Cisco
Secure Firewall
Platforms Deep Dive

Łukasz Bromirski

BRKSEC-2239

mr0vka@infosec.exchange

lukasz.bromirski.net

cisco Live!

Your Speaker

- CCIE #15929 (R&S/SP) & CCDE #2012::17
- running community projects:
  BGP Blackholing PL, BGP Free Full Feed,
  AS 112 cluster in Poland
- Co-founder of PLNOG and FreeBSD advocate
- MANRS Training Fellow
- https://lukasz.bromirski.net/
- Leading **Firewall Platform Team** at
  Cisco Security Business Group

CISCO CERTIFIED CCDE DESIGN

CISCO CERTIFIED CCIE SERVICE PROVIDER

CISCO CERTIFIED CCIE ENTERPRISE

cisco Live!

Agenda

- Cisco Secure Firewall platforms review
- Design considerations
  - Throughput
  - Scale
  - High Availability
  - Multi-Tenancy
  - Internet Edge
- Q&A

Cisco Secure Firewall
Platforms

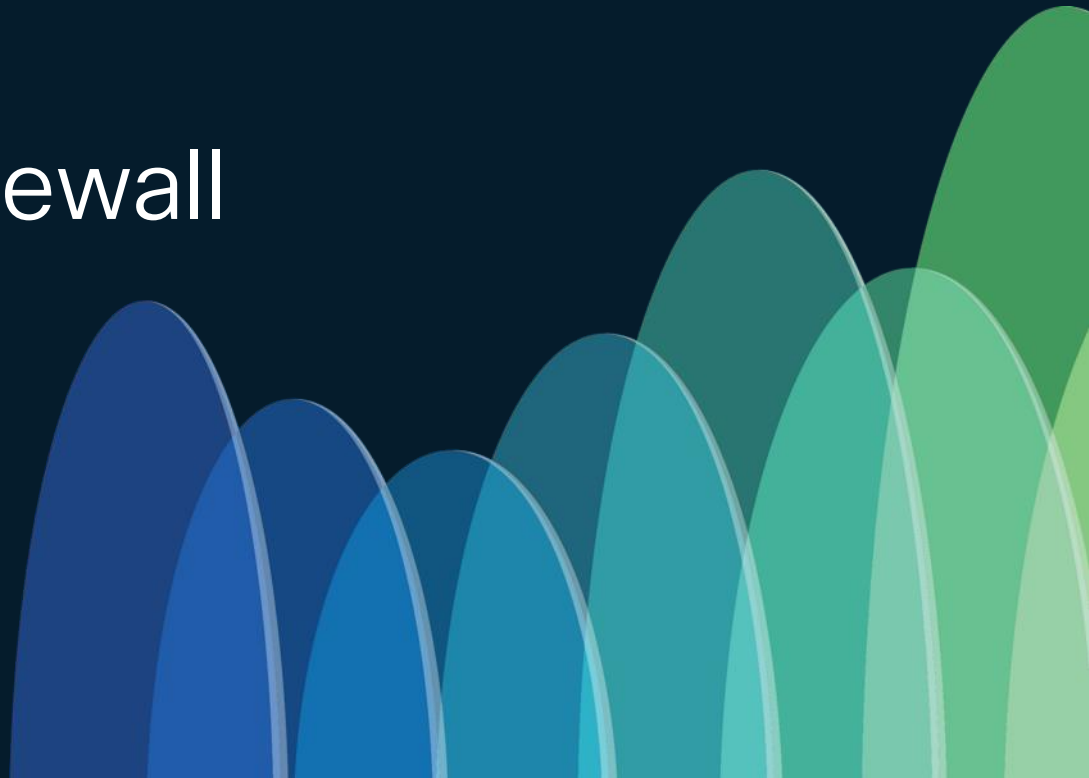

cisco Live!

