EX / QFX PRODUCT UPDATE

Juniper SE

2019.05

JUNIPER NETWORKS | Engineering Simplicity

AGENDA

- EX Series
- EX Architecture
- QFX Series
- QFX Software Architecture

EX SERIES









EX2300-C /EX2300	EX2300-MP	EX3400	EX4300	EX4300-MP	EX4600	EX4650	EX9200	EX9250
Access	**Multigigabit access**	Access	Access	**Multigigabit access**	Aggregation	Core & Aggregation	Core & Aggregation	Compact Core & aggregation
Up to 176Gbps	**Up to 264Gbps**	**Up to 336Gbps**	**Up to 496Gbps**	**Up to 960Gbps**	**Up to 1.44Tbps**	**Up to 2Tbps**	**Up to 13.2Tbps**	**Up to 4.8 Tbps**
12-24-48 x GbE 2-4 x 10GbE SFP+	24-48 x 1GbE 8-16 x 2.5GbE 4-6 x 10GbE SFP+	24-48 x GbE 4 x 10GbE SFP+ 2 x 40GbE QSFP+	24-48 x GbE 4 x 10GbE 4 x 40GbE QSFP+	24 x 10/100/1000 24 x 1/2.5/5/10GbE 2 x 40GbE 1 x 100GbE 4 x 10/40GbE	72 x 10GbE SFP+ 12 x 40GbE QSFP+	48 x 1/10/25GbE 8 x 40/100GbE	Up to 320x10GbE Up to 60x40GbE Up to 20x100GbE	Up to 144x10GbE Up to 36x40GbE Up to 24x100GbE
PoE / PoE+	PoE / PoE+	PoE / PoE+	PoE / PoE+	PoE / PoE+ / PoE++	N/A			
Virtual Chassis (VC) Capable					**VC & MC-LAG**	**EVPN/VXLAN MC-LAG**		
Junos Fusion Satelite Device					**N/A**		**Junos Fusion Aggregation Device**	

EX SERIES









EX2300-C /EX2300	EX2300-MP	EX3400	EX4300	EX4300-MP	EX4600	EX4650	EX9200	EX9250
Single 1.25Ghz	Single 1.25Ghz	Dual 1Ghz	Dual 1.5Ghz	Dual 2.2Ghz	Dual 1.5Ghz	Quad 2.3Ghz	Six-core 2Ghz	8Core 1.6Ghz
DRAM 2GB	DRAM 2GB	DRAM 2GB	DRAM 2GB / 3GB	DRAM 8GB	DRAM 8GB	DRAM 16GB	DRAM 64GB	DRAM 32 / 64GB
FLASH 2GB	FLASH 8GB	FLASH 2GB	FLASH 2GB / 4GB	FLASH 64GB	FLASH 32GB	FLASH 64GB	FLASH 64GB	FLASH 100GB
VC 최대 4개 지원	VC 최대 4개 지원	VC 최대 10개 지원	VC 최대 10개 지원	VC 최대 10개 지원	VC 최대 10개 지원	-	-	-
BGP/MPLS 미지원	BGP/MPLS 미지원	MPLS 미지원	MPLS 미지원	MPLS 미지원	-	-	-	-
MAC 16,000	MAC 16,000	MAC 16,000	MAC 64,000	MAC 272,000	MAC 288,000	MAC 288,000	MAC 1,000,000	MAC 1,000,000
ARP 1,500	ARP 5,000	ARP 16,000	ARP 64,000	ARP 64,000	ARP 48,000	ARP 64,000	ARP 512,000	ARP 512,000
VLAN 4.093	VLAN 4.093	VLAN 4.093	VLAN 4.093	VLAN 4.093	VLAN 4.096	VLAN 4.096	VLAN 32,000	VLAN 32,000

EX2300

12 x 10/100/1000BASE-T

USB/Console/Mgmt

2 x 1GbE or 10GbE SFP/SFP+

EX2300 SKUs	10/100/1000BASE-T Ports	1/10 GbE SFP/SFP+ Ports	Cooling	PoE/PoE+
EX2300-C-12T	12	2	Fanless	0
EX2300-C-12P	12	2	Fanless	*12

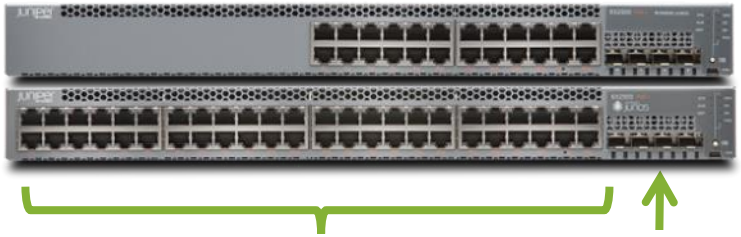

24, 48x10/100/1000BASE-T

4 x 1GbE or
4 x 10GbE SFP/SFP+

Console, Mgmt, USB and Fixed Power/Fan

EX2300 SKUs	10/100/1000BASE-T Ports	10GbE SFP+ Ports	PoE/PoE+ Ports	PoE/PoE+ Power Budget	Cooling	AC/DC
EX2300-24T	24	4	0	0	1x Fixed, AFO	Fixed AC
EX2300-24P	24	4	24	370W	2x Fixed, AFO	Fixed AC
EX2300-24T-DC	24	4	0	0	1x Fixed, AFO	Fixed DC
EX2300-48T	48	4	0	0	1x Fixed, AFO	Fixed AC
EX2300-48P	48	4	48	740W	2x Fixed, AFO	Fixed AC

EX2300-MP

16 or 32 10/100/1000Based-T , 8 or 16 1/2.5G , 4 or 6 10G 인터페이스 지원

WLAN 802.11AC 웨이브 2 AP 지원

Virtual Chassis 최대 4 개 지원

Juniper Enterprise Fusion 사용 가능 (Satellite Device)

POE / POE+ 지원

EX2300-MP

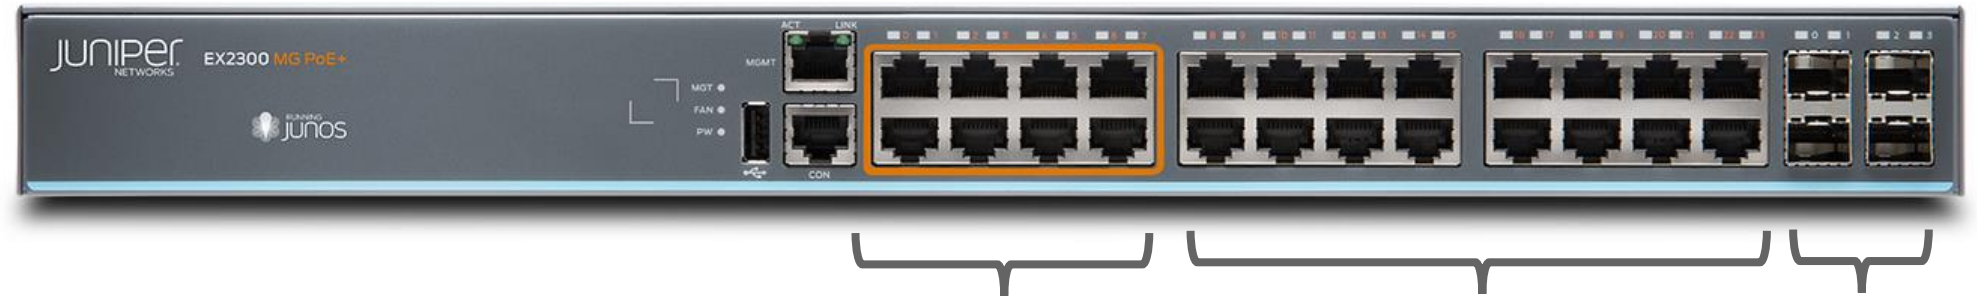

8 x 100M / 1G / 2.5G

16 x 10M /100M /1G

4 x 1G / 10G SFP/SFP+

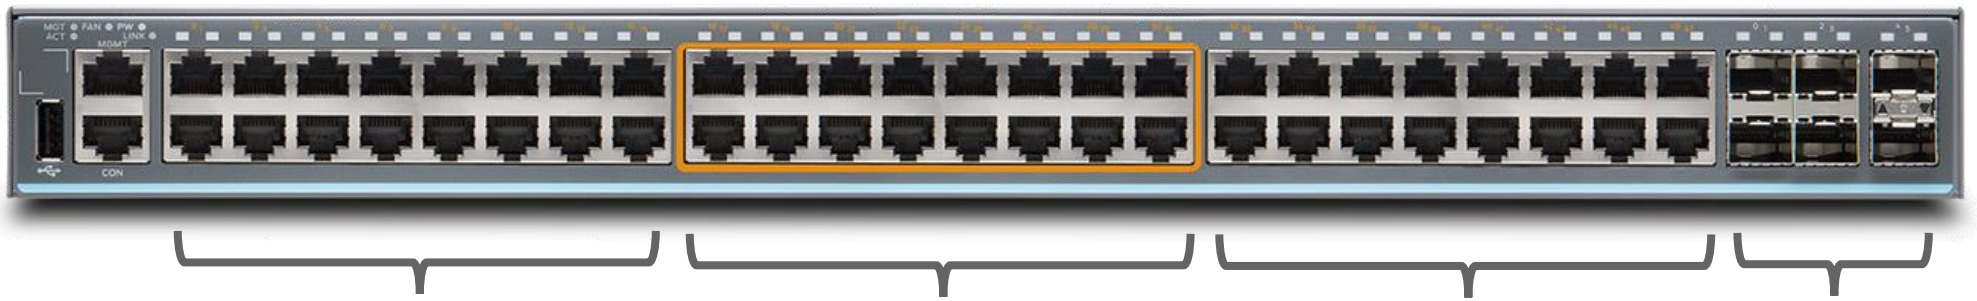

16 x 10M /100M /1G

16 x 10 / 100M / 1G / 2.5G

16 x 10M /100M /1G

6 x 1G / 10G SFP/SFP+

EX3400

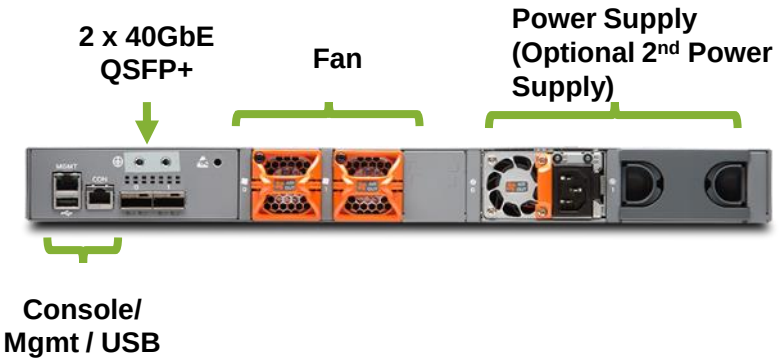

EX2300 SKUs	10/100/1000BASE-T Ports	10GbE SFP+ Ports	40Gbe QSFP+	PoE/PoE+ Power Budget	Cooling	AC/DC
EX3400-24T	24	4	2	0	2xFan	AC
EX3400-24P	24	4	2	MAX 720W	2xFan	AC
EX3400-24T-DC	24	4	2	0	2xFan	DC
EX3400-48T	48	4	2	0	2xFan	AC
EX3400-48P	48	4	2	MAX 1440W	2xFan	AC

EX3400

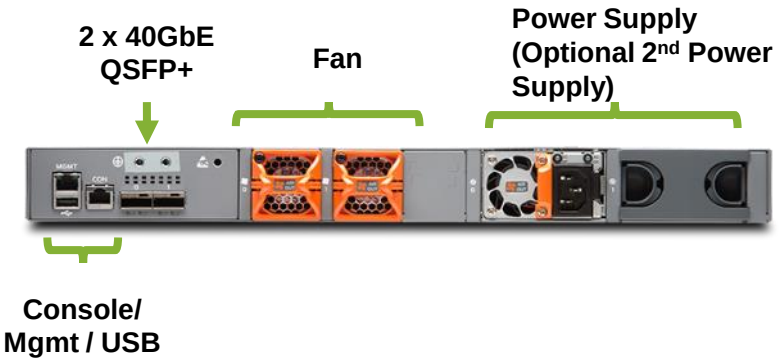

EX2300 SKUs	10/100/1000BASE-T Ports	10GbE SFP+ Ports	40Gbe QSFP+	PoE/PoE+ Power Budget	Cooling	AC/DC
EX3400-24T	24	4	2	0	2xFan	AC
EX3400-24P	24	4	2	MAX 720W	2xFan	AC
EX3400-24T-DC	24	4	2	0	2xFan	DC
EX3400-48T	48	4	2	0	2xFan	AC
EX3400-48P	48	4	2	MAX 1440W	2xFan	AC

EX4300

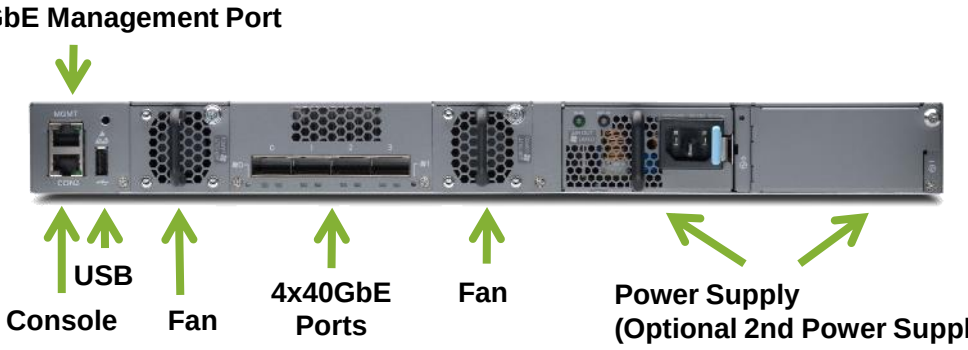

EX2300 SKUs	10/100/1000BASE-T Ports	10GbE SFP+ Ports	40Gbe QSFP+	PoE/PoE+ Power Budget	Cooling	AC/DC
EX4300-24T	24	4 (Option)	4	0	2xFan	AC
EX4300-24P	24	4 (Option)	4	550W	2xFan	AC
EX4300-48T	48	4 (Option)	4	0	2xFan	AC/DC
EX4300-48P	48	4 (Option)	4	900W	2xFan	AC

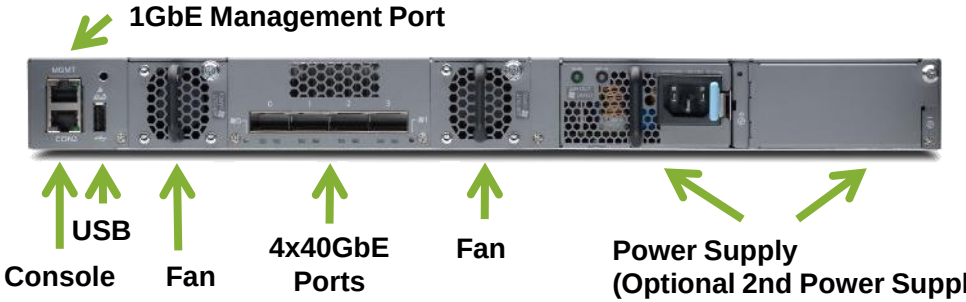

EX2300 SKUs	100/1000BASE-X Ports	10GbE SFP+ Ports	40Gbe QSFP+	Module	Cooling	AC/DC
EX4300-32F	32	4	4	8 x 1Gbe/10Gbe or 2 x 40Gbe	2xFan	AC/DC

EX4300-MP

48 포트 1 / 2.5 / 5 / 10 Gig 멀티 스피드 이더넷 스위치

10 / 40 / 100 Gig 업 링크 모듈 지원

EX4300 과 Virtual-Chassis 지원 (최대 10 개 장치)

Juniper Enterprise Fusion 및 Virtual-Chassis-Fabric 지원

POE / POE+ / POE++ 지원

EX4300-MP

EX4600

최대 40개 10Gbe 인터페이스 지원 (Fixed 24개, 8 x 10Gbe 모듈 2개 추가 장착 가능)

최대 12개의 40Gbe 인터페이스 지원 (Fixed 4개, 4 x 40Gbe 모듈 2개 추가 장착 가능)

Virtual-Chassis 지원 (최대 10 개 장치)

EX4600

Console/
Mgmt / USB

Redundant
Fan

Redundant
Power

24 x 10GbE SFP+

4 x 40GbE QSFP+

Module Slots (2) :
8 x 10GbE or 4 x 40GbE

EX4600 SOFTWARE ARCHITECTURE

Provides the foundation for advanced functions

- ISSU (In-Service Software Upgrade)
- Other Juniper applications for additional service in a single switch
- Third-party application
- Can bring up the system much faster

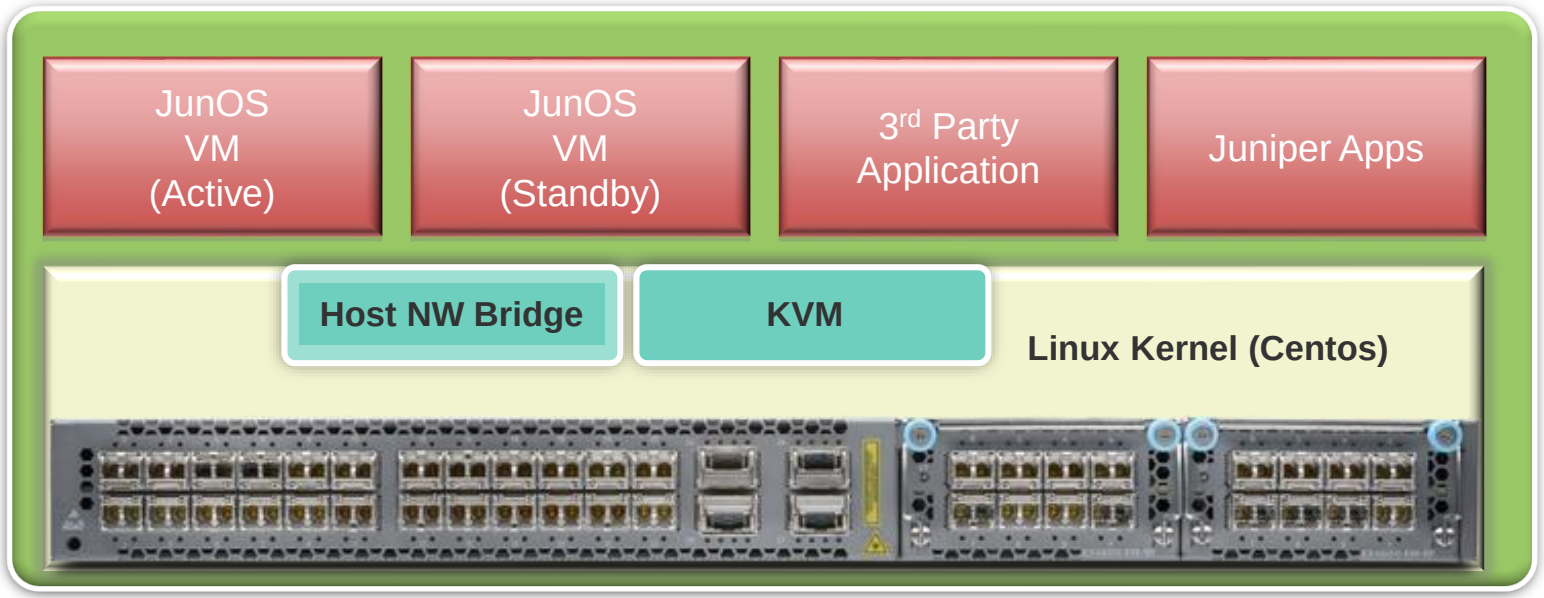

EX4600 SOFTWARE ARCHITECTURE

ISSU - In-service-software-upgrade

- All devices are managed by one RE
- Master issues upgrade command
- System launches a new Junos VM with new image as backup
- All states are synchronized to the new backup Junos
- “Hot-move” PFE control from current master to backup Junos
- New master will control the PFE forwarding
- Stop the new backup VM

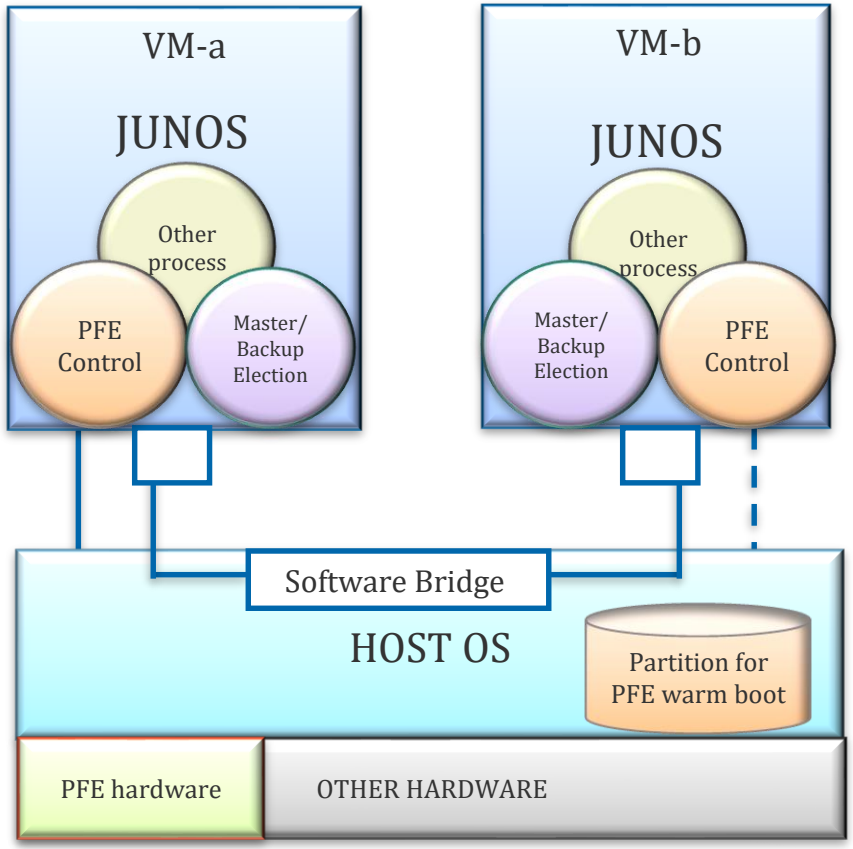

*** ISSU supported on Standalone switch; Not in VC

EX4600 SOFTWARE ARCHITECTURE

Full upgrade (upgrade junos & centos)

- Three ways to upgrade entire EX4600 chassis:
- a) USB install
- b) PXE install
- c) “request system software add jinstall-vjunos-<version>.tgz”
- Normally, only the Junos VM is upgraded
- Every jinstall-vjunos package contains a minimum-centos-compatible version #, currently it’s “13.2”. This means any version above 13.2 is compatible.
- ONLY when min-version compatibility checks fails will Centos OS be upgraded
- In addition, the “force” option can be specified with the “software add” command to force upgrade host Centos

EX4600 SOFTWARE ARCHITECTURE

reboot/halt/poweroff

System reboot commands:

> request system reboot <- reboots Junos VM

> request system reboot hypervisor <- reboots both Junos VM and hypervisor OS

Halt and poweroff the chassis:

