



*LET'S
BUILD
TOMORROW
TODAY*

Cisco Nexus 7000 / 7700 Switch Architecture



Tim Stevenson
Distinguished Engineer, Technical Marketing

BRKARC-3470

Cisco *live!*

Session Abstract

This session presents an in-depth study of the architecture of the latest generation of Nexus 7000 and Nexus 7700 data center switches. Topics include supervisors, fabrics, I/O modules, forwarding engines, and physical design elements, as well as a discussion of key hardware-enabled features that combine to implement high-performance data center network services.

Session Goal

- To provide a thorough understanding of the Nexus 7000 / Nexus 7700 switching architecture, chassis, supervisor, fabric, and I/O module design, key forwarding engine functions, and complete packet flows
- This session will examine the Nexus 7700 system, as well as the latest additions to the Nexus 7000
- This session will not examine NX-OS software architecture or other Nexus platform architectures



Agenda

- Introduction to Nexus 7000 / Nexus 7700
- Chassis Architecture
- Supervisor Engine and I/O Module Architecture
- Fabric Architecture
- Hardware Forwarding
- Packet Walks
- Conclusion

Introduction to Nexus 7000 / Nexus 7700 Platform

Data-center class Ethernet switches designed to deliver high performance, high availability, system scale, and investment protection

Designed for wide range of Data Center deployments, focused on feature-rich 10G/40G/100G density and performance

I/O Modules



Chassis



Supervisor Engines



Fabrics



Nexus 7000 / Nexus 7700 – Common Foundation

Nexus 7000

General purpose DC switching w/10/40/100G



Nexus 7700

Targeted at Dense 40G/100G deployments




Common Foundation




NX-OS

- Same release vehicles, versioning, feature-sets
- Common configuration model
- Common operational model



Fabrics

- Common fabric ASICs (Fab2) and architecture
- Same central arbitration model
- Same VOQ/QOS model







Fwding

- Identical forwarding ASICs (F2E, F3)
- Consistent hardware feature sets
- Consistent hardware scale

Agenda

- Introduction to Nexus 7000 / Nexus 7700
- Chassis Architecture
- Supervisor Engine and I/O Module Architecture
- Fabric Architecture
- Hardware Forwarding
- Packet Walks
- Conclusion


Nexus 7000 Chassis Family



Nexus 7700 Chassis Family


NX-OS 6.2(2) and later

Nexus 7718



NX-OS 6.2(2) and later

Nexus 7710




NX-OS 6.2(6) and later

Nexus 7706



Nexus 7702 Chassis



Agenda

- Introduction to Nexus 7000 / Nexus 7700
- Chassis Architecture
- Supervisor Engine and I/O Module Architecture
- Fabric Architecture
- Hardware Forwarding
- Packet Walks
- Conclusion


Supervisor Engine 2 / 2E

- Provides all control plane and management functions


Supervisor Engine 2 (Nexus 7000)	Supervisor Engine 2E (Nexus 7000 / Nexus 7700)
Base performance	High performance
One quad-core 2.1GHz CPU with 12GB DRAM	Two quad-core 2.1GHz CPU with 32GB DRAM

- Connects to fabric via 1G inband interface
- Interfaces with I/O modules via 1G switched EOBC
- Onboard central arbiter ASIC
Controls access to fabric bandwidth via dedicated arbitration path to I/O modules

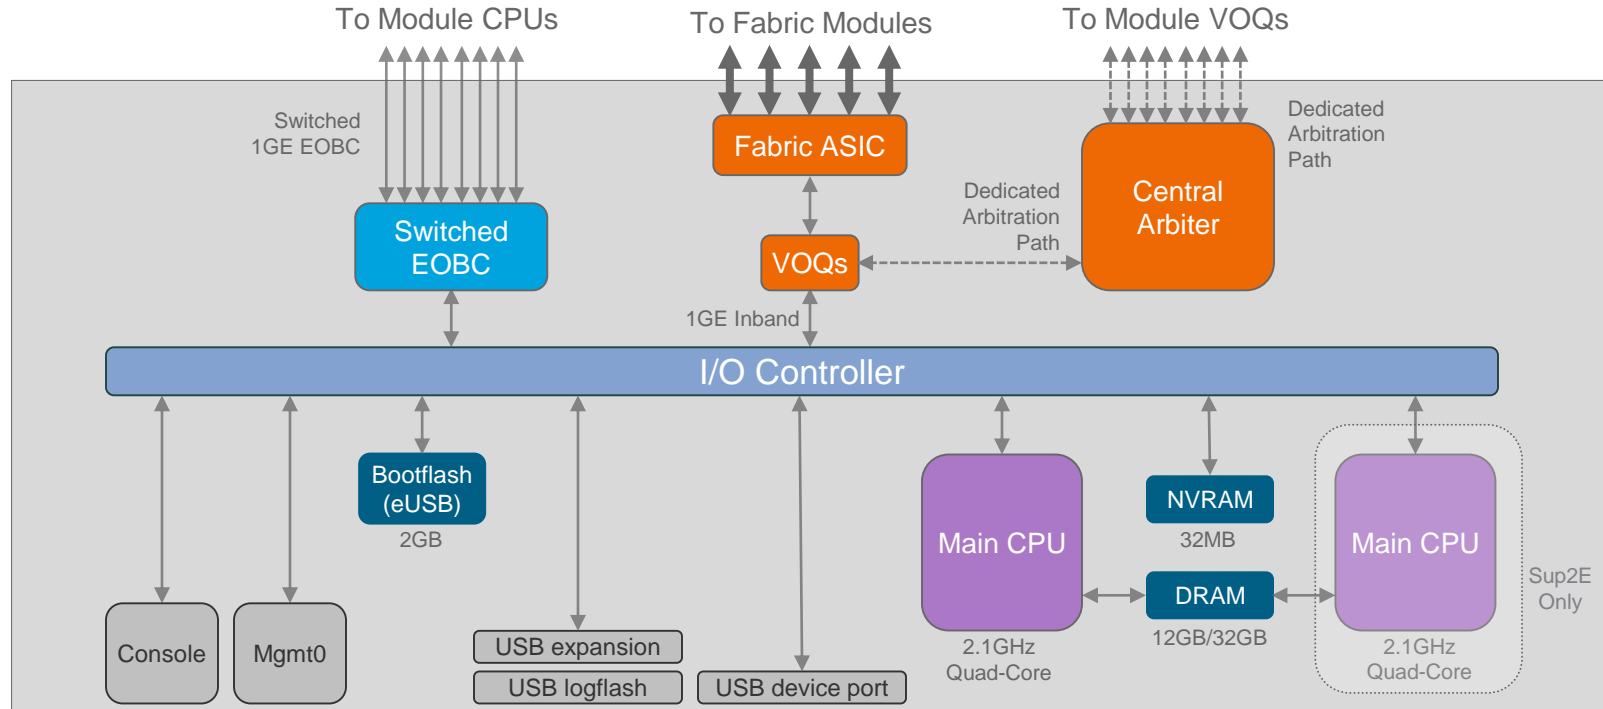
N7K-SUP2/N7K-SUP2E



N77-SUP2E




Supervisor Engine 2 / 2E Architecture



Reference: Component Functions – Supervisor Engines

- Main CPU(s) – Runs all system-level NX-OS processes and handles all control plane and management functions
- Switched EOBC – Provides switch 1G connections to each module CPU for internal system management and communication
- I/O Controller – Provides all I/O functions for supervisor components
- VOQs – Interface to central arbiter and local crossbar fabric, implements Virtual Output Queuing
- Fabric ASIC – Local fabric that provides first/third stage of three-stage crossbar
- Central Arbiter – Dedicated ASIC that controls access to fabric based on destination interface and priority of requests

Nexus 7000 / 7700 I/O Module Families



Nexus 7000 M2 I/O Modules

N7K-M224XP-23L / N7K-M206FQ-23L / N7K-M202CF-22L

- 10G / 40G / 100G M2 I/O modules
- Share common hardware architecture – multi-chipset
- Two integrated forwarding engines (120Mpps)
- Layer 2/Layer 3 forwarding with L3/L4 services (ACL/QOS) and advanced features (MPLS/OTV/GRE etc.)
- Large forwarding tables (900K FIB/128K ACL)
- 802.1AE LinkSec on all ports


N7K-M224XP-23L



N7K-M206FQ-23L



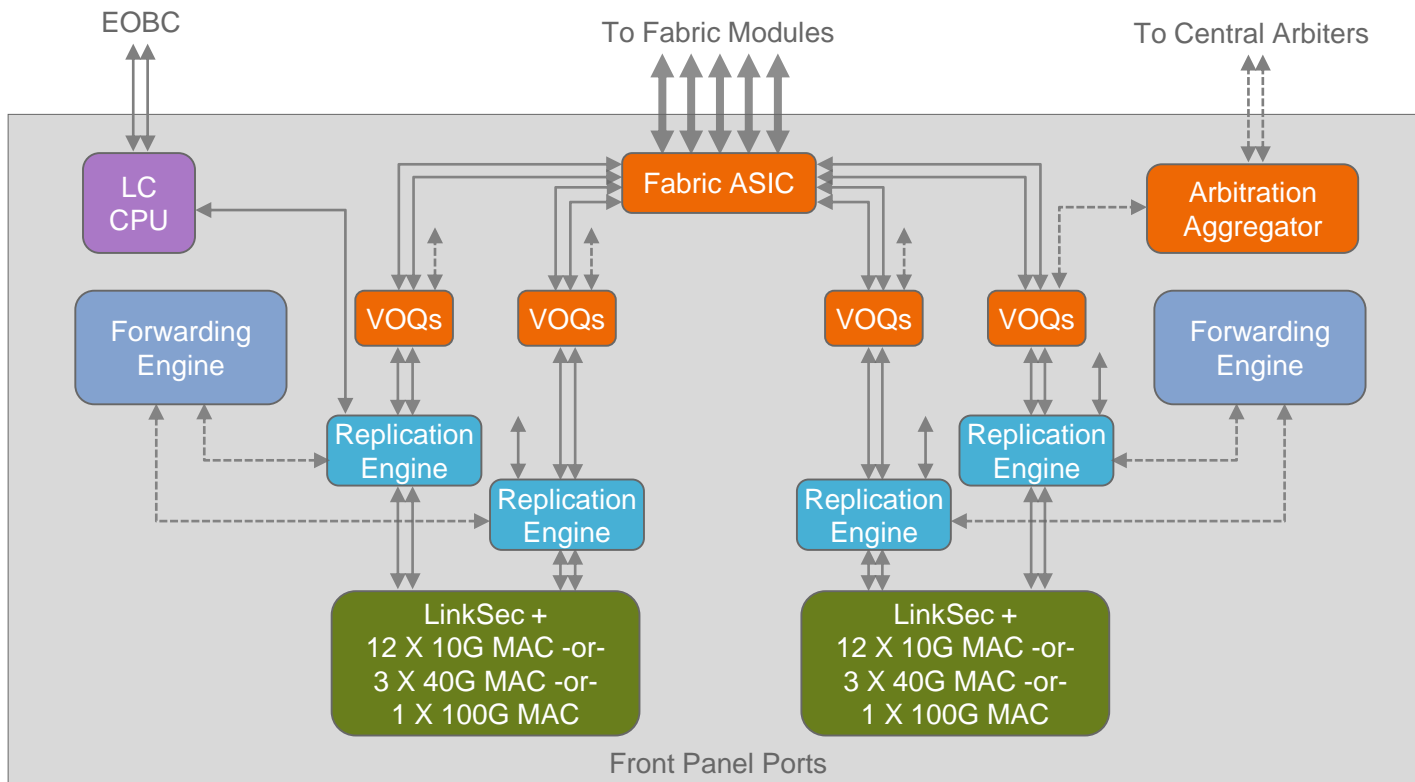
N7K-M202CF-22L



Module	Port Density	Optics	Bandwidth
M2 10G	24 x 10G (plus Nexus 2000 FEX support)	SFP+	240G
M2 40G	6 x 40G (or up to 24 x 10G via breakout)	QSFP+	240G
M2 100G	2 x 100G	CFP	200G

Nexus 7000 M2 I/O Module Architecture

N7K-M224XP-23L / N7K-M206FQ-23L / N7K-M202CF-22L




Reference: ASIC Functions – M2 Modules

- LinkSec + MAC – Provides port ASIC functions, including buffering/queuing, and performs 802.1ae encryption/decryption for front-panel ports
- Replication Engine – Bridge between front panel port, forwarding engine, and fabric; performs multicast and SPAN replication
- Forwarding Engine – Performs all Layer 2, Layer 3, and Layer 4 forwarding decisions and policy enforcement
- VOQs – Interface to central arbiter and local crossbar fabric, implements Virtual Output Queuing
- Arbitration Aggregator – Muxes arbitration requests from VOQs before sending to central arbiter on Supervisor Engine
- Fabric ASIC – Local fabric that provides first/third stage of three-stage crossbar
- (LC CPU – Linecard CPU, runs module-specific NX-OS processes and interfaces with Supervisor Engine over EOBC)

M2 Module 40G and 100G Flow Limits

- M2 modules use 10G Virtual Queuing Index (VQI)
- Each VQI sustains 10G traffic flow
- All packets in given 5-tuple flow hash to single VQI using port-channel load-balancing algorithm
- Single-flow limit is 10G



Nexus 7000 / Nexus 7700 F2E I/O Modules


N7K-F248XP-25E / N7K-F248XT-25E / N77-F248XP-23E

- 48-port 1G/10G with SFP/SFP+ transceivers
- 48-port 1G/10GBaseT with RJ-45 connectors (Nexus 7000)
- 480G full-duplex fabric connectivity
- System-on-chip (SOC) forwarding engine design
 - 12 independent SOC ASICs
- Layer 2/Layer 3 forwarding with L3/L4 services (ACL/QOS)
- Interoperability with M1/M2, in Layer 2 mode on Nexus 7000
 - Proxy routing for inter-VLAN/L3 traffic