Cisco Secure Firewall

Full coverage, from IoT/OT & Branch / SASE to Enterprise/Carrier Class modular chassis

Physical appliances

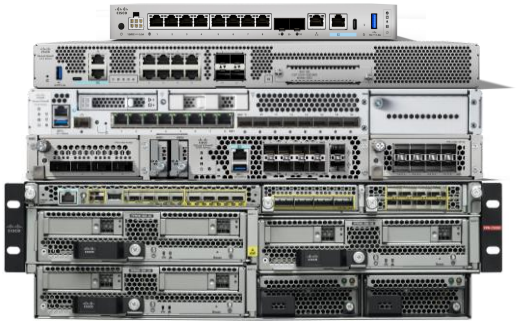

Cisco Secure Firewall hardware appliances
running either ASA or FTD application

Private & Public cloud

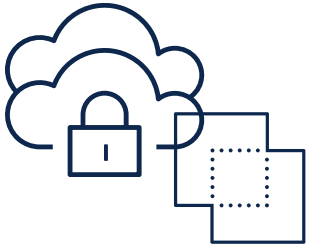

Cisco Multicloud Defense, ASAv and FTDv application
Running on all major public cloud and private cloud hypervisors

IoT and integrations

ISA 3000
Running either ASA or FTD application

Catalyst 9300
ASAc running as a container

Meraki MX and Catalyst 8000
Snort 3 running in container

Cisco Secure Firewall Hardware

Full coverage, from IoT/OT & Branch / SASE to Enterprise/Carrier Class modular chassis

ISA 3000
<0.7 Gbps

1010
<1 Gbps

11xx
2-5 Gbps

21xx
2.5-10 Gbps

1200 Series Compact
6-9 Gbps

1200 Series
9-18 Gbps

3100 Series
10-45 Gbps
up to 0.57Tbps in 16x cluster

4200 Series
65-145 Gbps
up to 1.79Tbps in 16x cluster

41xx
19-53 Gbps

93xx
55-68 Gbps

OT/IoT | Branch / SASE | Campus / Enterprise / Data Center / SP

* all performance values for 1024B avg. packet size with NGFW traffic profile

Secure Firewall 4200 Series

FTD 7.4

ASA 9.20

- 3 models – 4215/4225/4245
  - 32-128 (64-256) cores (4245 has two CPUs)
  - 8x1/10/25G SFP/SFP+ and two Network Module bays
  - 256GB-1TB of RAM
  - Two NVMe slots, 1.8TB of RAID1 protected space
  - AC redundant PS
- Advanced FPGA and one to four VPN crypto hardware accelerators
- Clustering support on all models, up to 16x nodes
- Up to 145Gbps for NGFW traffic profiles (~3x over 4100)
  - up to 45Gbps with 50% of TLS 1.2/1.3 mix
  - up to 140Gbps for IPsec traffic
- Up to 190Gbps for ASA traffic profiles (>2x over 4100)

Secure Firewall 4200 Series Overview

FTD 7.4

ASA 9.20

Appliance-Mode Security Platform for FTD or ASA Application

- Fixed configurations: 4215, 4225, 4245
- Lightweight virtual Supervisor module w/Multi-Instance (7.6) and Clustering
- Integrated Datapath FPGA w/Flow Offload and Crypto Engines
- Rear dual redundant power supplies and triple fan trays

SFP Data Interfaces

- 8x1/10/25GE

1RU

NVMe Drives

- Up to 2x900GB in RAID1 on 4215/4225 (SED)
- Up to 2x1.8TB in RAID1 on 4245 (SED)

Expansion Network Modules

- Standard: 8x1/10GE, 8x1/10/25/**50**GE, 4x10/40GE, 2x100GE, 4x40/100/200GE, **2x200/400GE** SFP+ (with 7.6)
- Fail-to-Wire: 8x1GE Copper; 6x10GE or 6x25GE SFP+ (SR and LR variants)