> request system halt

> request system power-off

EX4650

1/10/25 Gbe 인터페이스 48 포트 지원

40/100 Gbe Uplink 인터페이스 지원

EVPN/VXLAN L2/L3 GATEWAY 지원

EX4650

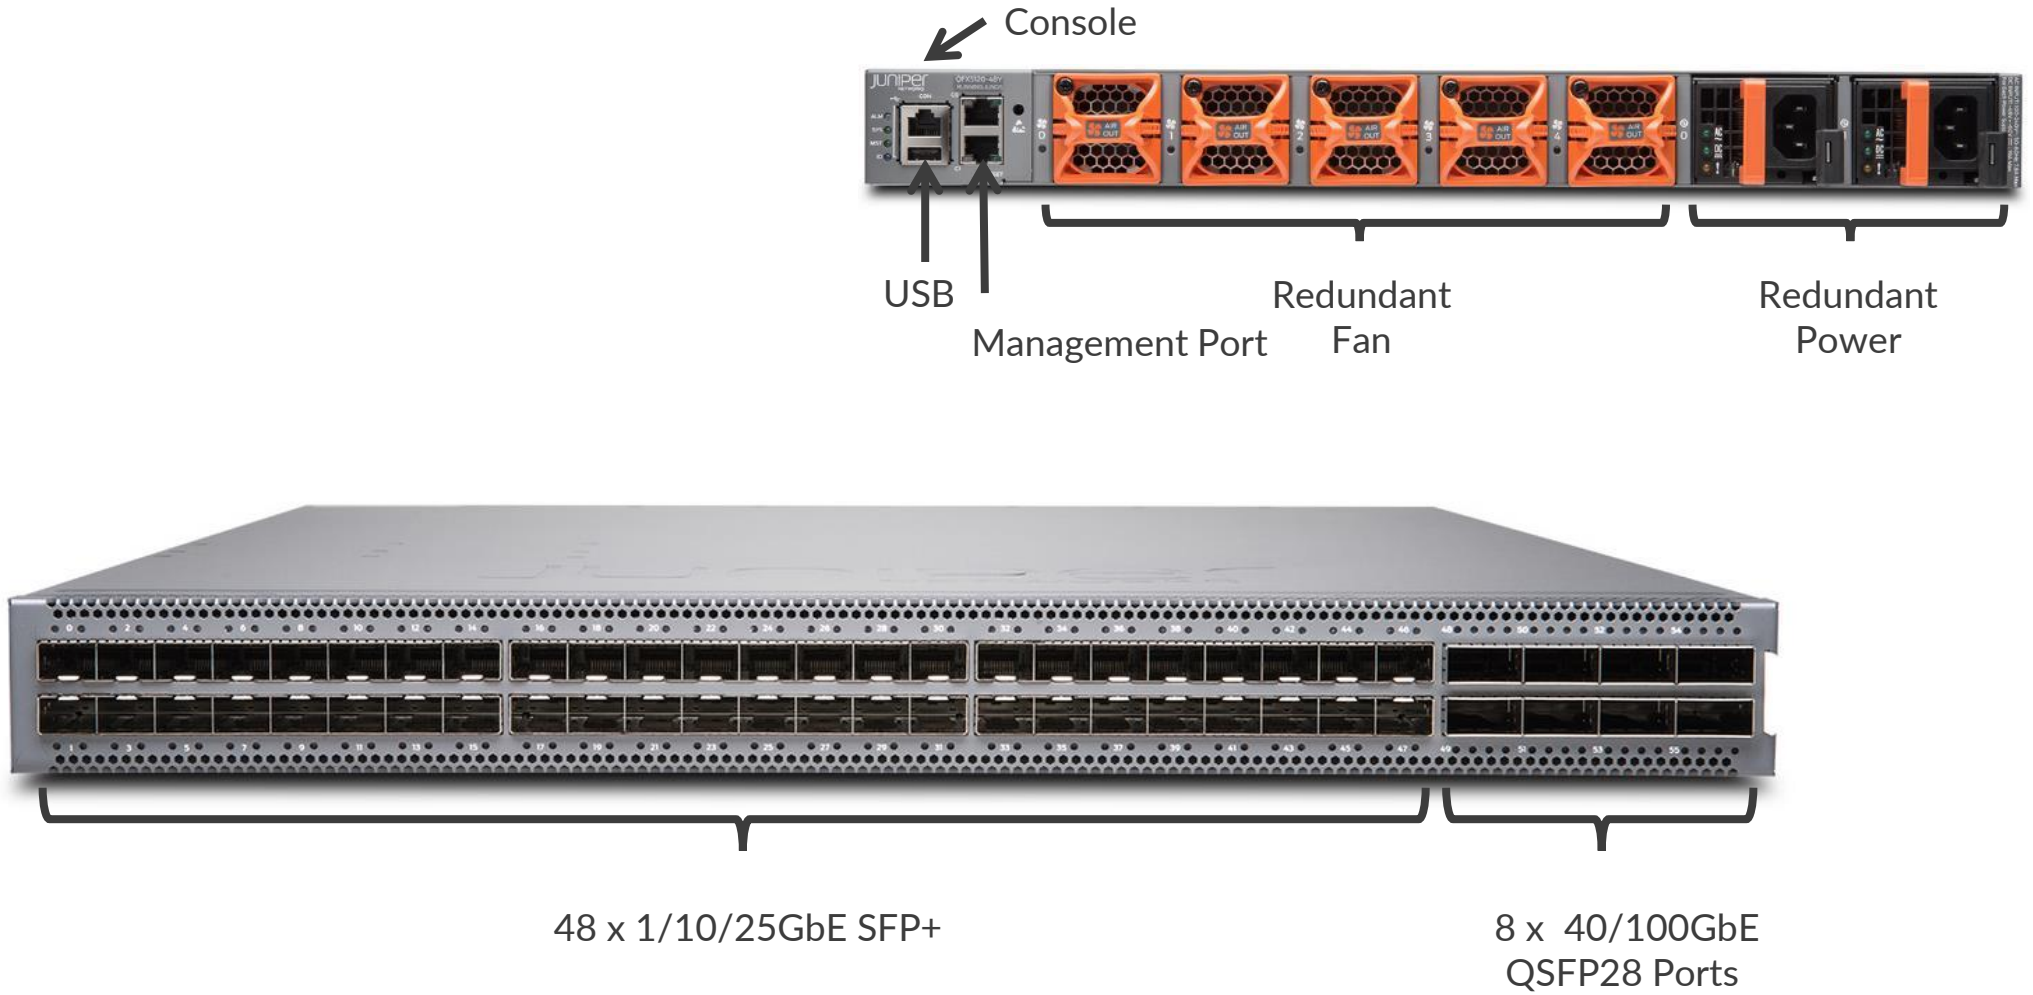

EX4650

- First 48 ports support 1G/10G/25G port speeds
- They are in 10G mode by default – interface will NOT get created automatically on inserting 1G/25G transceivers.
- CLI used to configure the port speed to 1G/25G mode manually

  set chassis fpc 0 pic 0 port <> speed <1G/25g>
- Configuration applies to a quad (set of 4 ports) at a time, and configuration is only allowed on the first port of the quad. For example, in order to configure ports 4-7 in 25G mode, we need to configure below CLI only

  set chassis fpc 0 pic 0 port 4 speed 25g
- 25G/40G/100G interfaces naming convention used is “et”

1 RU PLATFORM: EX9251

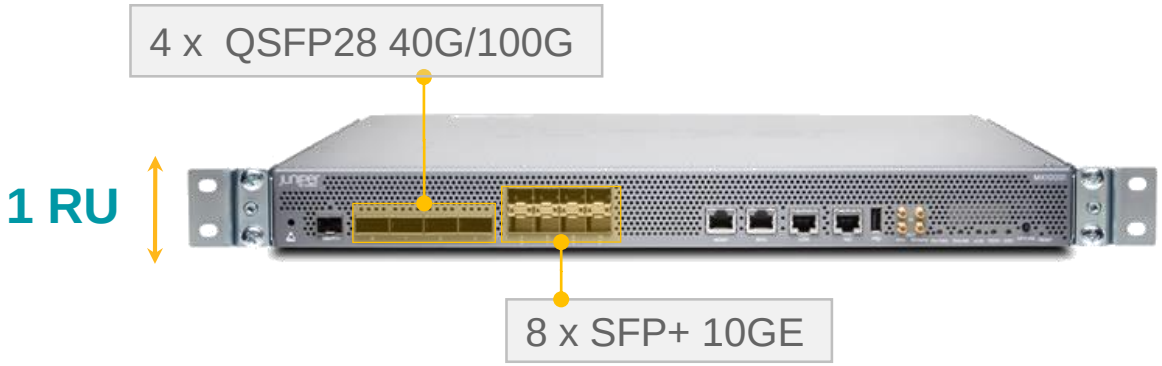

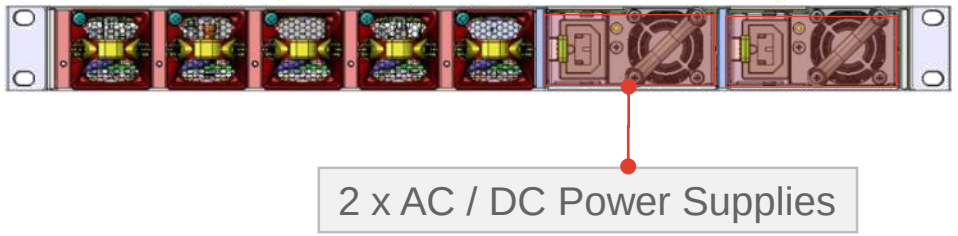

Compact 1RU Modular System

- **Fixed form factor 1RU platform**
- **Shallow 19” depth**
- **Multi-rate (4)100GE, (4)40GE and (24)10/1GE ports**
- **Junos Fusion Enterprise aggregation device**
- **PTP support for timing requirements**

X86 Based Routing Engine

- **X86 based single RE**

Power

- **Power efficient ~ 0.9W/GbE**
- **AC/DC PSMs with N+N redundancy**

Cooling and NEBS

- **Redundant, Front to back air-cooling**
- **NEBS Compliant**

3 RU PLATFORM: EX9253

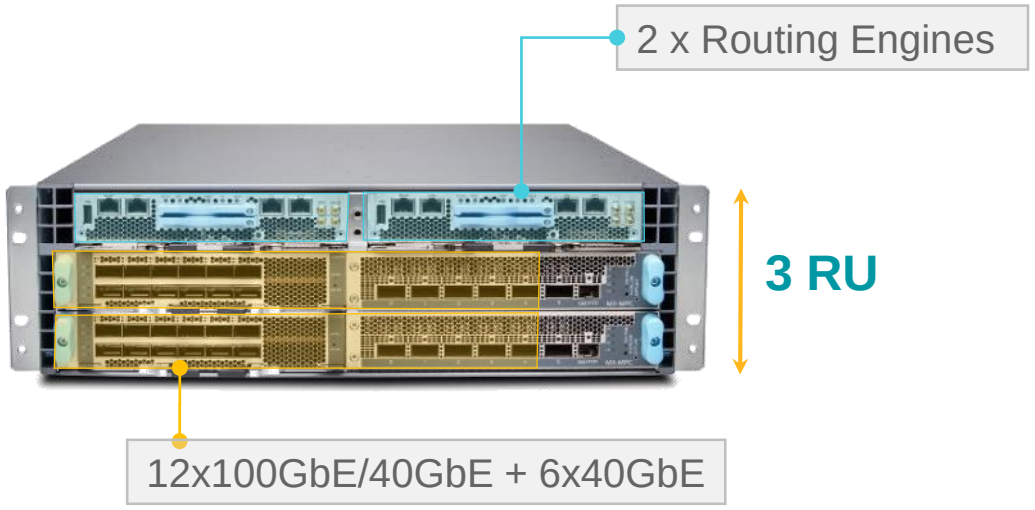

Compact 3RU Modular System

- **Compact high density core and distribution**
- **Junos Fusion Enterprise aggregation device**
- **2 line card slots**
- **PTP support for timing requirements**

Line Card

- **Multi-rate 100GbE, 40GbE and 10GbE ports**
  - 12x100GbE/40GbE + 6x40GbE
  - 12x100GbE/40GbE with MACsec + 6x40GbE
- **12x100GE, 18x40GE or 72x10GE/1GE per linecard**

X86 Based Routing Engine

- **Modular, redundant and upgradable**

Power

- **Power efficient ~ 0.9W/GbE**
- **AC/DC PSMs with N+N redundancy**

Cooling and NEBS

- **Redundant, front to back air-cooling, NEBS compliant**

EX9200

14- Slots

ROUTING ENGINE (1+1)

Linecard Slots (11+1)

8- Slots

ROUTING ENGINE (1+1)

Linecard Slots (6)

4- Slots

ROUTING ENGINE (1+1)

Linecard Slots (2+1)

System Feature

- MC-LAG
- Logical Systems
- Junos Fusion for Enterprise
- 1/10/40/100G Interface

Key Differentiations

- Comprehensive Layer 2
- Full featured IPv4/IPv6
- Scalable full featured multicast
- MPLS/VPLS/EVPN

EX9200

EX9214

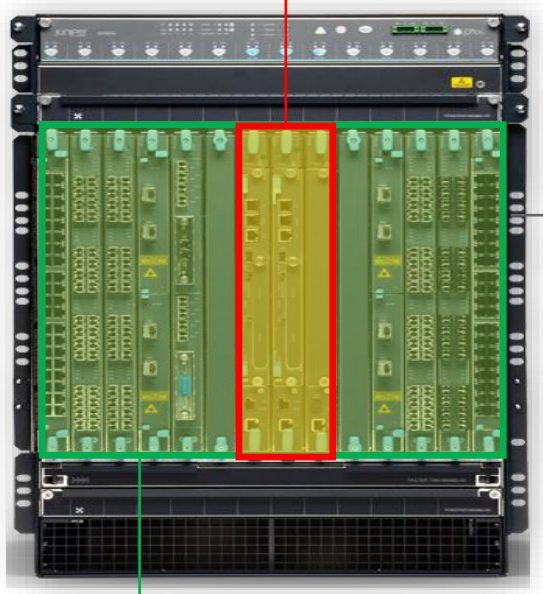

SWICH FABRIC (2+1)
ROUTING ENGINE (1+1)

4 INTERNAL POWER
4 EXTERNAL POWER

LINE CARDS
40P 10/100/1000
40P 100/1000 SFP
32P, 40P 10G SFP+
12P 40G QSFP OR
4P 100G

- The default mode is 2+1 on EX9214 and 1+1 on EX9208/EX9204
- Supports up to 480Gbps per slot

