Routed Access (Stub)
- OSPFv2 for Routed Access (Limited\*)
- HSRP
- VRRP
- IGMPv1, v2, v3
- PIM v2 (sparse)
- Routed ACLs
- uRPF

**Base Enterprise**

- BGPv4
- EIGRP
- OSPFv2
- VRF-Lite

**Future:**

- v6 Routing (OSPFv3, IS-ISv6, BGPv6)
- HSRPv6/VRRPv3

**Enterprise LAN**

**Storage Services:**

- 8-port FC/FCoE License
- Chassis license for 5548 and 5596
- FcoE NPV (available also as standalone license)

**FCoE License**

- FabricPath

**Enhanced L2**

**Base** ISSU vPC Port Profile LACP ACL QoS STP LLDP XML SNMP 1588

(\*Limited to 256 network routes)

# NX-OS Releases

## NX-OS Strategy and Recommended Nexus 7000 Releases

- NX-OS on Nexus 7000 Minimum software Recommendation

[http://www.cisco.com/en/US/docs/switches/datacenter/sw/nx-os/recommended\\_releases/recommended\\_nx-os\\_releases.html](http://www.cisco.com/en/US/docs/switches/datacenter/sw/nx-os/recommended_releases/recommended_nx-os_releases.html)

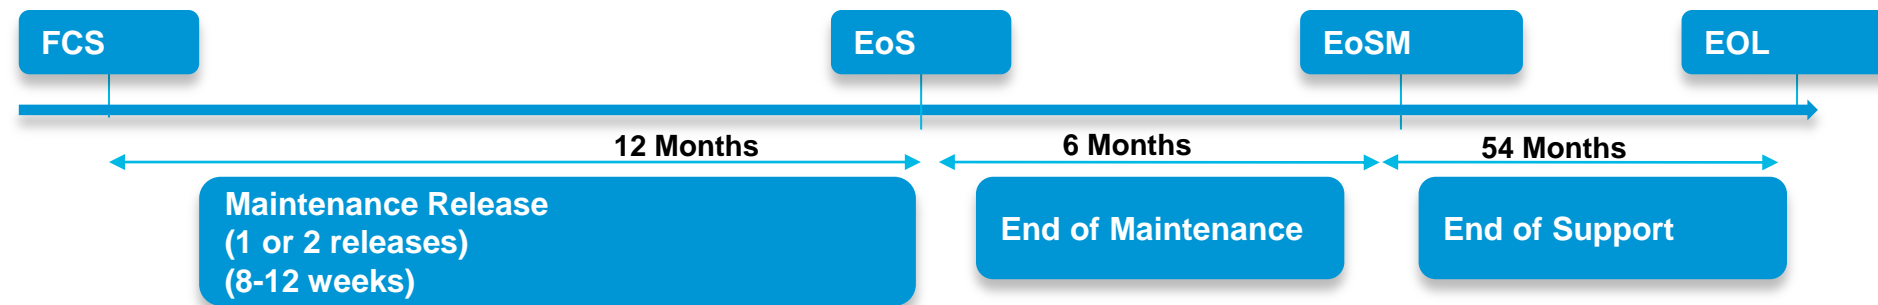
The screenshot displays two browser windows. The left window shows the Cisco website page titled "Minimum Recommended Cisco NX-OS Releases for Cisco Nexus 7000 Series Switches". The page includes a navigation menu, a table of contents, and a main content area with the date "April 21, 2011". The right window shows the "Cisco NX-OS Software Release Strategy" document, which includes a diagram of the NX-OS 4 and 5 trains and a section on "Long-Lived and Short-Lived Releases".

- NX-OS Software Release Strategy Document


[http://www.cisco.com/en/US/prod/collateral/iosswrel/ps9494/ps9372/guide\\_c07-658595.html](http://www.cisco.com/en/US/prod/collateral/iosswrel/ps9494/ps9372/guide_c07-658595.html)



# NX-OS Software Life Cycle



## Short Lived Release



## Long Lived Release

- 4.2 is a long-lived release for NX-OS on N7K
- 5.2 is a long-lived NX-OS train on N7K and N5K
- 6.2 is expected to be the long-lived release for the 6.x train



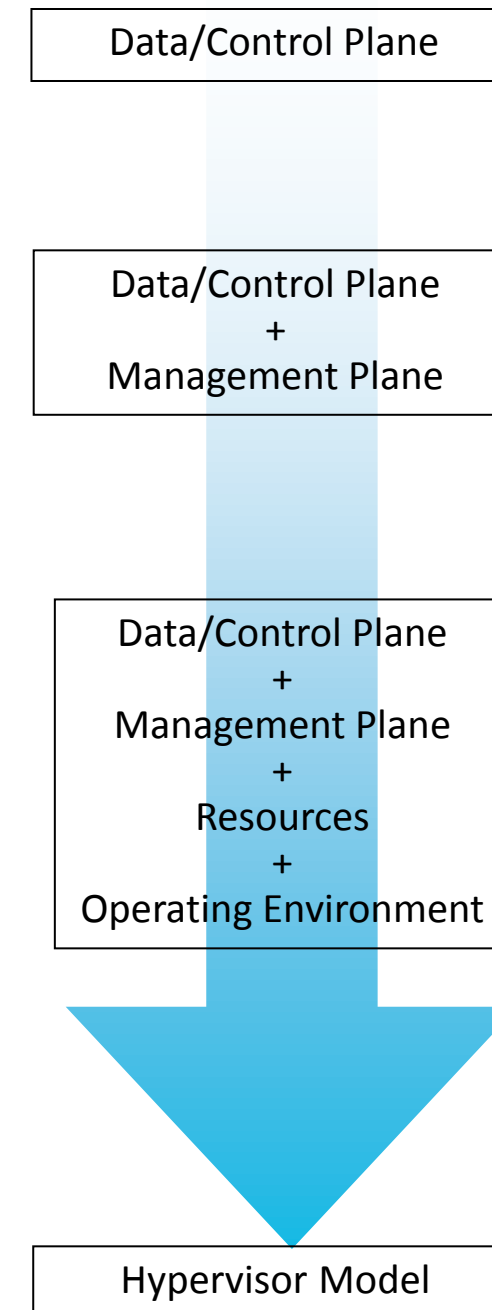
# Session Agenda

- NX-OS Origins & Overview
- NX-OS Modular Architecture
- High-Availability Infrastructure
- High-Availability Features & Capabilities
- Innovation
- Conclusion
- Command Line Interface
- Operational & Management Features
- Licensing & Lifecycle