Secure Firewall 4200 Series Architecture

FTD 7.4

ASA 9.20

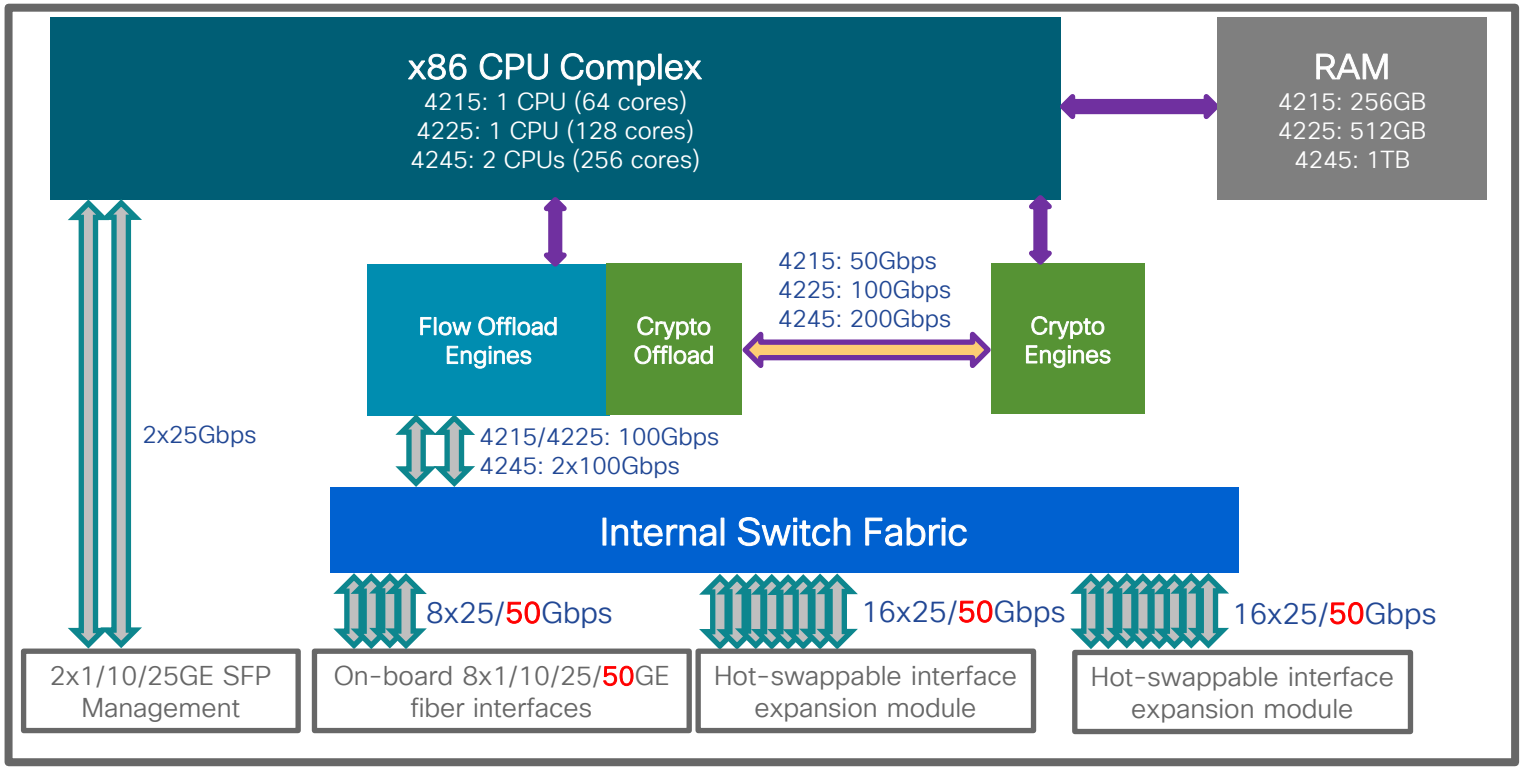

Secure Firewall 4200 Series Performance

FTD 7.4

ASA 9.20

	4215	4225	4245
FW+AVC+IPS HTTP 1024B Avg Packet	65Gbps	85Gbps	145Gbps
IPsec VPN HTTP 1024B Avg Packet	45Gbps (45Gbps per tunnel)	80Gbps (57Gbps per tunnel)	140Gbps (57Gbps per tunnel)
TLS Decryption HTTP 1024B Avg Packet 50% Flows Decrypted	20Gbps	30Gbps	45Gbps