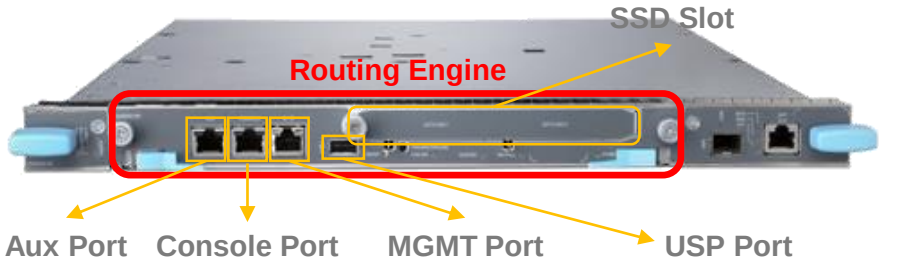

- Up to 64G DRAM, 64G SSD
- Six-core, 2 GHz Intel processor

EX9200

EX9200-40XS

10G 40P

EX9200-12QS

40G 12P or 100G 4P

EX9200-32XS

10G 32P

EX9200-40F

1G SFP 40P

EX9200-40T

1G-T 40P

EX9200-2C-8XS

2P 100G, 8P 10G

EX9200-6QS

6P 40G

AGENDA

- EX Series
- **EX Architecture**
- QFX Series
- QFX Software Architecture

CONTROL PLANE AND DATA PLANE BASIC

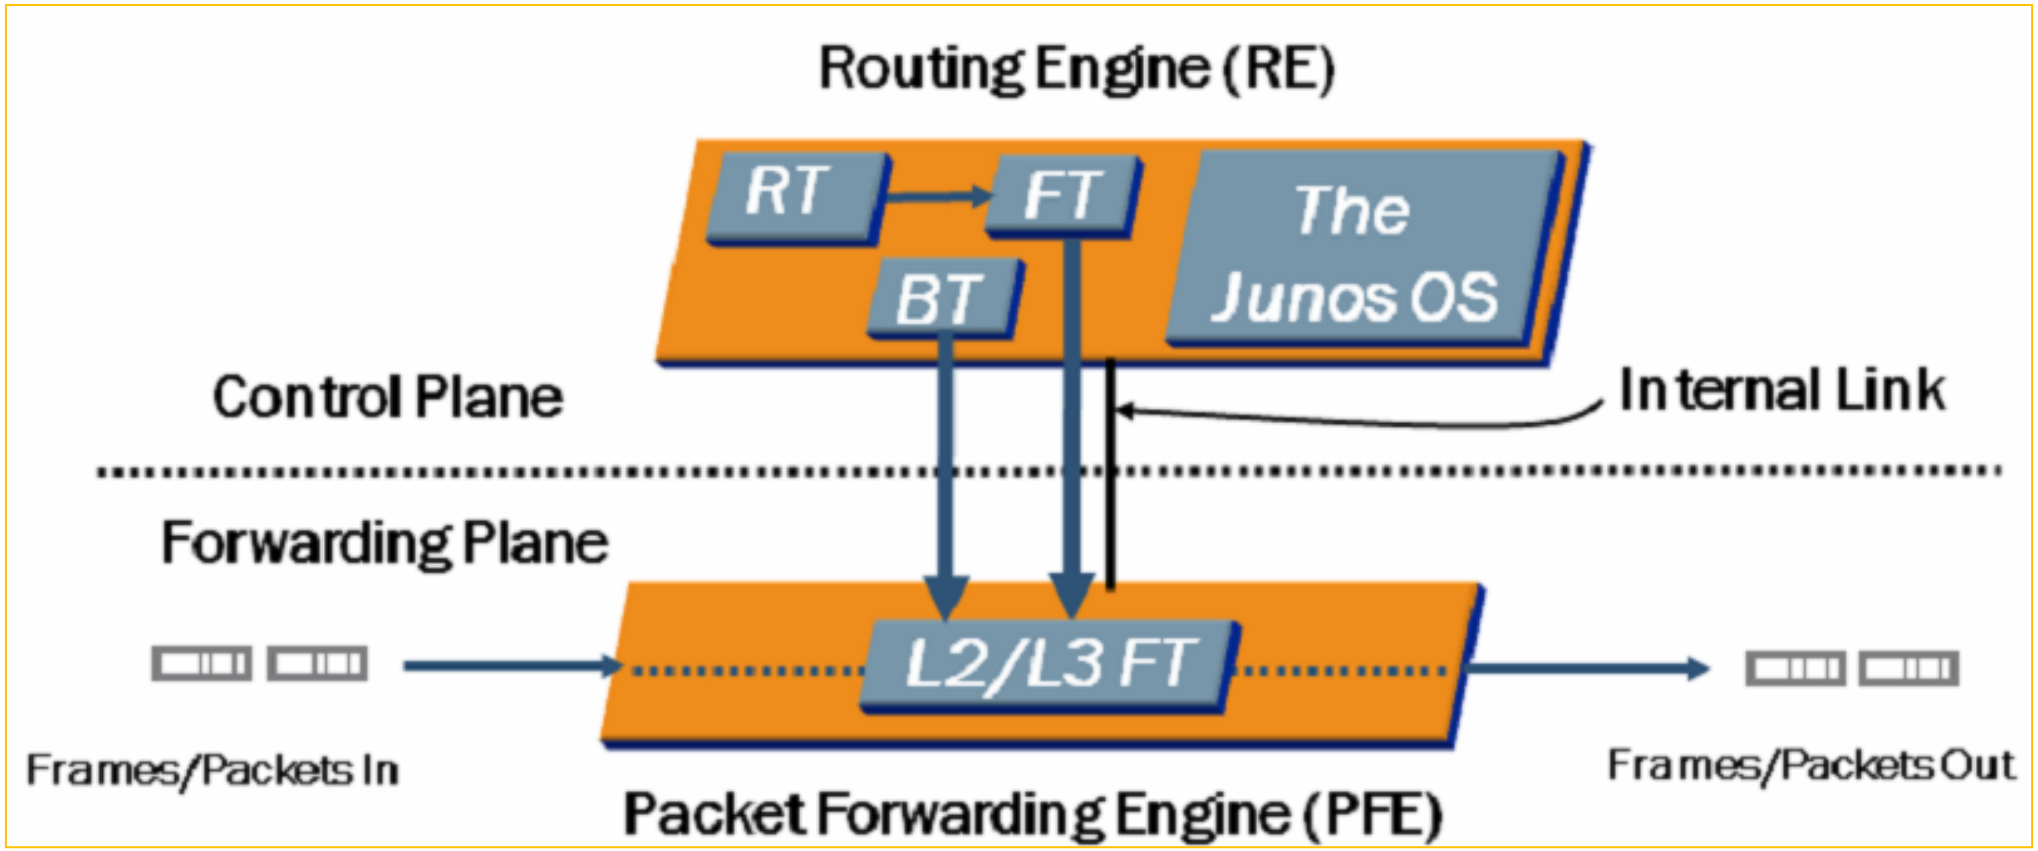

PACKET WALK - KNOWN

1. Packet enters a port and its attached PFE

2. PFE performs packet lookup and ingress processing

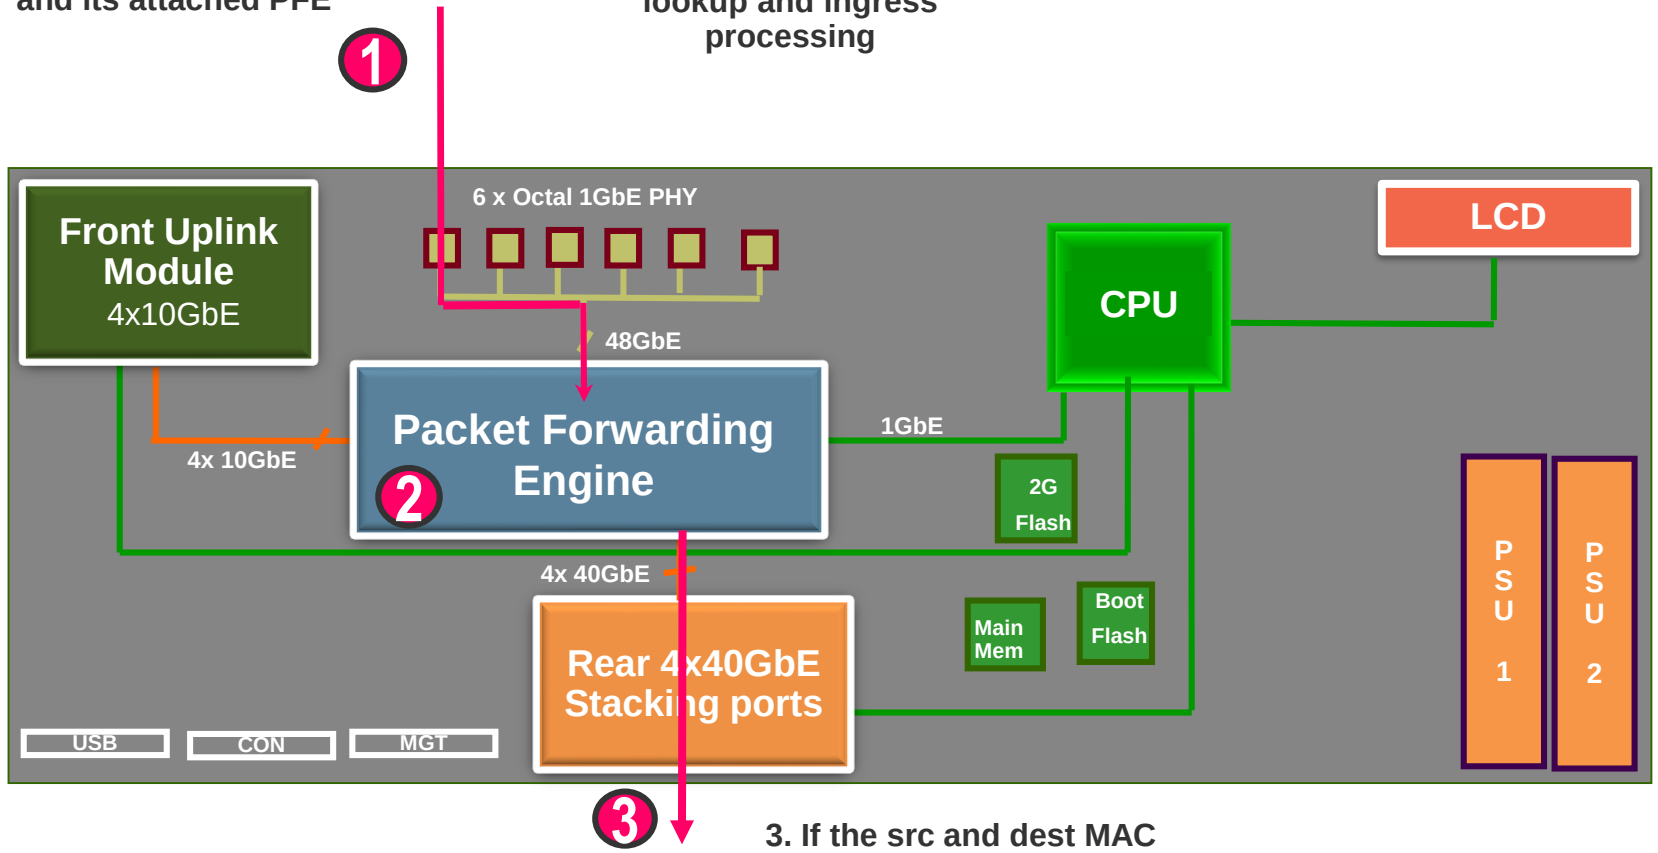

3. If the src and dest MAC is known, the PFE forwards the packet to the egress port

PACKET WALK - UNKNOWN

1. Packet enters a port and its attached PFE

2. PFE performs packet lookup and ingress processing

3. If src and dest MAC is unknown, PFE sends packet to CPU for learning

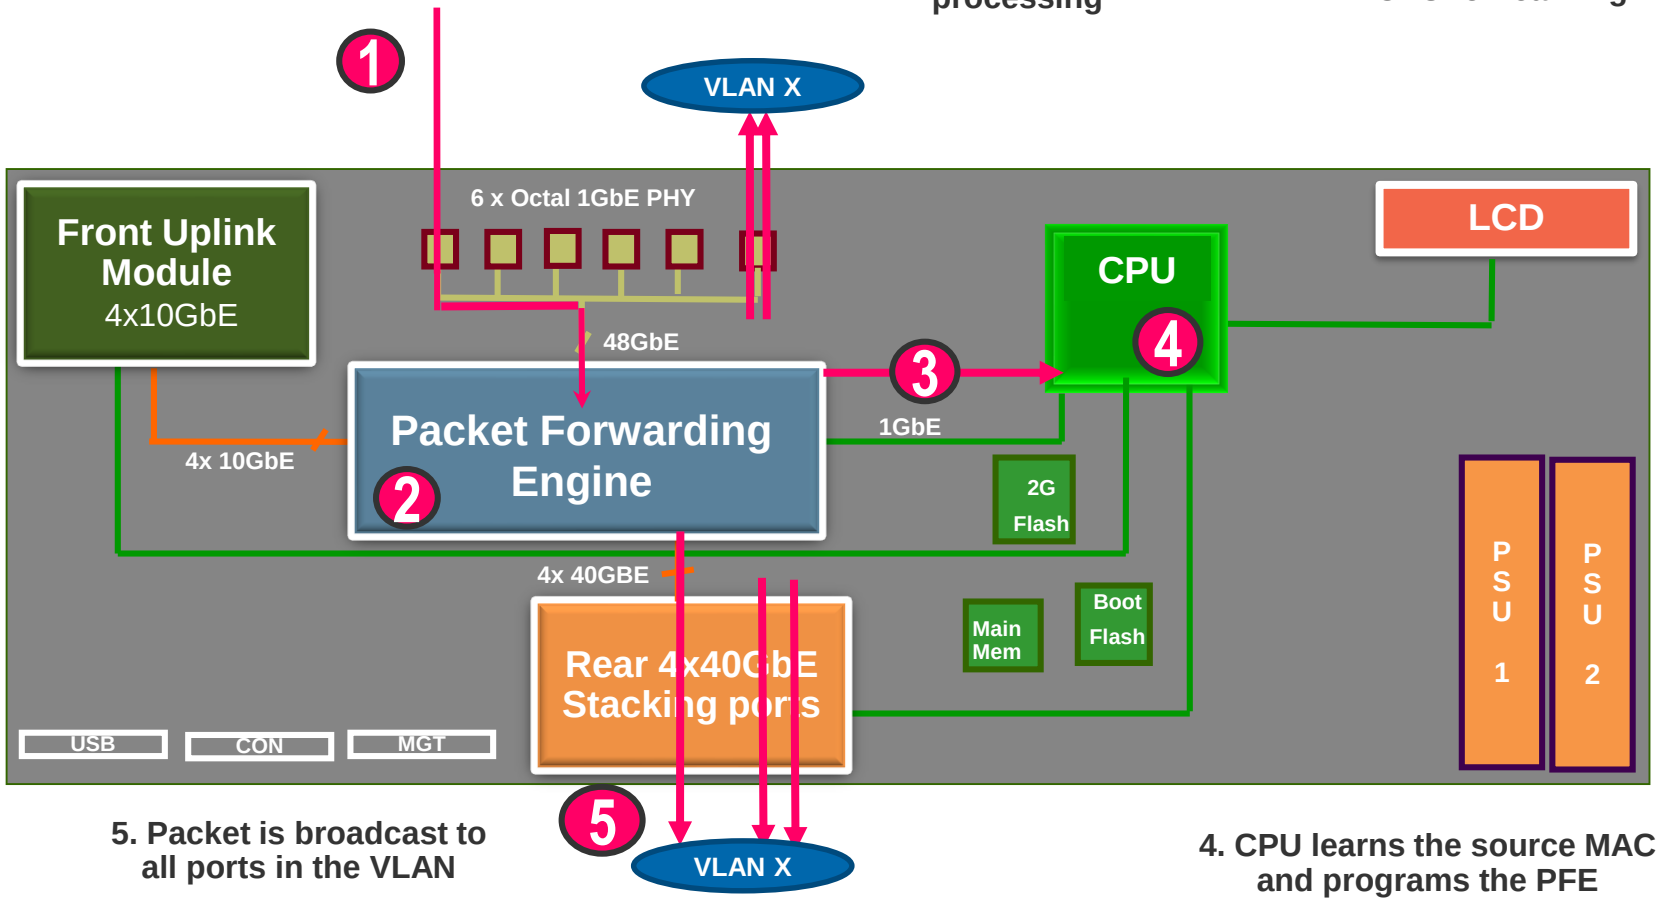

5. Packet is broadcast to all ports in the VLAN

4. CPU learns the source MAC and programs the PFE

PACKET WALK IN VIRTUAL CHASSIS - UNKNOWN

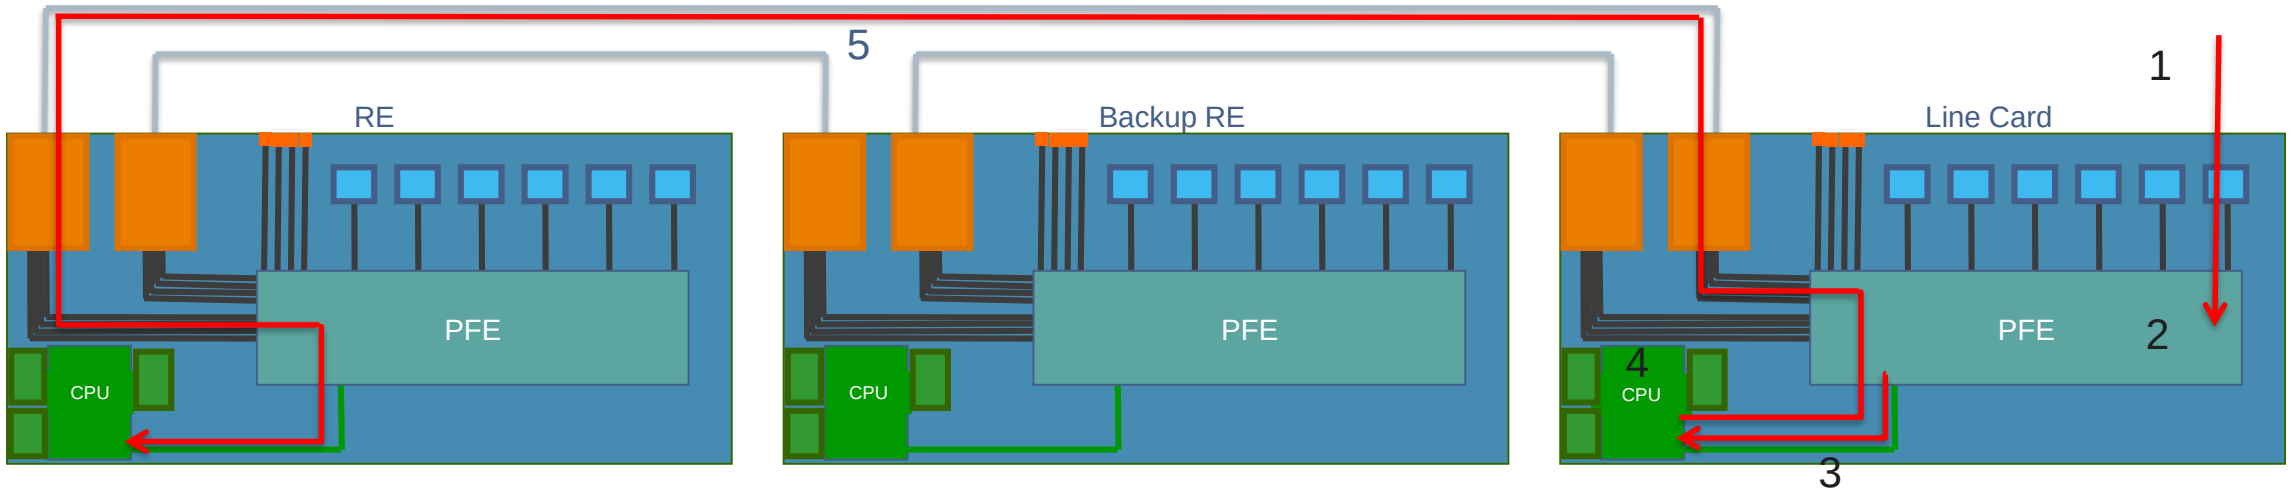

1. Packet enters port and the PFE on the line card

2. PFE performs source MAC lookup and the source MAC is unknown

3. A notification is sent to the local CPU

4. LC CPU encapsulates the packet header info (MAC, VLAN, port #) in an IPC message to send to the Master's CPU for learning

5. The IPC message is routed to the Master's CPU via the shortest path, first via the PCI-e to the first PFE en-route, then via the Virtual Chassis backplane arriving at the Master's CPU through its PCI-e connection

PACKET WALK IN VIRTUAL CHASSIS - UNKNOWN

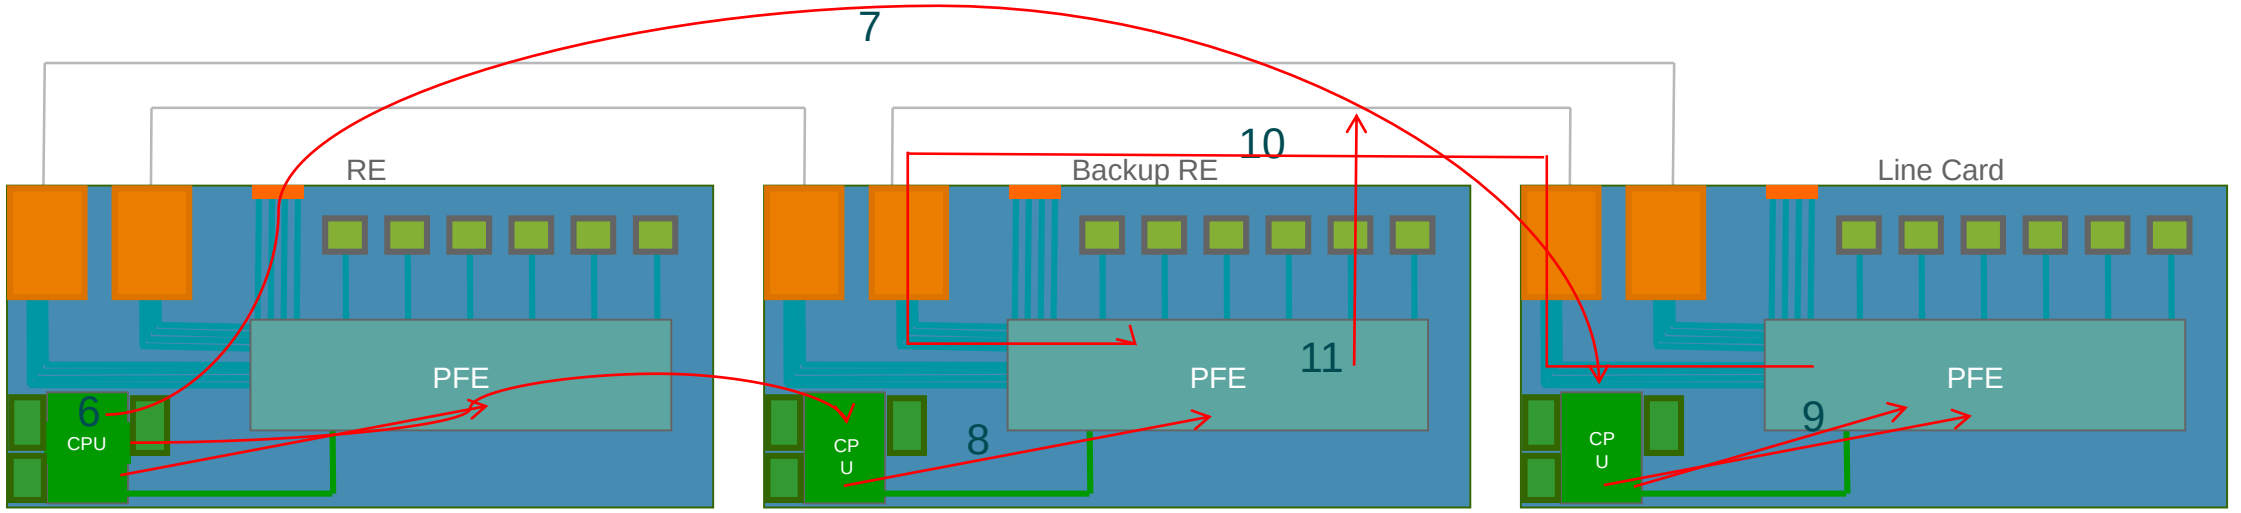

6. Master's CPU decides whether to learn the MAC address or discard it based on configuration (e.g. MAC-limiting)
7. If Master's CPU learns the MAC address, it sends IPC messages telling all other CPUs to program the new MAC address into their PFE's TCAM. As in Step 5, these messages traverse both the PCI-e connections and the Virtual Chassis backplane to reach CPUs on every switch in the Virtual Chassis.
8. Each CPU programs the new learned MAC address in their respective PFE's MAC table
9. The ingress CPU (LC CPU), re-injects the stored original packet back out to the ingress PFE
10. The ingress PFE forwards the packet to the egress PFE based on the lookup result of the packet's destination MAC address
11. The egress PFE forwards the packet to the egress port based on the same lookup result, no additional lookup needed

PACKET WALK IN VIRTUAL CHASSIS - L3 UNICAST

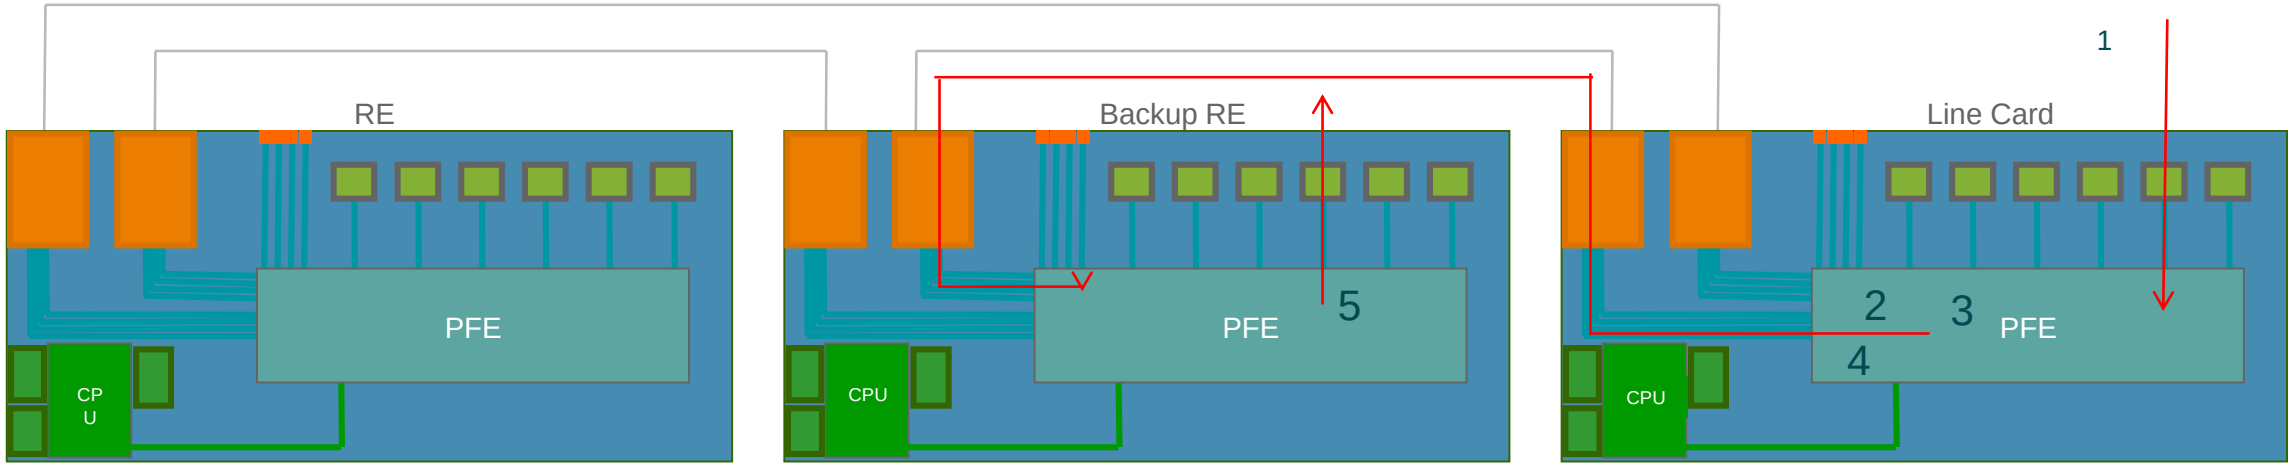

1. Packets enter a port on the “LC” switch
2. Ingress PFE performs MAC address lookup, finds source MAC known, and subsequently performs lookup on the destination MAC address
3. Destination MAC address matches that of the switch itself, which triggers L3 lookup
4. Ingress PFE forwards packet based on L3 lookup result – if no entry found in L3 TCAM or if next-hop missing, then packet is dropped in hardware
5. The egress PFE forwards the packet to the egress port based on the same L3 lookup result; no additional lookup needed

PACKET WALK IN VIRTUAL CHASSIS - CONTROL PACKET

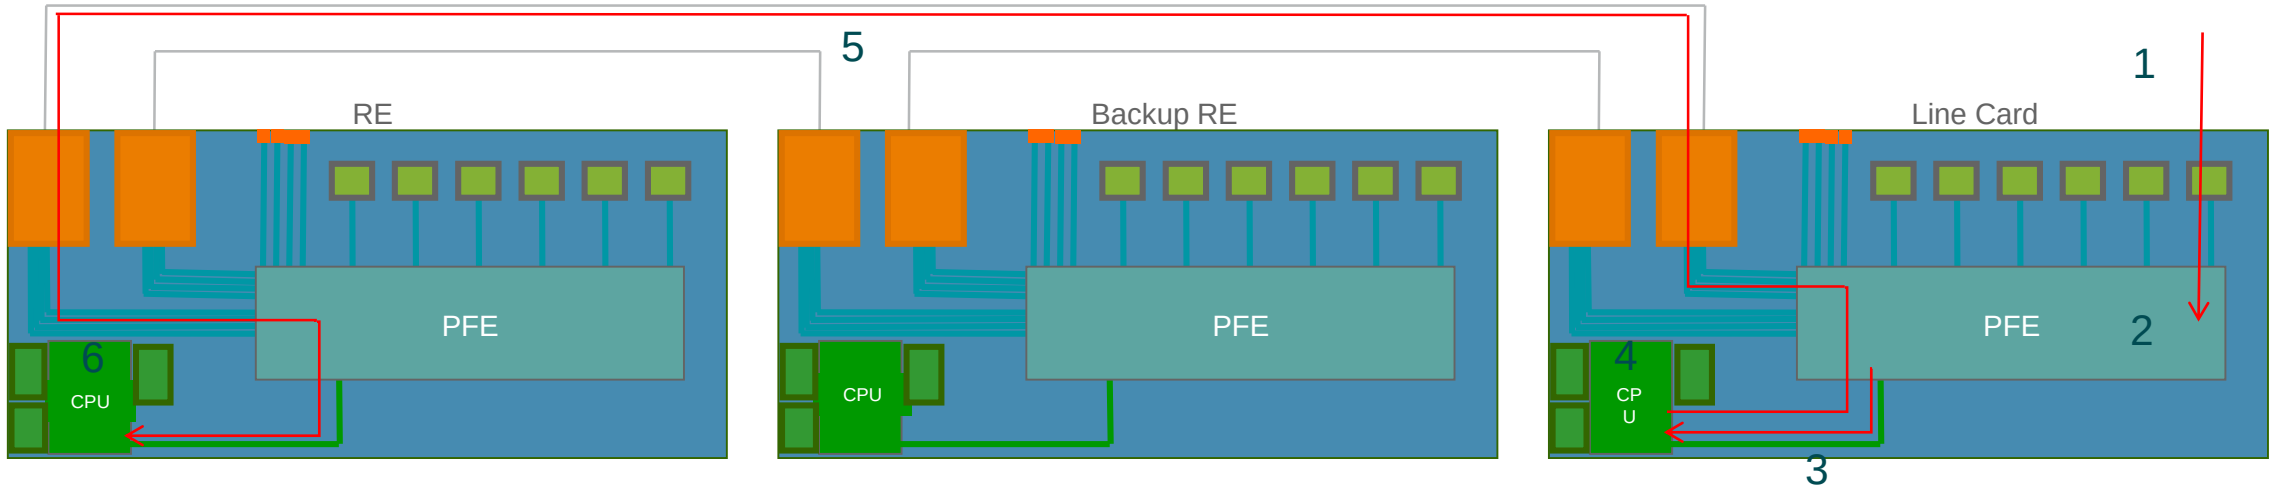

1. Packets enter a port on the “LC” switch
2. PFE performs MAC address lookup; realizes destination is RE CPU
3. PFE forwards packet to local CPU via PCI-e connection
4. LC CPU does NOT store the packet but relays its entirety in an IPC message to the Master’s CPU for processing
5. The IPC message is routed to the Master’s CPU via the shortest path, first via the PCI-e to the first PFE en-route, then via the Virtual Chassis backplane to reach one of the PFEs on the Master switch, eventually arriving at the Master’s CPU through its PCI-e connection
6. Master’s RE CPU processes the packet

AGENDA

- EX Series
- EX Architecture
- **QFX Series**
- QFX Software Architecture

QFX SERIES DATACENTER SWITCH PRODUCT PORTFOLIO (1/2)

LEAF, SPINE, DCI DC EDGE

QFX5100				QFX5110		QFX5120
QFX5100-48S	**QFX5100-48T**	**QFX5100-24Q**	**QFX5100-96S**	**QFX5110-48S**	**QFX5110-32Q**	**QFX5120-48Y**
Leaf	**Leaf**	**Leaf, Spine**	**Leaf**	**Leaf**	**Spine**	**Leaf**
Up to 1.44Tbps	**Up to 1.44Tbps**	**Up to 2.56Tbps**	**Up to 2.56Tbps**	**Up to 1.76Tbps**	**Up to 2.56Tbps**	**Up to 4Tbps**
1/10G 48P, 40G 6P	1/10G-T 48P, 40G 6P	Up to 40G 32P	1/10G 96P, 40G 8P	1/10G 48P, 40/100G 4P	40G 32P or 40G 20P + 100G 4P	1G/10G/25G 48P, 40/100G 8P
		2 Expansion Slot	2RU	Higher Scale, 100G Support		**Low Latency 550n**
VC, VCF, BGP, MPLS, VXLAN, EVPN				VC, VCF, BGP, MPLS, EVPN, VXLAN Inter Routing, PTP		BGP, MPLS, EVPN, VXLAN

QFX SERIES DATACENTER SWITCH PRODUCT PORTFOLIO (2/2)

LEAF, SPINE, DCI DC EDGE

QFX5200		QFX10002			QFX10000	
QFX5200-48Y	**QFX5200-32C**	**QFX10002-36Q**	**QFX10002-72Q**	**QFX10002-60C**	**QFX10008**	**QFX100016**
Leaf	Spine	Leaf, Spine	Leaf, Spine	Leaf, Spine	Spine	Spine
Up to 3.6Tbps	Up to 6.4Tbps	Up to 2.88Tbps	Up to 5.76Tbps	Up to 12Tbps	Up to 48Tbps	Up to 96Tbps
10/25G 48P, 40/100G 6P	50G 64P or 40G/100G 32P	40G 36P or 100G 12P	40G 72P or 100G 24P	40/100G 60P	Up to 10G 1152P, 40G 288P, 100G 240P	Up to 10G 2304P, 40G 576P, 100G 480P
Multispeed, High Throughput PFE		**Juniper Q5 ASIC** 12GB Buffer, VoQ	**Juniper Q5 ASIC** 24GB Buffer, VoQ	**Juniper Q5 ASIC** 24GB Buffer, VoQ	**Juniper Q5 ASIC** 12GB Buffer per LC (60S 8GB), VoQ	
BGP, MPLS, VXLAN, EVPN		BGP, MPLS, EVPN, VXLAN Inter Routing, PTP			BGP, MPLS, EVPN, VXLAN Inter Routing, PTP MACSEC QFX10000-30C-M	