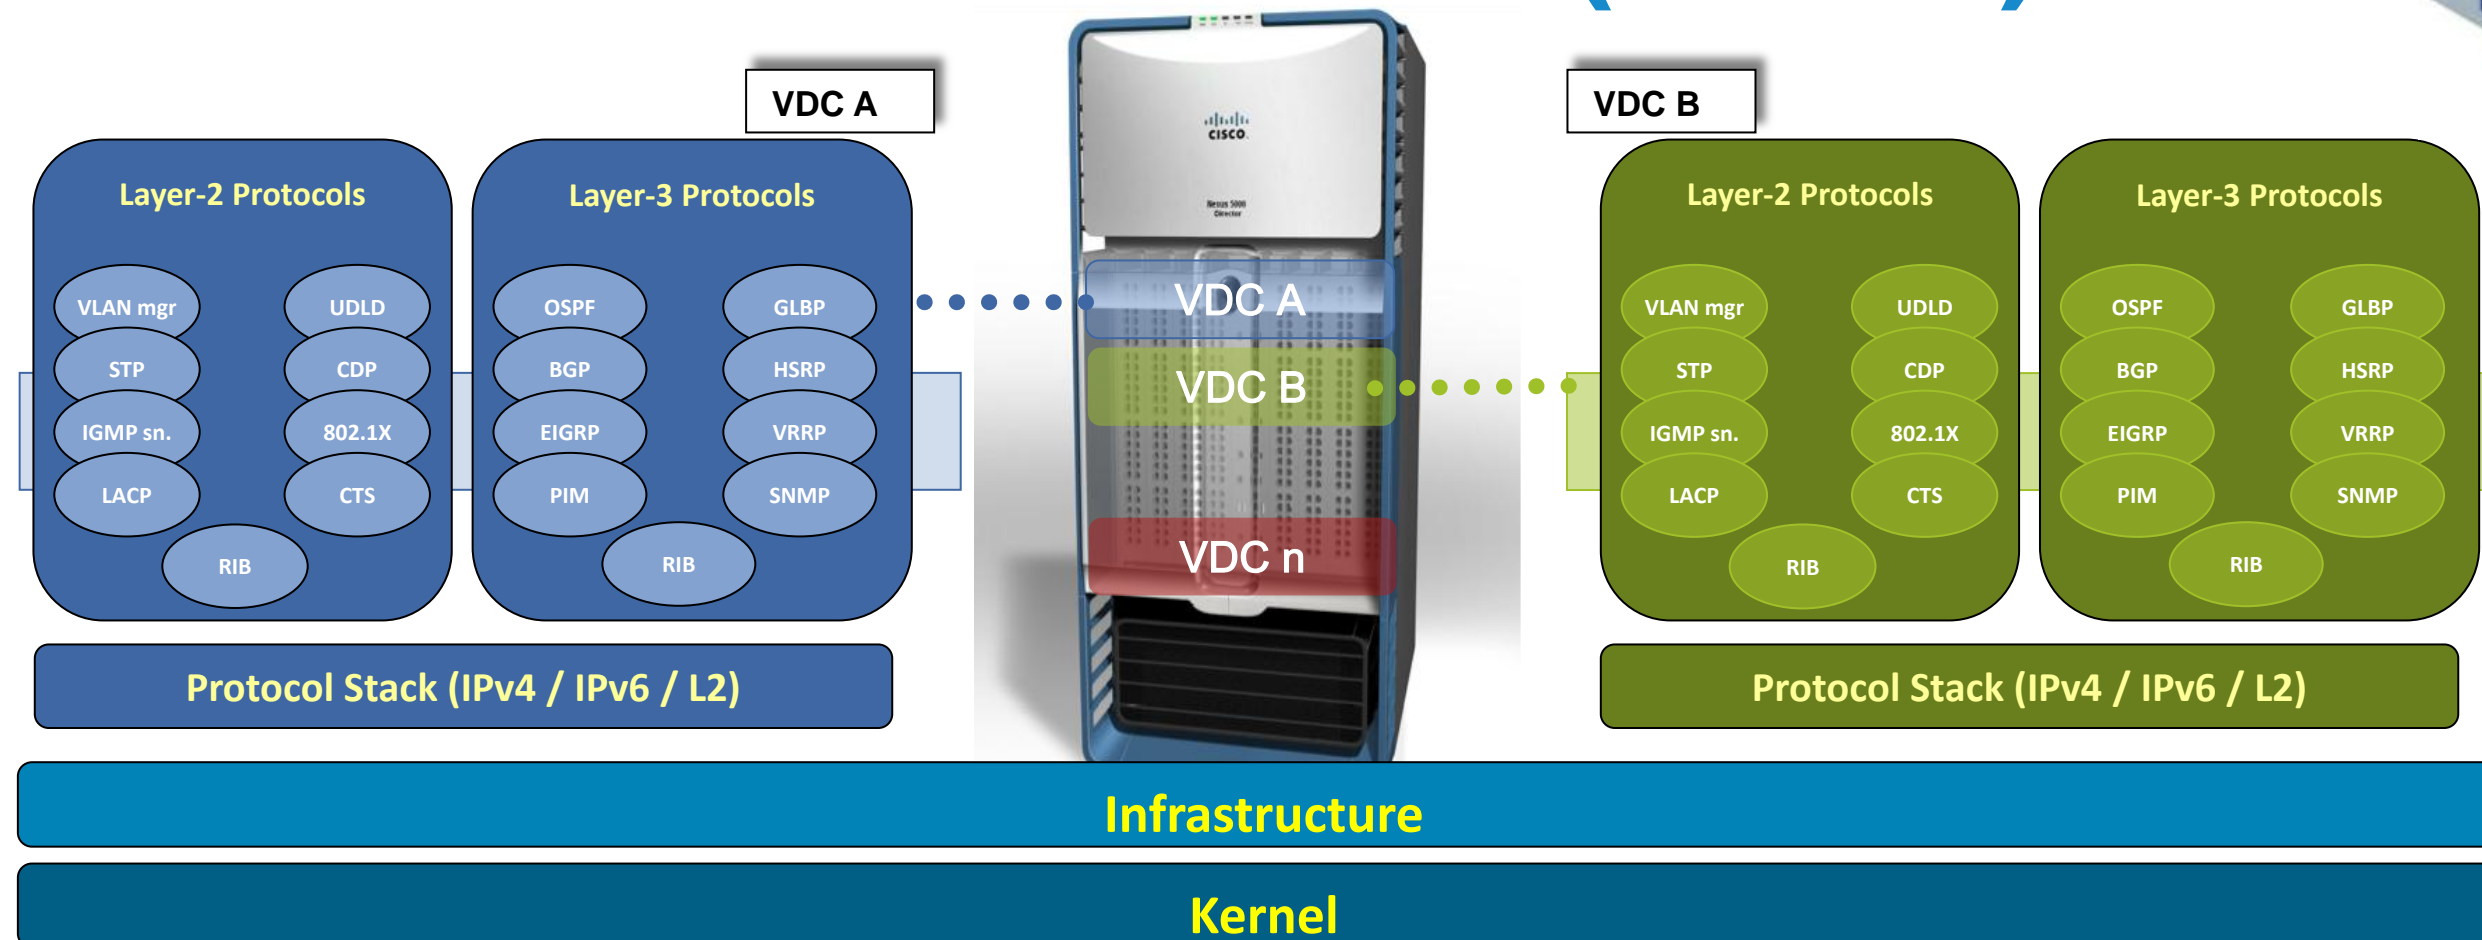


# Various Degrees of Virtualisation

- VRFs and VLANs
  - Logical separation of data-plane (and some control-plane) functionality
- Virtual Contexts (i.e. Firewalls, ACE, etc.)
  - Logical separation of configuration or management and data-plane
- Virtual Device Contexts
  - Logical separation of control-plane, data-plane, management, resources, and system processes