Nexus 7000
N7K-F248XP-25E



Nexus 7000
N7K-F248XT-25E

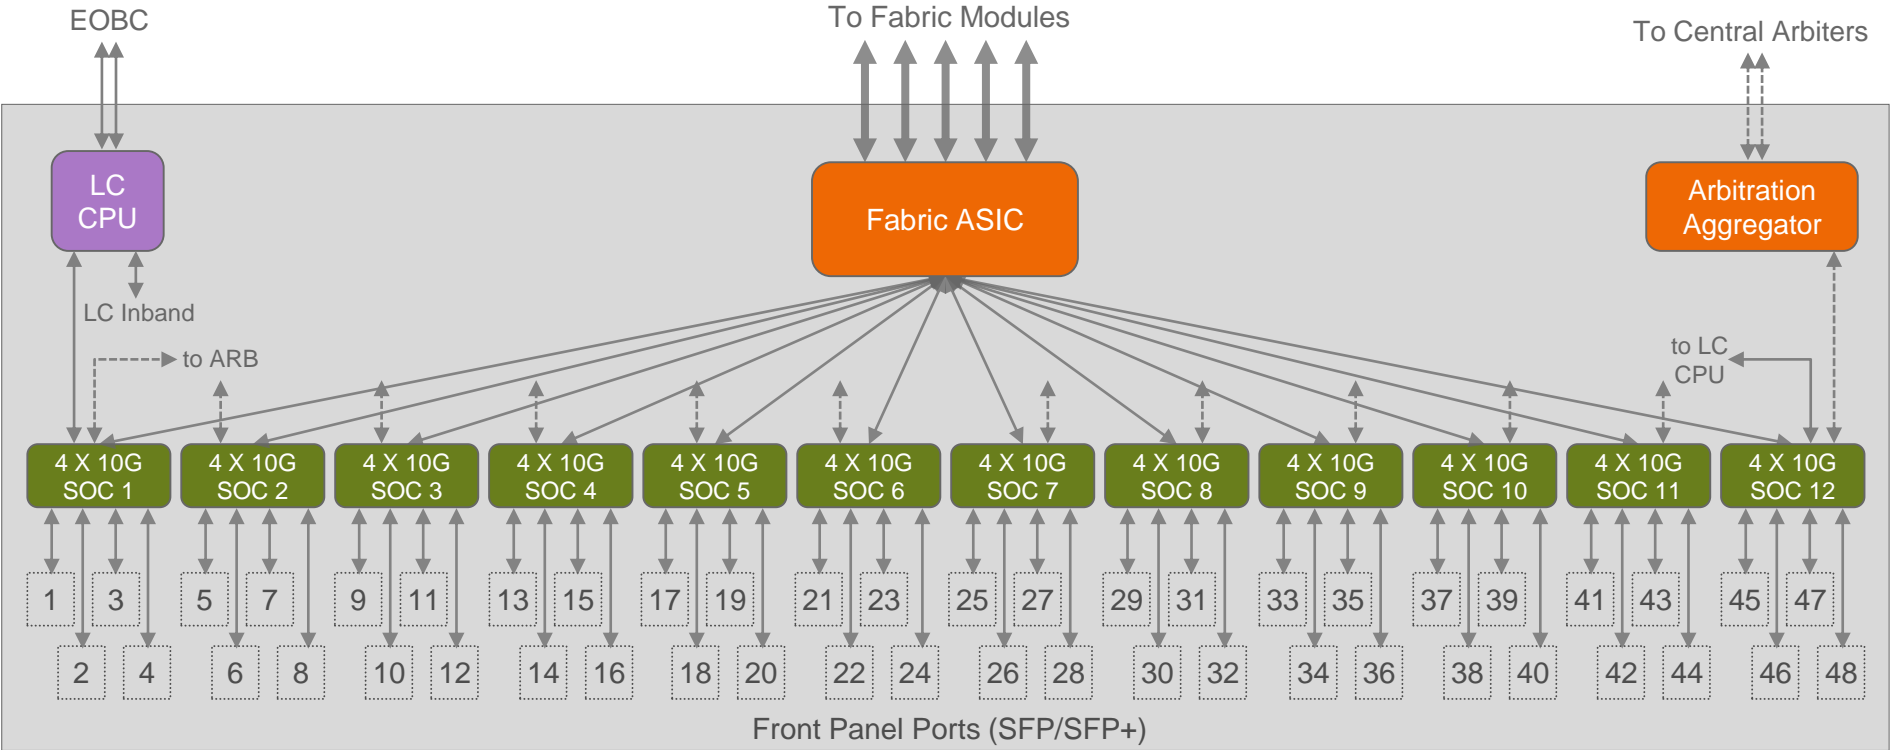


Nexus 7700
N77-F248XP-23E



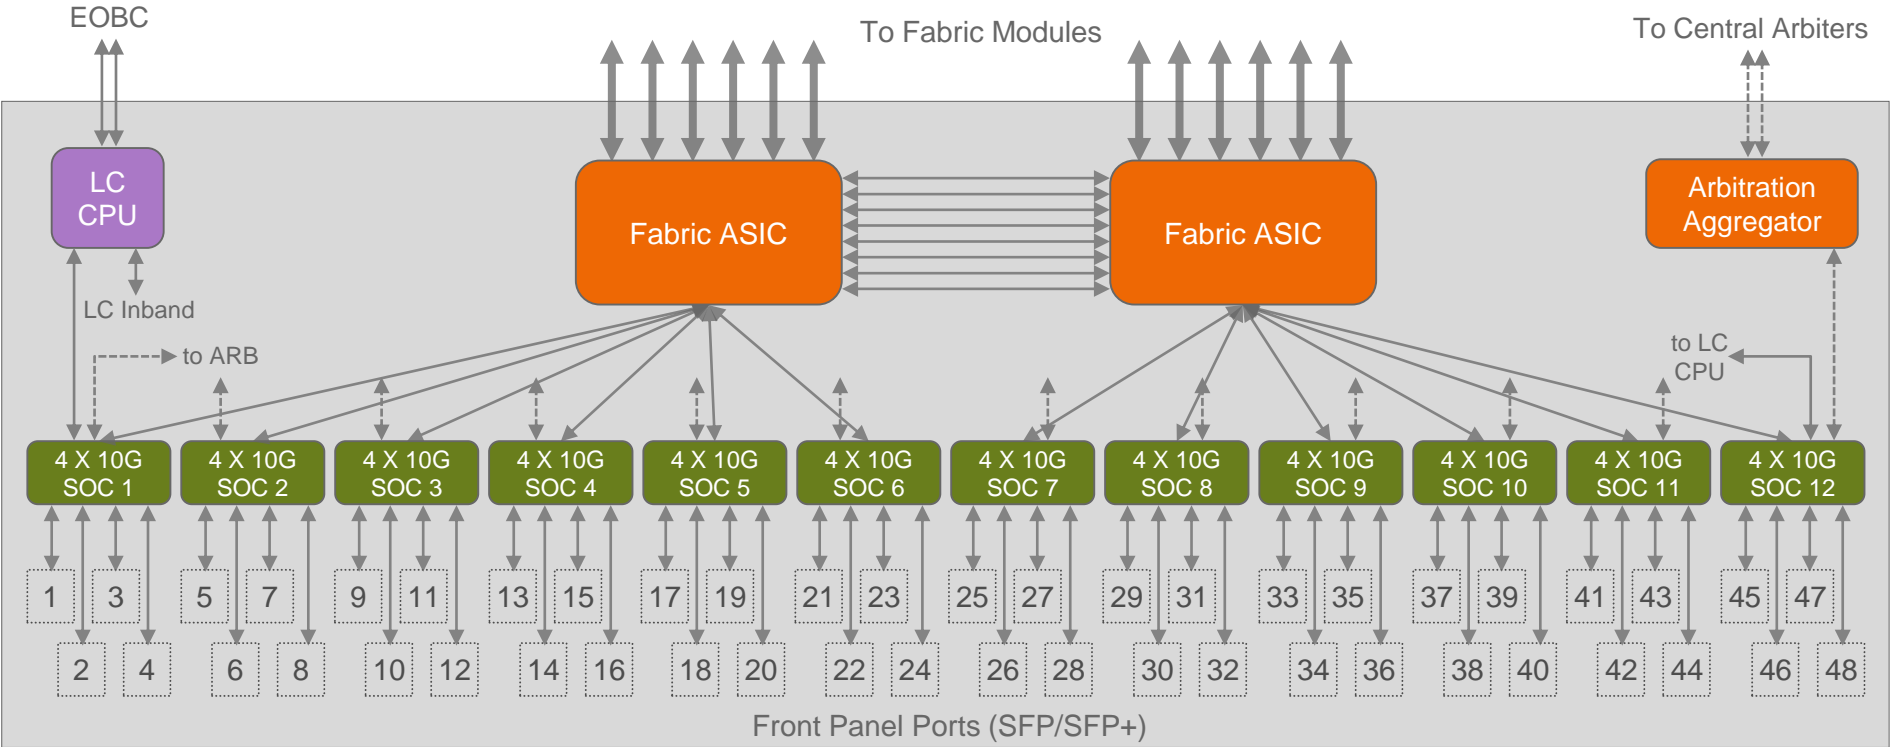
Nexus 7000 F2E Module Architecture

N7K-F248XP-25E / N7K-F248XT-25E



Nexus 7700 F2E Module Architecture

N77-F248XP-23E



Reference: ASIC Functions – F2E Modules

- 4 X 10G SoC – Four-port 10G system-on-chip; provides Port ASIC, Replication Engine, Forwarding Engine, and VOQ functions
- Arbitration Aggregator – Muxes arbitration requests from SoCs before sending to central arbiter on Supervisor Engine
- Fabric ASIC – Local fabric that provides first/third stage of three-stage crossbar
- (LC CPU – Linecard CPU, runs module-specific NX-OS processes and interfaces with Supervisor Engine over EOBC)


Nexus 7000 F3 I/O Modules

N7K-F348XP-25 / N7K-F312FQ-25 / N7K-F306CK-25


- 10G / 40G / 100G F3 I/O modules
- Share common hardware architecture
- SOC-based forwarding engine design
 - 6 independent SOC ASICs per module
- Layer 2/Layer 3 forwarding with L3/L4 services (ACL/QOS) and advanced features (FP/MPLS/OTV/GRE/VXLAN etc.)
- **Require Supervisor Engine 2 / 2E**



N7K-F348XP-25



N7K-F312FQ-25




N7K-F306CK-25

Module	Port Density	Optics	Bandwidth
F3 10G	48 x 1/10G (plus Nexus 2000 FEX support)	SFP+	480G
F3 40G	12 x 40G (or up to 48 x 10G via breakout)	QSFP+	480G
F3 100G	6 x 100G	CPAK	550G

Nexus 7700 F3 I/O Modules

N7K-F348XP-25 / N7K-F312FQ-25 / N7K-F306CK-25

- 10G / 40G / 100G F3 I/O modules
- Share common hardware architecture
- SOC-based forwarding engine design
 - 6 independent SOC ASICs per 10G module
 - 12 independent SOC ASICs per 40G/100G module
- Layer 2/Layer 3 forwarding with L3/L4 services (ACL/QOS) and advanced features (FP/MPLS/OTV/GRE/VXLAN etc.)



N77-F348XP-23



N77-F324FQ-25

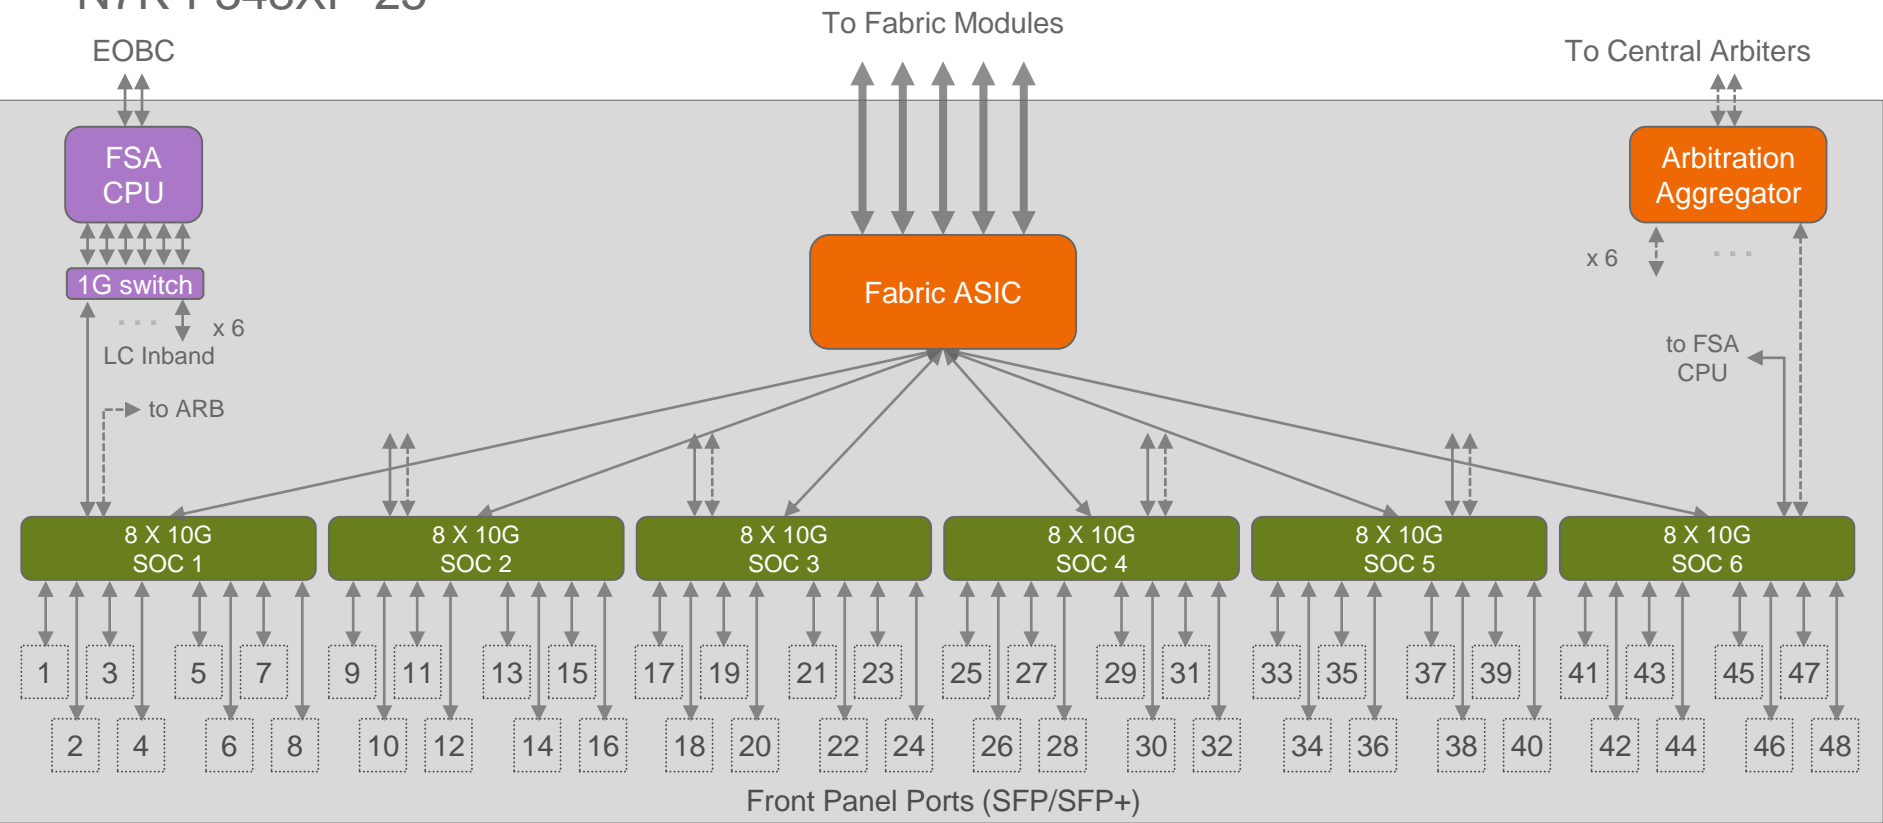


N77-F312CK-26

Module	Port Density	Optics	Bandwidth
F3 10G	48 x 1/10G (plus Nexus 2000 FEX support)	SFP+	480G
F3 40G	24 x 40G (or up to 76 x 10G + 5 x 40G via breakout)	QSFP+	960G
F3 100G	12 x 100G	CPAK	1.2T

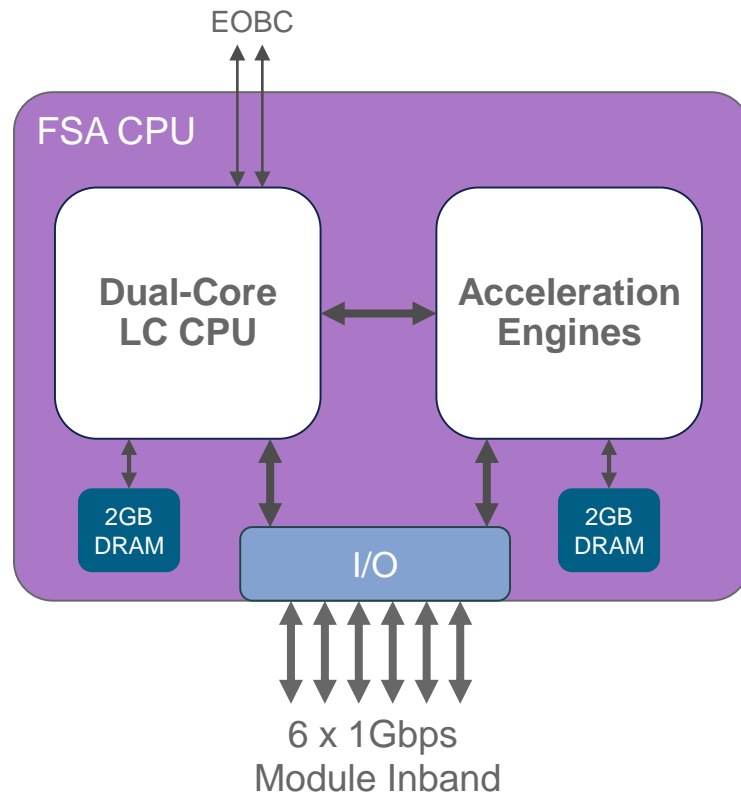
Nexus 7000 F3 48-Port 1G/10G Module Architecture

N7K-F348XP-25



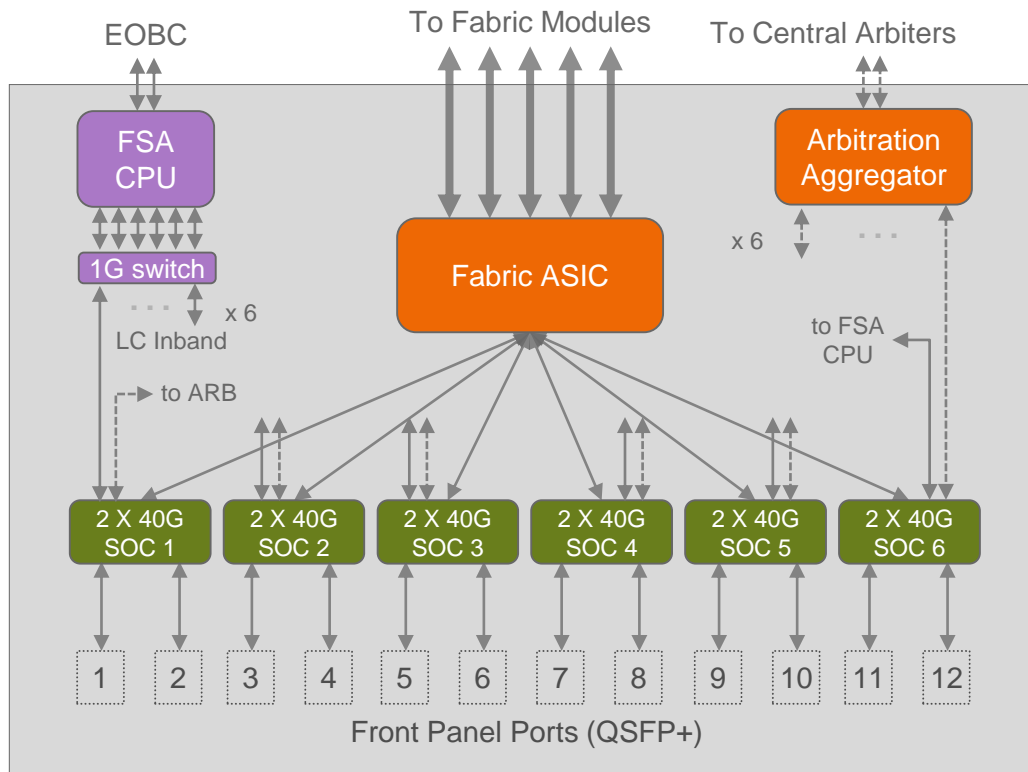
Fabric Services Accelerator (FSA) for F3

- High-performance module CPU with on-board acceleration engines
 - 6Gbps inband connectivity from SOC's to FSA
 - Multi-Mpps packet processing
 - 2 X 2GB dedicated DRAM
- Performance/scale boost for distributed fabric services, including sampled Netflow and BFD (roadmap)
- Other potential applications include distributed ARP/ping processing, data plane packet analysis (wireshark), network probing, etc.



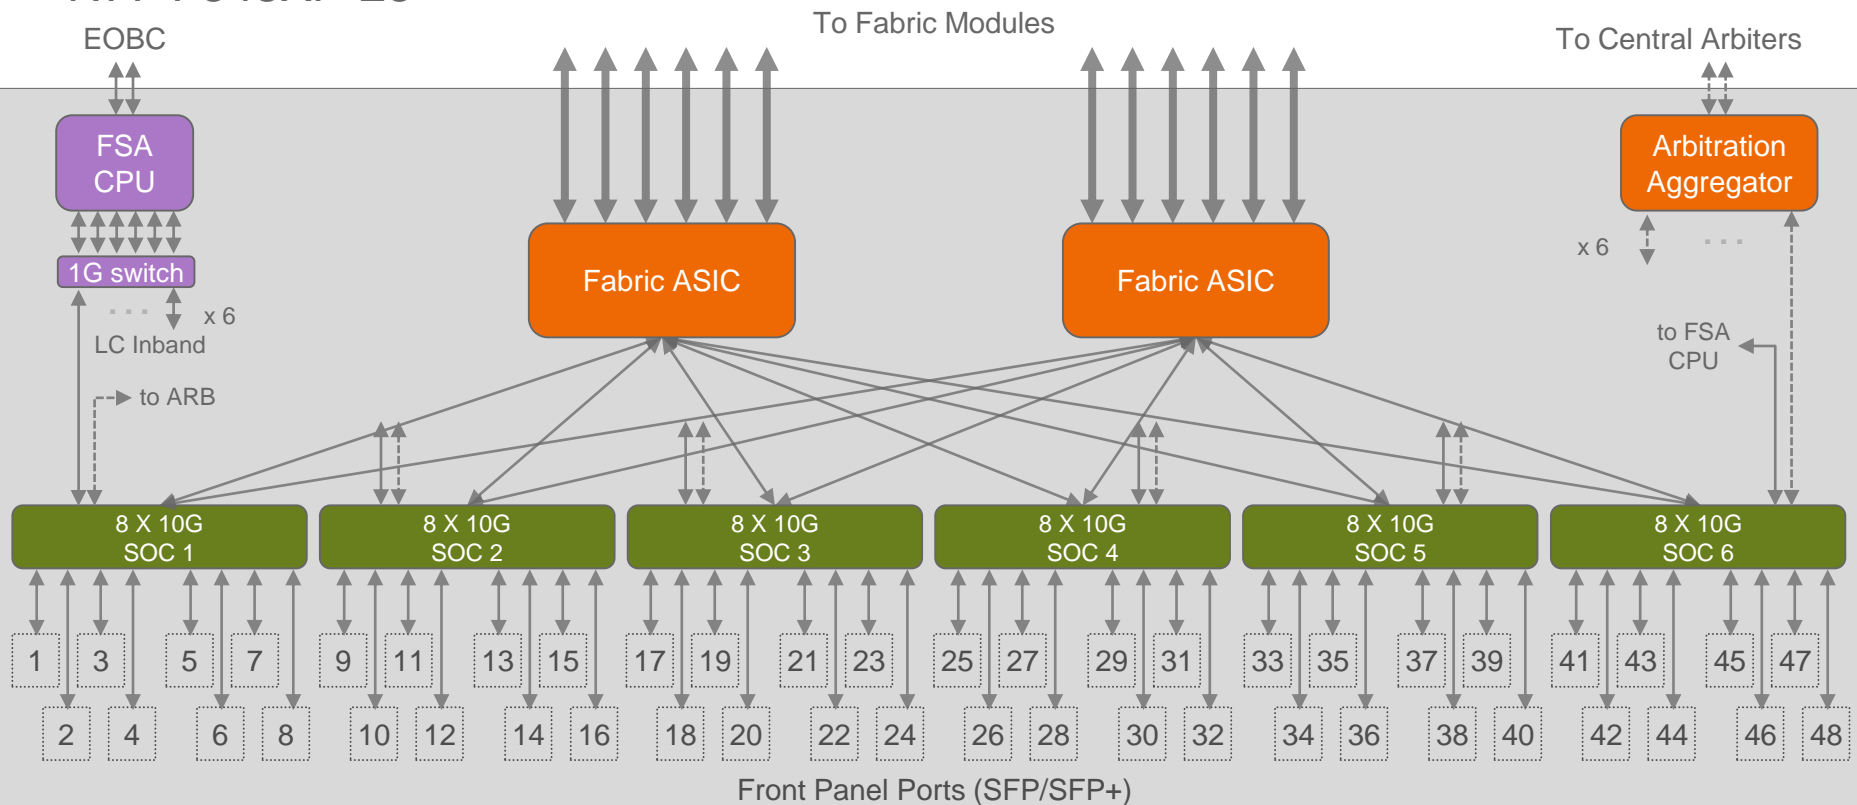
Nexus 7000 F3 12-Port 40G Module Architecture