QFX5100

10/40/100Gbe 지원 장비

Spine-and-Leaf, Layer 3 또는 OpenClos IP Fabric

Virtual-Chassis / Virtual-Chassis-Fabric 지원

QFX5100-24Q

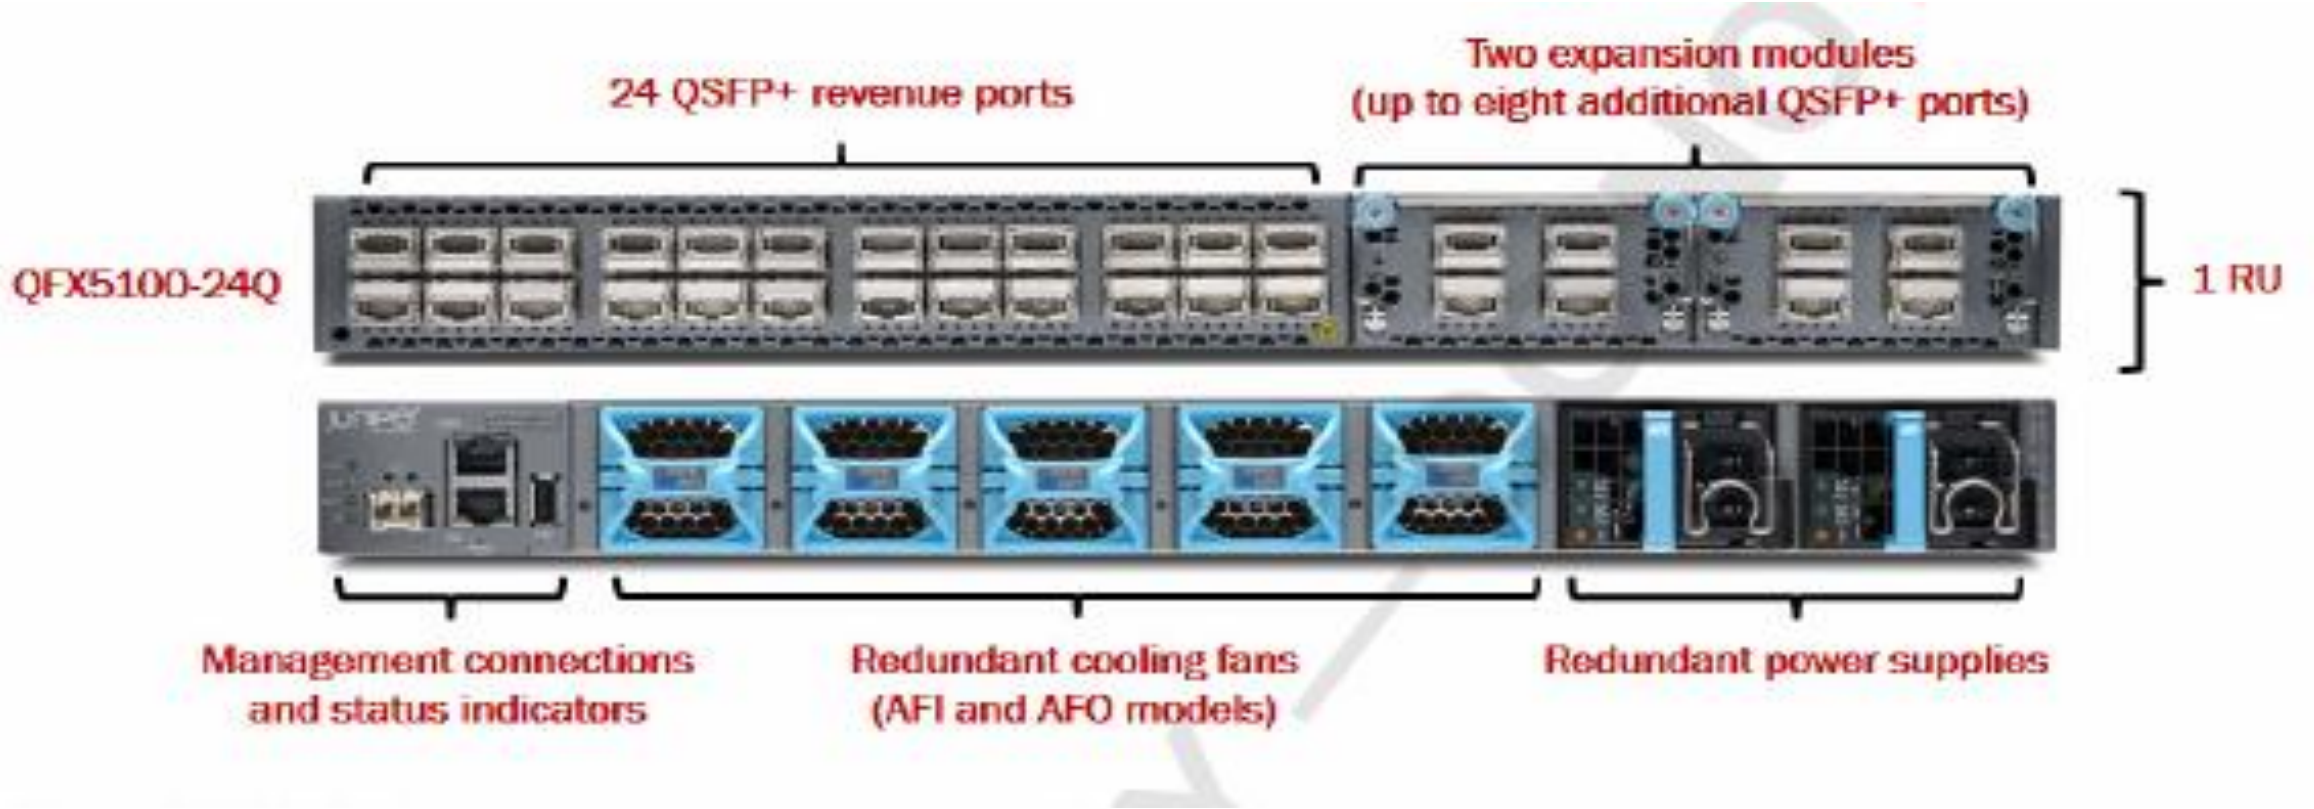

QFX5100-48S AND QFX5100-48T

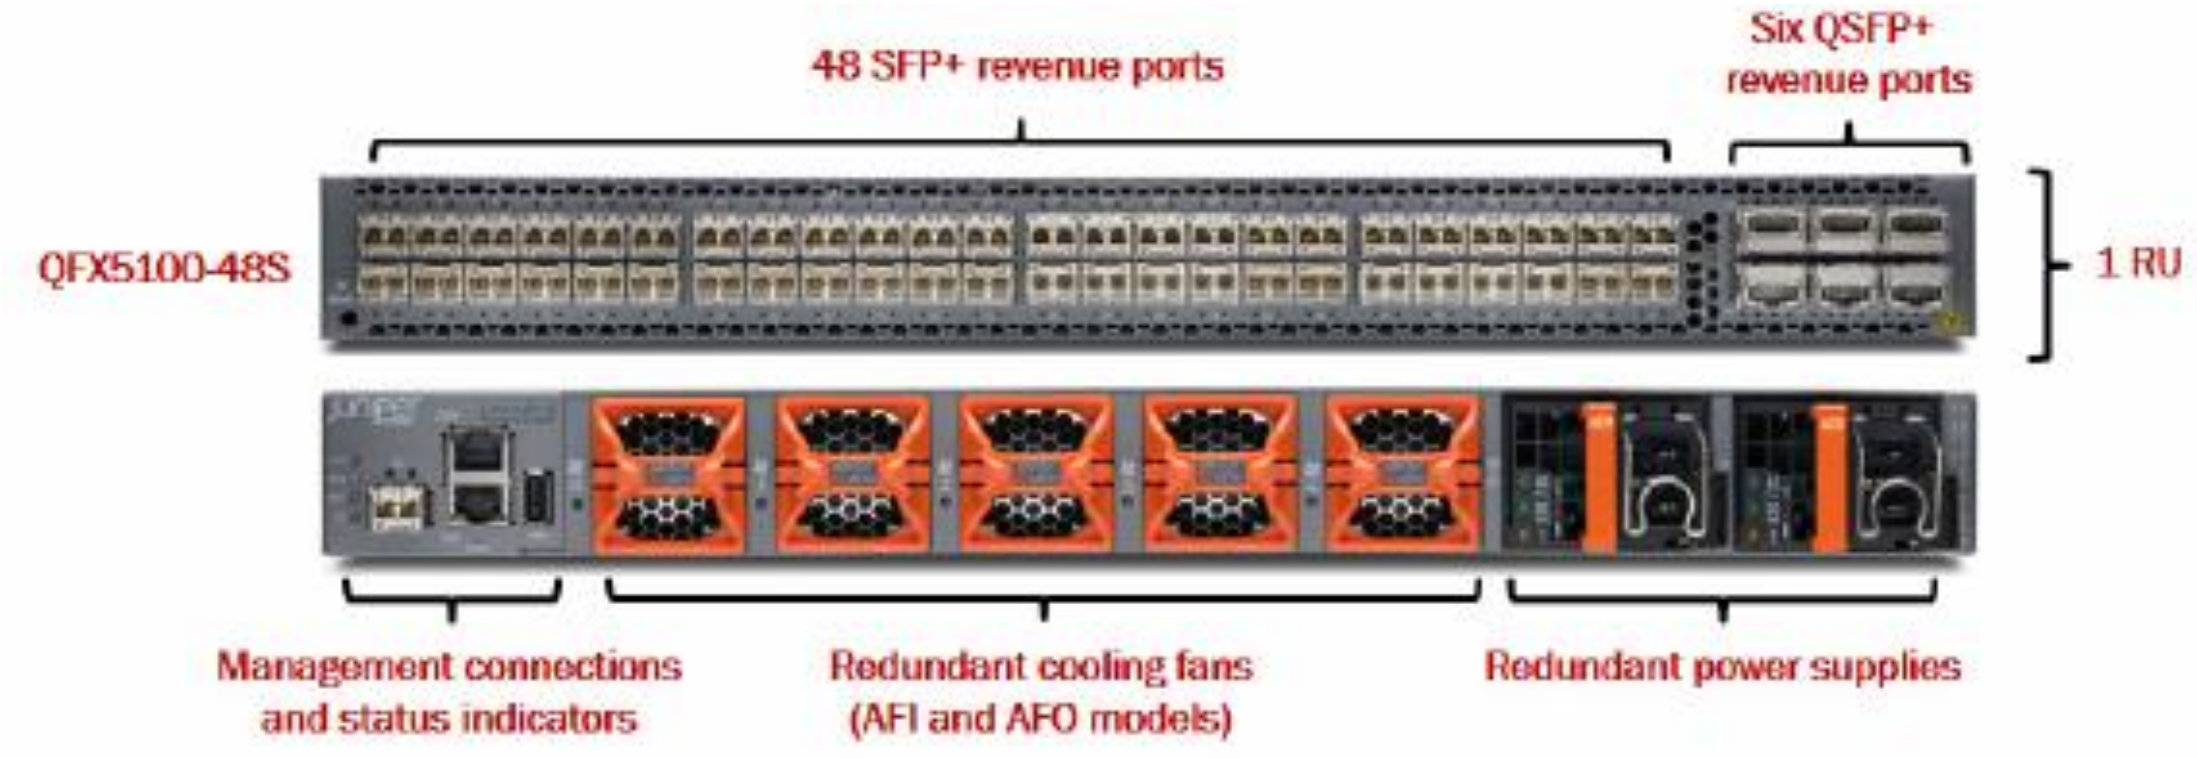

QFX5100-96S

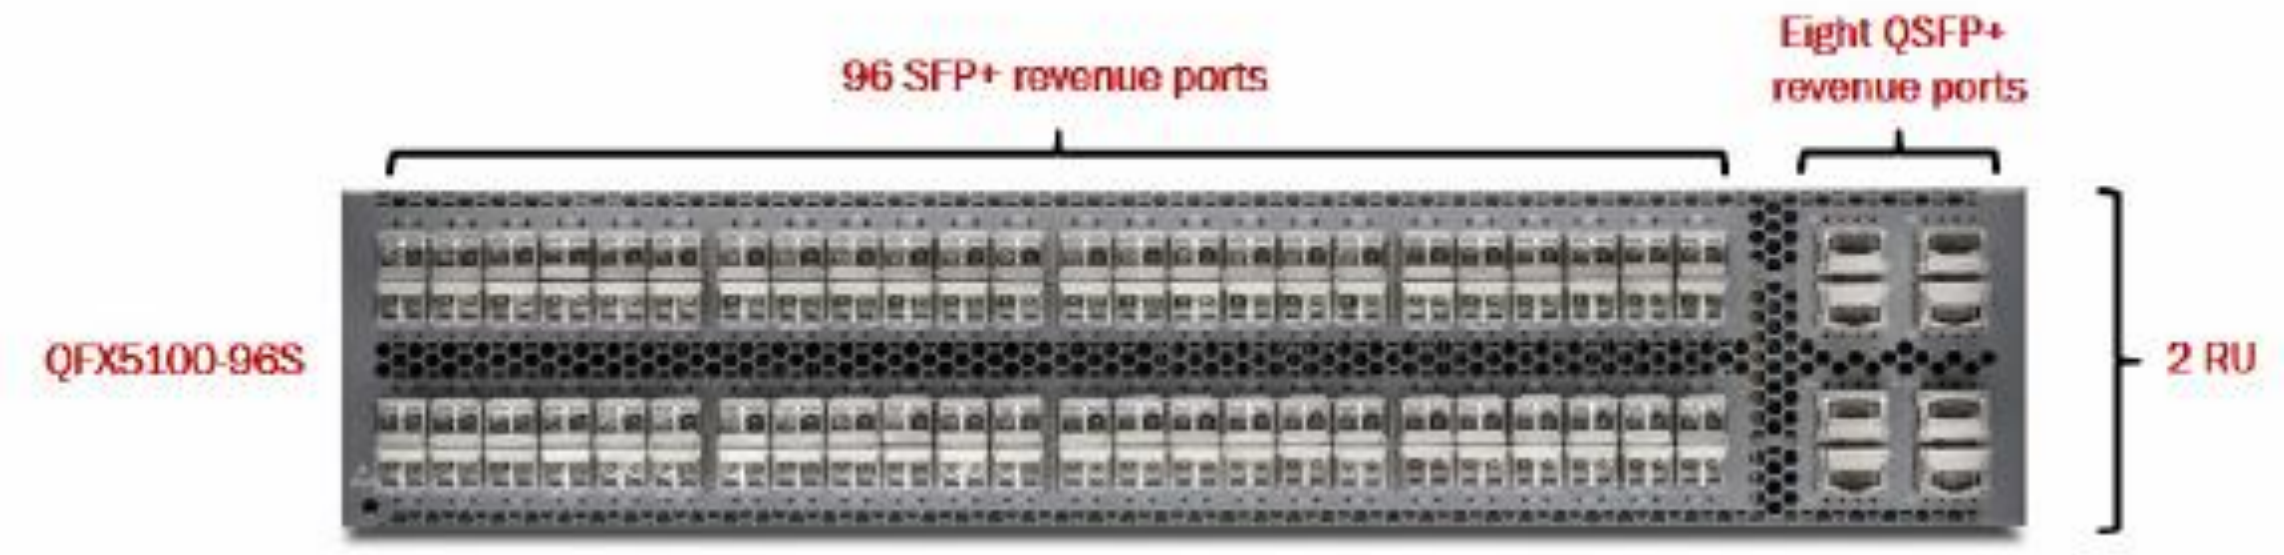

QFX5110

10/40/100Gbe 지원 장비

Spine-and-Leaf, Layer 3 또는 OpenClos IP Fabric

Virtual-Chassis / Virtual-Chassis-Fabric 지원

QFX5110-48S

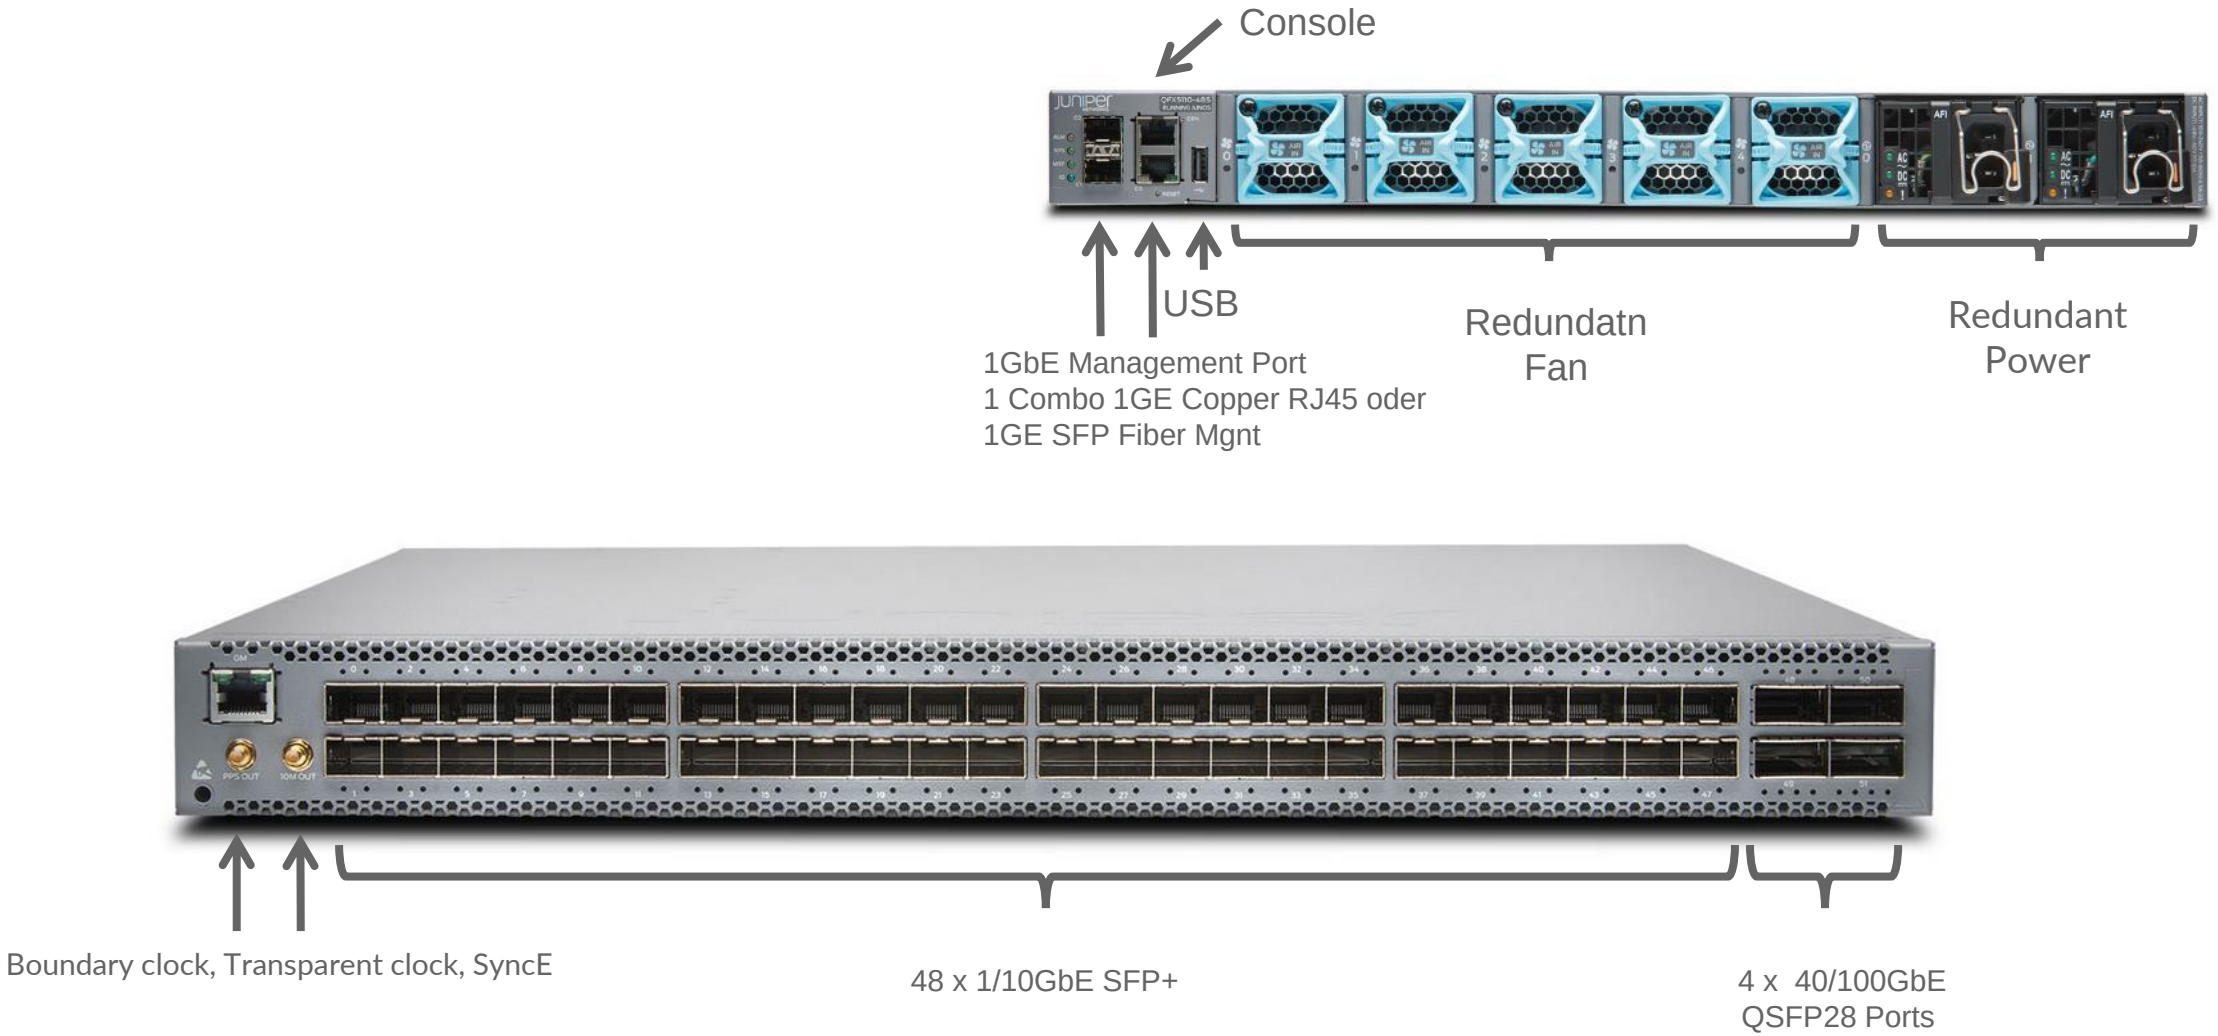

QFX5110-32Q

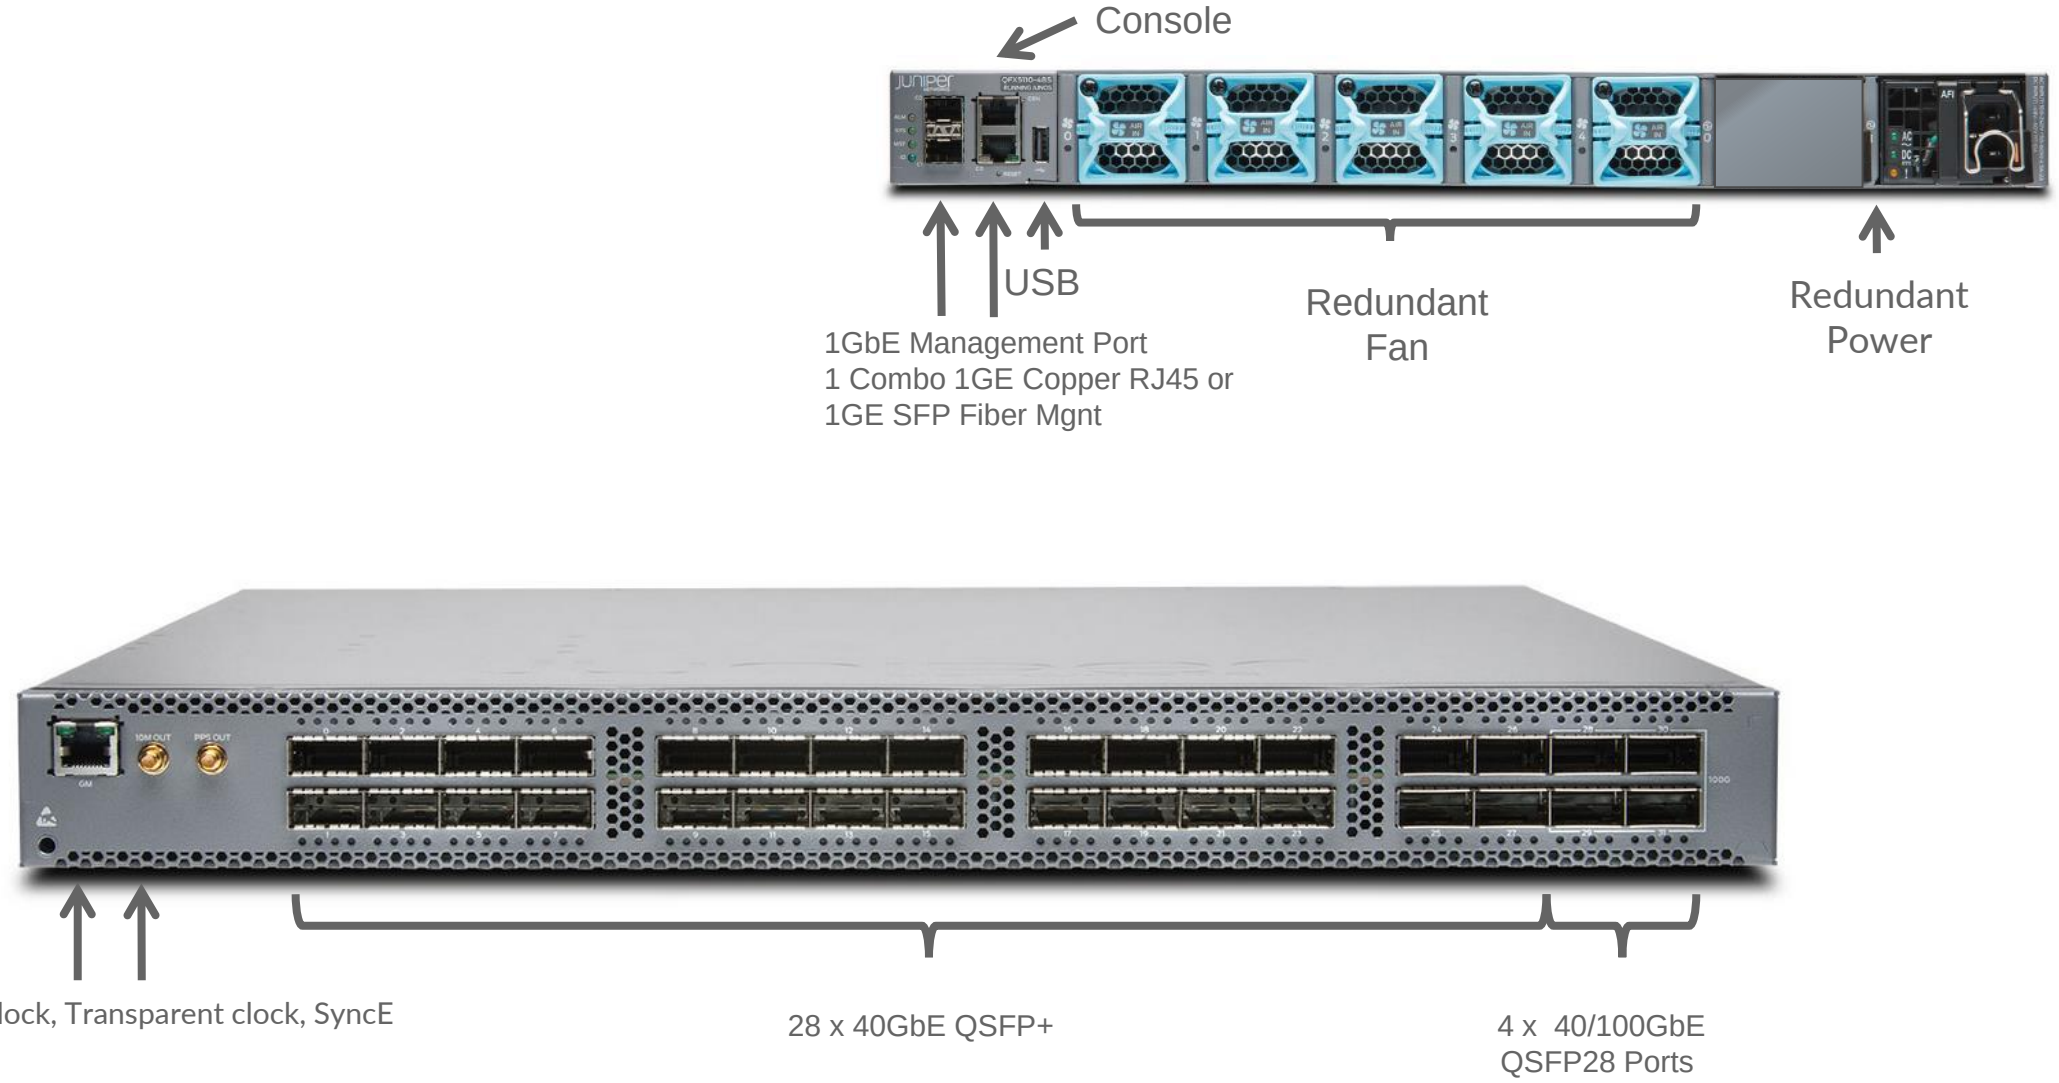

QFX5110-32Q System Modes

- **Default Mode**: All the 32 ports function as 40G only. None of the ports can be channelized in this mode.