# Virtual Device Contexts (VDCs)



## ■ VDC—Virtual Device Context

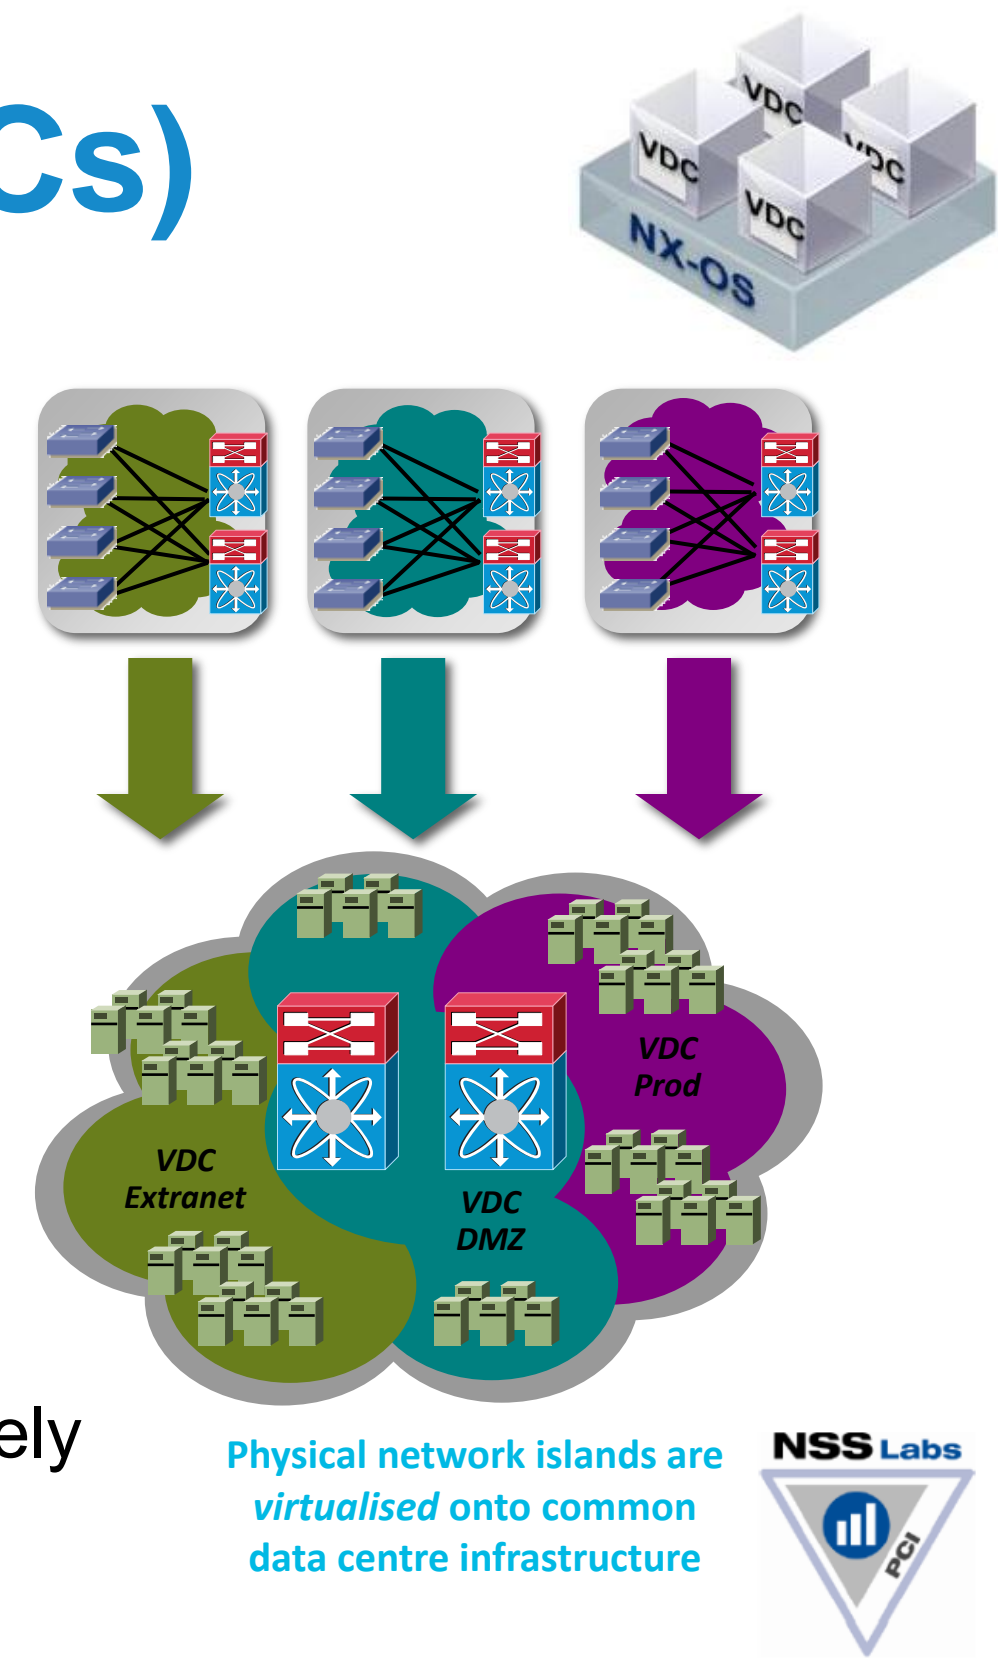
- Flexible separation/distribution of **Software Components**
- Flexible separation/distribution of **Hardware Resources**
- Securely delineated **Administrative Contexts**

## VDCs are not...

- The ability to run different OS levels on the same box at the same time
- based on a **hypervisor** model; there is a single 'infrastructure' layer that handles h/w programming...