N7K-F312FQ-25



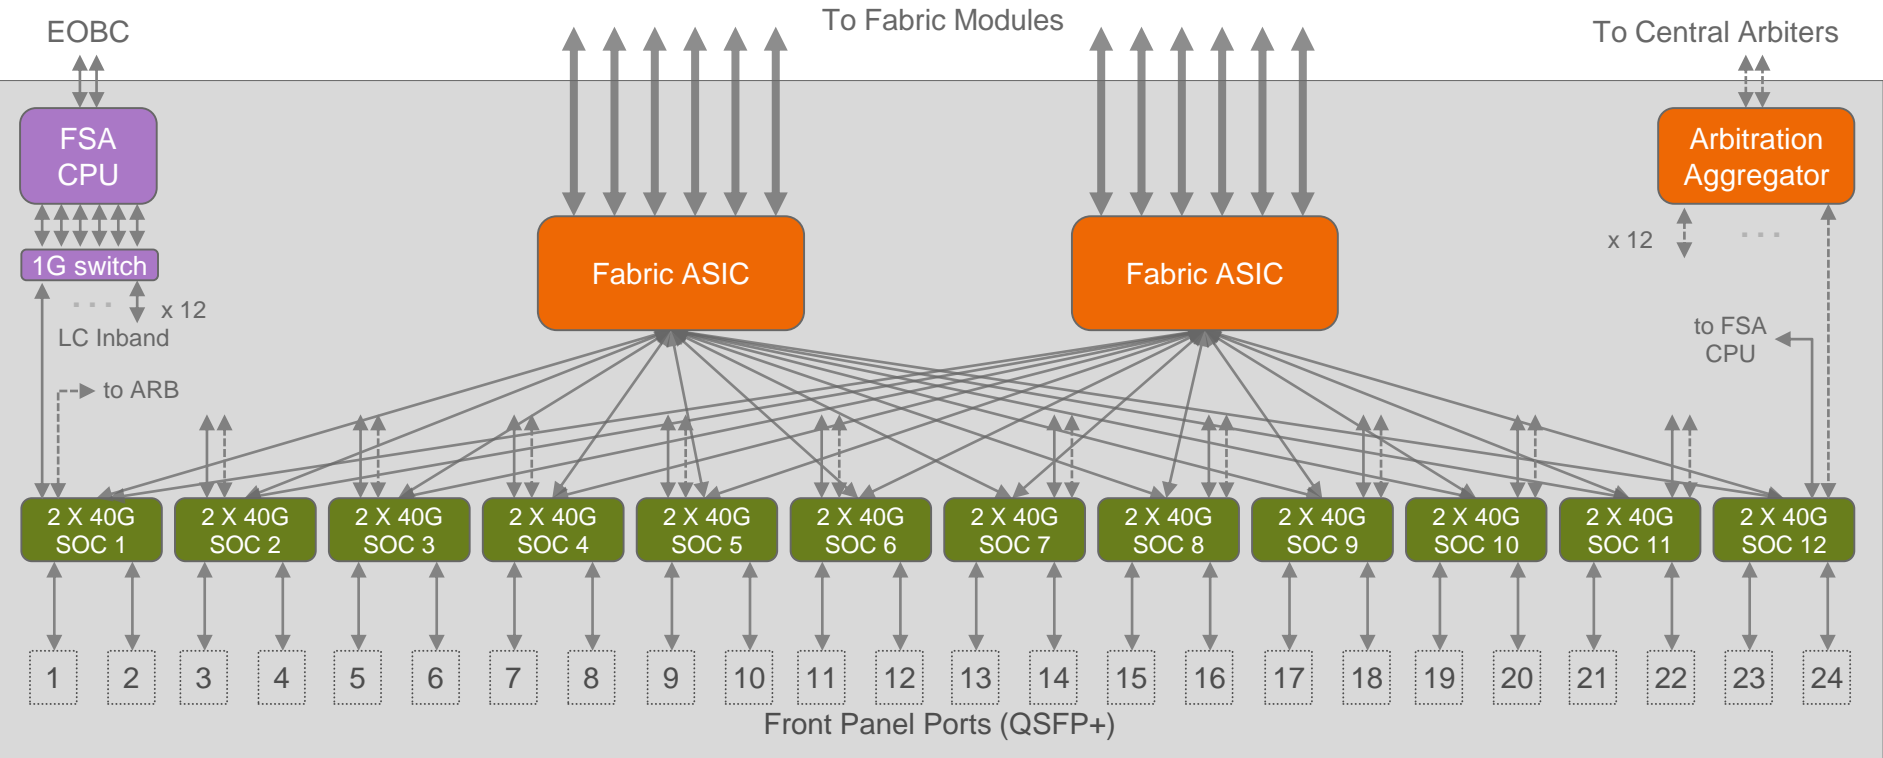
Nexus 7700 F3 48-Port 1G/10G Module Architecture

N77-F348XP-23



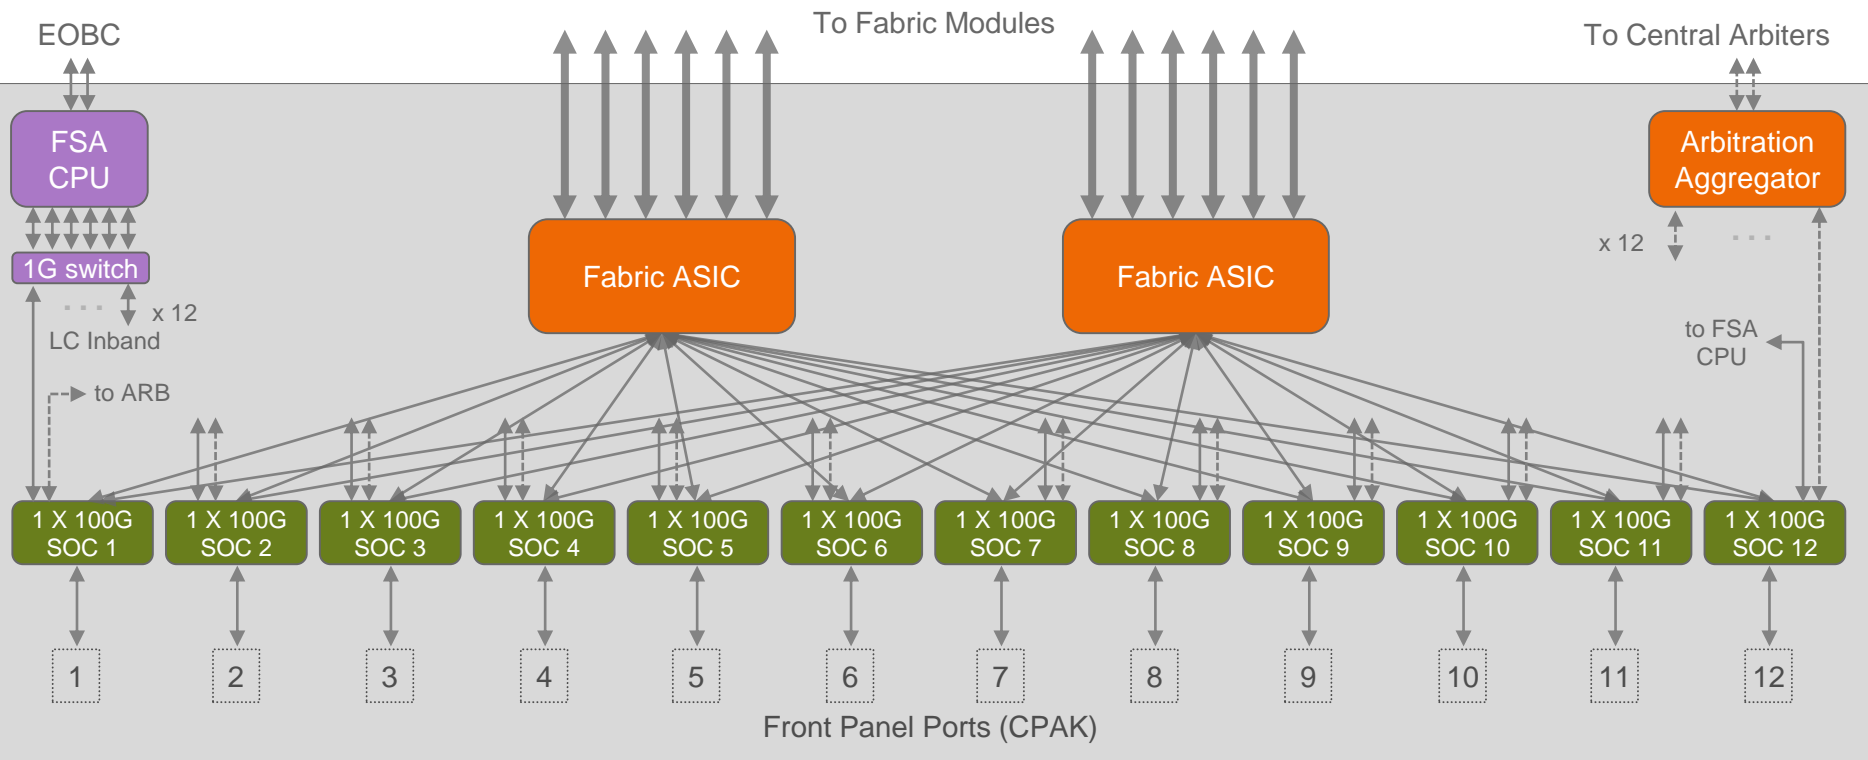
Nexus 7700 F3 24-Port 40G Module Architecture

N77-F324FQ-25



Nexus 7700 F3 12-Port 100G Module Architecture

N77-F312CK-26




Reference: ASIC Functions – F3 Modules

- 8 X 10G / 2 X 40G / 1 X 100G SOC – 10/40/100G capable system-on-chip; provides Port ASIC, Replication Engine, Forwarding Engine, and VOQ functions
- Arbitration Aggregator – Muxes arbitration requests from SoCs before sending to central arbiter on Supervisor Engine
- Fabric ASIC – Local fabric that provides first/third stage of three-stage crossbar
- (FSA CPU – Fabric Services Accelerator, a linecard CPU with built-in application acceleration for higher performance BFD, sampled Netflow, and other functions; runs module-specific NX-OS processes and interfaces with Supervisor Engine over EOBC)

F3 Module 40G and 100G Flows

- Virtual Queuing Index (VQI) sustains 10G, 40G, or 100G traffic flow based on destination interface type
- No single-flow limit – full 40G/100G flow support



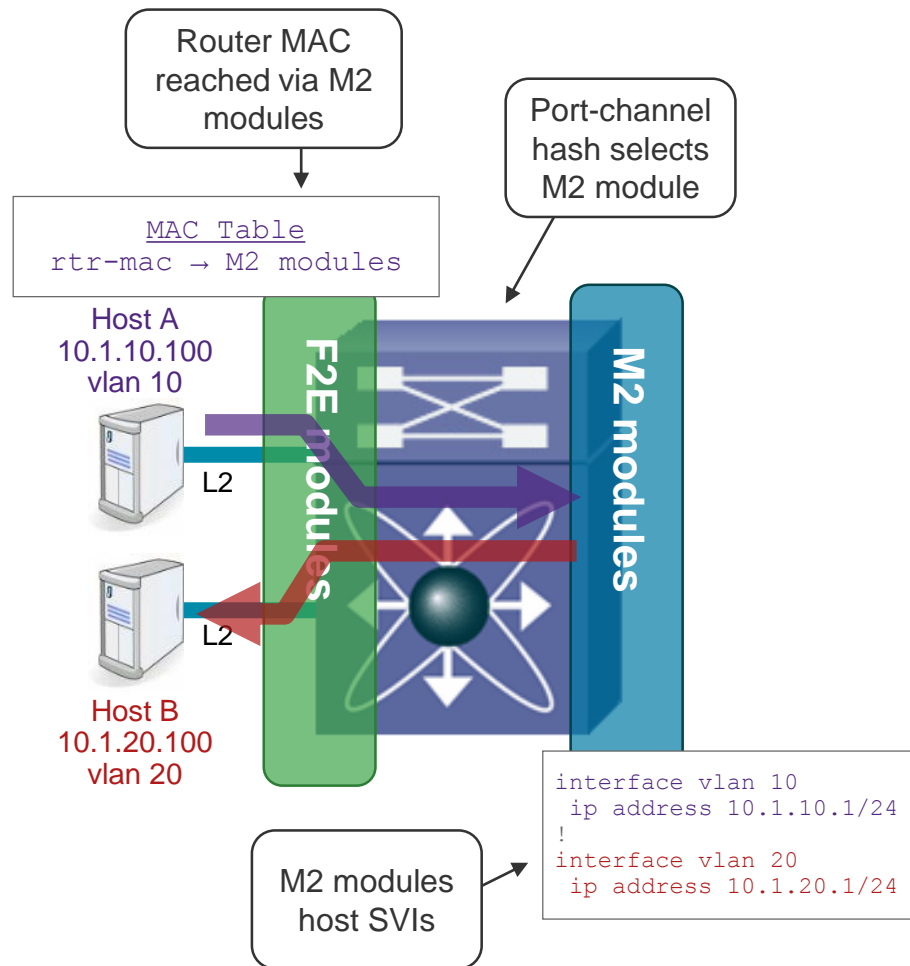
I/O Module Interoperability

- General module interoperability rule is: “+/-1 generation” in same Virtual Device Context (VDC)
- Two main module interoperability models:
 - “Proxy Forwarding”
 - “Ingress Forwarding” with Lowest Common Denominator

Proxy Forwarding Model

M2 + F2E VDC

- F2E modules run in pure Layer 2 mode – all L3 functions disabled
- M2 modules host SVIs and other L3 functions
- From F2E perspective, Router MAC reachable via M2 modules
- All packets destined to Router MAC forwarded through fabric toward one M2 module, selected via port-channel hash
- M2 module(s) perform all L3 forwarding and policy, pass packets back over fabric to output port
- **Key consideration:** M-series L3 routing capacity versus F-series front-panel port count – How much Layer 3 routing is required?



Ingress Forwarding with Lowest Common Denominator Model

F3 + M2 VDC -or- F3 + F2E VDC

- F3 module interoperability always “Ingress Forwarding” – NO proxy forwarding
 - Ingress module receiving packet makes all forwarding decisions for that packet
- Supported feature set and scale based on Lowest Common Denominator
 - Feature available if all modules support the feature
 - Table sizes based on lowest capacity

Not all features supported by software today...

Module Types in VDC	Layer 2	Layer 3	VPC	MPLS	OTV	Fabric Path	VXLAN	Table Sizes
F3	✓	✓	✓	✓	✓	✓	✓	F3 size
F3 + M2	✓	✓	✓	✓	✓	✗	✗	F3 size
F3 + F2E	✓	✓	✓	✗	✗	✓	✗	F2E size
M2 + F2E + F3	Not supported							

Interoperability Forwarding Model Matrix

Module Combination	Interoperability Model
M1 + M2	Lowest Common Denominator
M + F1	Proxy Forwarding
M + F2E	Proxy Forwarding
F2 + F2E	Lowest Common Denominator
F2 + F2E + F3	Lowest Common Denominator
M2 + F3	Lowest Common Denominator
M2 + F2/F2E + F3	NOT SUPPORTED

Agenda

- Introduction to Nexus 7000 / Nexus 7700
- Chassis Architecture
- Supervisor Engine and I/O Module Architecture
- Fabric Architecture
- Hardware Forwarding
- Packet Walks
- Conclusion

Crossbar Switch Fabric Modules

- Provide interconnection of I/O modules
- Nexus 7000 and Nexus 7700 fabrics based on Fabric 2 ASIC
- Each installed fabric increases available per-payload slot bandwidth

Fabric Module	Supported Chassis	Per-fabric module bandwidth	Max fabric modules	Total bandwidth per slot
Nexus 7000 Fabric 2	7009 / 7010 / 7018	110Gbps per slot	5	550Gbps per slot
Nexus 7700 Fabric 2	7706 / 7710 / 7718	220Gbps per slot	6	1.32Tbps per slot

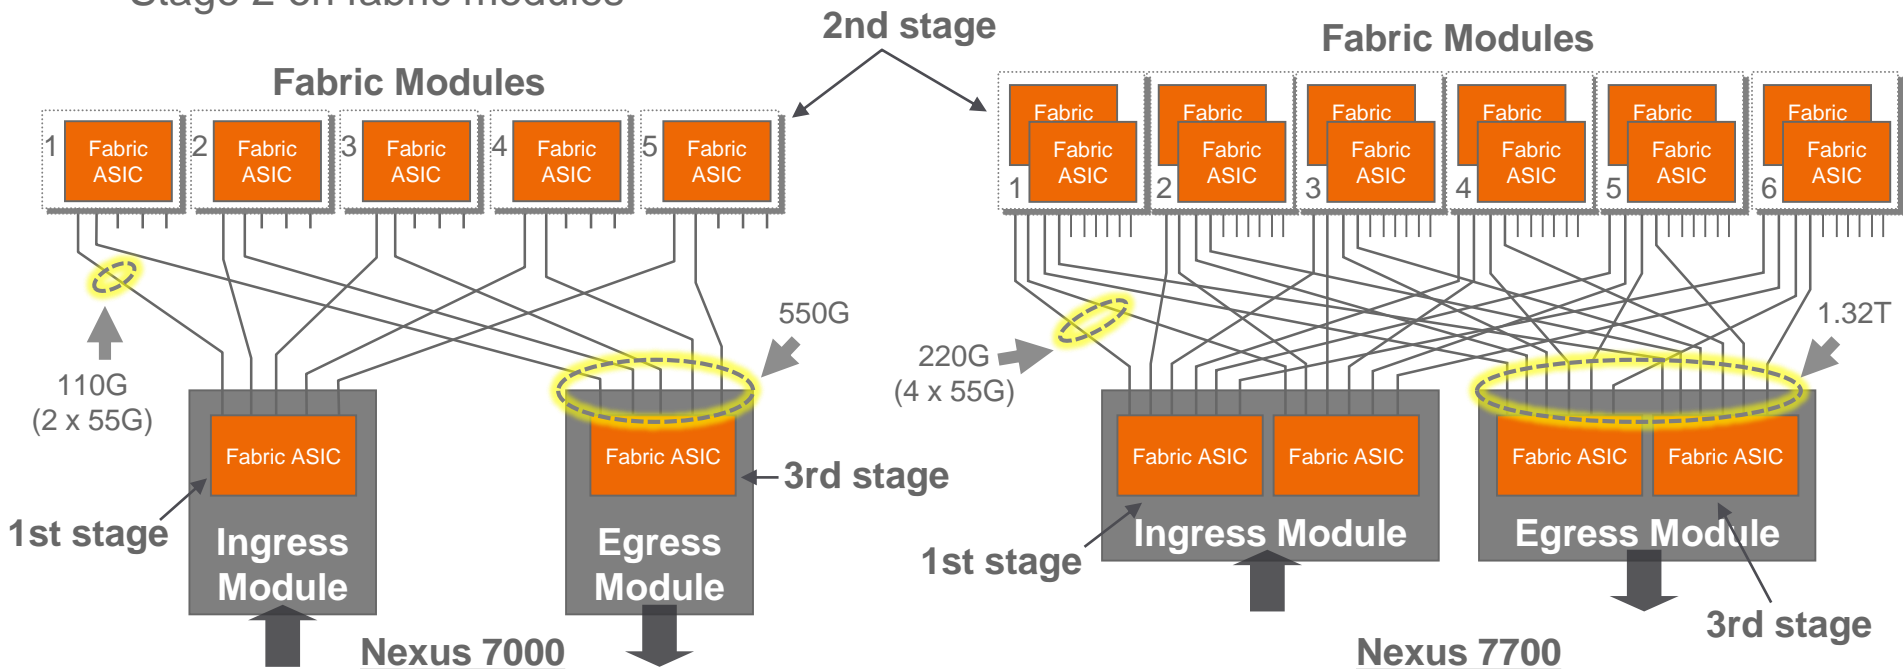
- Different I/O modules leverage different amount of available fabric bandwidth
- Access to fabric bandwidth controlled using QOS-aware central arbitration with VOQ



Multistage Crossbar

Nexus 7000 / Nexus 7700 implement 3-stage crossbar switch fabric

- Stages 1 and 3 on I/O modules
- Stage 2 on fabric modules



I/O Module Capacity – Nexus 7000

550Gbps

per slot bandwidth

One fabric:

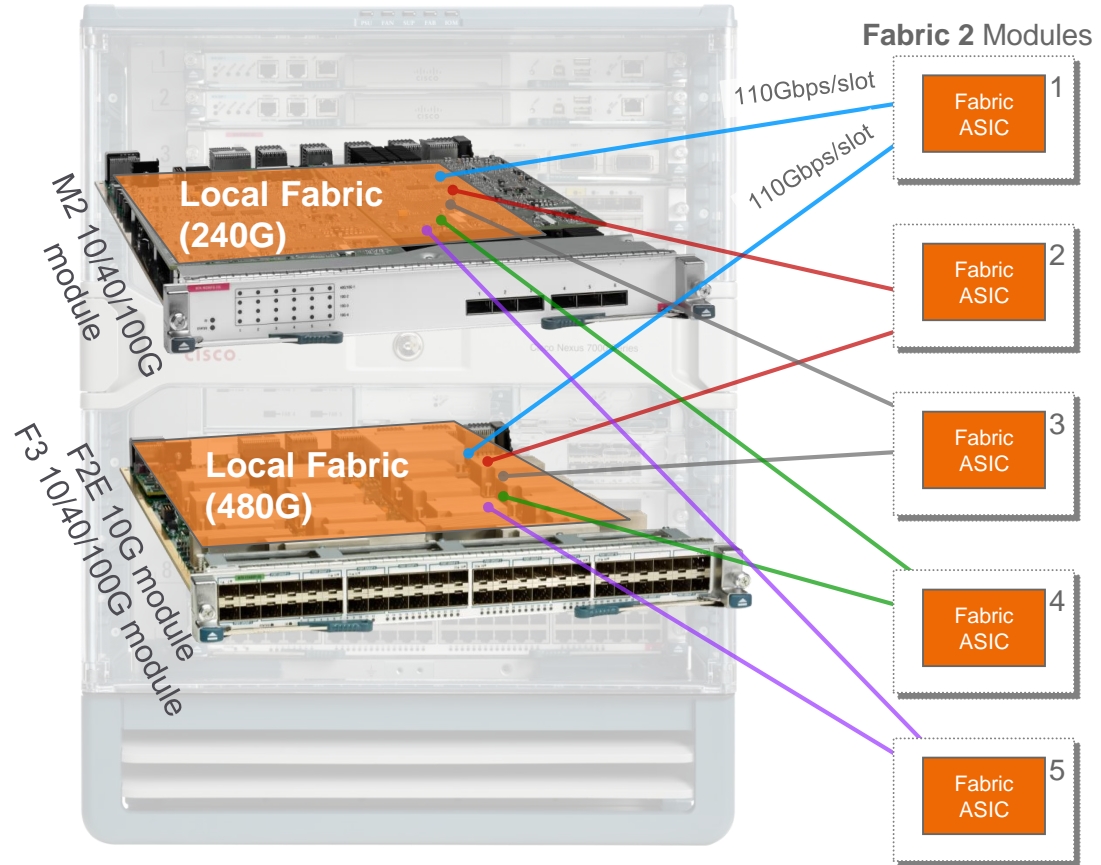
- Any port can pass traffic to any other port in VDC

Three fabrics:

- 240G M2 module has maximum bandwidth

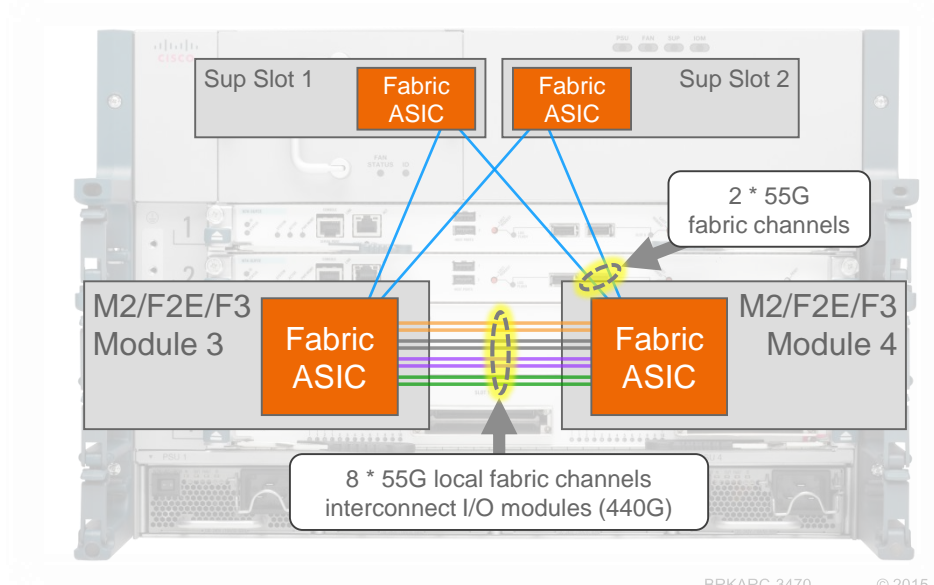
Five fabrics:

- 480G F2E/F3 module has maximum bandwidth



What About Nexus 7004?

- Nexus 7004 has no fabric modules
- Each I/O module has local fabric with 10 available fabric channels
 - I/O modules connect “back-to-back” via 8 fabric channels
 - Two fabric channels “borrowed” to connect supervisor engines



I/O Module Capacity – Nexus 7700

1320Gbps
per slot bandwidth

One fabric:

- **Any port** can pass traffic to **any other port** in VDC

Three fabrics:

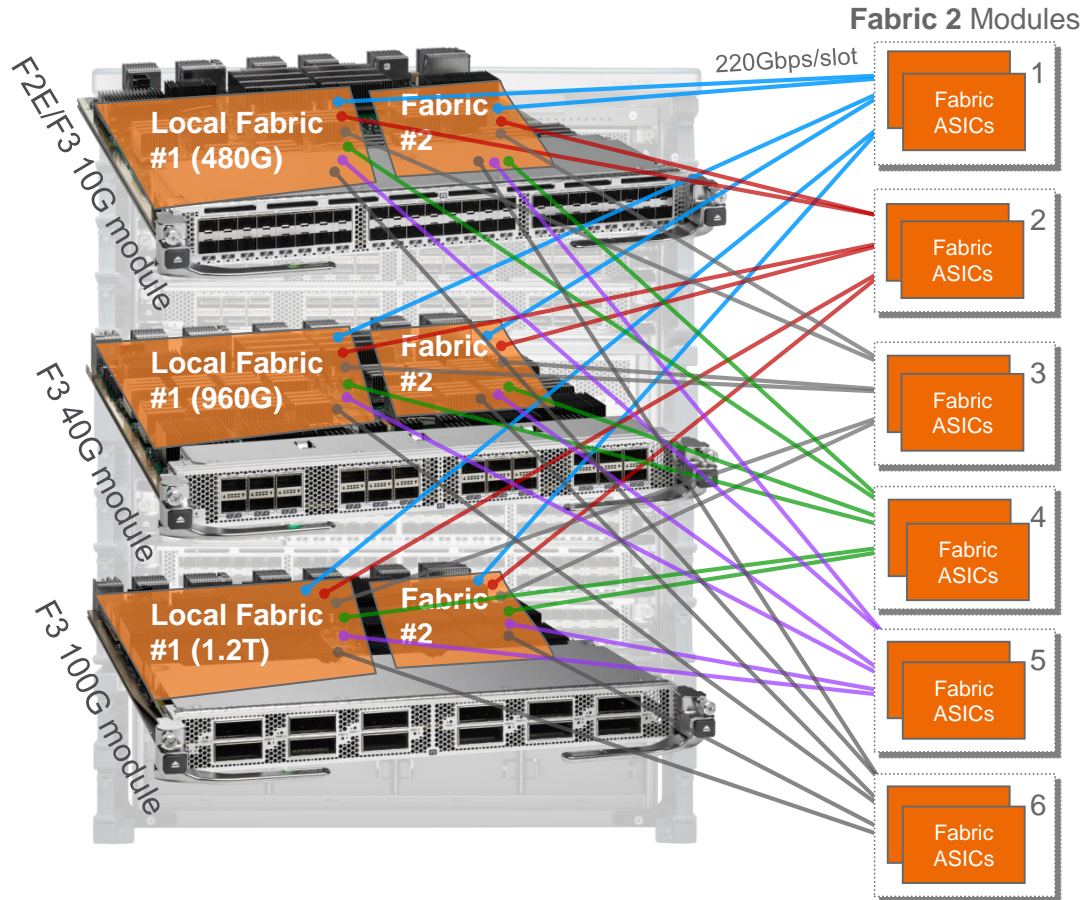
- 480G F2E/F3 10G module has maximum bandwidth

Five fabrics:

- 960G F3 40G module has maximum bandwidth

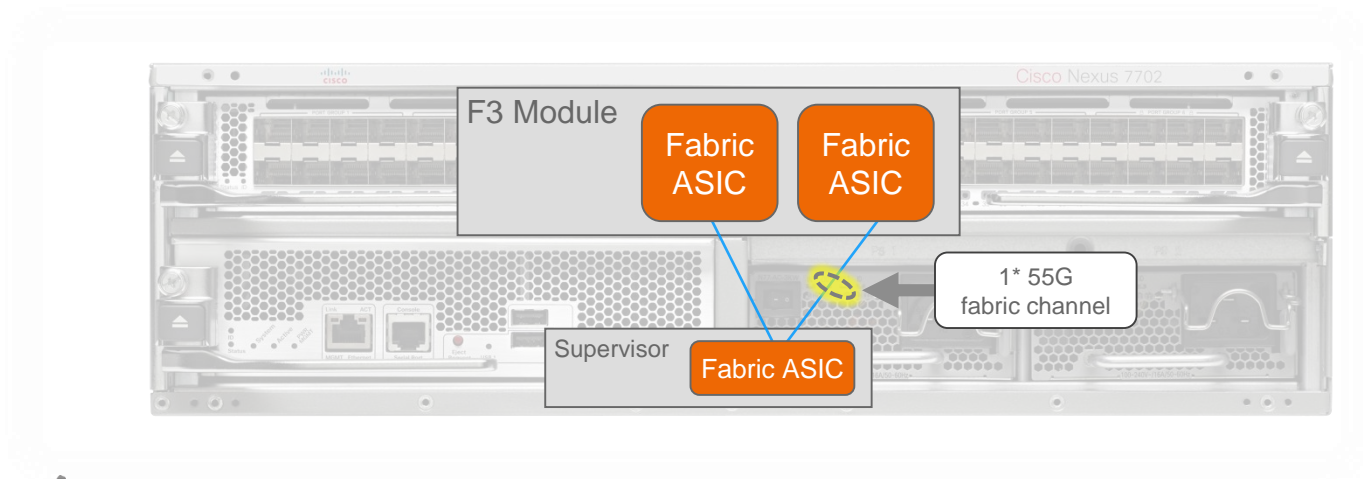
Six fabrics:

- 1.2T F3 100G module has maximum bandwidth



What About Nexus 7702?

- Nexus 7702 has no fabric modules
- Single I/O module – all traffic locally switched
- Two fabric channels connect to supervisor engine



Agenda

- Introduction to Nexus 7000 / Nexus 7700
- Chassis Architecture
- Supervisor Engine and I/O Module Architecture
- Fabric Architecture
- Hardware Forwarding
- Packet Walks
- Conclusion

Hardware Forwarding Process

- Ingress queuing and scheduling
- Perform forwarding lookups
- Forward through fabric
- Egress queuing and scheduling



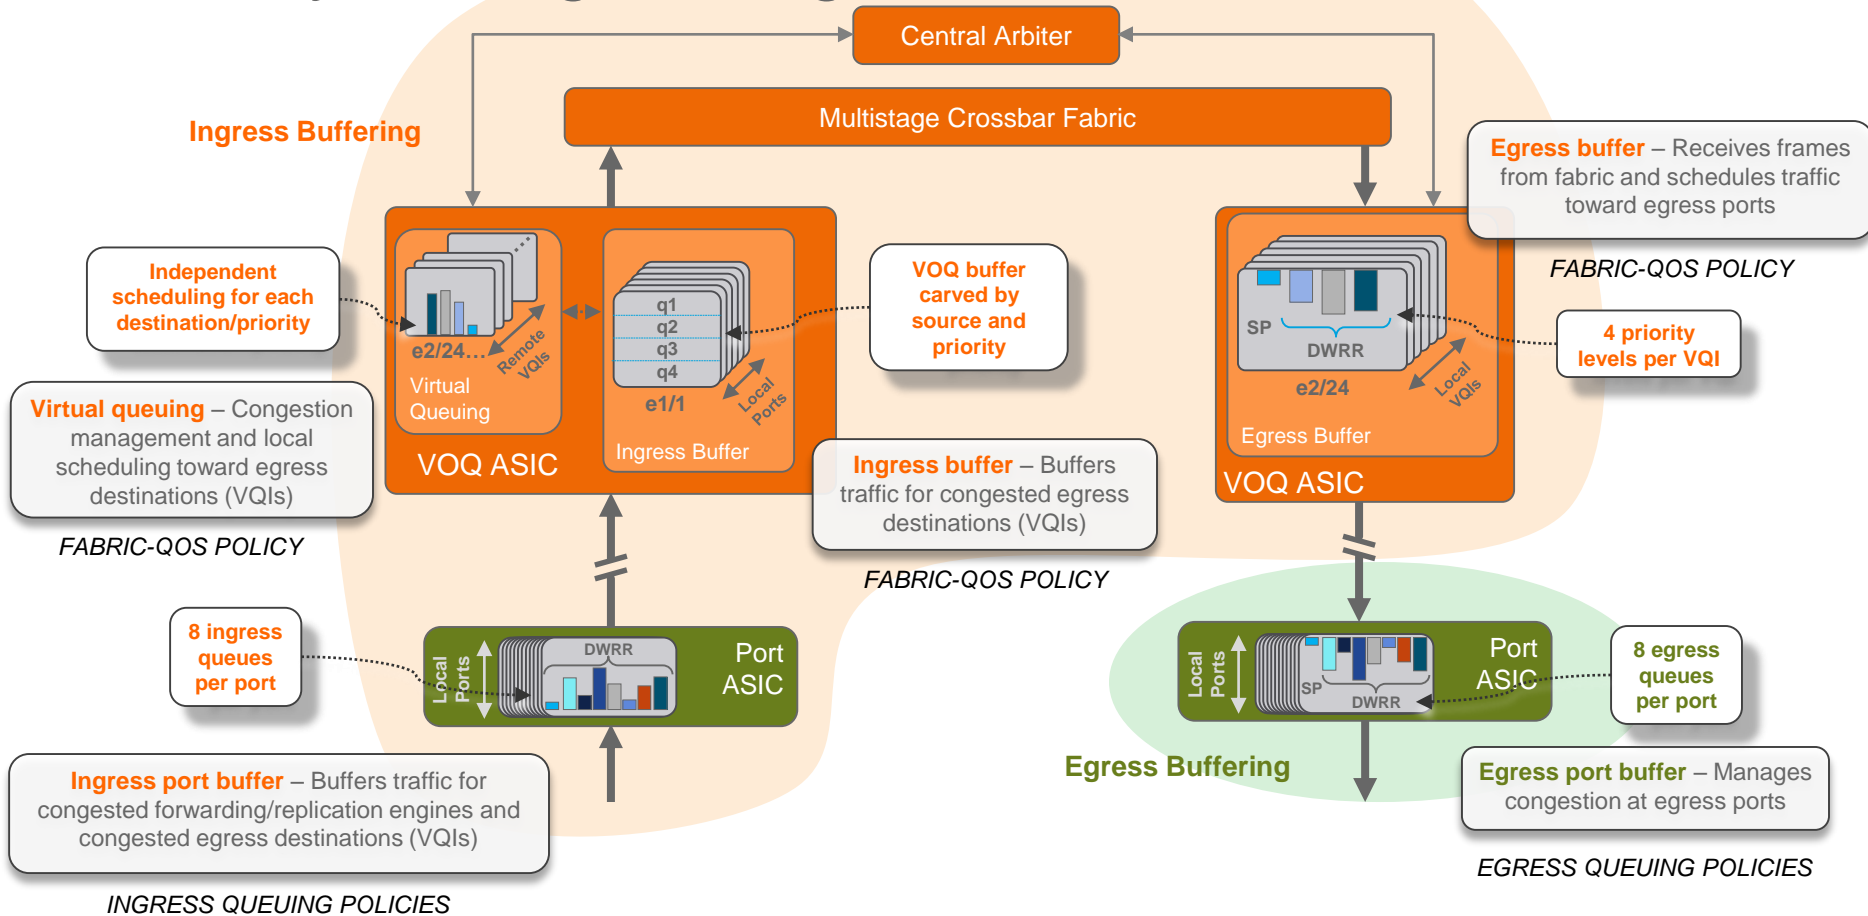
Ingress / Egress Queuing and Scheduling

- Nexus 7000 / Nexus 7700 use **queuing policies** and **network-QOS policies** to define queuing and scheduling behavior
- Default policies always in effect in absence of any user configuration
- Queuing model varies by I/O module family
 - **M-series modules**: hybrid model combining egress-buffered architecture with ingress-buffered architecture
 - **F-series modules**: pure ingress-buffered architecture