- **Flexi-pic Mode** : The ports 0 to 19 come up as 40G and can be channelized as 4x10G. Ports 28 to 31 are 100G only and cannot be channelized. Ports 20 to 27 are rendered inactive and cannot be used.

QFX5110-32Q System Modes

- The CLI to change the system mode is as follows:
  “request chassis system-mode default-mode”
  “request chassis system-mode flexi-pic-mode”
- With the change of system mode, dcpfe process is restarted which causes the ports to go down and come up in the requested system mode.
- The existing mode in the system can be checked with the CLI:
  “show chassis system-mode”

QFX5110-32Q System Modes

```
root> show chassis system-mode
localre:
--------------------------------------------------------------------------
Current System-Mode Configuration:
Flexi-pic mode
{master:0}
root> request chassis system-mode default-mode
System-mode has changed. This will restart PFE.
root> show chassis system-mode
localre:
--------------------------------------------------------------------------
Current System-Mode Configuration:
Default-mode
```

QFX5120

10/25/40/100Gbe 지원 장비

Spine-and-Leaf, Layer 3 또는 OpenClos IP Fabric

Virtual-Chassis / Virtual-Chassis-Fabric 지원

QFX5120-48Y

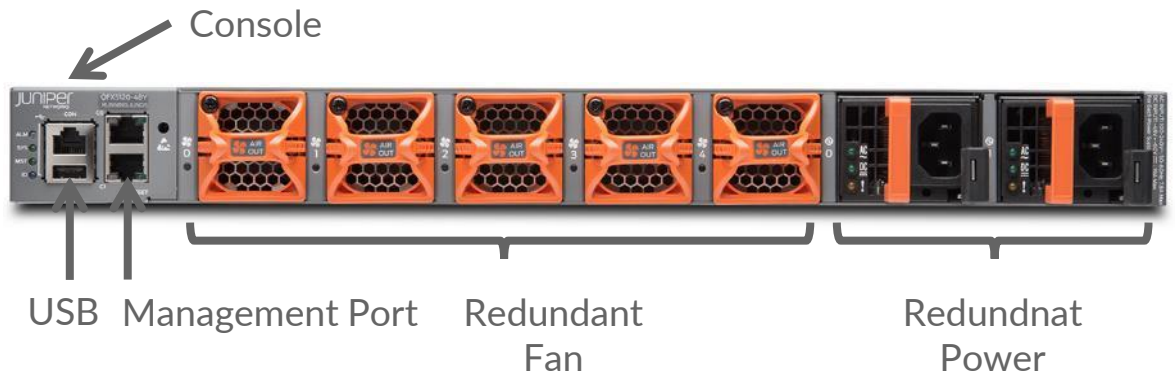

EX4650

- First 48 ports support 1G/10G/25G port speeds
- They are in 10G mode by default – interface will NOT get created automatically on inserting 1G/25G transceivers.
- CLI used to configure the port speed to 1G/25G mode manually

  set chassis fpc 0 pic 0 port <> speed <1G/25g>
- Configuration applies to a quad (set of 4 ports) at a time, and configuration is only allowed on the first port of the quad. For example, in order to configure ports 4-7 in 25G mode, we need to configure below CLI only

  set chassis fpc 0 pic 0 port 4 speed 25g
- 25G/40G/100G interfaces naming convention used is “et”

QFX5120-32C

32 x 40/100Gbe QSFP+ / QSFP28

QFX5120-32C

- First 32 ports support 100G/40G port speeds and 10G/25G/50G channelization speeds
- 25G/40G/100G interfaces naming convention used is “et”

- Limitations:
  - ✓ Port 31 does not support 4x10G and 4x25G channelization mode
  - ✓ 2x50G channelization not supported at FRS

QFX5200

10/25/40/50 and 100GbE 지원

Spine-and-Leaf, Layer 3 또는 OpenClos IP Fabric

중형 및 대형 데이터 센터 용 고밀도 패브릭 설계에 최적화

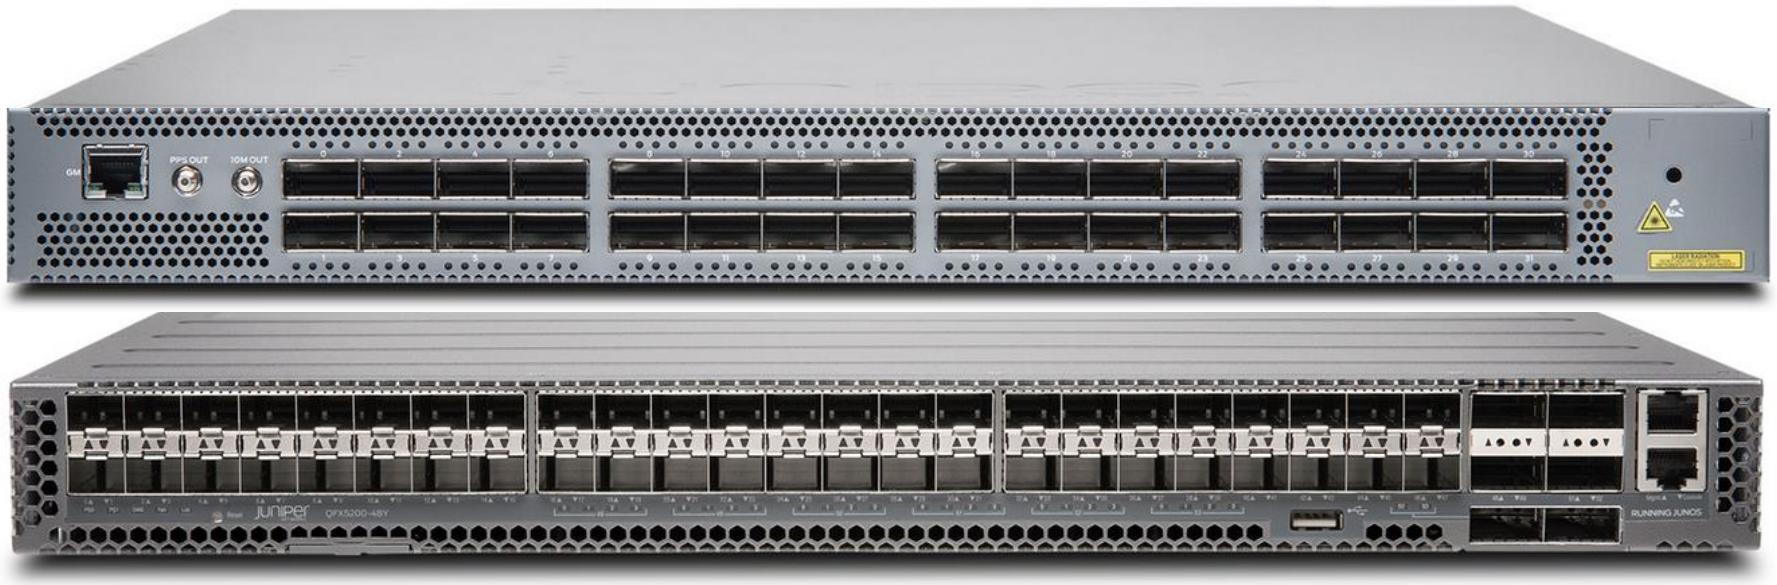

QFX5200-48Y

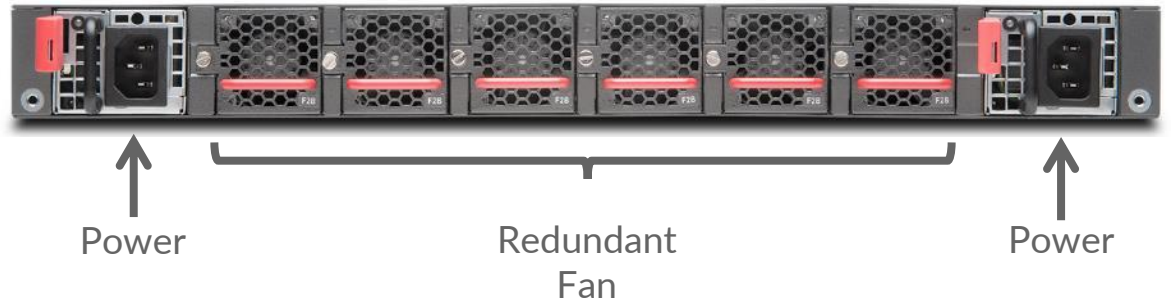

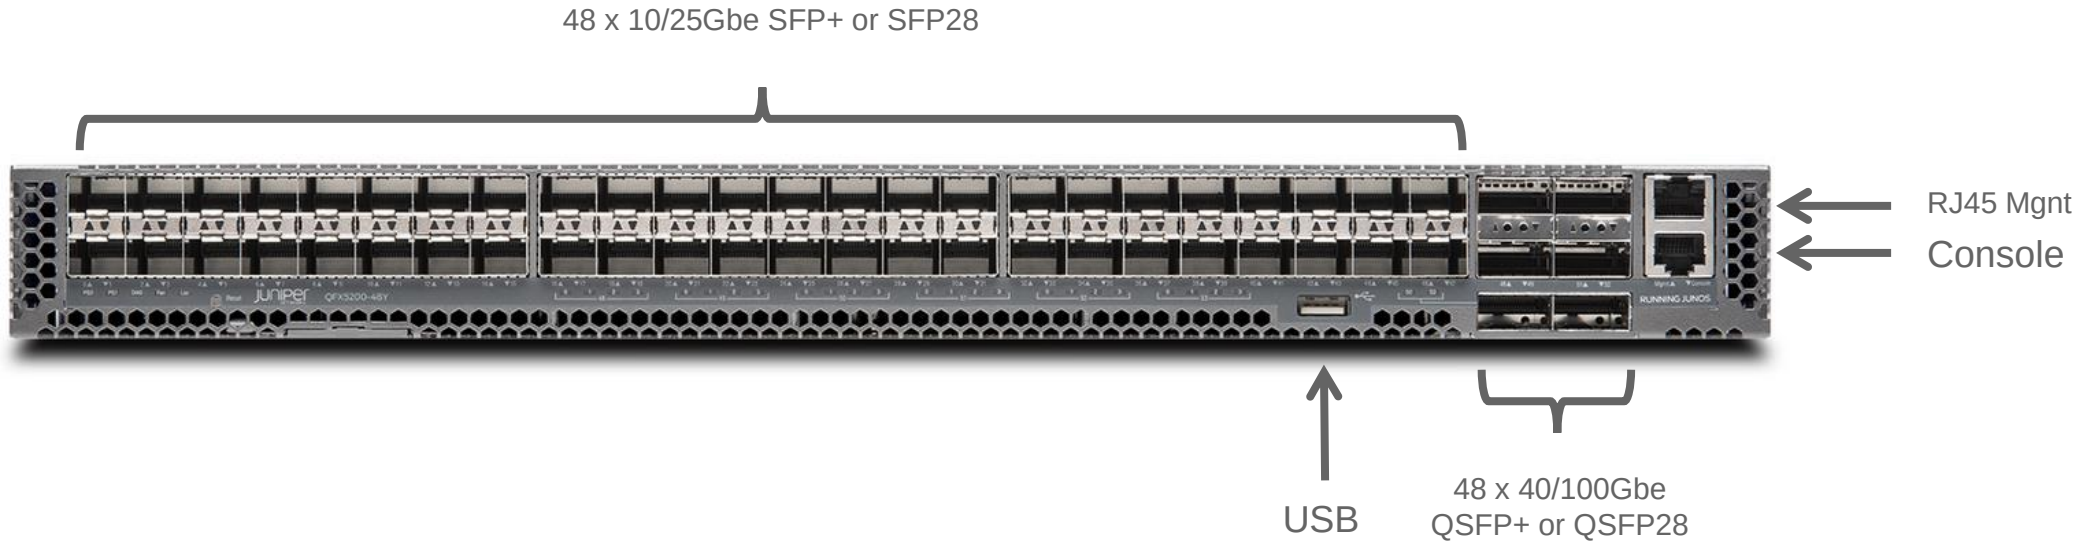

QFX5200-32C

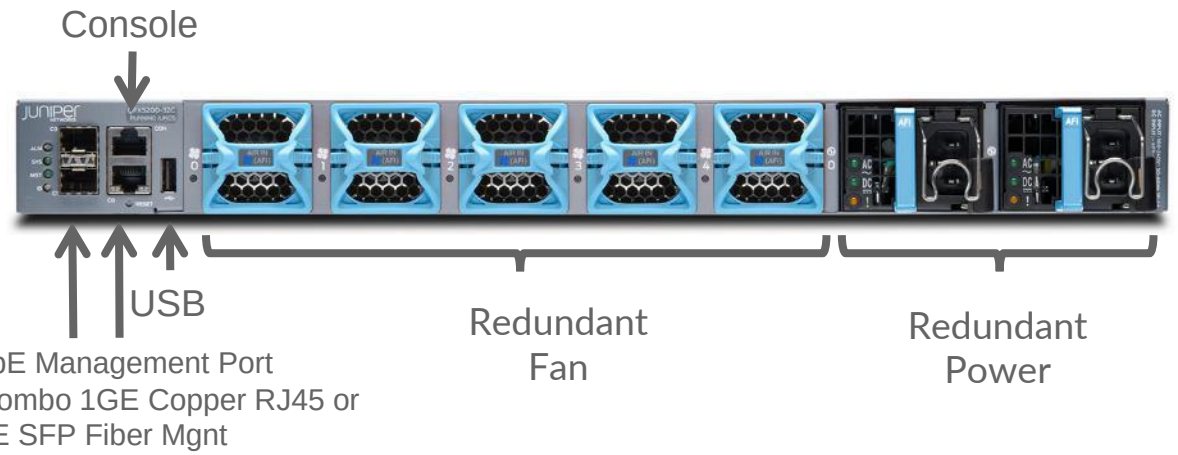

QFX10002-36Q AND QFX10002-72Q

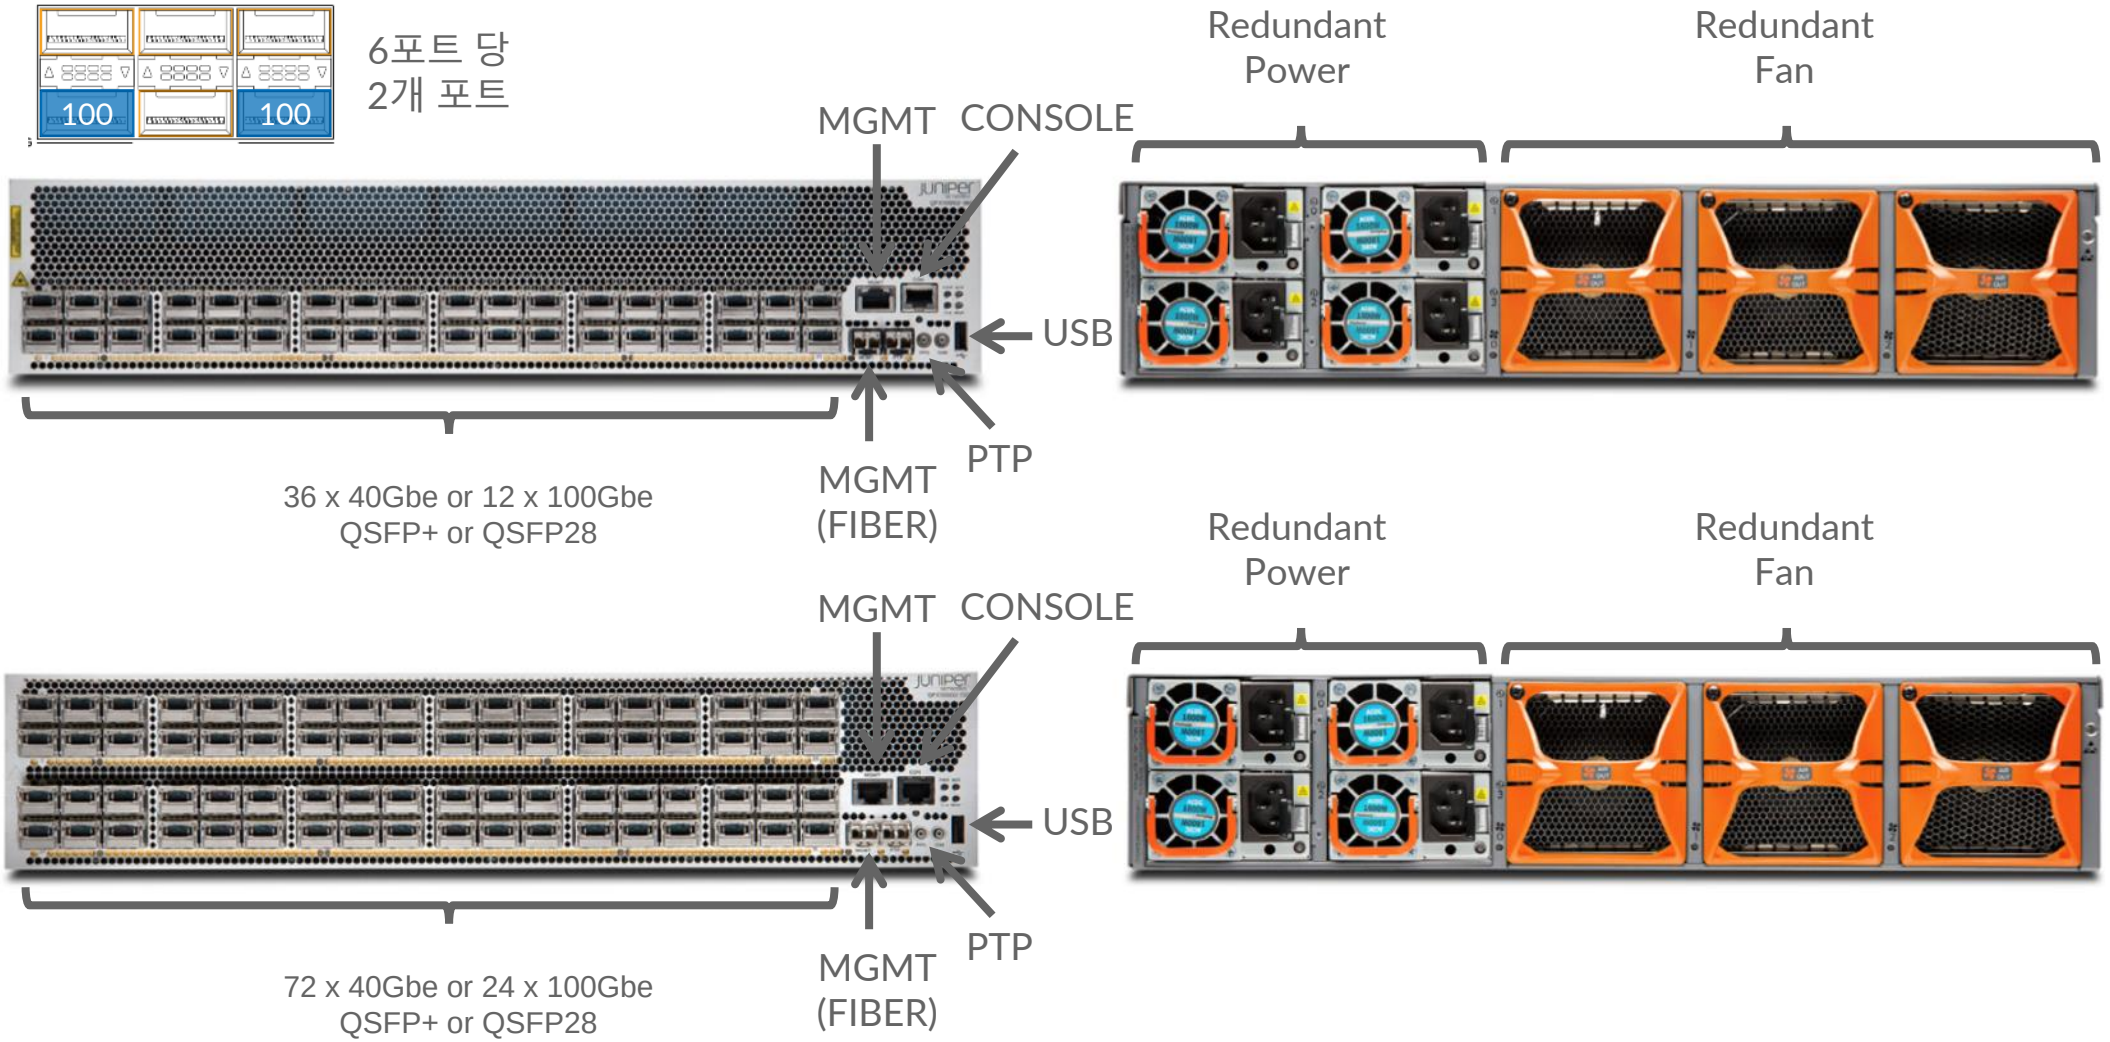

QFX10002-60C

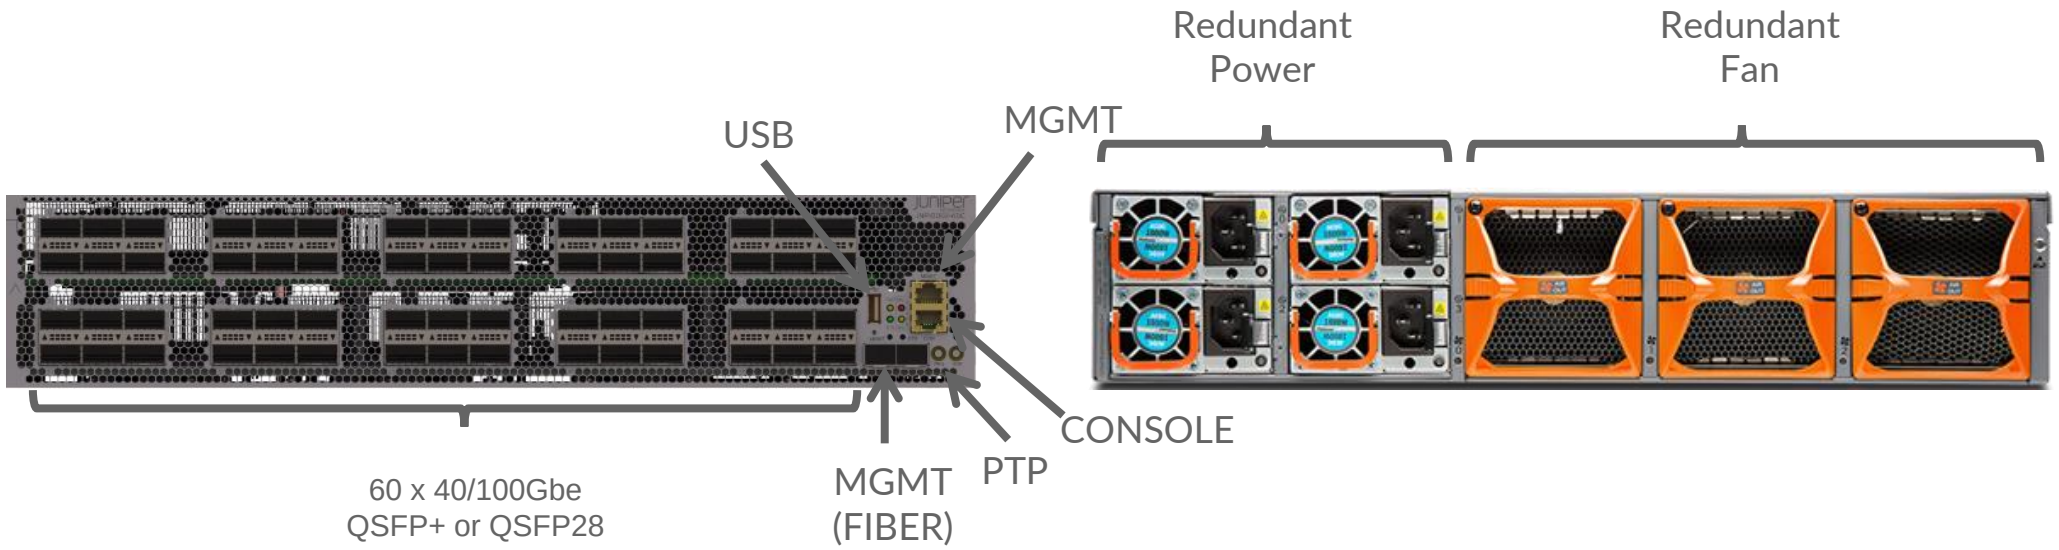

QFX10002-60C

PFE	Ports
0	30,32,34,36,38
1	31,33,35,37,39
2	40,42,44,46,48
3	41,43,45,47,49
4	50,52,54,56,58
5	51,53,55,57,59