# Virtual Device Contexts (VDCs)

- Typical silo/stovepipe design
  - Production, Development, Test
  - Intranet, Internet, DMZ, Extranet
  - Application A, Application B, Application C
  - Customer A, Customer B, Customer C
  - Cluster A, Cluster B, Cluster C
- VDCs enable collapsing of physical infrastructure into logical infrastructure
- Preserves security, administration, and organisational boundaries, & fault isolation
- “The results clearly demonstrate that VDCs can be effectively deployed as though they are physically separate devices”  
**Source: NSS Labs\***



\*Ref: <http://nsslabs.com/research/network-security/virtualization/cisco-nexus-7000-q2-2010.html>

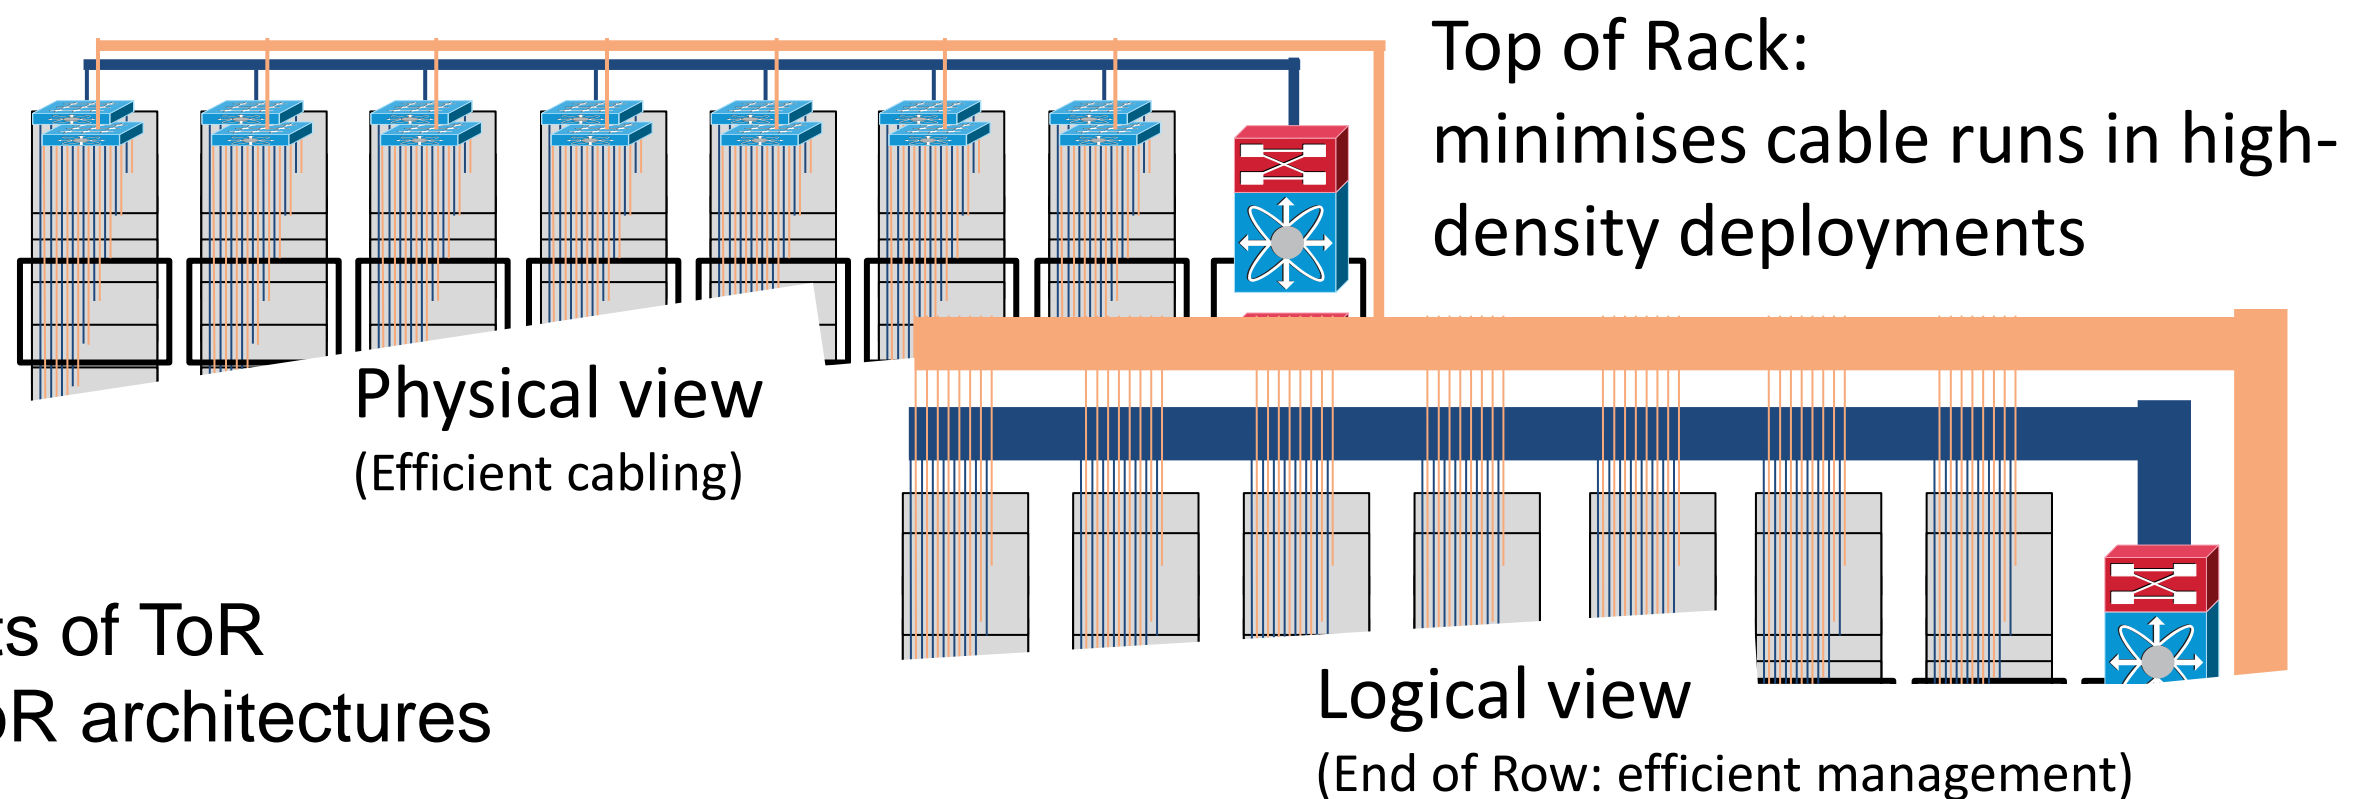
# FEX-Link



- Nexus 7000/5x00 + FEX is like a “Virtual Chassis”
- Nexus 2000 FEX is a “Virtual Line Card” to its “parents”
- No Spanning Tree between the FEX and its “parent”