M2 – Hybrid Ingress/Egress Buffered

M2 10G module used as example

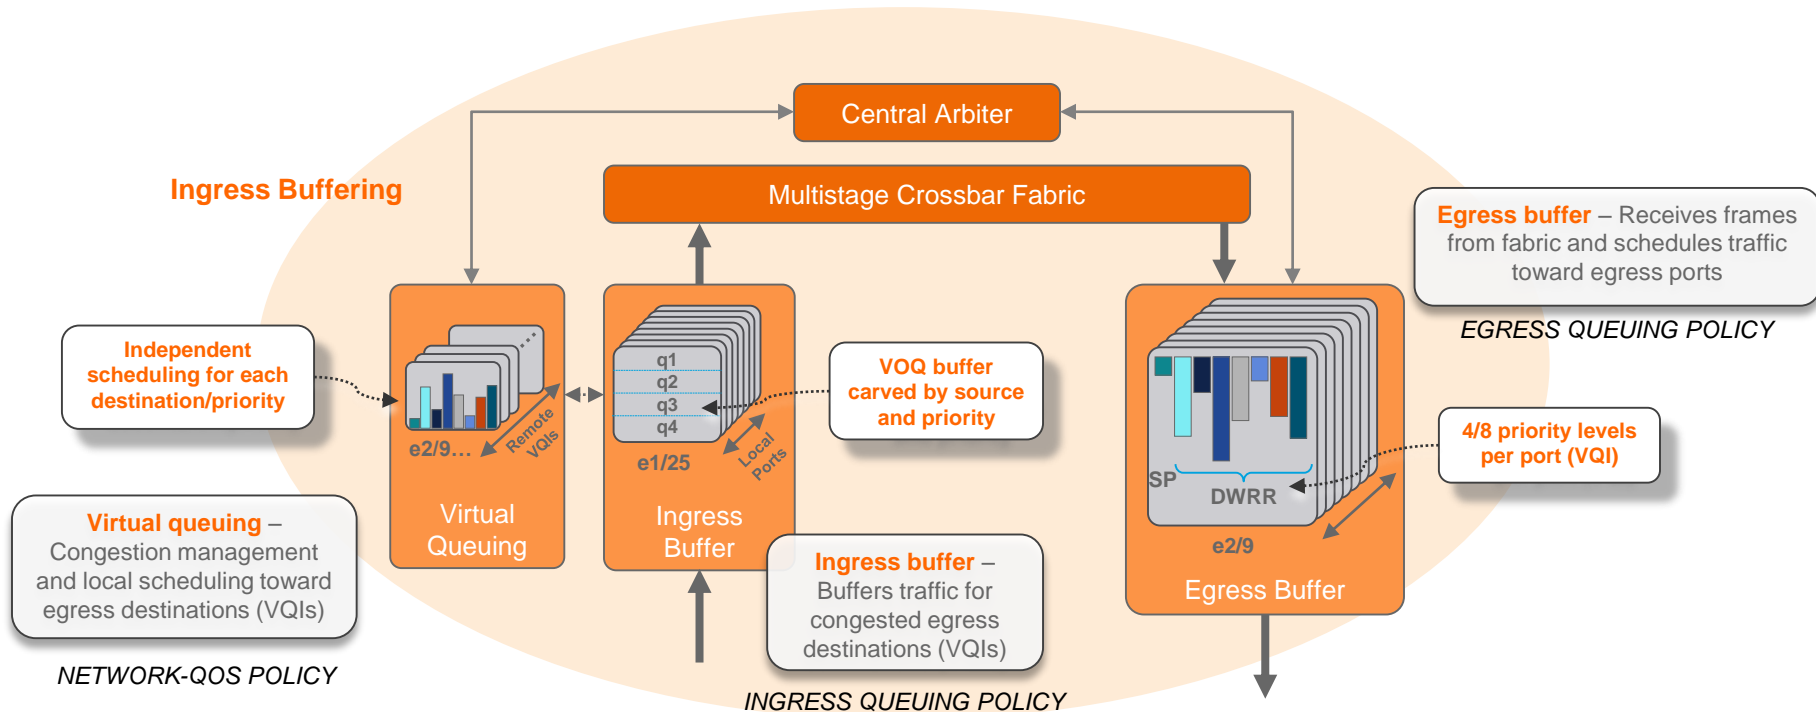
M Series



F2E/F3 – Ingress Buffered

N7700 10G F3 module used as example

F Series

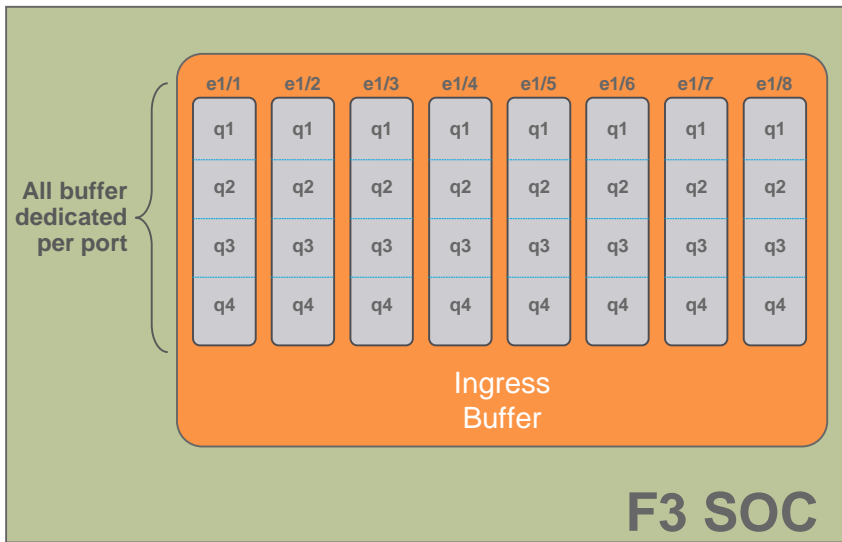


F3 Buffer Sharing

10G F3 used as example

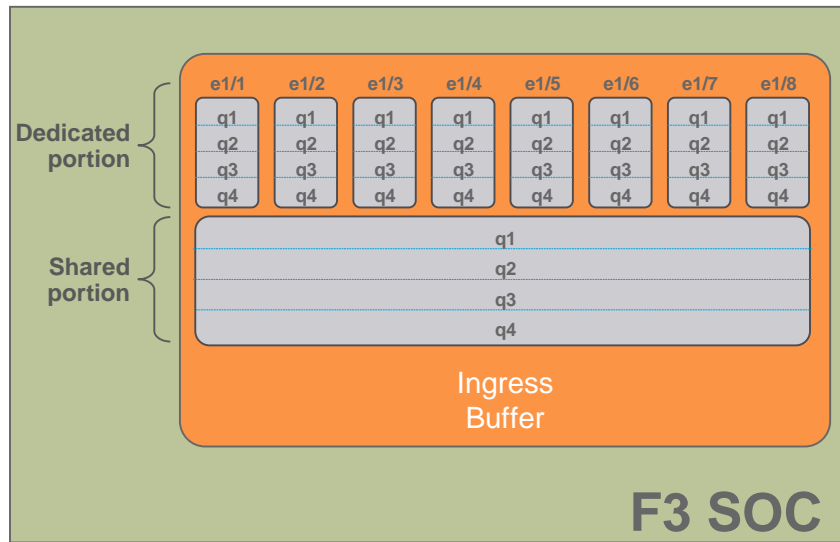
F Series

Default – Dedicated per Port Buffer



- All ingress buffer carved equally among ports
- One port cannot use another port's unused buffer

Optional – Shared Buffer + Dedicated per Port Buffer



- Total buffer split into dedicated per-port portion and shared portion
- Any port can consume buffer in shared portion
- Increases burst absorption, assuming not all ports burst simultaneously

Hardware Forwarding Lookups

- Layer 2 and Layer 3 packet flow virtually identical in hardware
- Forwarding engine / decision engine pipeline provides consistent L2 and L3 lookup performance
- Pipelined architecture also performs ingress and egress ACL, QOS, and Netflow lookups, affecting final forwarding result

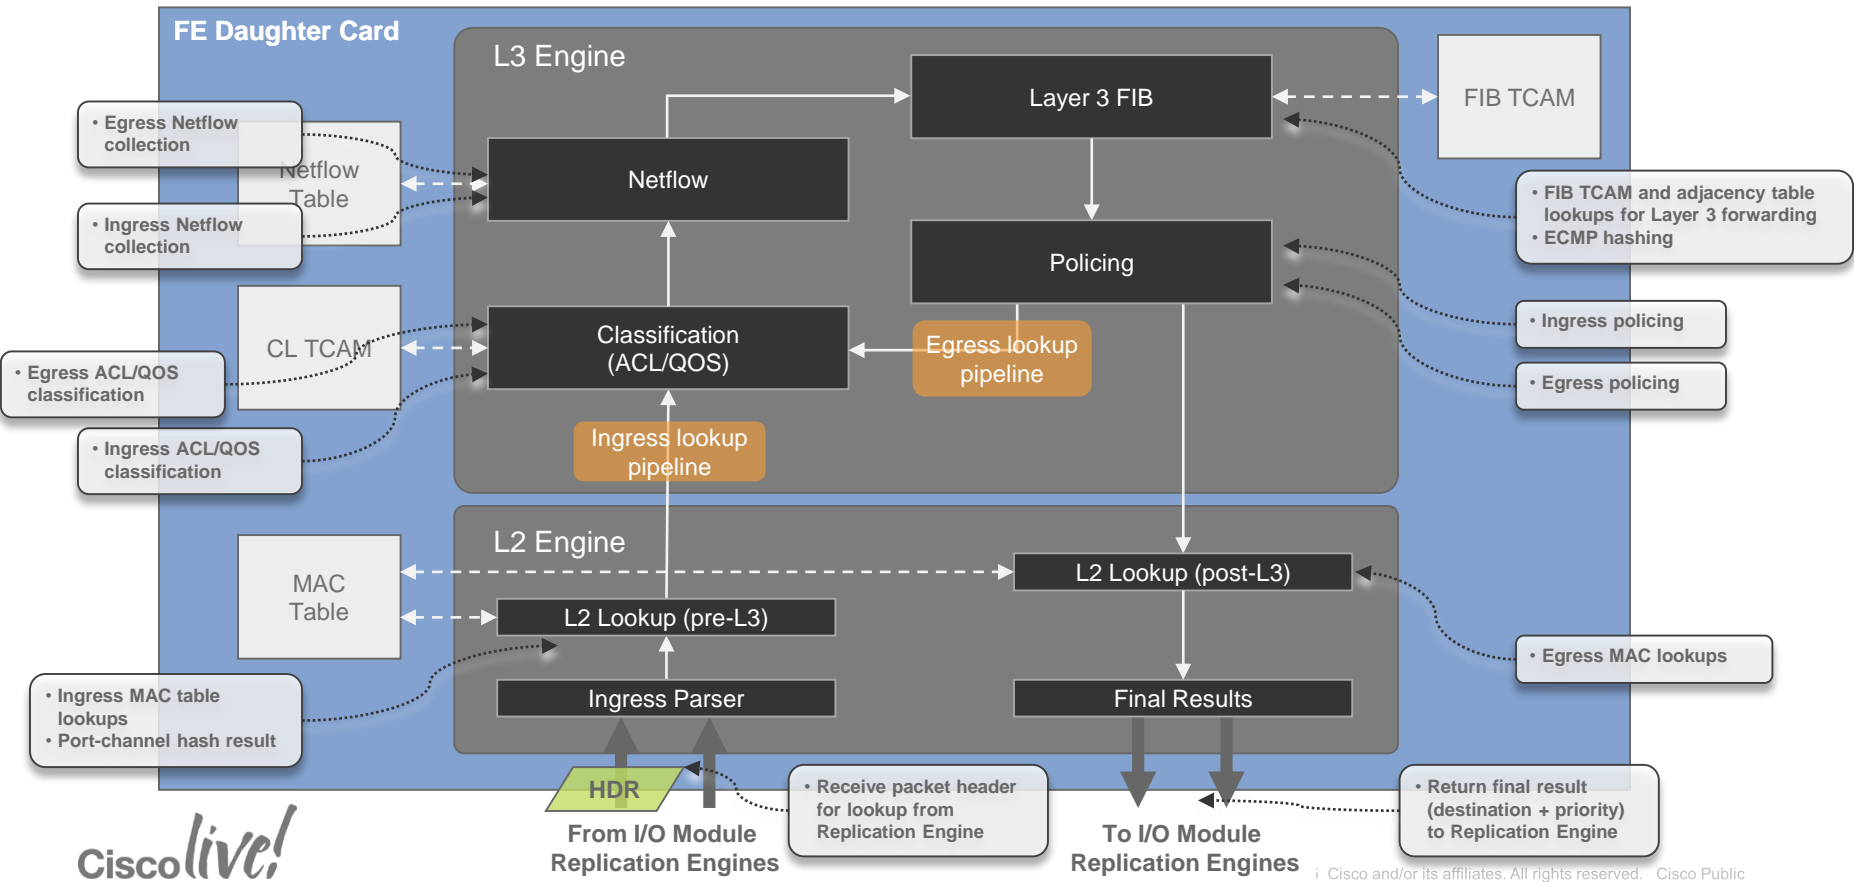


M2 Forwarding Engine Hardware

- Two hardware forwarding engines integrated on every M2 I/O module
- Layer 2 switching (with hardware MAC learning)
- Layer 3 IPv4/IPv6 unicast and multicast
- MPLS/VPLS/EoMPLS
- OTV / GRE
- RACL/VACL/PACL
- QOS remarking and policing policies
- Ingress and egress Netflow (full and sampled)

Hardware Table	M-Series Modules without Scale License	M-Series Modules with Scale License
MAC Address Table	128K	128K
FIB TCAM	128K IPv4 / 64K IPv6	900K IPv4 / 350K IPv6
Classification TCAM (ACL/QOS)	64K	128K
Netflow Table	1M	1M

M-Series Forwarding Engine Architecture



F2E Forwarding Engine Hardware

- 4 x 10G SOC with decision engine
- Layer 2 switching (with hardware MAC learning)
- Layer 3 IPv4/ IPv6 unicast and multicast
- FabricPath forwarding
- RACL/VACL/PACL
- QOS remarking and policing policies
- Ingress sampled Netflow

Hardware Table	F2E Capacity
MAC Address Table	16K
FIB TCAM	32K IPv4/16K IPv6
Classification TCAM (ACL/QOS)	16K

F3 Forwarding Engine Hardware

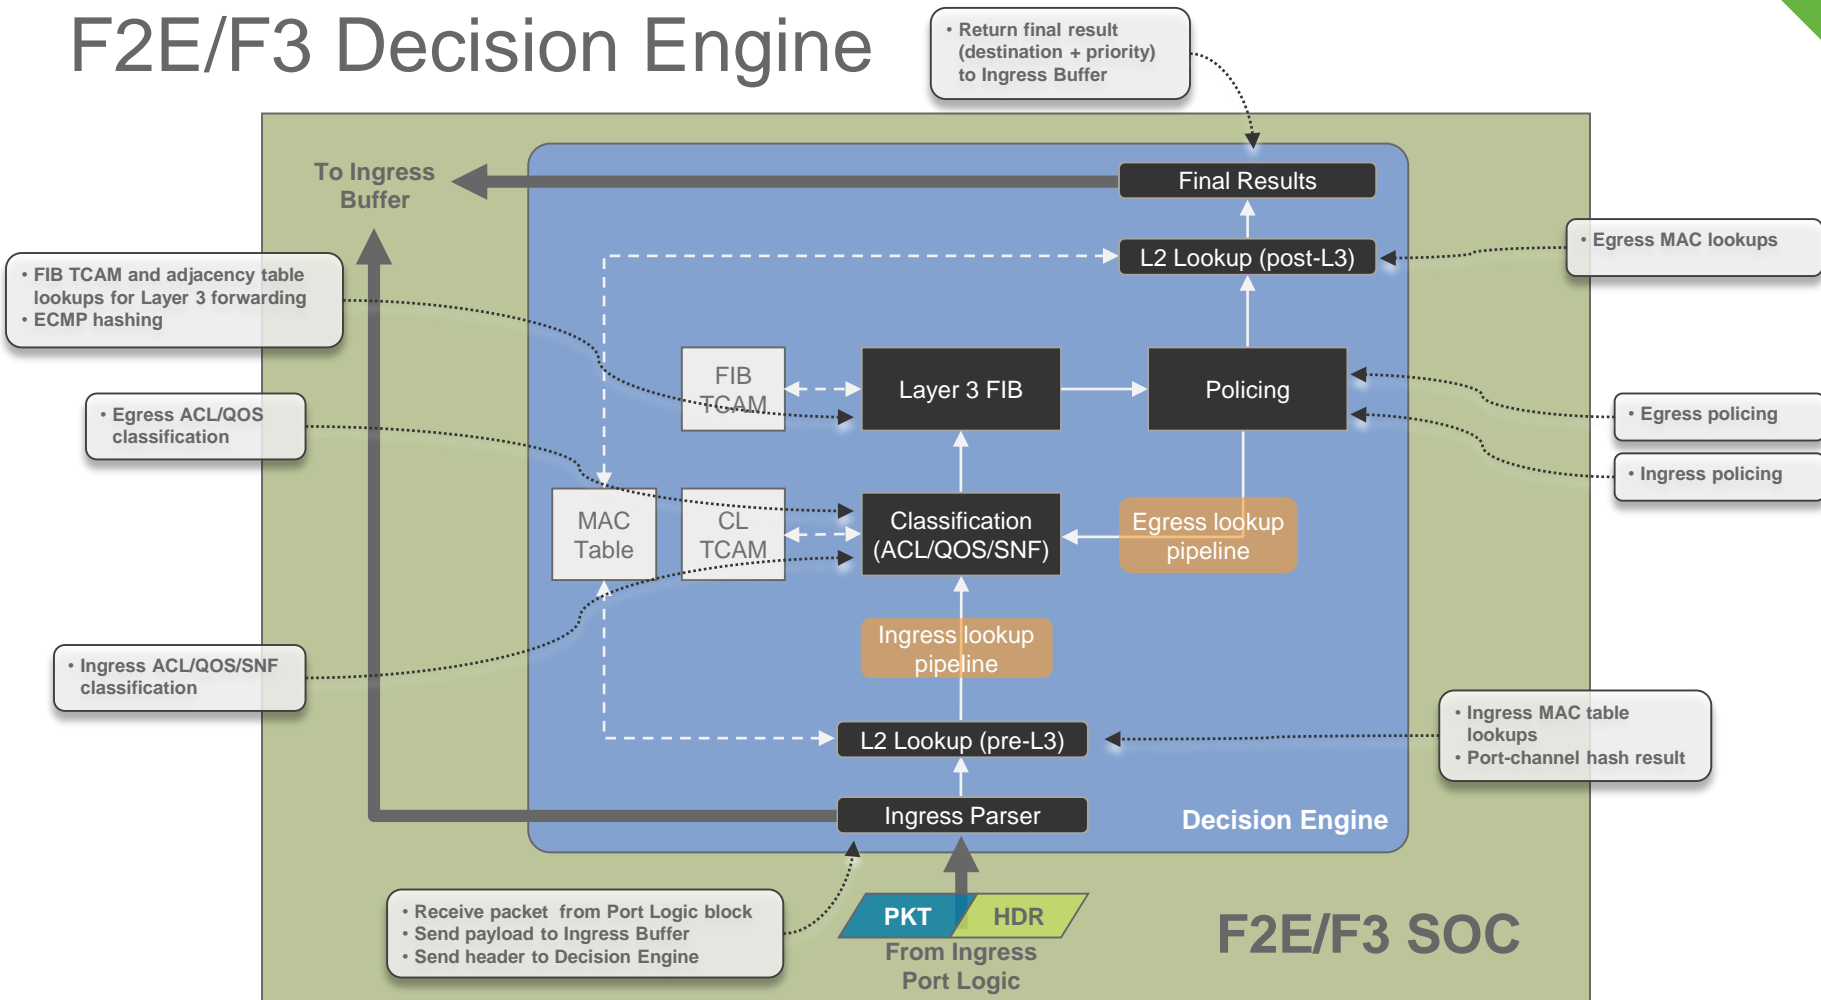
- 8 x 10G, 2 x 40G, or 1 x 100G SOC with decision engine
- Layer 2 switching (with hardware MAC learning)
- Layer 3 IPv4/ IPv6 unicast and multicast
- FabricPath forwarding

- RACL/VACL/PACL
- QOS remarking and policing policies
- Ingress sampled Netflow

- MPLS/VPLS/EoMPLS
- OTV / GRE tunnels
- LISP
- VXLAN

Hardware Table	F3 Capacity
MAC Address Table	64K
FIB TCAM	64K IPv4/32K IPv6
Classification TCAM (ACL/QOS)	16K

F2E/F3 Decision Engine



Layer 2 Hardware Forwarding

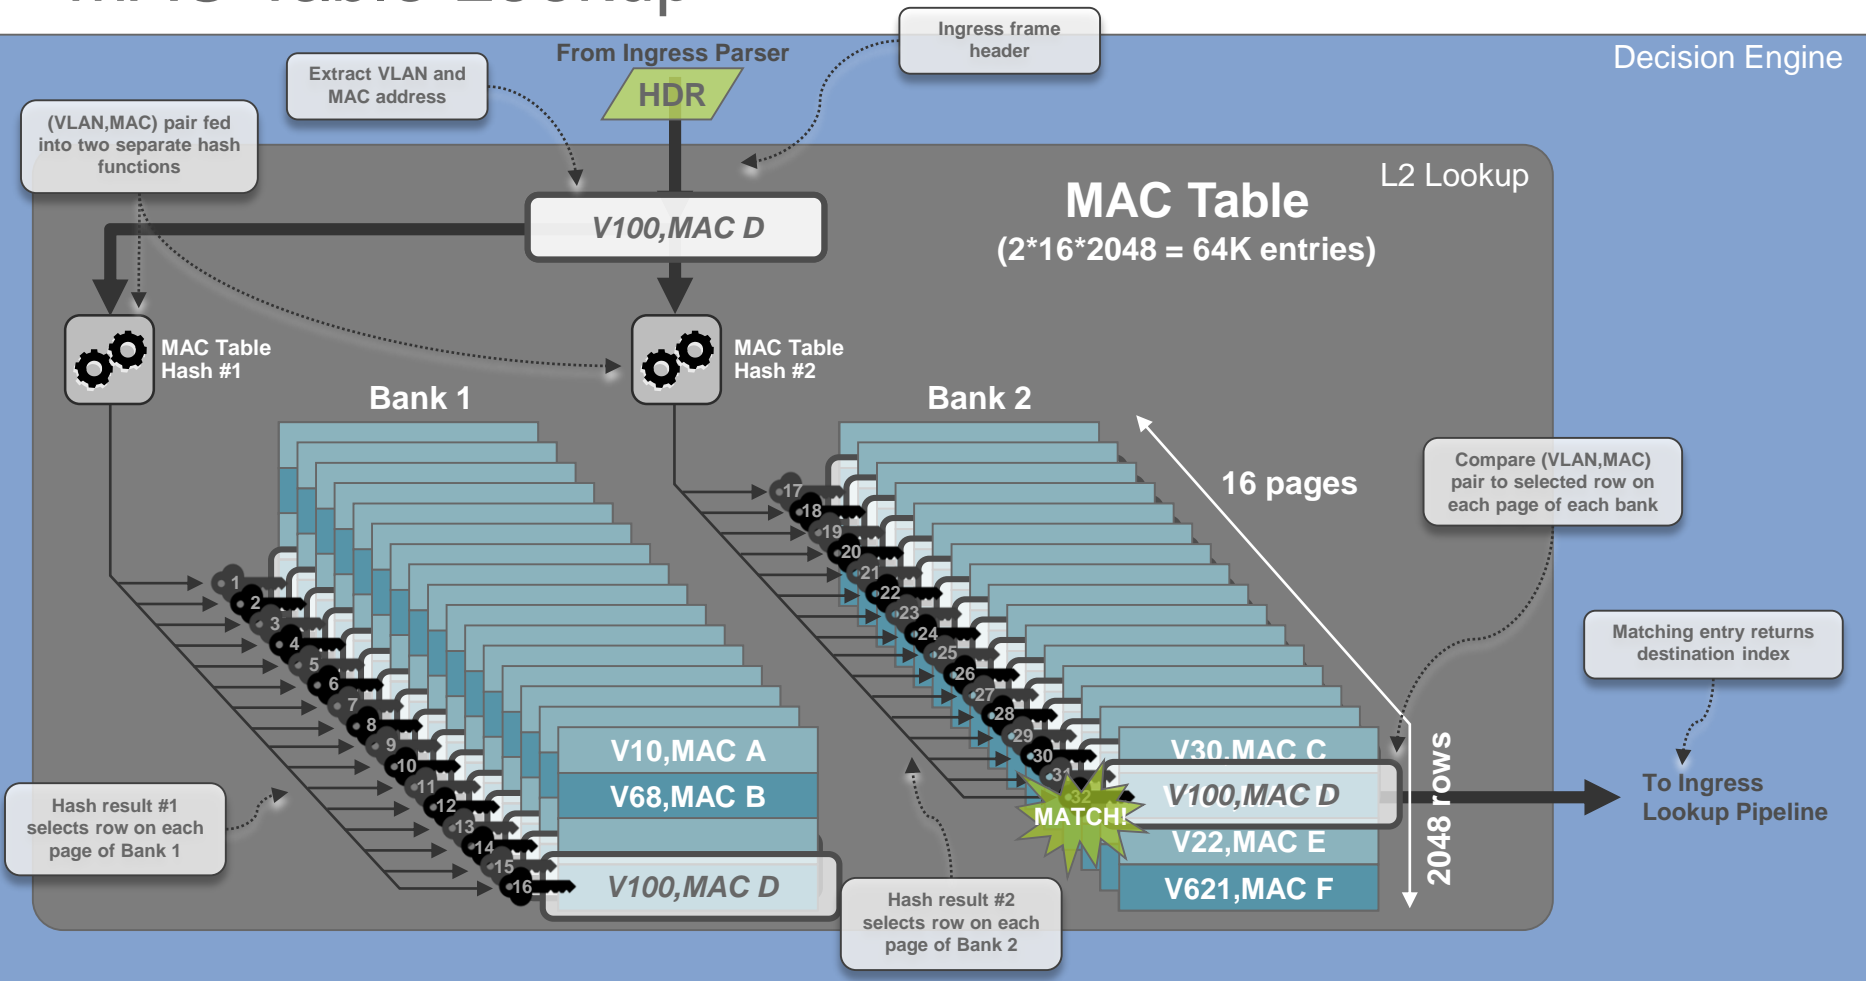
- Layer 2 forwarding driven by MAC table lookups
- Source and destination MAC lookups performed for each frame, based on {VLAN,MAC} pairs
- Source MAC lookup drives new learns and refreshes aging timers
- Destination MAC lookup returns destination port