PFE	Ports
6	0,2,4,6,8
7	1,3,5,7,9
8	10,12,14,16,18
9	11,13,15,17,19
10	20,22,24,26,28
11	21,23,25,27,29

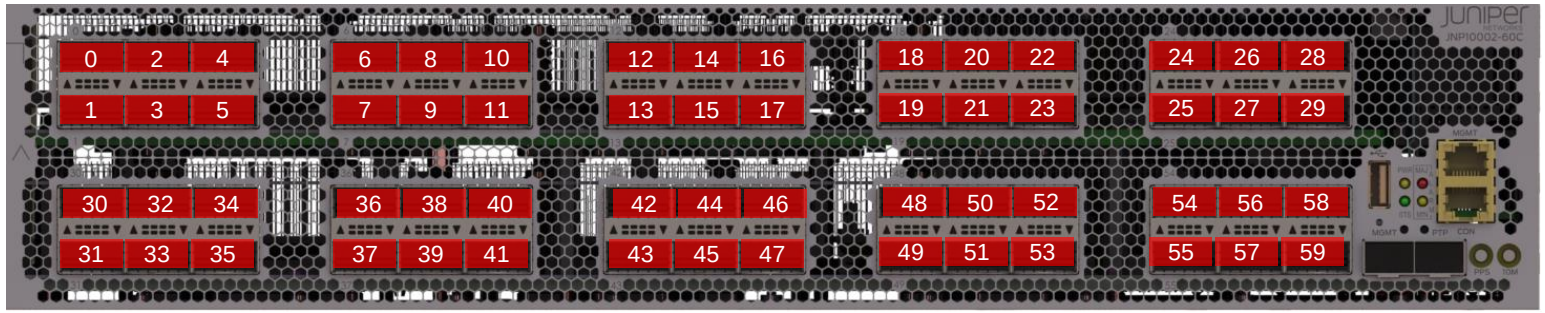

QFX10002-60C

Switches Operates in Non-Channelized (Mode D) and Channelized (Mode A) mode

By default any PE comes up with Mode D

Port Speed	Non-Channelized Mode (Mode D)	Channelized Mode (Mode A)
100 Gbps	60	48
40 Gbps	60	60
10 Gbps	0	192

Port Behavior is tied to the ASIC mode associated with the port

Must configure each port individually in order to channelize a 40G Port

- Any mode changes A to D or D to A at PE, will require cold-reboot of the PFE complex

QFX10002-60C

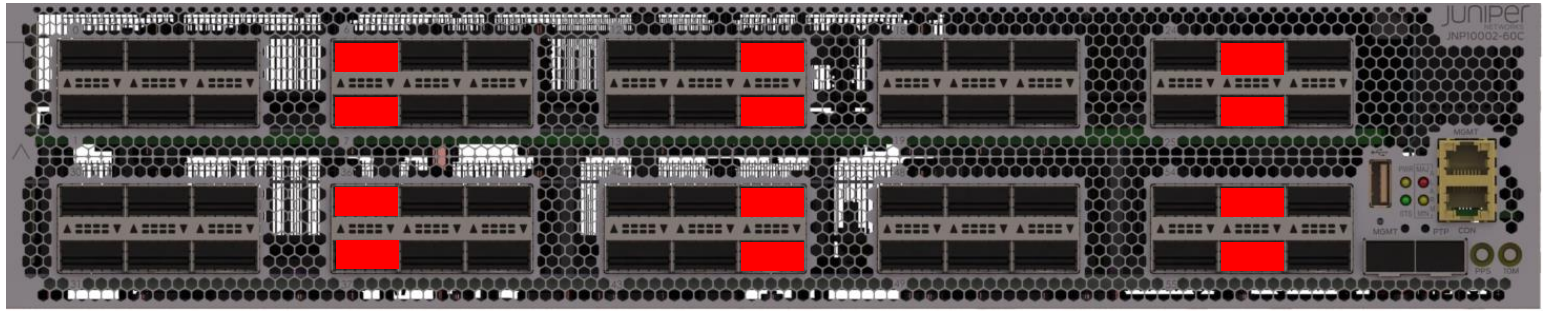

Disabled Ports

- Once PE is put into mode-A, others ports in the same can operate in 40G, 100G or 10G speeds
- Interfaces will come up in 40G or 100G mode depend on the optics plugged in

CLI:

- set chassis fpc <slot> pic <> <type – 100G/other type> pe <> mode <A/D>
- set chassis fpc <slot> pic <> port <> channelization-speed 10g

QFX10002-60C

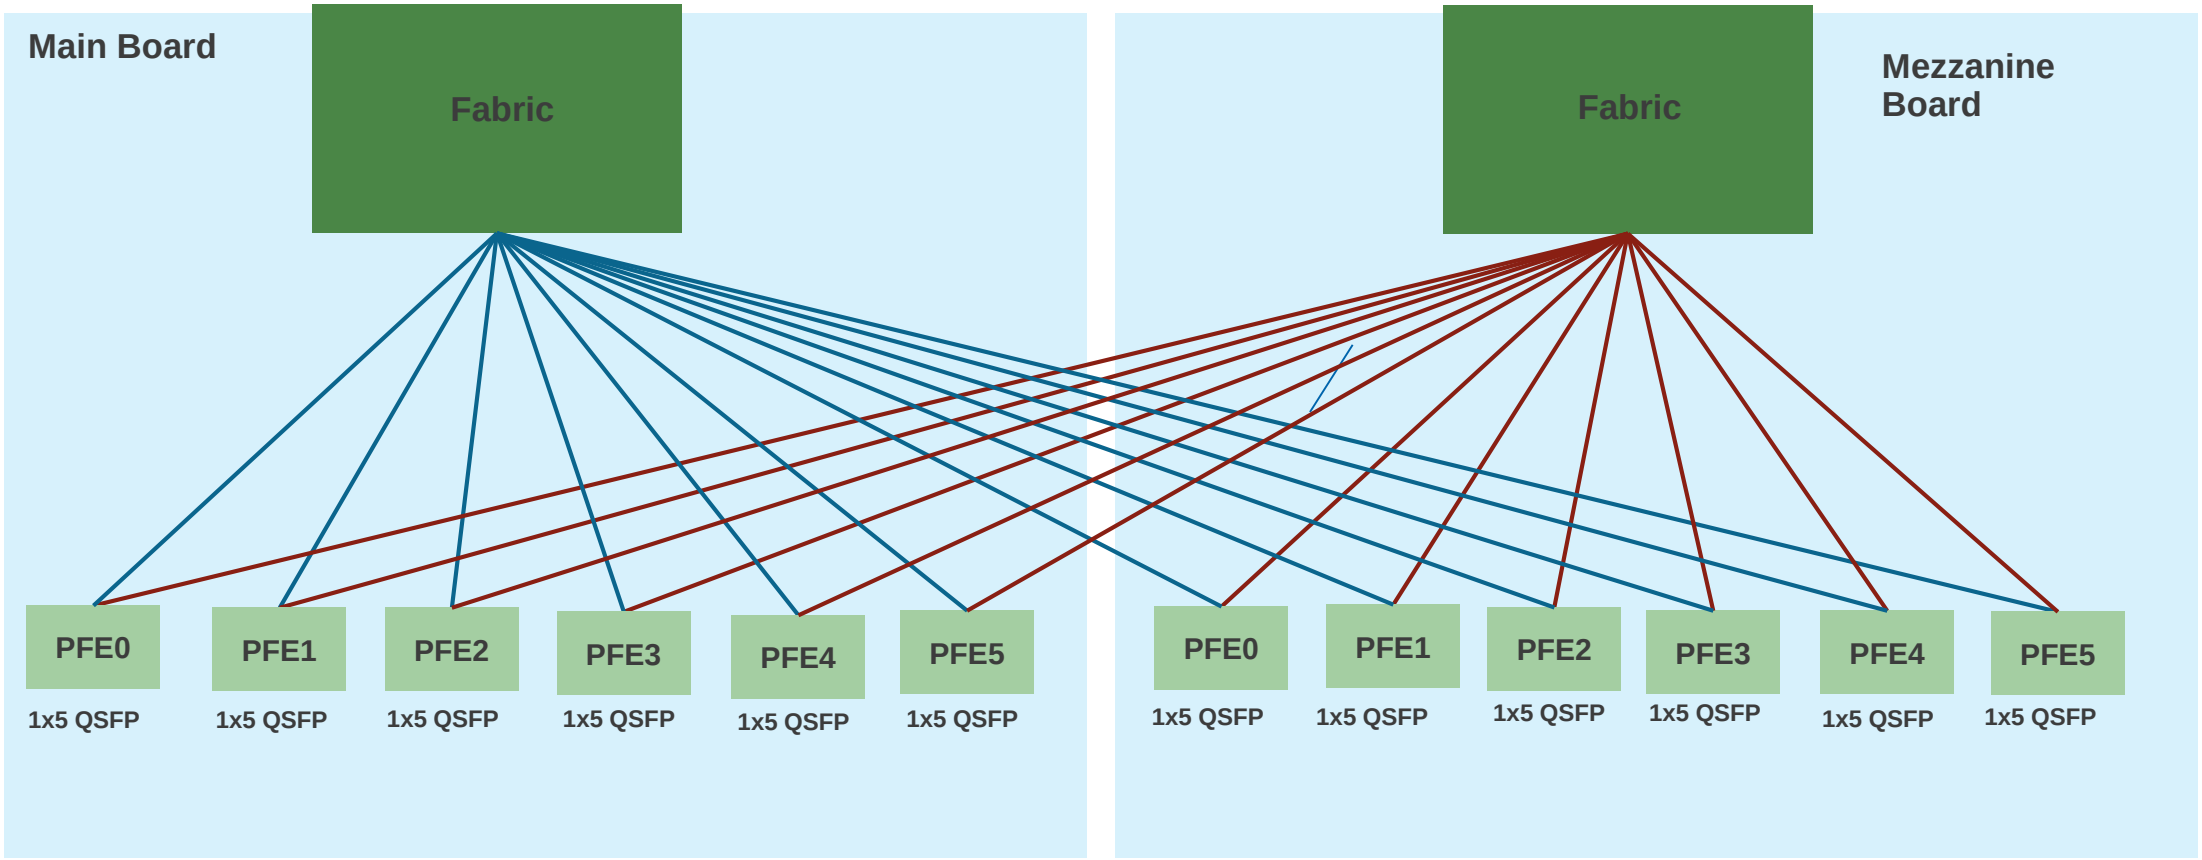

Main Board: Gladden RE, 6x PE and One PF and 30 X QSFP
Mezz Board: 6xPE, one PF and 30 X 100GE QSFP

QFX10002-60C

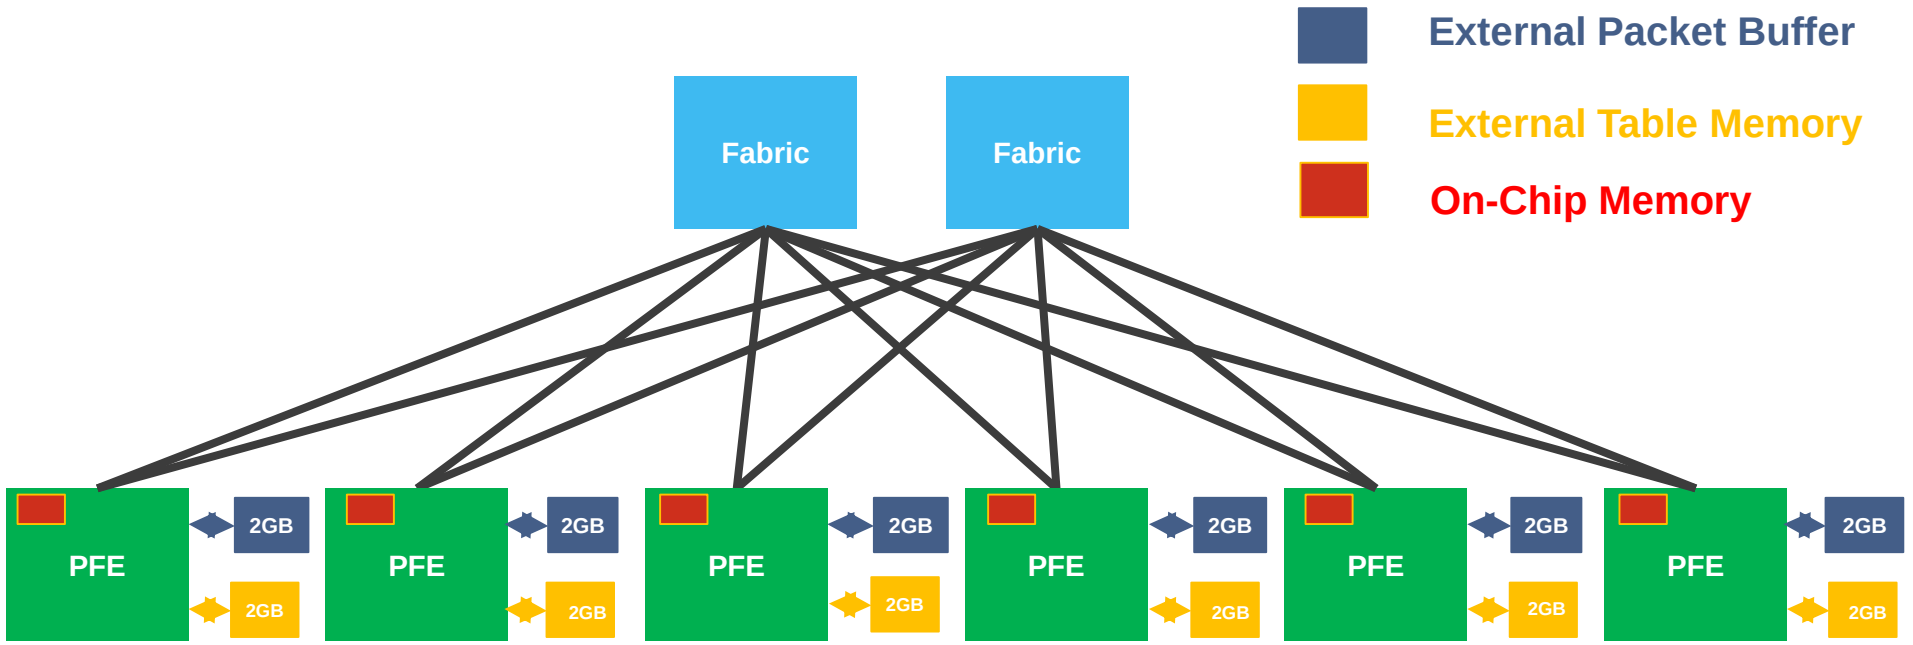

Table Memory Summary

- Each PFE Chip has dedicated memory for storing forwarding tables
- Small on chip memory per PFE for MAC & ACL
- External Memory is 4GB total – 2GB for FIB(External table memory) and 2GB for buffering
- Buffering average reduces to 34ms compared to 67ms on Elit, each port can burst 100ms

QFX10008 AND QFX10016

ROUTING ENGINE (1+1)

Linecard Slots (16)

ROUTING ENGINE (1+1)

Linecard Slots (8)

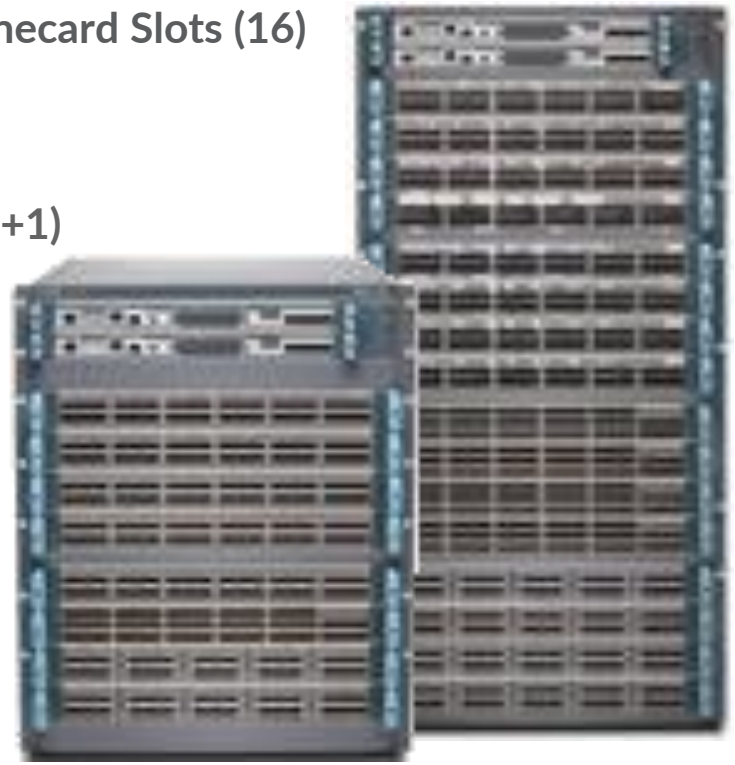

System Feature

- MC-LAG
- Enabling Cloud Infrastructure
- Juniper Q5 Silicon
- 1/10/40/100G Interface

Key Differentiations

- High Multi-dimensional scale and features
- 4M FIB capable
- Up to 100ms Buffer per Port
- High Logical Scale

QFX10008 AND QFX10016

QFX10008

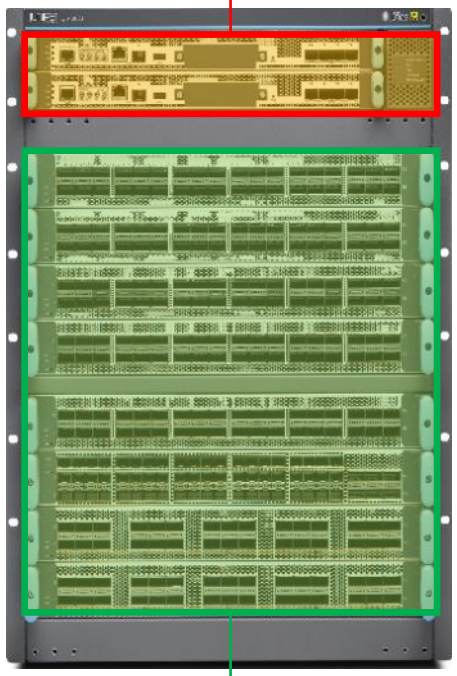

ROUTING ENGINE (1+1)

6 FABRIC CARDS

6 POWER SUPPLY

MID Plane이 없는 구조

LINE CARDS
60P 1/10G + 6P 40G
36P 40G OR 12P 100G
30P 40/100G

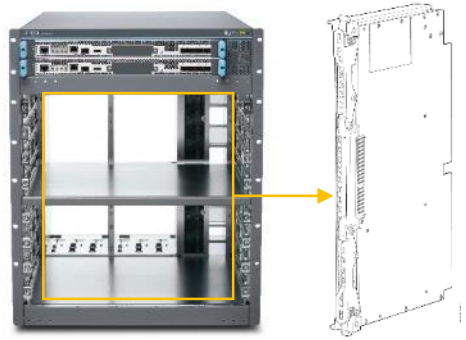

- Up to 6 SIB (5+1)
- Up to 42Tbps
- Hot-removable and hot-insertable
- Contains Two Paradise Fabric (PF) ASICs
- I2C bus from CB to SIB for basic card control

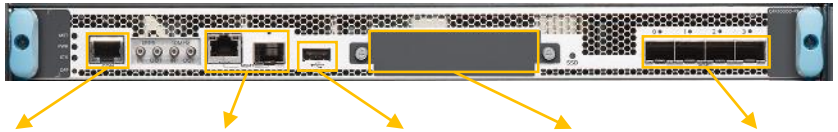

- Up to 32G DRAM, 50G SSD + 50G SSD (Extra Slot)
- 4 core, 2.5 GHz Intel processor
- Control Plane 연동을 위한 4 Port 10G Ethernet