- No local switching on the FEX
- NX-OS Linecard code runs on the 2148/2248/2232



# Access Layer with Nexus 2000



## Benefits of ToR and EoR architectures

- Physically on the top of each server rack, Logically an end of access row device
- Reduces cable runs: Majority of physical cabling is within the rack, <2m cable
- Reduce management points in the network
  - 1,500 server HA network with 48-port ToR access switches: 34 management points
  - 1,500 server HA network with Nexus 7000 and Nexus 2000 : 2 management points
  - 1,500 server HA network with Nexus 5500 and Nexus 2000 : 4 management points
- Easier to ensure feature consistency across hundreds or thousands of ports


# Adapter FEX and VM-FEX

## Features:

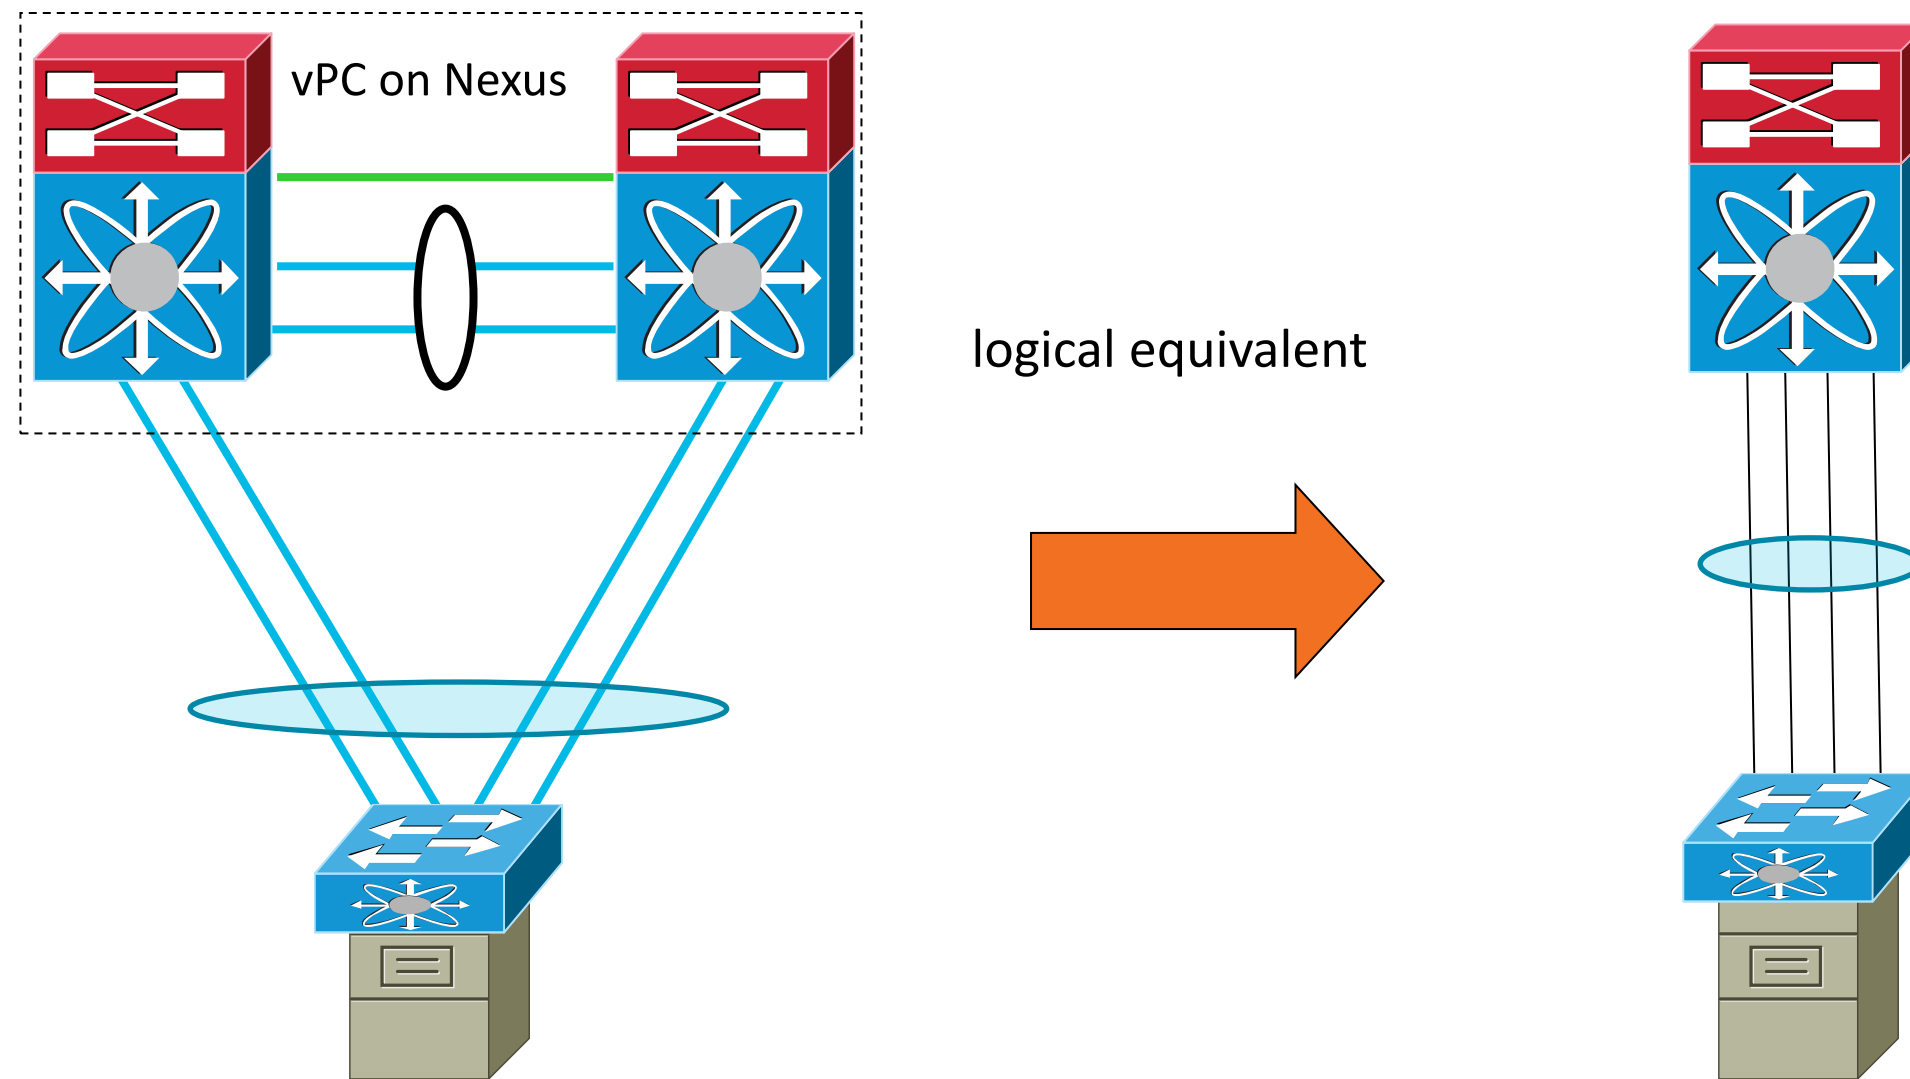
- Adapter FEX split a physical NIC into multiple logical NICs
- VM-FEX extends Adapter FEX technology to virtual machine
- Pre-Standard IEEE 802.1BR