MAC Table Lookup

F3 module used as example

Decision Engine




MAC Table Details

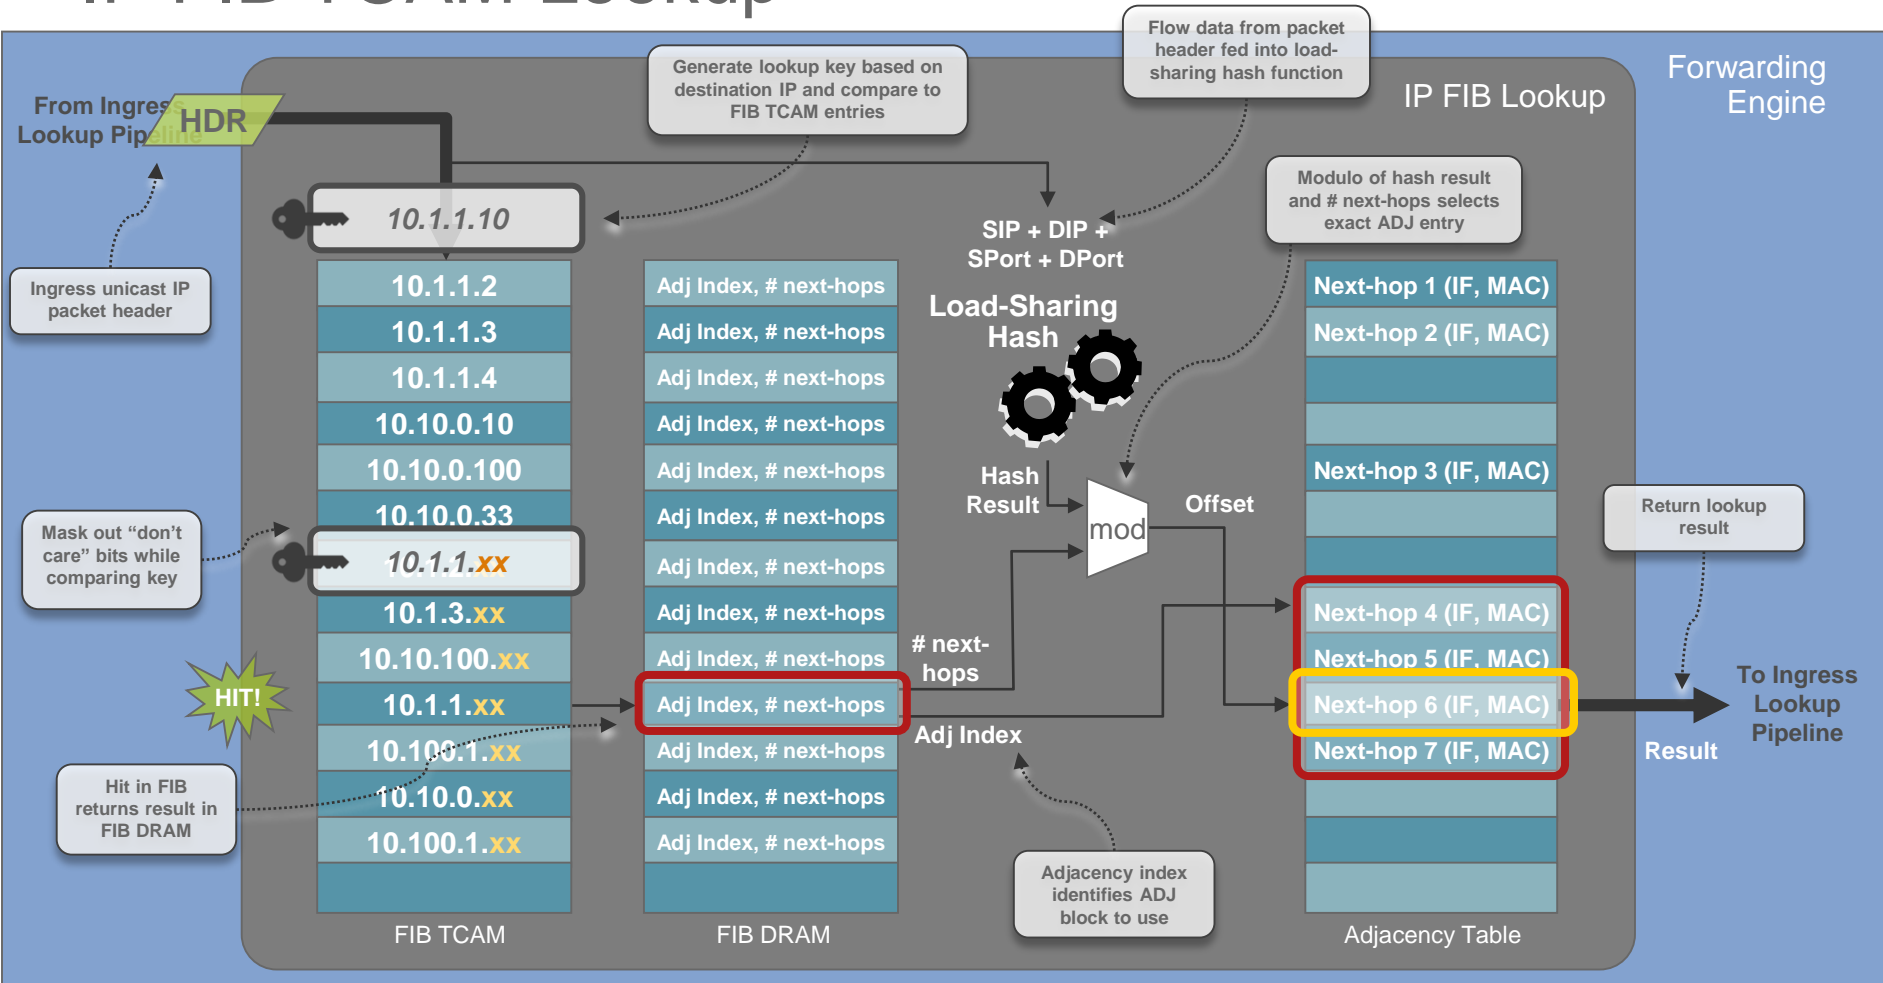
- MAC table is a hash table
- Effective capacity of any hash table dependent on “quality” of input to hash
- Hash collisions can occur and 100% utilization may not be possible
- Hash collisions more likely when many MACs are similar, e.g.:
 - Performance testers (all ports sending incrementing MAC addresses)
 - Virtualized environments (new VMs deployed en masse with incrementing MACs)
 - Homogenous server environments (all hosts using the same NICs from the same vendor and all purchased at roughly the same time)
 - IP multicast deployments (many L2 group MACs which start with 01005E and likely to have sequential destination IP group addresses)

Layer 3 Hardware Forwarding

- Layer 3 forwarding driven by FIB table lookups
- Forwarding tables built by control plane on supervisor engine
 - OSPF, EIGRP, IS-IS, BGP, statics, etc.
- Tables downloaded to forwarding engine hardware for data-plane forwarding
- FIB TCAM lookup based on longest-match destination IP prefix lookup
- FIB match returns rewrite (next-hop) information in adjacency table



IP FIB TCAM Lookup



Classification Lookups



- Matching packets
 - Layer 2, Layer 3, and/or Layer 4 information
- Used to decide whether to apply a particular policy to a packet
 - Enforce security, QOS, or other policies
- Some examples:
 - Match TCP/UDP source/destination port numbers to enforce security policy
 - Match source IP addresses to apply policy-based routing (PBR)
 - Match 5-tuple to apply QOS marking policy
 - Match protocol-type to apply Control Plane Policing (CoPP)
 - etc.

CL TCAM Lookup – ACL

Security ACL

```
ip access-list example
  permit ip any host 10.1.2.100
  deny ip any host 10.1.68.44
  deny ip any host 10.33.2.25
  deny ip any host 10.24.77.7
  deny ip any host 10.24.77.1
  permit tcp any any eq 22
  deny tcp any any eq 23
  deny udp any any eq 514
  permit tcp any any eq 80
  permit udp any any eq 161
```

Decision Engine

Classification

From Ingress/Egress
Lookup Pipeline

Packet header

HDR

Generate lookup key
based on packet fields
and compare to CL
TCAM entries

Mask out "don't
care" bits while
comparing key

ACEs from
security ACL
(x = don't care)

10.1.1.1 / 10.2.2.2 / tcp / 33992 / 80

XXXXXXXX / 10.2.2.2 / xx / xxx / xxx

XXXXXXXX / 10.1.68.44 / xx / xxx / xxx

XXXXXXXX / 10.33.2.25 / xx / xxx / xxx

XXXXXXXX / 10.24.77.7 / xx / xxx / xxx

XXXXXXXX / 10.24.77.1 / xx / xxx / xxx

XXXXXXXX / XXXXXXXX / tcp / xxx / 80

XXXXXXXX / XXXXXXXX / tcp / xxx / 23

XXXXXXXX / XXXXXXXX / udp / xxx / 514

XXXXXXXX / XXXXXXXX / tcp / xxx / 80

XXXXXXXX / XXXXXXXX / udp / xxx / 161

Hit in CL TCAM
returns result in
CL SRAM

Permit

Deny

Deny

Deny

Deny

Permit

Deny

Deny

Permit

Permit

Actions from
security ACL

Return lookup result,
affecting final packet
handling (forward or drop)

To Ingress/
Egress Lookup
Pipeline

CL TCAM

CL SRAM

Fields to match:
src IP | dst IP | protocol | src port | dst port

HIT!

CL TCAM Lookup – QOS

QOS Policy ACLs

```
ip access-list police
 permit ip any 10.3.3.0/24
 permit ip any 10.4.12.0/24
 permit ip any 10.10.1.0/24
 permit ip any 10.20.2.0/24
ip access-list remark-dscp-32
 permit udp 10.0.1.0/24 any
 permit udp 10.1.1.0/24 any
ip access-list remark-dscp-40
 permit tcp 10.0.1.0/24 any
 permit tcp 10.1.1.0/24 any
ip access-list remark-prec-3
 permit tcp any 10.2.3.0/24 eq 23
 permit tcp any 10.5.5.0/24 eq 23
```

Decision Engine

Classification

From Ingress/Egress
Lookup Pipeline

Packet header

HDR

Generate lookup key
based on packet fields
and compare to CL
TCAM entries

Mask out "don't
care" bits while
comparing key

ACEs from QOS
policy ACLs
(x = don't care)

10.1.1.1 / 10.2.2.2 / tcp / 33992 / 80

xxxxxxx / 10.2.2.xx / xx / xxx / xx

xxxxxxx / 10.4.12.xx / xx / xxx / xxx

xxxxxxx / 10.10.1.xx / xx / xxx / xxx

xxxxxxx / 10.20.2.xx / xx / xxx / xxx

10.1.1.xx / xxxxxxx / tcp / xxx / xxx

10.1.1.xx / xxxxxxx / udp / xxx / xxx

10.0.1.xx / xxxxxxx / tcp / xxx / xxx

10.1.1.xx / xxxxxxx / tcp / xxx / xxx

xxxxxxx / 10.2.3.xx / tcp / xxx / 23

xxxxxxx / 10.5.5.xx / tcp / xxx / 23

Hit in CL TCAM
returns result in
CL SRAM

Policer ID 1

Policer ID 1

Policer ID 1

Policer ID 1

Remark DSCP 32

Remark DSCP 32

Remark DSCP 40

Remark DSCP 40

Remark IP Prec 3

Remark IP Prec 3

Actions from
QOS policy

To Ingress/
Egress Lookup
Pipeline

Return lookup result,
affecting final packet
handling (police or remark)

CL TCAM

CL SRAM

Fields to match:
src IP | dst IP | protocol | src port | dst port

HIT!

Full and Sampled Netflow

Netflow collects flow data for export to collector(s)

Full Netflow: Accounts for every packet of every flow on interface, up to capacity of hardware Netflow table

- Available on **M-Series** modules only

Sampled Netflow: Accounts for M in N packets on interface using random packet-based sampling

- **M2**: Accounts sampled flows, up to capacity of hardware Netflow table
- **F2E/F3**: Accounts hardware sampled flows in software

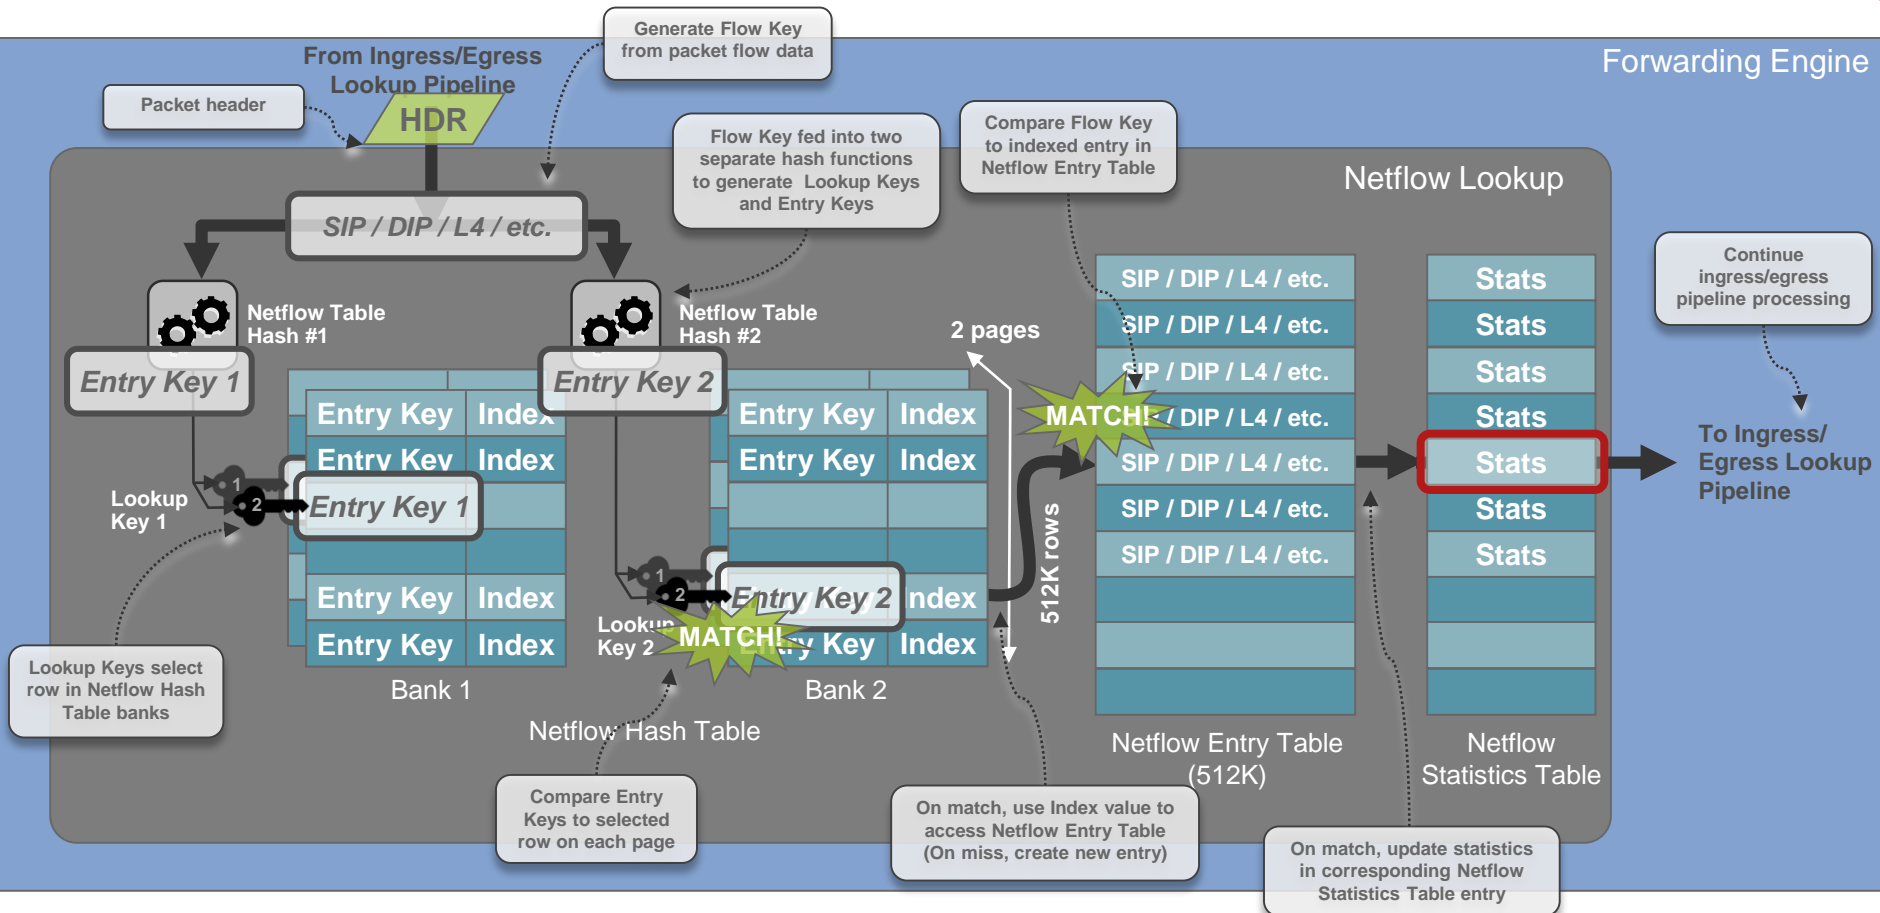
M2 Netflow Table

Netflow “Table” actually consists of three hardware components in M2 forwarding engine:

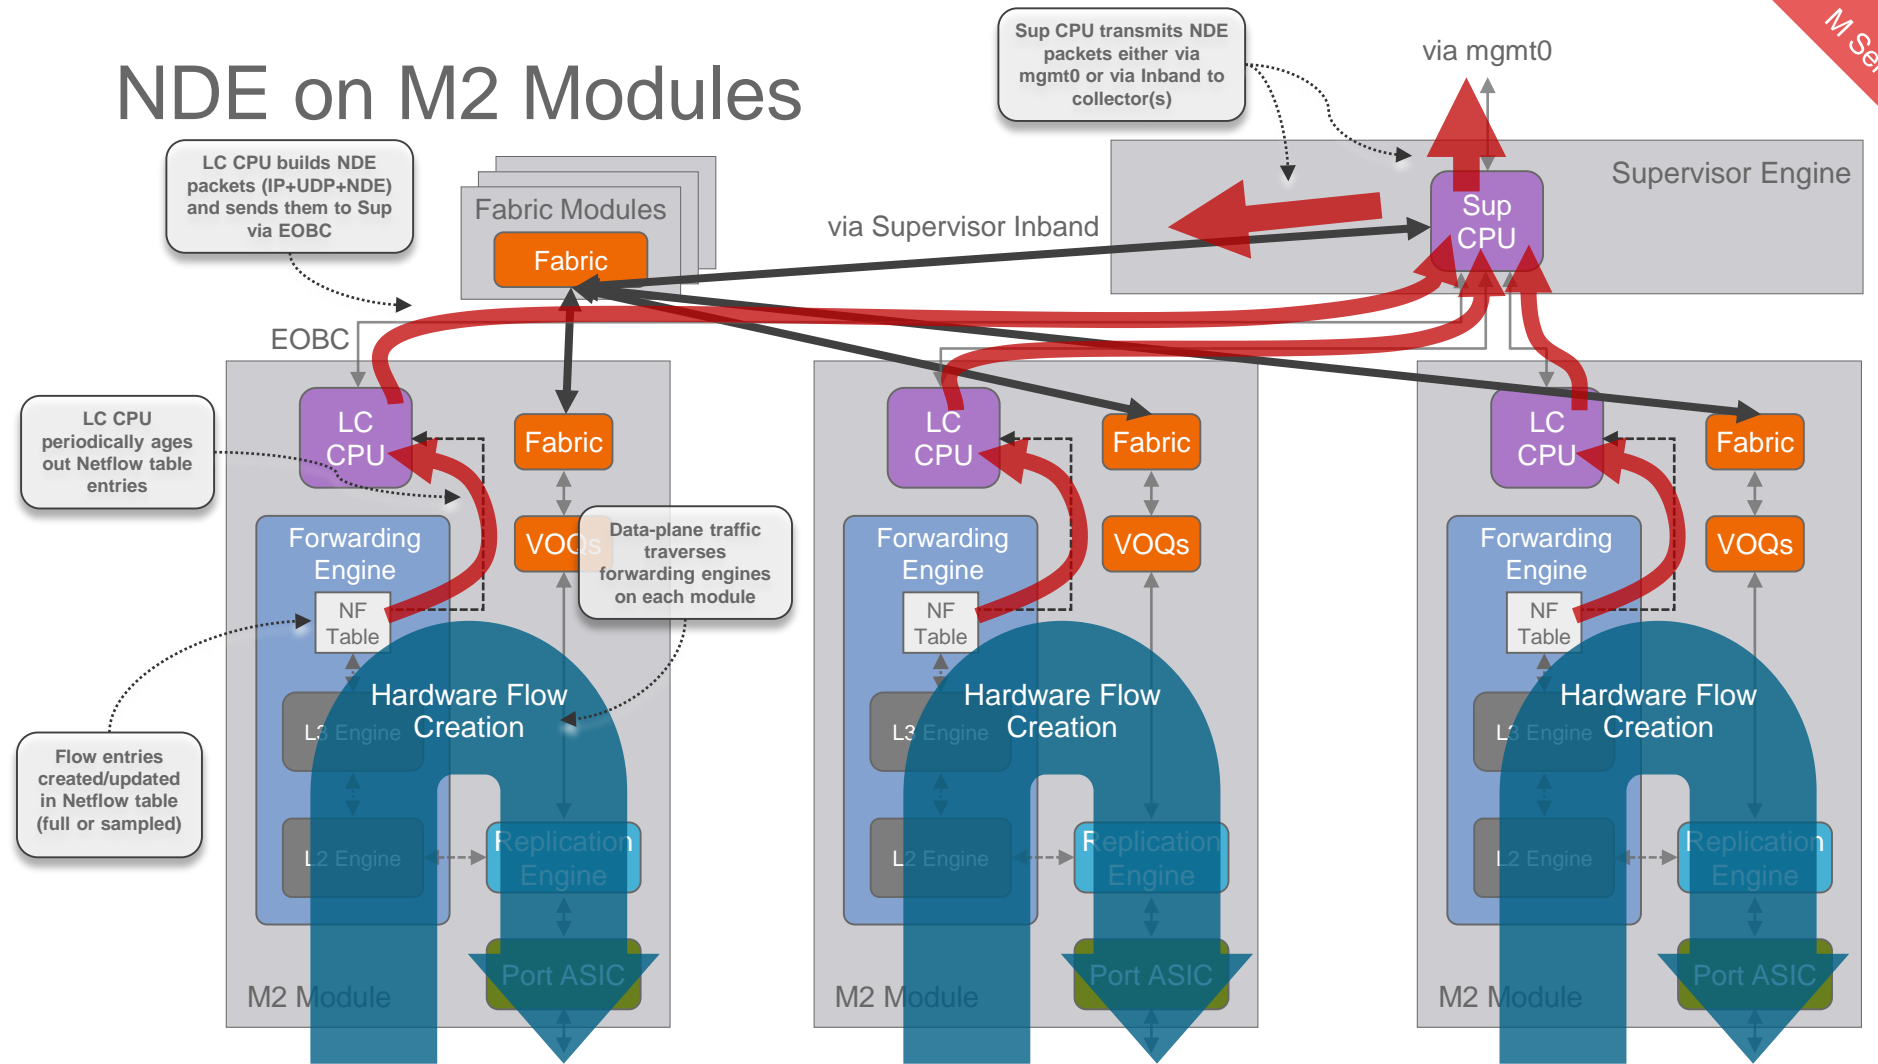
- **Netflow Hash Table:** Contains Netflow Entry Keys and corresponding indexes to Netflow Entry Table (speeds lookups and minimizes hash collisions)
- **Netflow Entry Table:** Contains actual Netflow flow data
- **Netflow Statistics Table:** Contains statistics for corresponding flow entries

M2 Netflow Lookup

Forwarding Engine



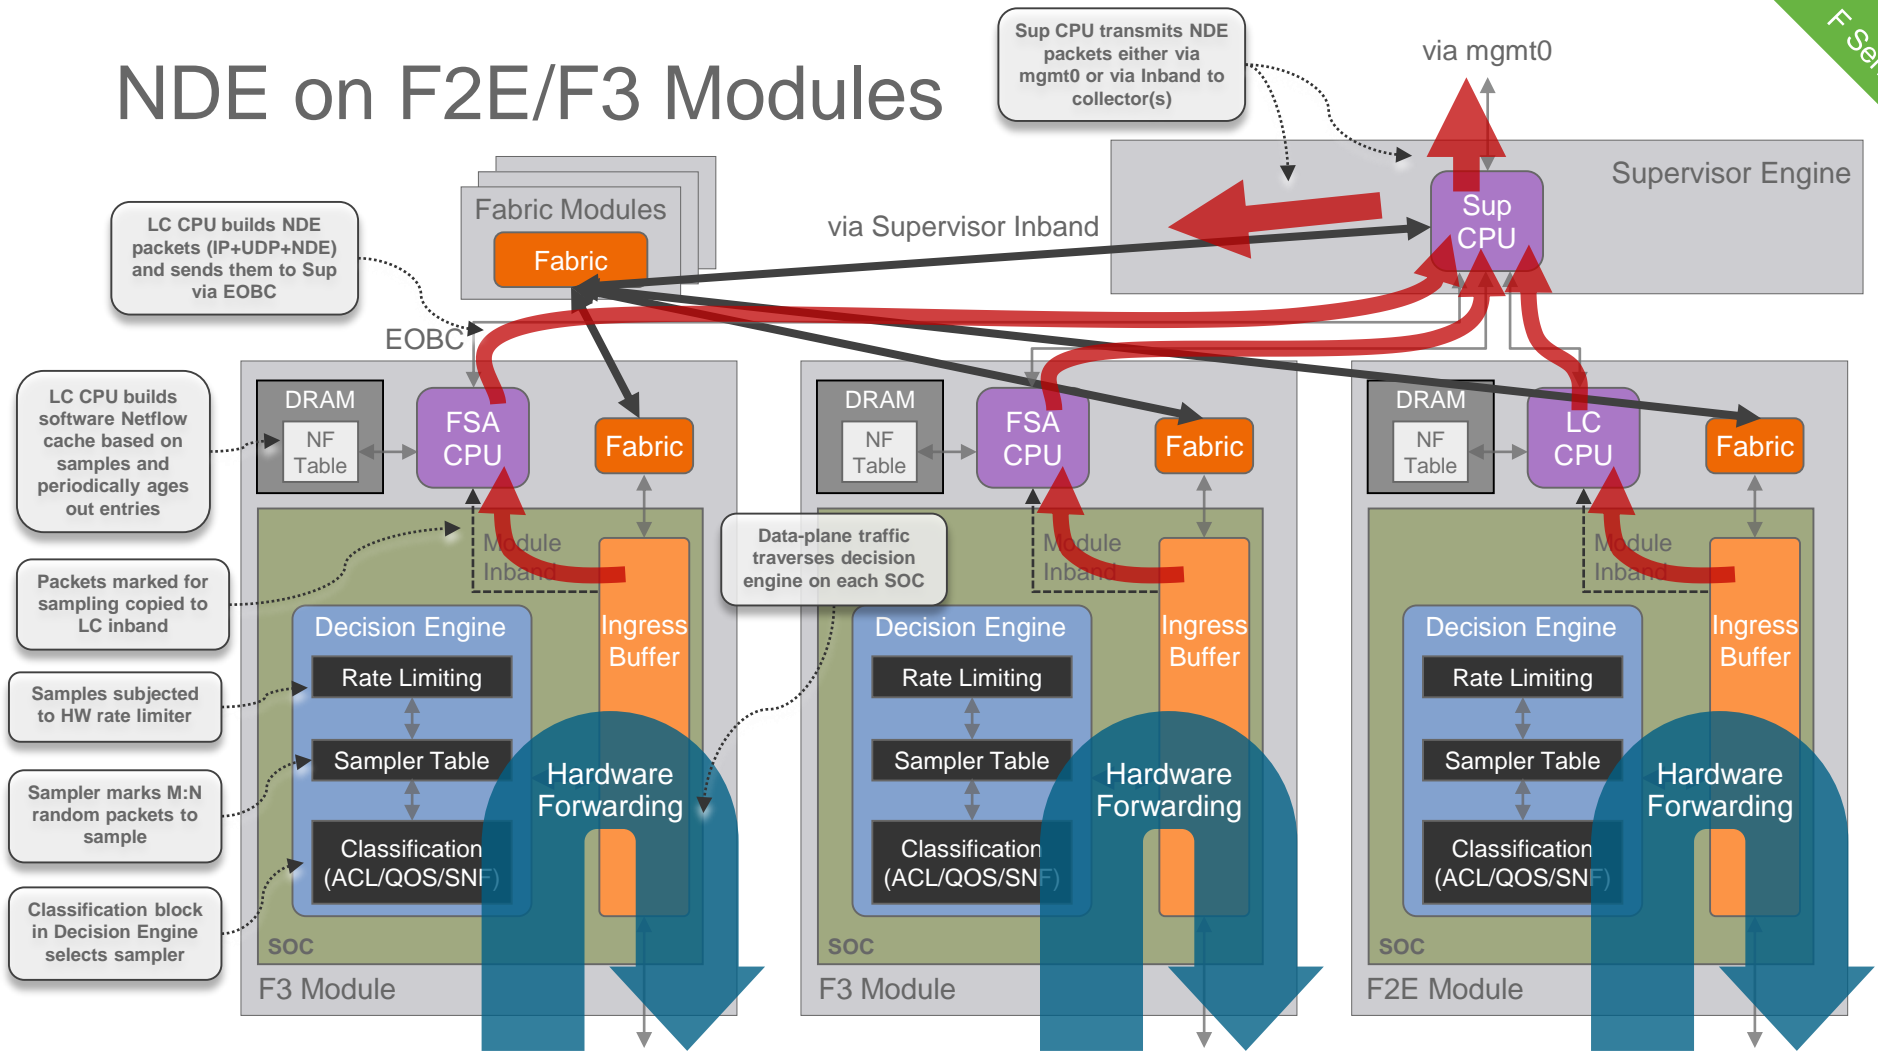
NDE on M2 Modules



F2E/F3 Sampled Netflow

- Hardware-based sampling with software-based Netflow cache
- Classification lookup selects Netflow sampler-table entry
- Sampler table defines which sampler to use (defines M:N)
- Copy of randomly sampled packets sent to LC CPU/FSA via module inband
 - Sampled copies “sliced” to reduce bandwidth consumption
- Sampled copies rate-limited to protect LC CPU
 - F2E to ~1000pps per module
 - F3 to ~3000pps per module
 - Roadmap for ~50Kpps per F3 module using FSA

NDE on F2E/F3 Modules



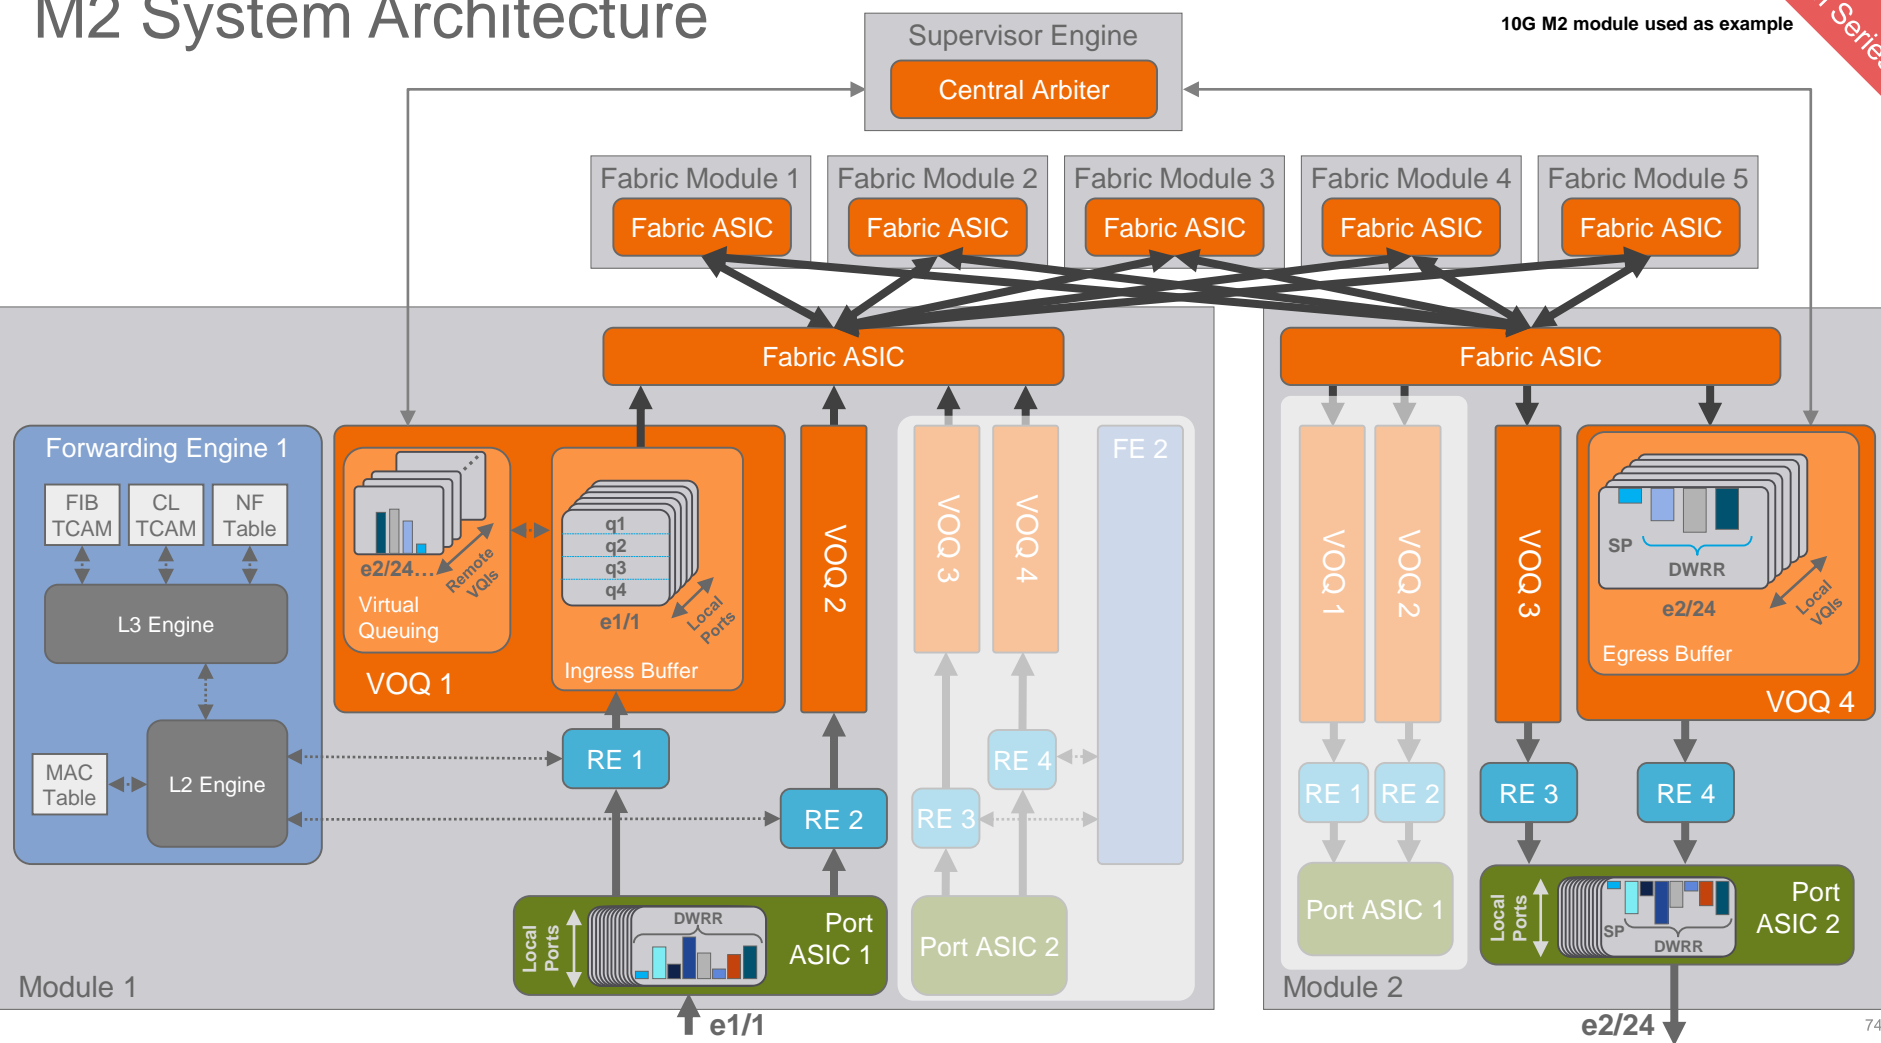
Agenda

- Introduction to Nexus 7000 / Nexus 7700
- Chassis Architecture
- Supervisor Engine and I/O Module Architecture
- Fabric Architecture
- Hardware Forwarding
- Packet Walks
- Conclusion