QFX10008 AND QFX10016

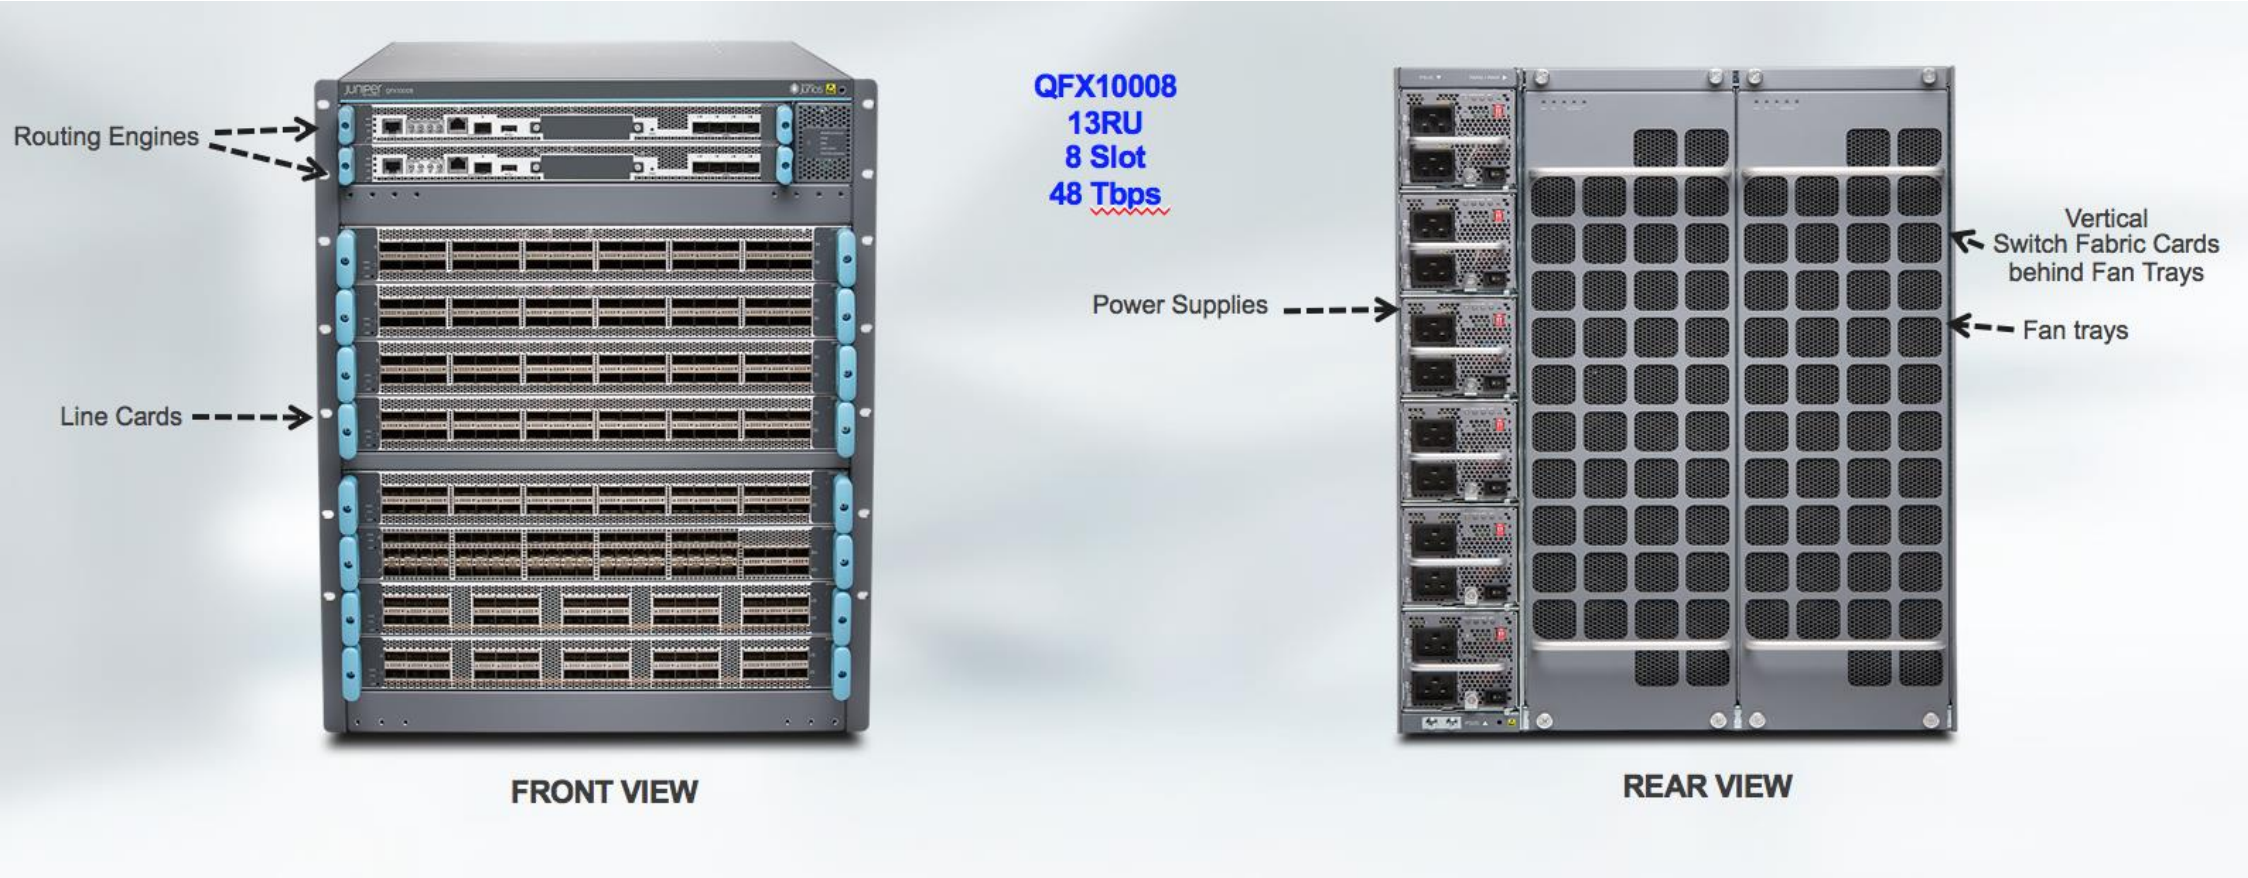

QFX10008 AND QFX10016

QFX10008 AC PSU

- 220V Dual Input, 12V DC Output
- 2850W Output Power
- 80% Minimum Efficiency (Gold)
- All PSUs share the load
- PSU includes internal Fan for cooling

QFX10008 DC PSU

- 40-72VDC Input Range, Nominal 48V DC Input, 12V DC Output
- 2650W Output Power
- 91% Minimum Efficiency
- All PSUs share the load
- PSU includes internal Fan for cooling

QFX10008 Power Distribution

- 12V Power is Distributed to all FRUs in the chassis
- Copper Bus Bars used for distribution

QFX10008 AND QFX10016

QFX10008 Cooling Subsystem

- Two Fan Trays, each with 11 Fans
- Two Fan Tray Controllers, each driving one Fan Tray
- System can operate continuously in case of a single fan failure (Minor Alarm)
- System can operate for a limited time in case of a failure of a fan tray (Major Alarm condition)
- Fan Speed is controlled by JunOS through I2C interface between Control board and Fan Tray Controller
- Fan Power split into 4 rails to avoid single point of failure
- Each Fan Tray has LEDs for indicating status of that Fan Tray and corresponding Fan Tray Controller
- Each Fan Tray also has LEDs indicating status of Three SIBs behind that Fan Tray

QFX10008 AND QFX10016

QFX10008 AND QFX10016

- Easy migration to the next level of performance without having to upgrade the chassis to accommodate multi-generational I/O card upgrades.
- Allows **direct connection between the line cards and fabric cards.**
- Significantly reduces overall system power consumption and optimizes cooling and airflow efficiency through the system while providing an easy upgrade path to next-generation interconnect technology.

QFX10008 AND QFX10016

QFX10008 CB Features

- Intel IVY Bridge 4 core 2.5GHz CPU
- 32GB DDR3 SDRAM – Four 8GB DIMMs
- 50 GB Internal SSD Storage
- One 2.5” External SSD slot
- 10GB Ethernet Switch for Control Plane connectivity with Line Cards
- PCI Express Switch for Control Plane Connectivity with SIBs
- I2C bus segments from CB FPGA to all FRUs
- RS232 Console Port & USB Port
- RJ45 and SFP Management Ethernet
- Includes PTP logic, and SMB connectors for PTP

QFX10008 AND QFX10016

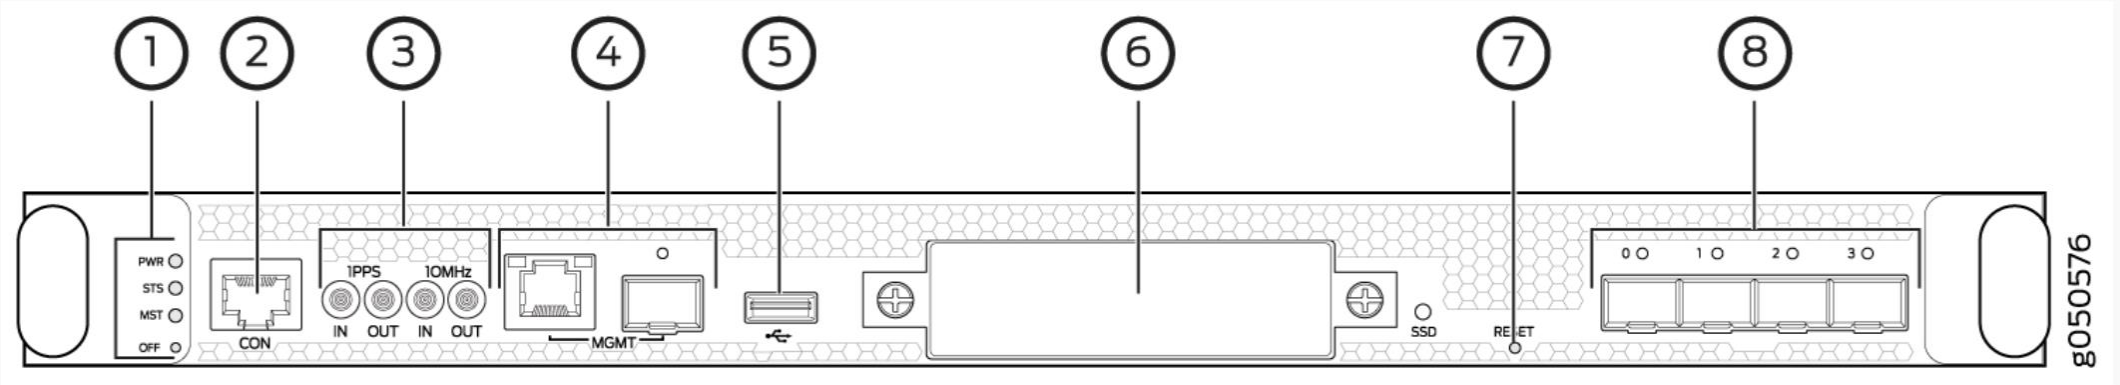

1—Control Board status LEDs	5—USB 2.0 port
2—Console port	6—SATA SSD slot
3—PTP-capable connections: SMB In, SMB Out, 10 MHz In, 10 MHz Out	7—Reset button
4—Ethernet management ports: RJ-45 for 10/100/1000 BASE-T and small-form factor pluggable (SFP) for fiber. If both copper and fiber cables are installed, the RJ-45 is the default.	8—Four SFP+ ports (reserved for future use)

QFX10008 AND QFX10016

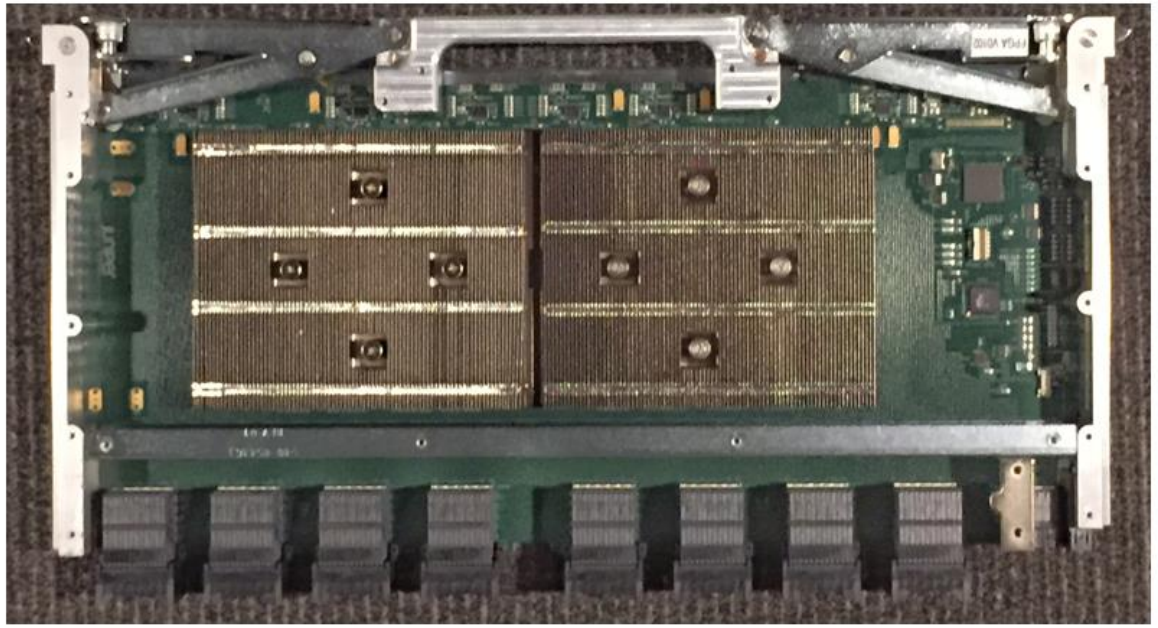

QFX10008 SIB Features

- Contains Two Paradise Fabric (PF) ASICs
- Provides Data Path Connection to all Eight line cards in the chassis
- PCI Express control path from CB to SIB
- I2C bus from CB to SIB for basic card control
- Ejector levers designed to handle large insertion/extraction force

- Creates the switch fabric for the QFX10008.
- Six SIBs that are installed vertically, mid-chassis, between the line cards and the control boards in the front and the fan trays in the rear.
- Five SIBs are required for operation with the sixth providing N+1 redundancy.
- Each SIB has eight connectors that match and connect to a connector on one of the eight line cards.
- Hot-removable and hot-insertable field-replaceable units (FRUs).
- They are not visible from the outside of the switch chassis. You must remove one of the fan trays in order to view the SIBs. The SIBs are numbered from left to right SIB0 to SIB5.