## Benefits:

- Single point of management
- Increased 10 G bandwidth utilisation – less power, cabling adaptors with adapter FEX
- Dynamic network & security policy mobility during VM migration with VM-FEX



# Virtual Port Channel (vPC) Objectives




- Provides a loop-free topology
- Maximises bandwidth / lower over-subscription
- Improved convergence & availability

# vPC Topology Example


## Back to Back

- Two layers of vPC peers can be connected back-to-back e.g. N7k to N5k
- Opportunity for very high bandwidth using an evolutionary development of STP
- Up to 32-way port-channel



# FabricPath: an Ethernet Fabric

Shipping on Nexus 7000 and Nexus 5500



- Connect a group of switches using an **arbitrary** topology
- With a simple CLI, aggregate them into a Fabric:

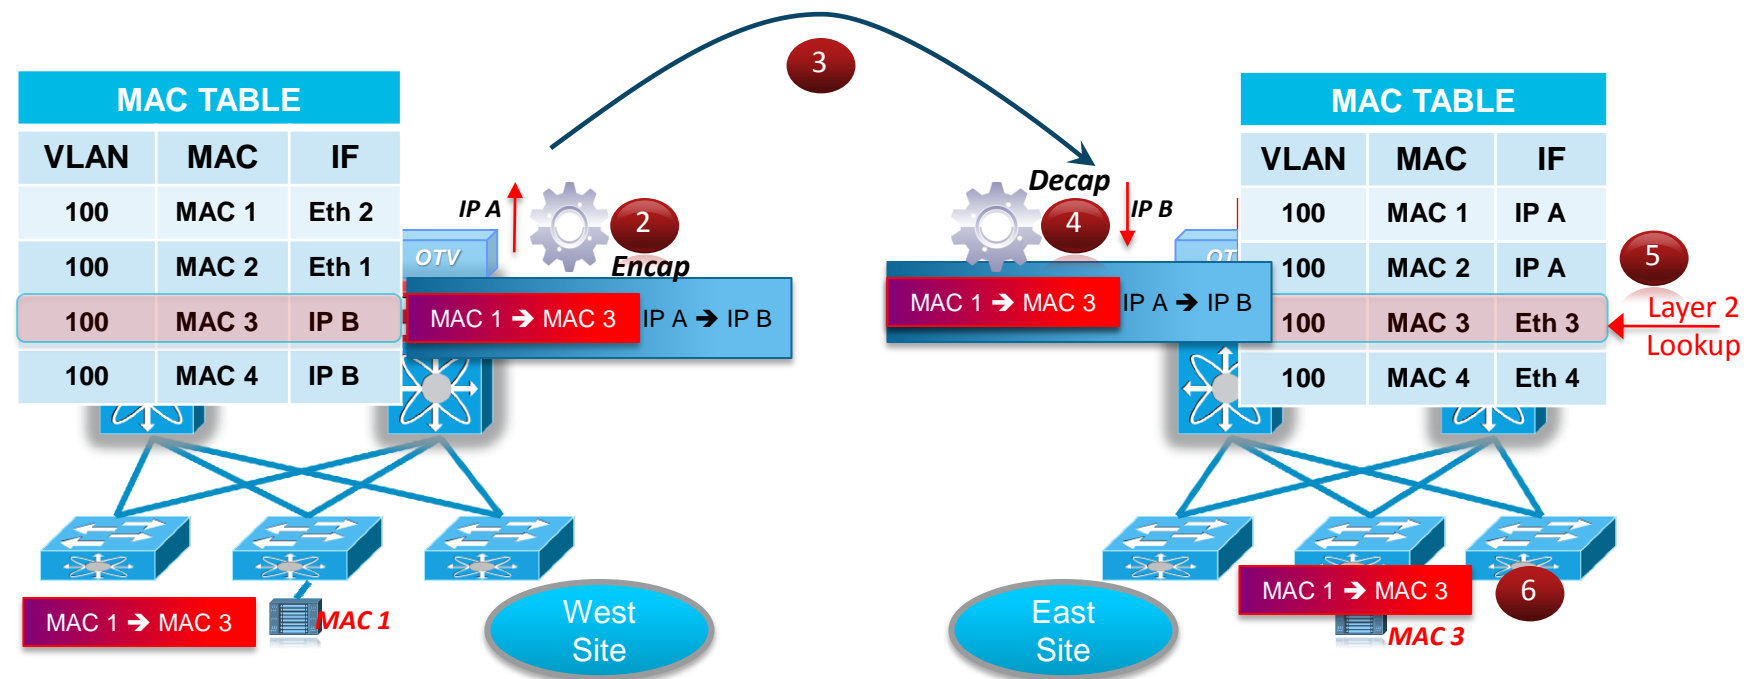

```
N7K(config)# interface ethernet 1/1  
N7K(config-if)# switchport mode fabricpath
```