M2 System Architecture

10G M2 module used as example

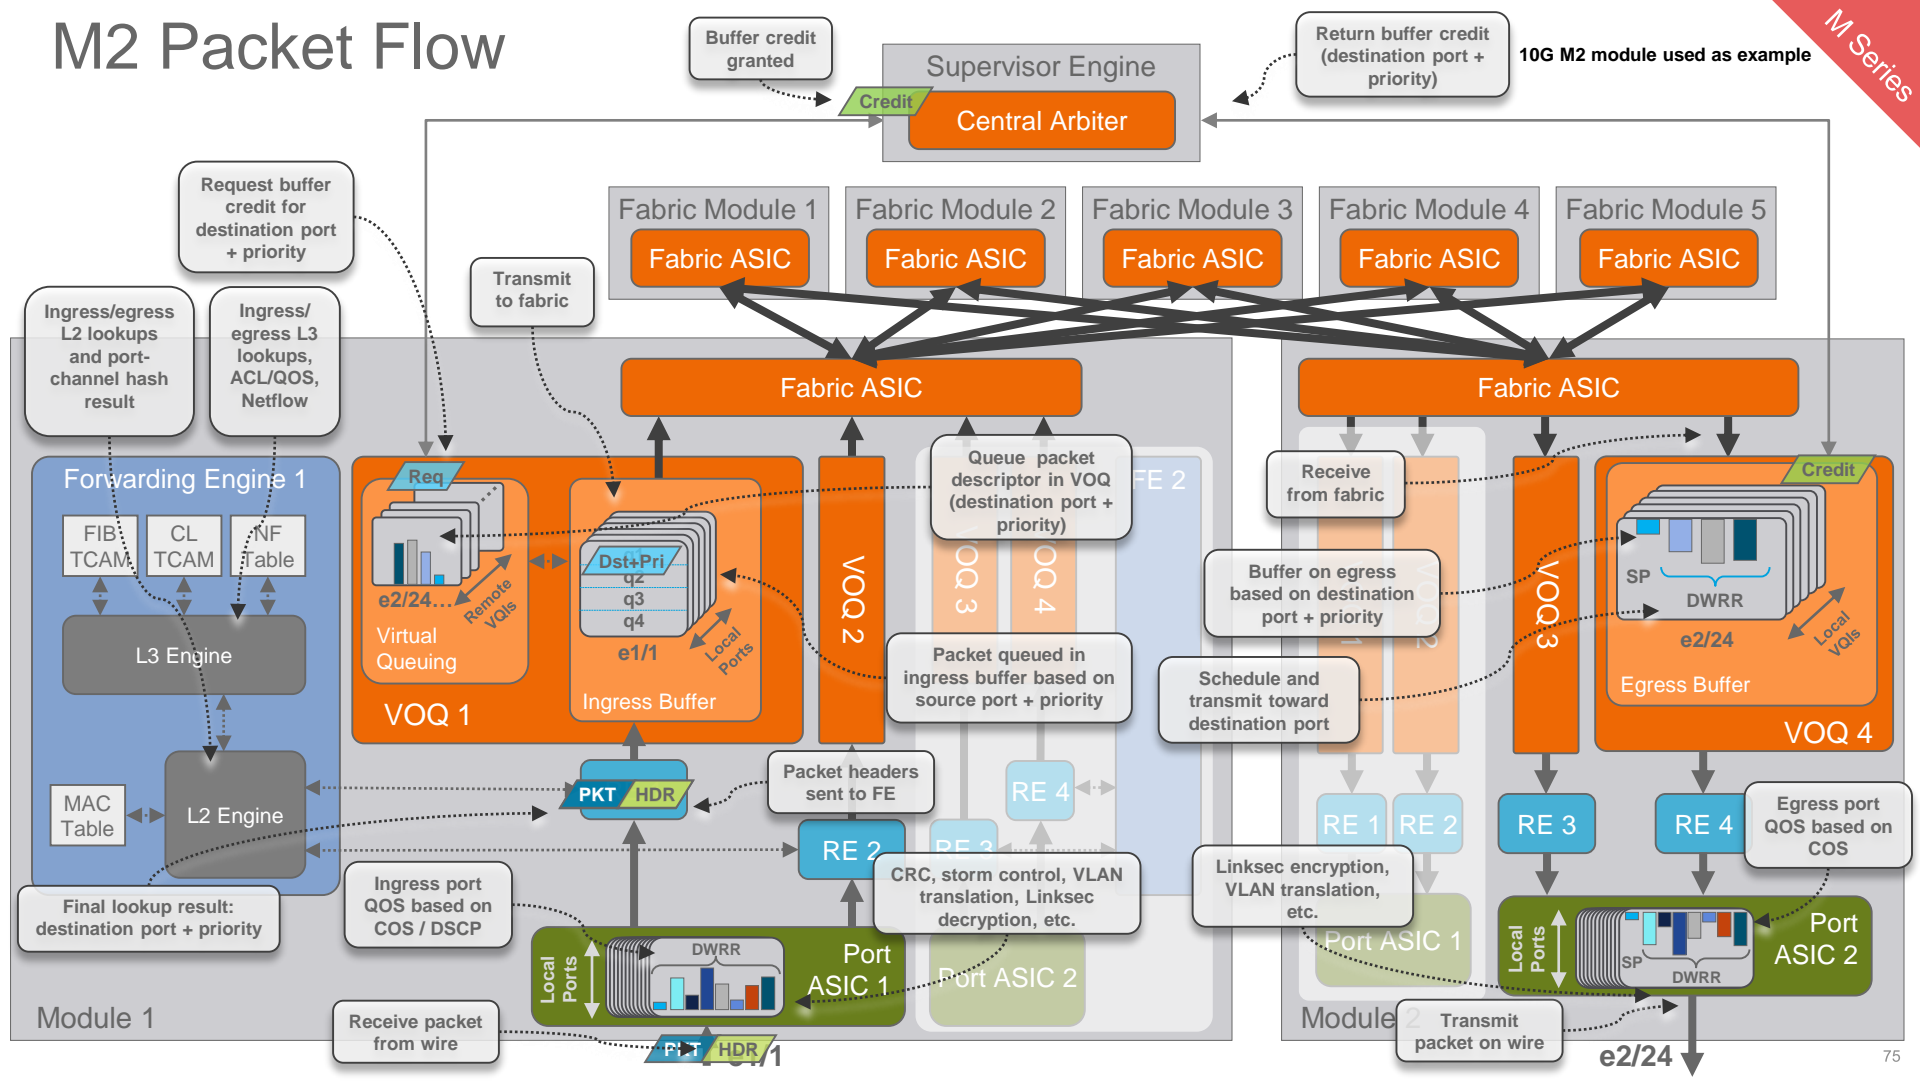
M Series



M2 Packet Flow

M Series

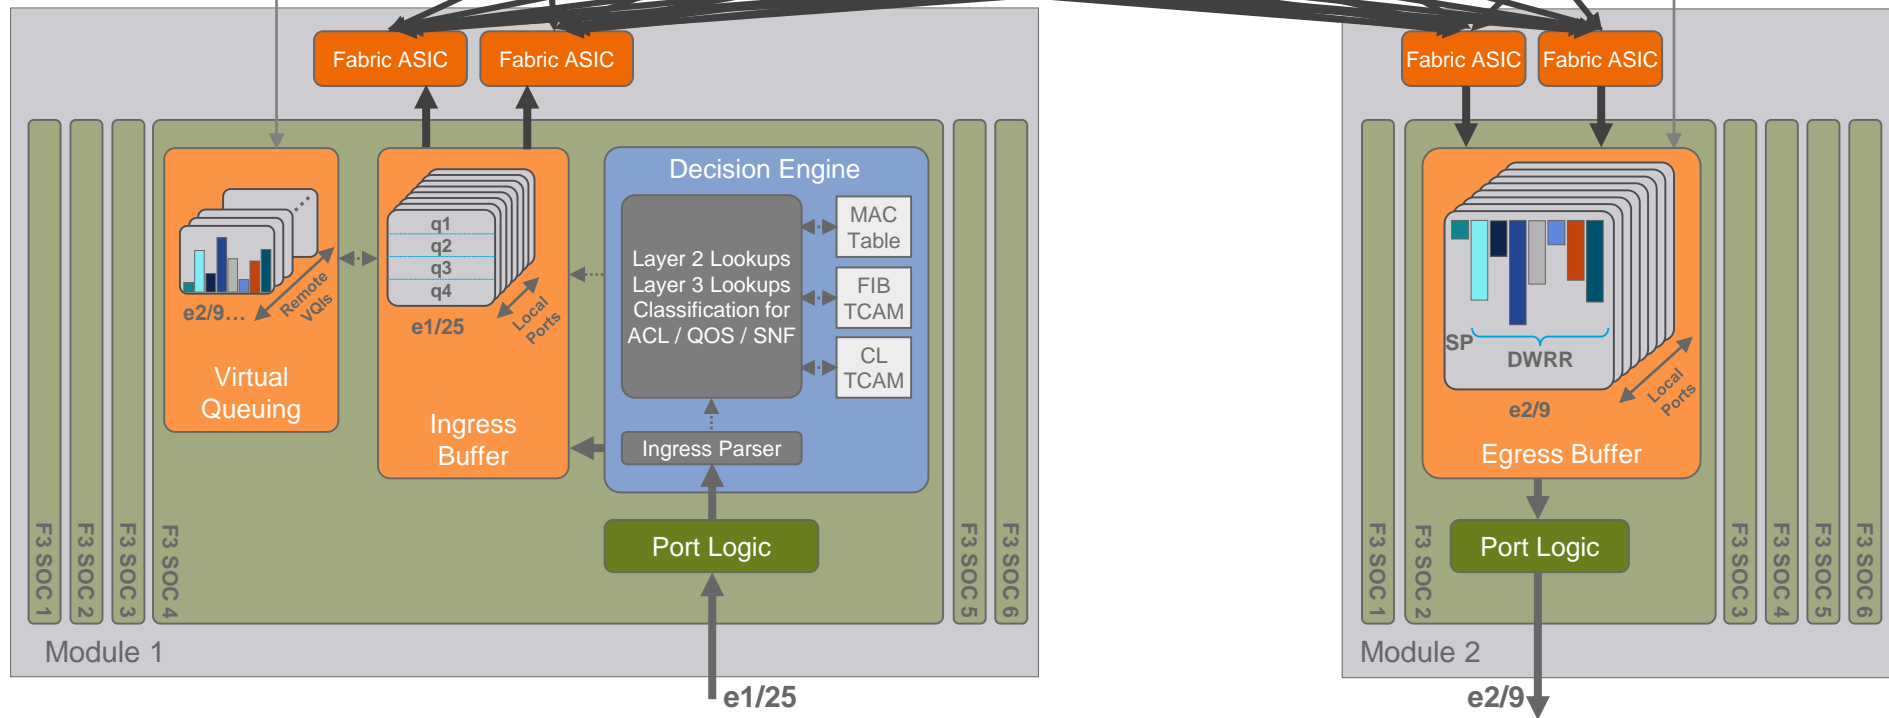
10G M2 module used as example



F2E/F3 System Architecture

N7700 10G F3 module used as example

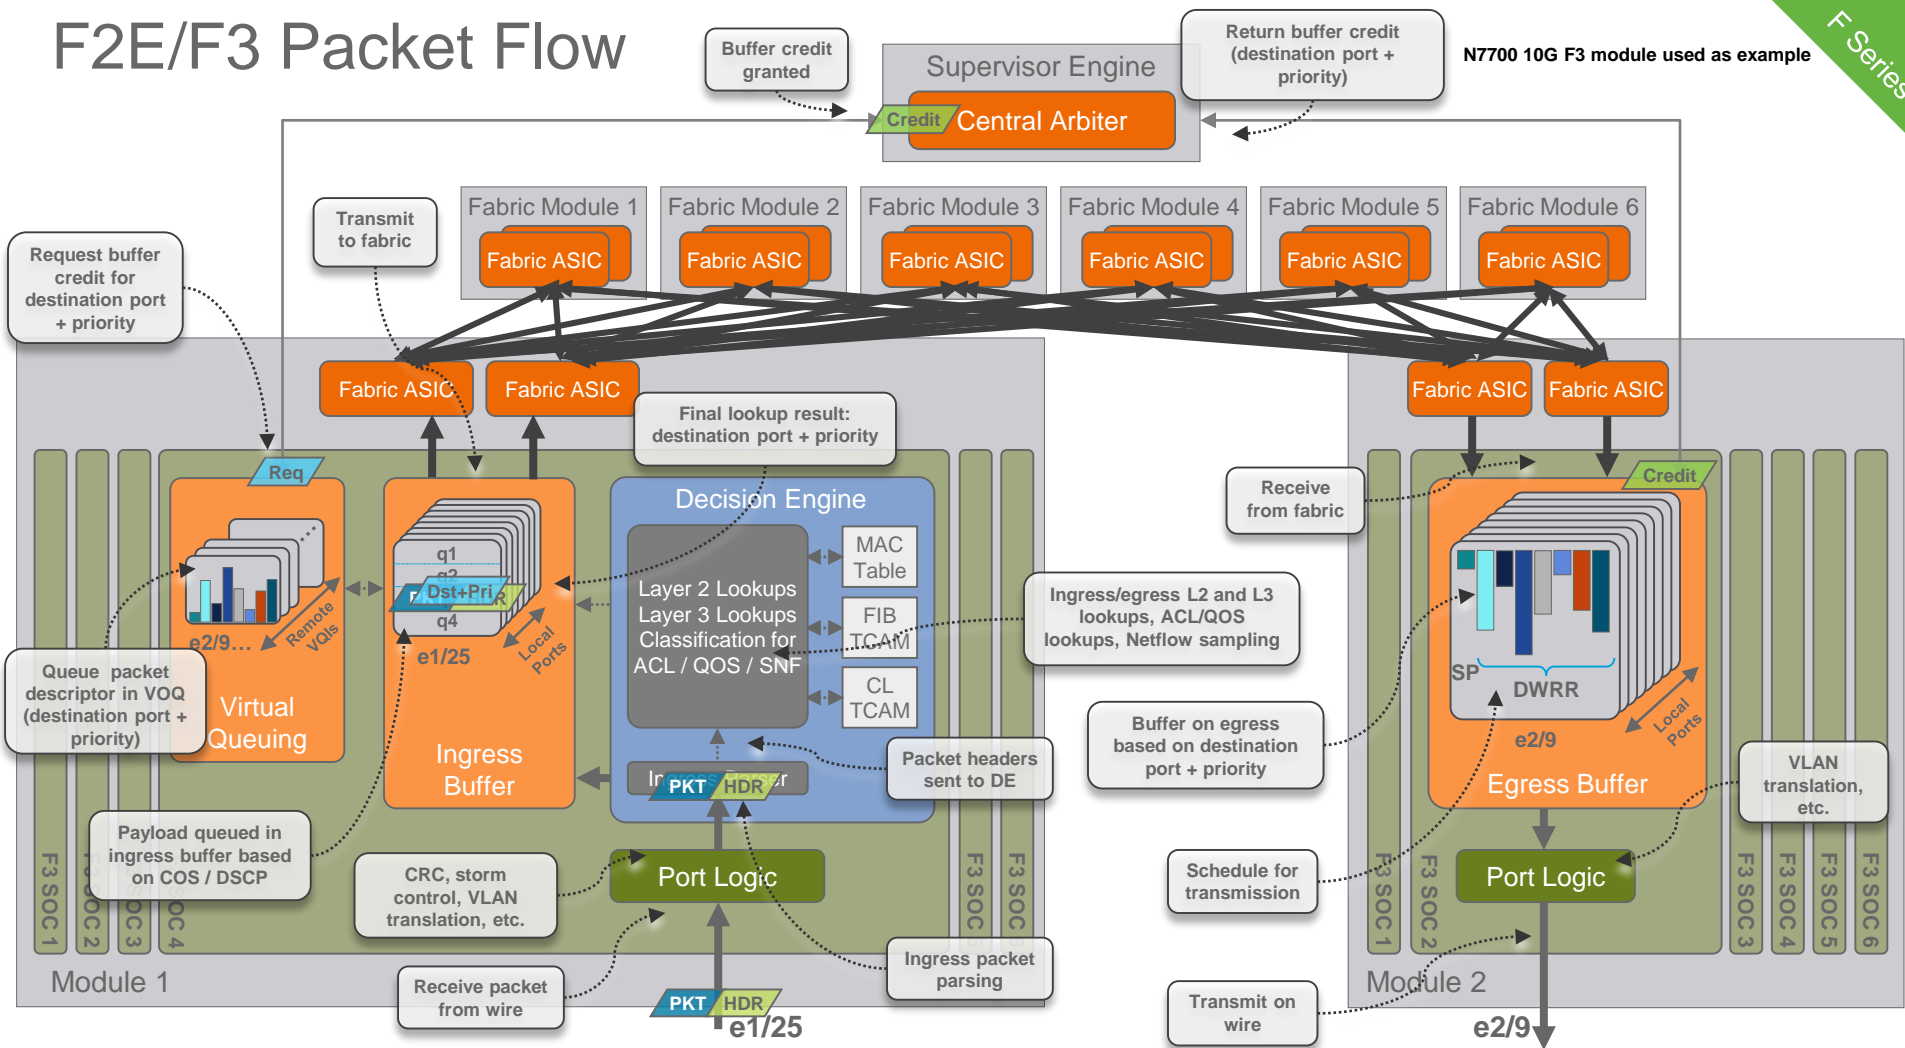
F Series



F2E/F3 Packet Flow

F Series

N7700 10G F3 module used as example



Agenda

- Introduction to Nexus 7000 / Nexus 7700
- Chassis Architecture
- Supervisor Engine and I/O Module Architecture
- Fabric Architecture
- Hardware Forwarding
- Packet Walks
- Conclusion

Conclusion

- You should now have a thorough understanding of the Nexus 7000 / Nexus 7700 switching architecture – including chassis, supervisors, I/O modules, and fabrics, as well as forwarding lookups and complete system-level packet flows...
- **Any questions?**



Reference: Acronym Decoder

- ACL—Access Control List
- ADJ—Adjacency
- ASIC—Application Specific Integrated Circuit
- CFP—C Formfactor Pluggable
- CoPP—Control Plane Policing
- COS—Class of Service
- DE—Decision Engine
- DSCP—Differentiated Services Code Point
- DWRR—Deficit Weighted Round Robin
- ECMP—Equal Cost Multi Path
- EOBC—Ethernet Out-of-Band Channel
- FCoE—Fiber Channel over Ethernet
- FE—Forwarding Engine
- FEX—Fabric Extender (Nexus 2000 family)
- FIB—Forwarding Information Base
- GRE—Generic Route Encapsulation
- MPLS—Multiprotocol Label Switching
- NDE—Netflow Data Export
- OTV—Overlay Transport Virtualization
- PACL—Port ACL
- PBR—Policy-Based Routing
- QOS—Quality of Service
- QSFP+—40G Quad Small-Formfactor Pluggable
- RACL—Router ACL
- RE—Replication Engine
- RPF—Reverse Path Forwarding
- RU—Rack Unit
- SFP+—10G Small-Formfactor Pluggable
- SNF—Sampled Netflow
- SOC—System-on-chip/switch-on-chip
- SP—Strict priority (queue)
- TCAM—Ternary CAM
- VACL—VLAN ACL
- VDC—Virtual Device Context
- VOQ—Virtual Output Queuing
- VQI—Virtual Queuing Index
- VXLAN—Virtual Extensible LAN

Participate in the “My Favorite Speaker” Contest

Promote Your Favorite Speaker and You Could Be a Winner

- Promote your favorite speaker through Twitter and you could win \$200 of Cisco Press products (@CiscoPress)
- Send a tweet and include
 - Your favorite speaker’s Twitter handle
 - Two hashtags: #CLUS #MyFavoriteSpeaker
- You can submit an entry for more than one of your “favorite” speakers
- Don’t forget to follow @CiscoLive and @CiscoPress
- View the official rules at <http://bit.ly/CLUSwin>

Complete Your Online Session Evaluation

- Give us your feedback to be entered into a Daily Survey Drawing. A daily winner will receive a \$750 Amazon gift card.
- Complete your session surveys though the Cisco Live mobile app or your computer on Cisco Live Connect.



Don't forget: Cisco Live sessions will be available for viewing on-demand after the event at [CiscoLive.com/Online](https://cislive.com/online)

Continue Your Education

- Demos in the Cisco campus
- Walk-in Self-Paced Labs
- Table Topics
- Meet the Engineer 1:1 meetings
- Related sessions



Thank you



TOMORROW starts here.