QFX10008 AND QFX10016

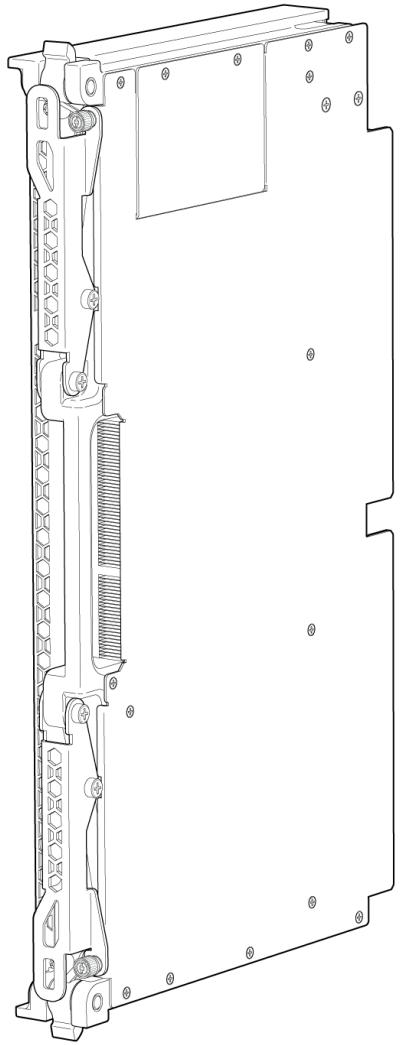

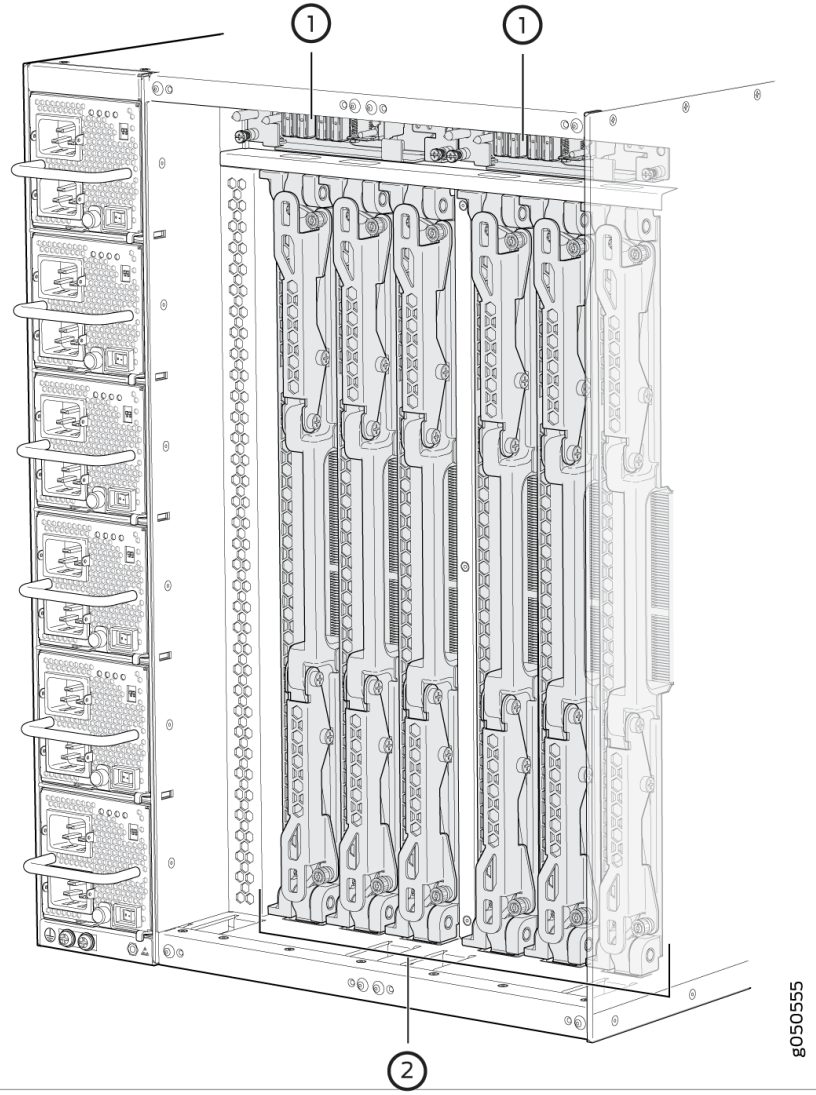

QFX10008 AND QFX10016

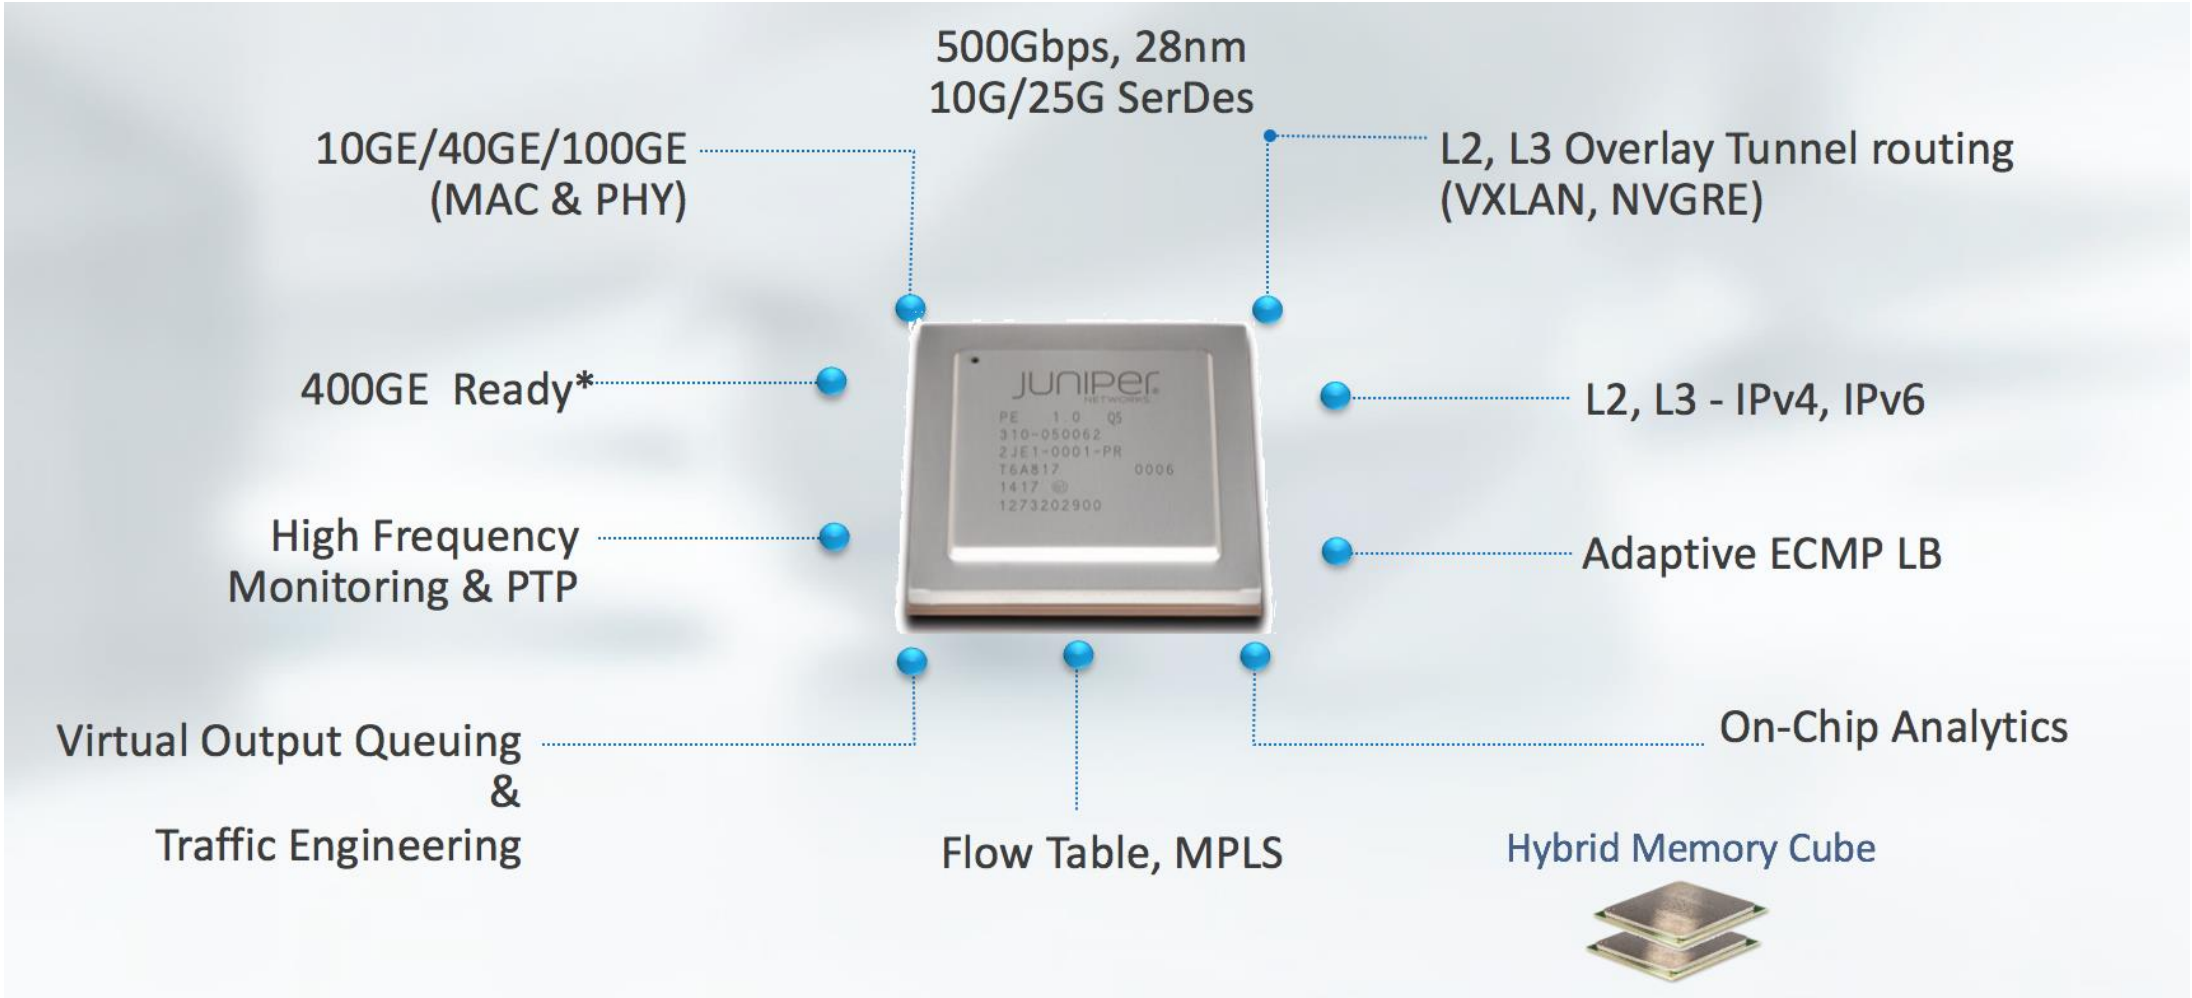

QFX10008 AND QFX10016

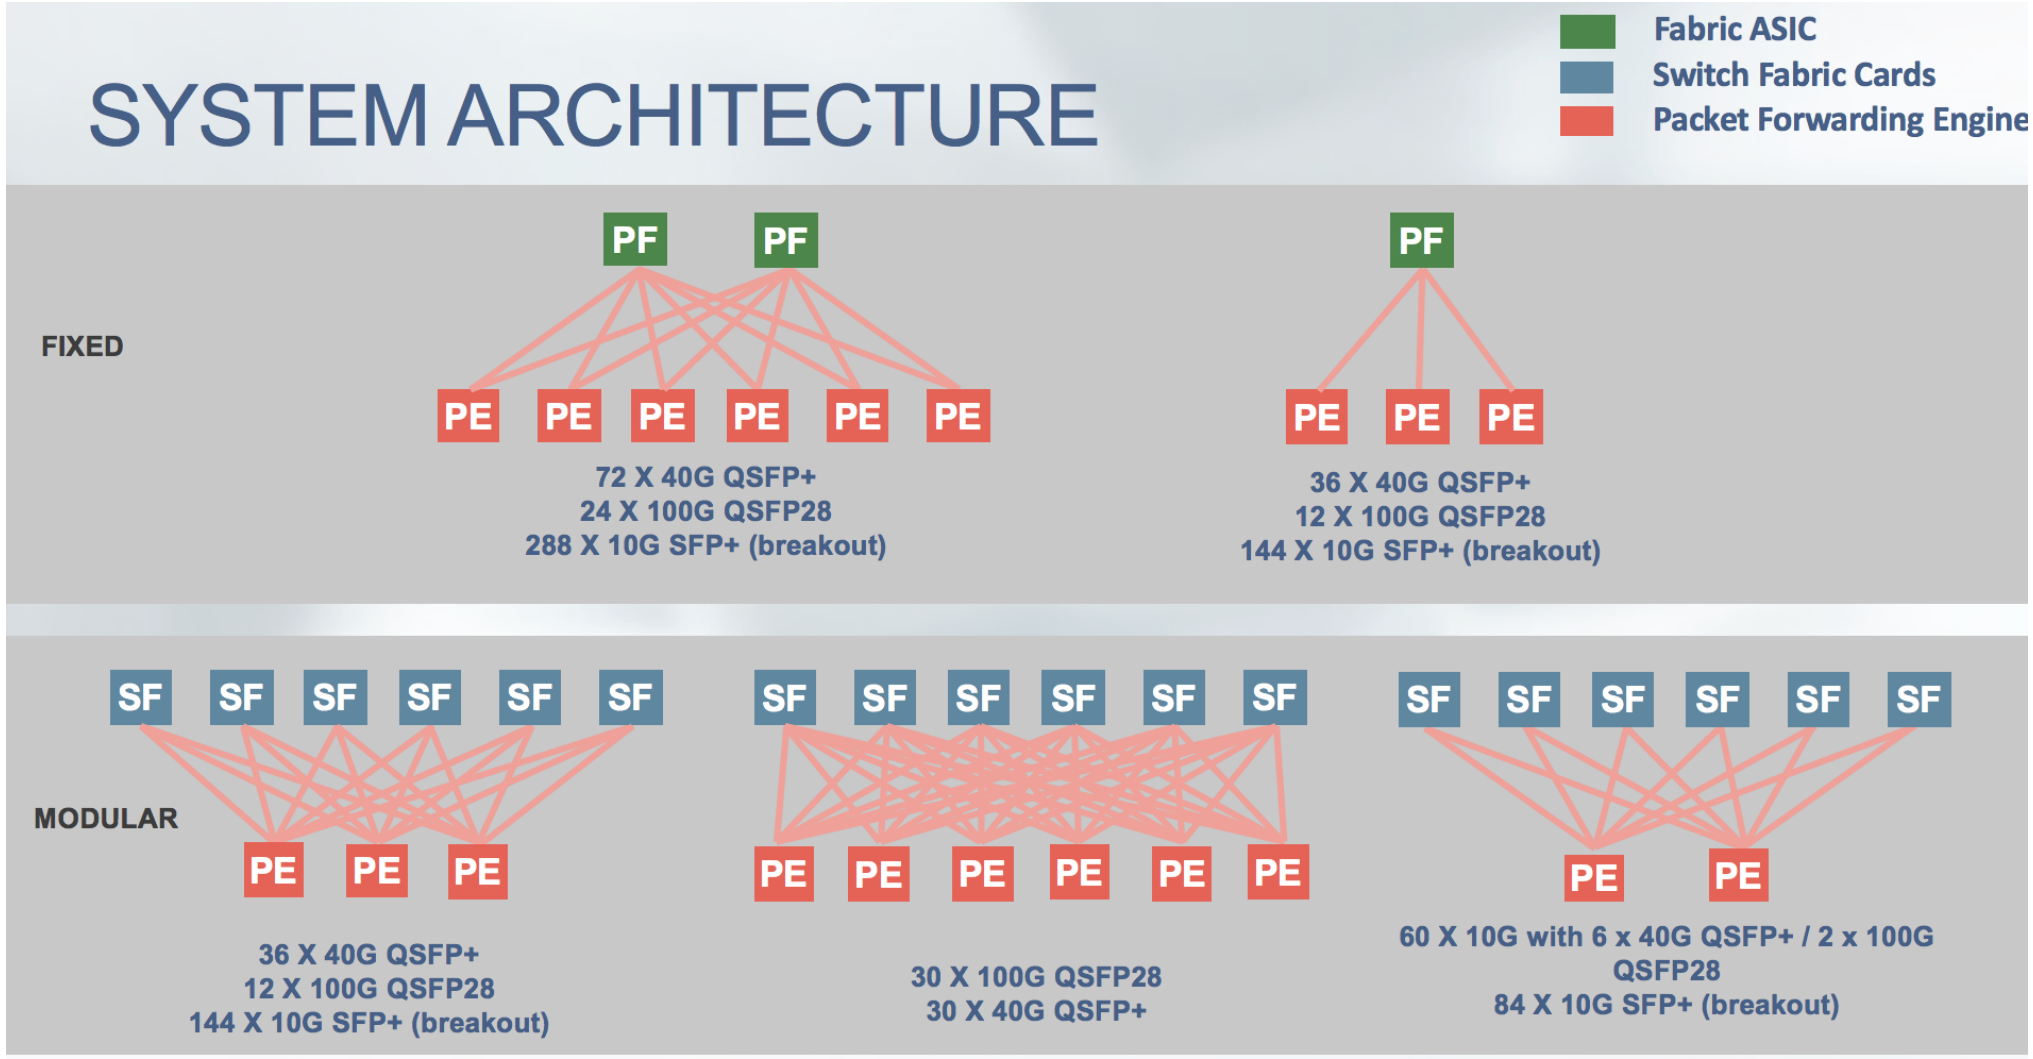

QFX10008 AND QFX10016

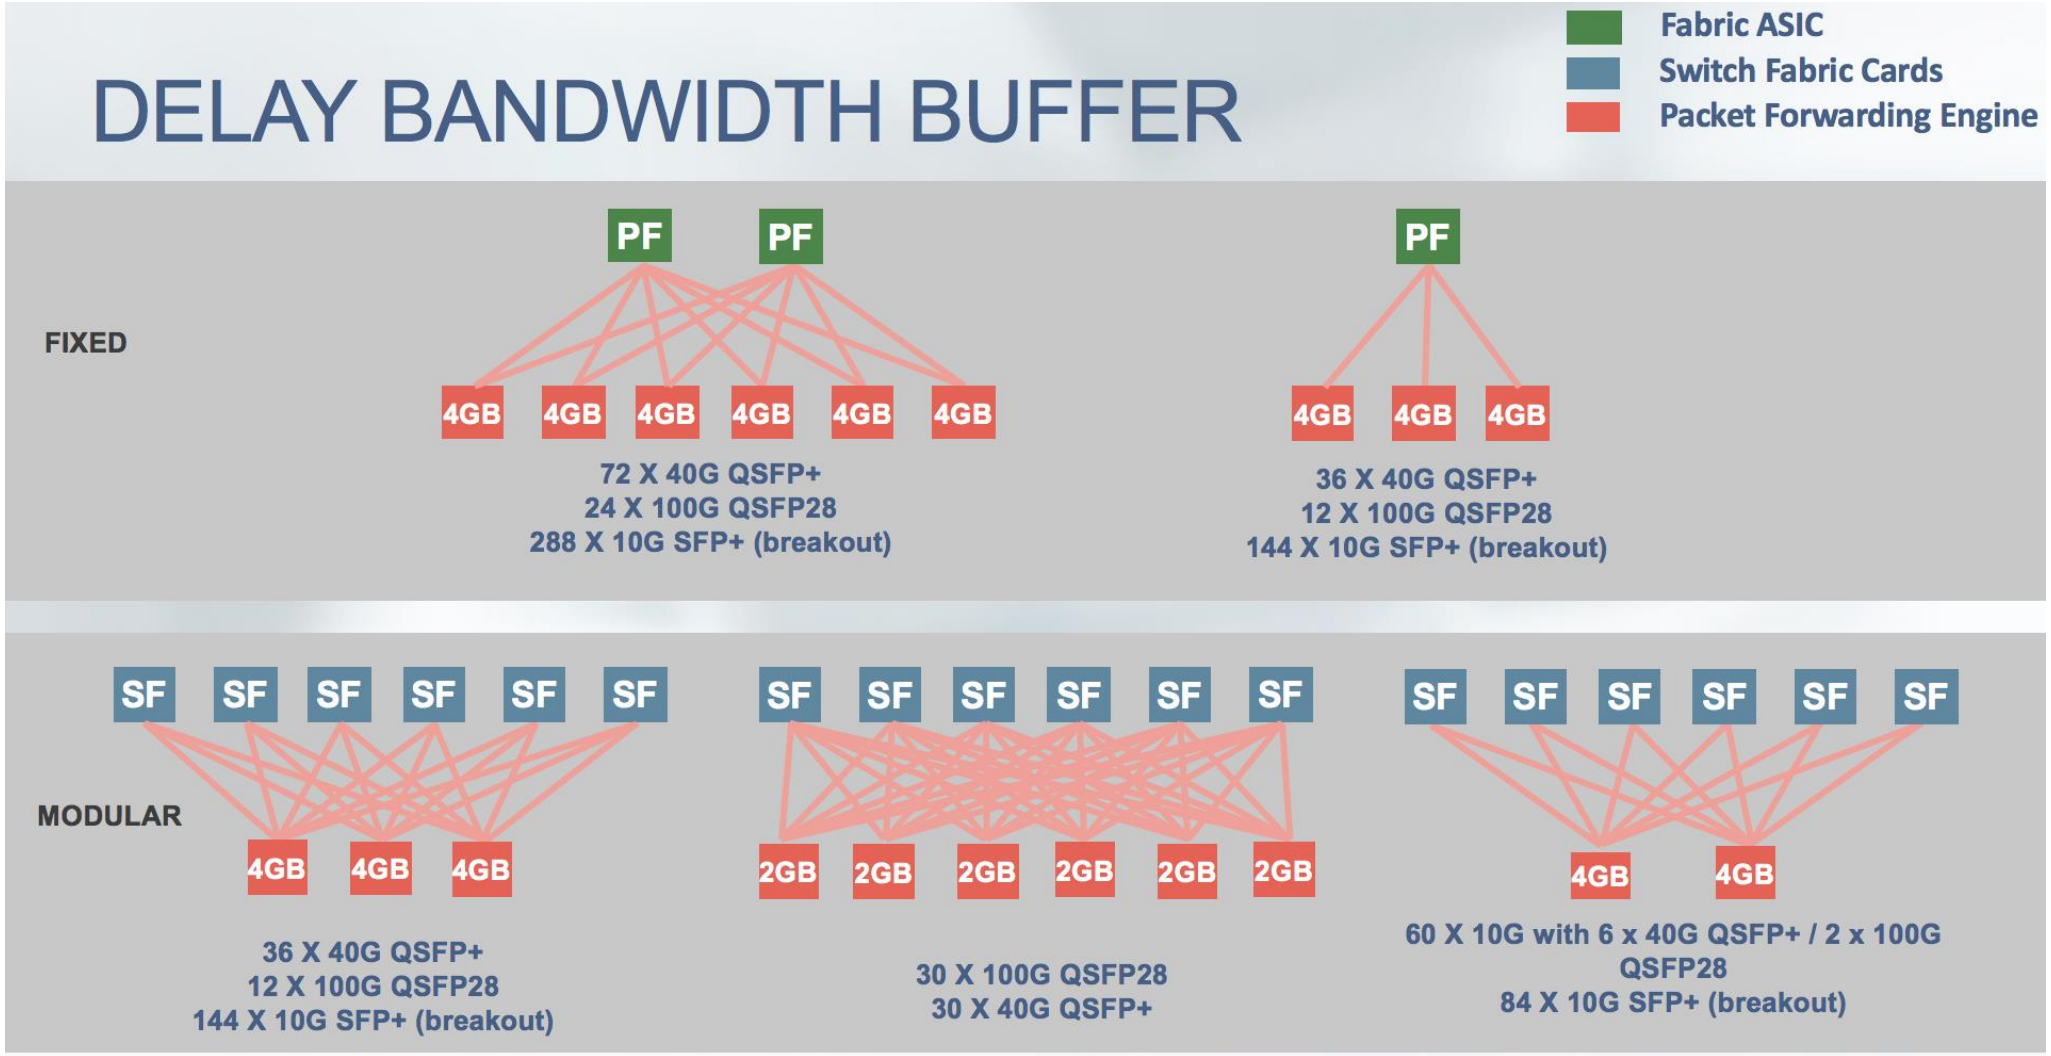

QFX10008 AND QFX10016

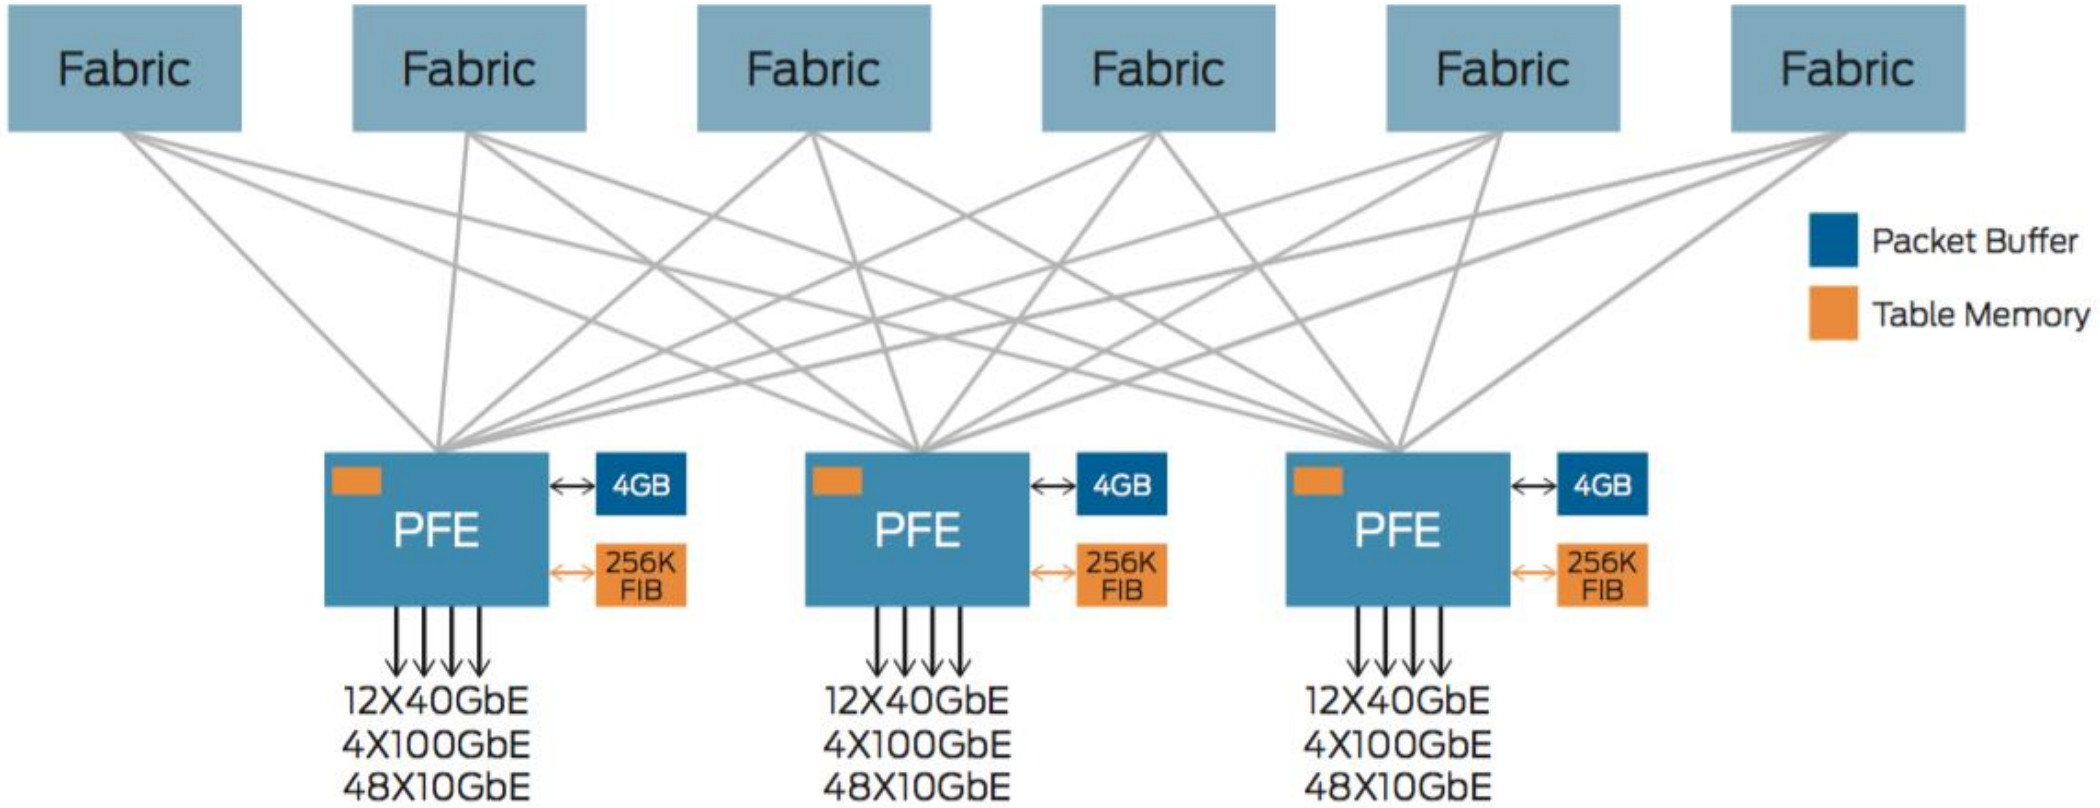

Figure 13: System architecture of the QFX10000-36Q I/O card for modular QFX10000 switches

AGENDA

- EX Series
- EX Architecture
- QFX Series
- QFX Software Architecture

ADVANCED JUNOS SOFTWARE ARCHITECTURE

Provides the foundation for advanced functions

- ISSU (In-Service Software Upgrade)
- Other Juniper applications for additional service in a single switch
- Third-party application
- Can bring up the system much faster

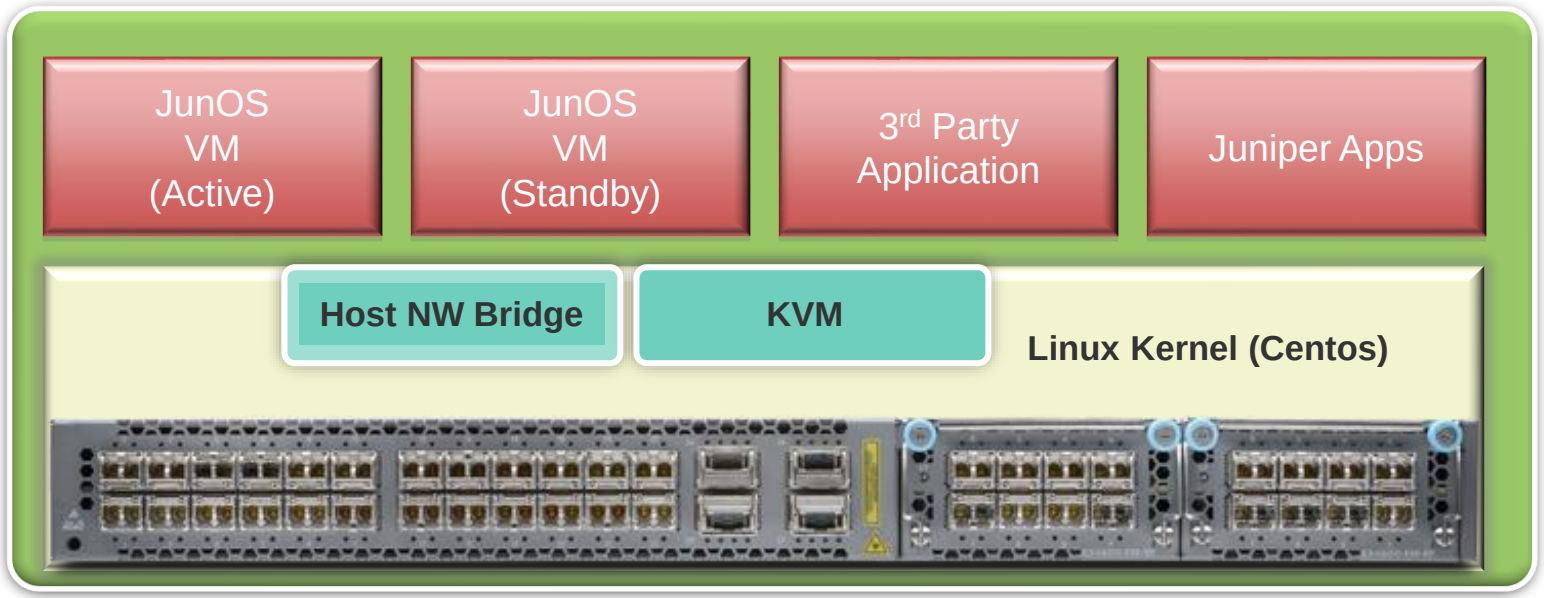

ADVANCED JUNOS SOFTWARE ARCHITECTURE

ISSU - In-service-software-upgrade

- Master Junos VM controls the hardware– PFE and FRU on the system
- Master issues upgrade command
- System launches a new Junos VM with new image as backup
- All states are synchronized to the new backup Junos
- Detach PFE from current master, then attach to backup Junos (hot move)
- The PFE control component in new master will control the forwarding
- Stop the new backup VM

*ISSU will be available post-FRS

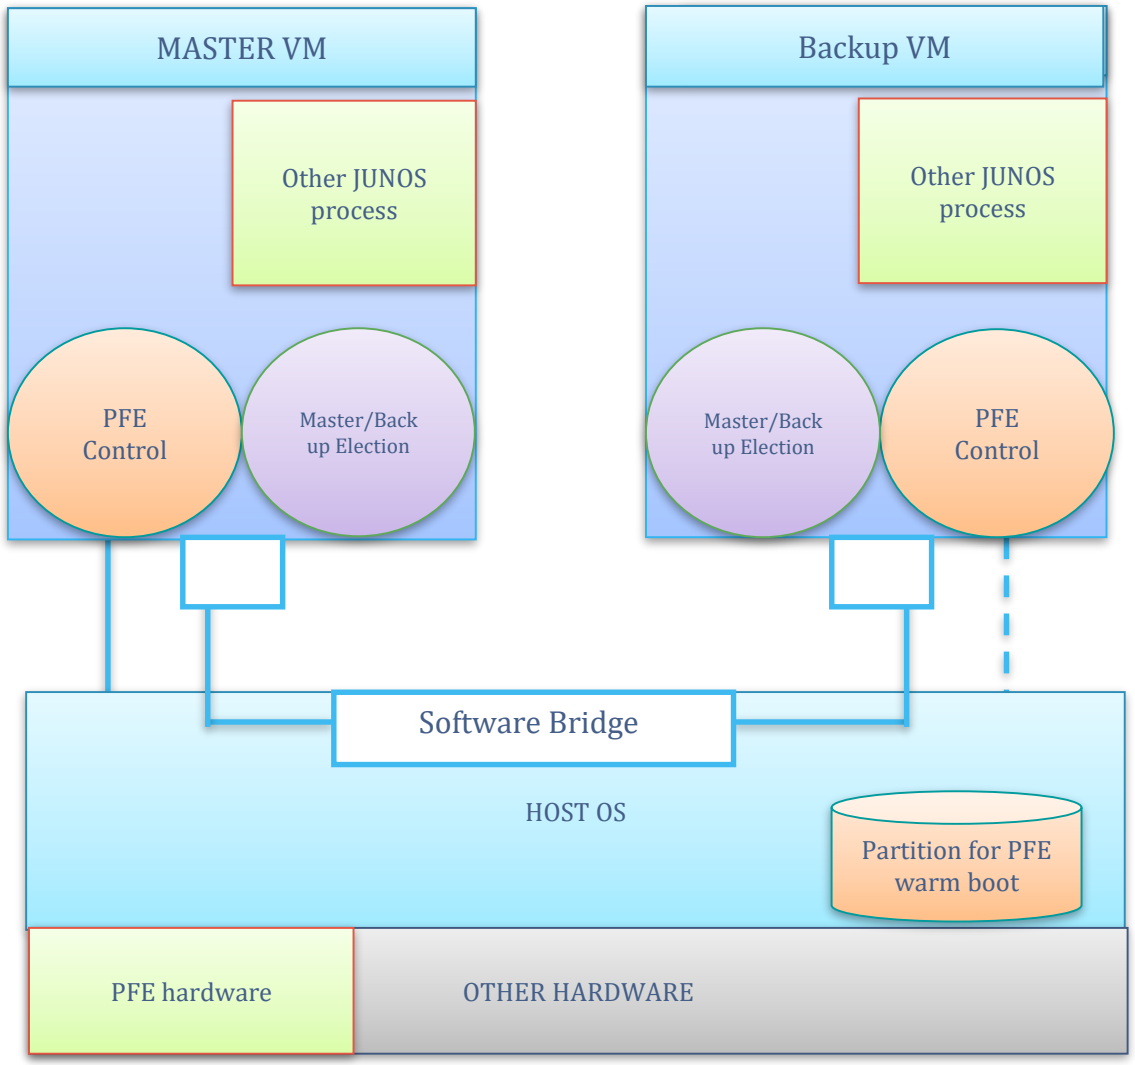


THANK YOU
JUNIPER NETWORKS
Engineering Simplicity