- An open protocol based on L3 technology provides Fabric-wide intelligence and ties the elements together
- Eliminates Spanning tree limitations
- High resiliency, fast network re-convergence
- Any VLAN Anywhere in the Fabric




# Overlay Transport Virtualisation

## Virtual Machine Mobility



# Overlay Transport Virtualisation


## Virtual Machine Mobility



MAC TABLE		
VLAN	MAC	IF
100	MAC X	<b>IP B</b>
100	MAC 2	Eth 1
100	MAC 3	Eth 1
100	MAC 4	Eth 2

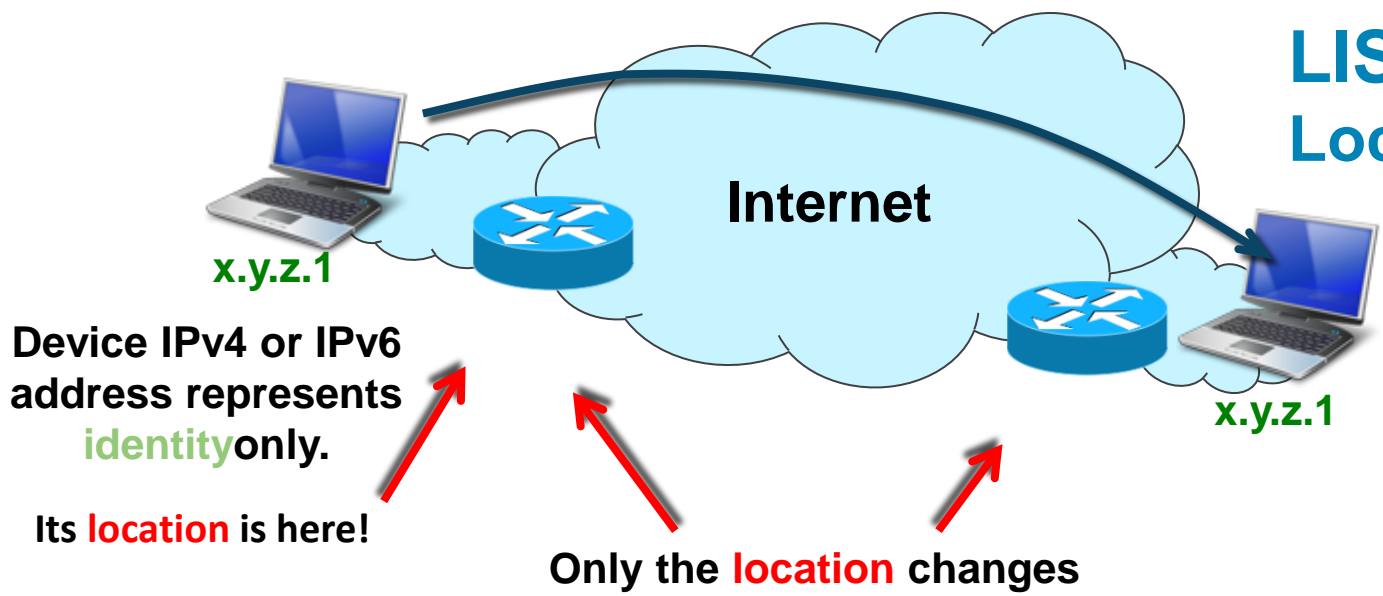
# Location ID/Separation Protocol (LISP)

End-Point “location” and “identity”



Today's Internet Behaviour  
Loc/ID “overloaded” semantic

When the device moves, it gets a new IPv4 or IPv6 address for its new identity and location



LISP Behaviour  
Loc/ID “split”

When the device moves, keeps its IPv4 or IPv6 address. It has the same identity

# Unified Fabric

Fibre Channel Over Ethernet (FCoE), Introduced with Nexus 5000

## FCoE

### Fewer Cables

–Both block I/O & Ethernet traffic co-exist on same cable

### Fewer adapters needed

### Overall less power


### Interoperates with existing SAN's

–Management of SAN's remains constant

### No Gateway

## Benefits

- Mapping of FC Frames over Ethernet
- Enables FC to Run on a Lossless Ethernet Network



# Nexus 7000 FCoE

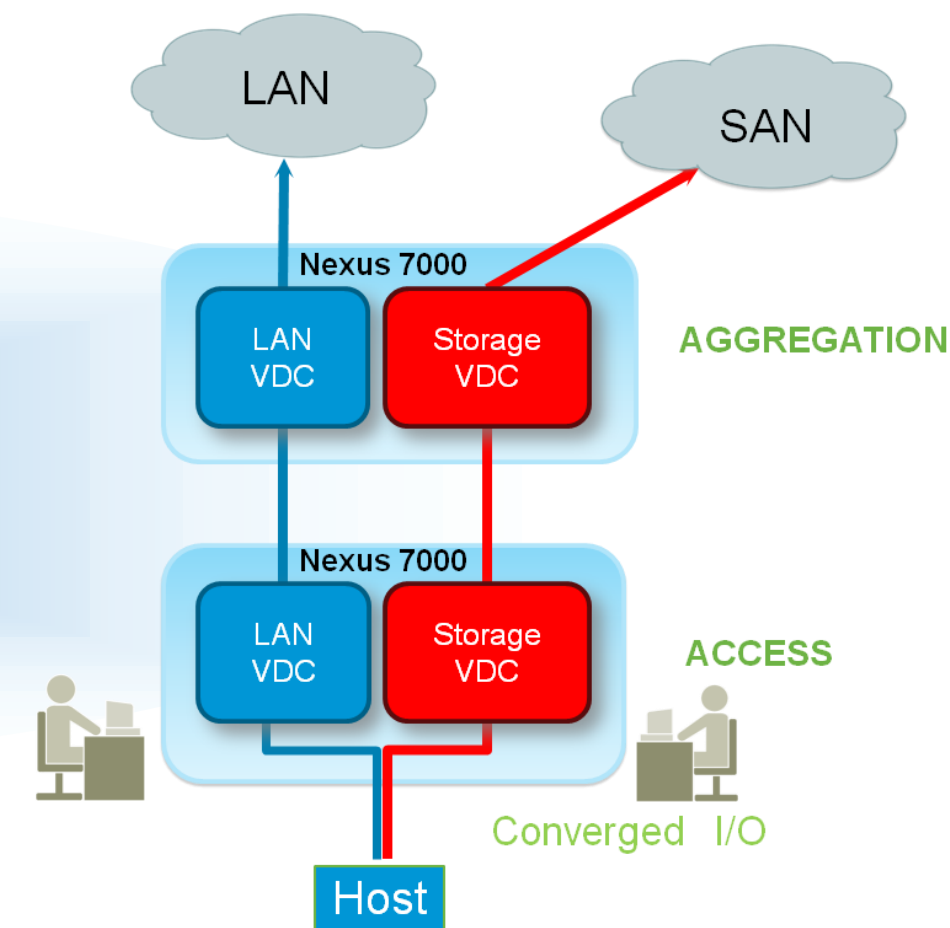
## Director-Class Convergence

- Industry's first Ethernet Storage Director
- Highest performance storage networking platform



### Benefits

- Meets the high availability requirements for mission critical SANs
- Flexibility to support iSCSI, NAS, and FCoE
- Designed to scale beyond 15Tbps, 40G and 100G ready



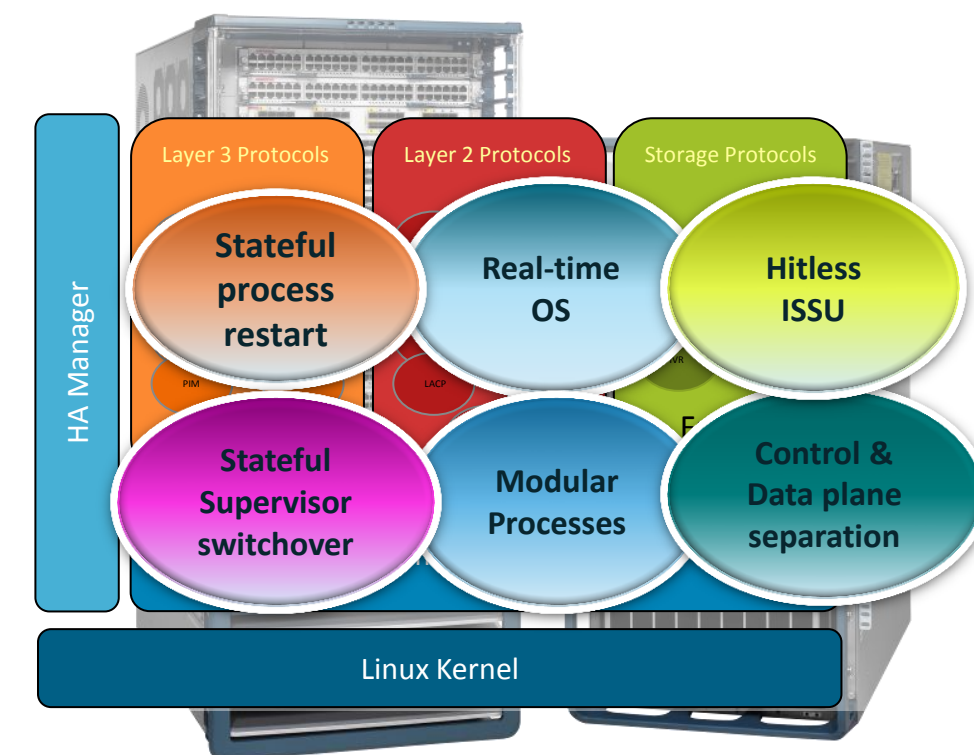
### Use-cases

- Access
- Aggregation
- Core



# MPLS High Availability in NX-OS

- Rich MPLS Layer 3 VPN functionality
- MPLS supported on existing Nexus 7000 M-series I/O modules (including 40/100G)
- Industry's first MPLS switch to implement hitless software upgrade with NSF/SSO:



- Stateful process restart and Graceful restart
- Stateful switchover and ISSU
- BFD for fast convergence (BGP and TE-FRR)

Service	Stateful Restart	ISSU
Layer-3 VPN	✓	✓
Traffic Engineering	✓	✓
mVPN	✓	✓

Protocol	NSR	Graceful Restart	ISSU
LDP	-	✓	✓
MP-BGP	-	✓	✓
RSVP	✓	✓	✓

# Session Agenda

- NX-OS Origins & Overview
- NX-OS Modular Architecture
- High-Availability Infrastructure
- High-Availability Features & Capabilities
- Command Line Interface
- Operational & Management Features
- Licensing & Lifecycle
- Innovation
- Conclusion



# NX-OS Software Architecture

## Top Things to Remember

- NX-OS built around **High-Availability** as a core principle
- NX-OS highly-granular **modularity** for improved efficiency and fault isolation
- NX-OS built to compartmentalise, **scale** (up or down), be **portable**, and **extendable**
- Based on proven SAN-OS/IOS & **secure/standard** features implementation
- Enabling **virtual mobility** and **cloud services**

# NX-OS Software Architecture

## Summary

**N**

**X**

**O**

**S**

**Highly Available and Secure**

**Modular and Efficient**

**Full-Featured and Cloud Ready**

# Q & A





# Complete Your Online Session Evaluation

## Give us your feedback and receive a Cisco Live 2013 Polo Shirt!

Complete your Overall Event Survey and 5 Session Evaluations.

- Directly from your mobile device on the Cisco Live Mobile App
- By visiting the Cisco Live Mobile Site [www.ciscoliveaustralia.com/mobile](http://www.ciscoliveaustralia.com/mobile)
- Visit any Cisco Live Internet Station located throughout the venue

Polo Shirts can be collected in the World of Solutions on Friday 8 March 12:00pm-2:00pm



Cisco *live!* 365

Don't forget to activate your Cisco Live 365 account for access to all session material,

communities, and on-demand and live activities throughout the year. Log into your Cisco Live portal and click the "Enter Cisco Live 365" button.

[www.ciscoliveaustralia.com/portal/login.wv](http://www.ciscoliveaustralia.com/portal/login.wv)

Cisco *live!*