Up to 3x ↑ Boost in FW+AVC+IPS

Up to 6x ↑ Boost in IPsec VPN

Up to 5x ↑ Boost in TLS Decrypt

Secure Firewall 3100 Series

FTD 7.1

ASA 9.17

- 5 models – 3105 & 3110/20/30/40
  - single CPU, 12-32 cores
  - 8x1G TX
  - 8x1/10G or 8x1/10/25G plus NetMod bay
  - 64-256GB of RAM
  - two SSD slots
  - AC/DC redundant PS (400W)
- Advanced NPU and VPN crypto hardware
- Clustering support on 3110-3140, up to 16x nodes
- 17-45 Gbps for FW+AVC+IPS with 1024 bytes average packet size
- 11-39.4 Gbps for IPsec with 1024 bytes average packet size with release 7.2

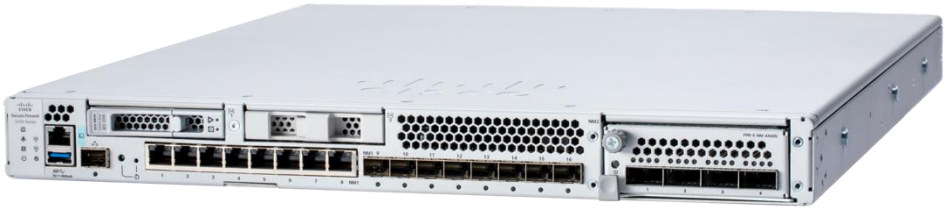

Secure Firewall 3100 Series

Overview

FTD 7.1

ASA 9.17

Appliance-Mode Security Platform for FTD or ASA Application

- Fixed configurations: 3105, 3110, 3120, 3130, 3140
- Lightweight virtual Supervisor module w/Multi-Instance and Clustering
- Integrated Datapath FPGA w/Flow Offload and Crypto Engine
- Rear dual redundant power supplies and fan trays

SFP Data Interfaces

- 8x1/10GE on 3105-3120
- 8x1/10/25GE on 3130-3140

1RU

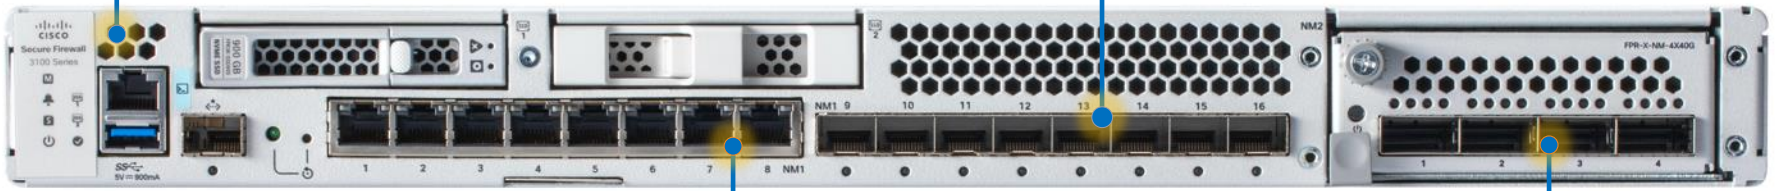

Copper Data Interfaces

- 8x10/100/1000BaseT

Network Module

- 8x1/10/25GE or 6x10/25GE FTW on 3105-3120
- 4x40GE, 2x40GE FTW and 2x100GE on 3130-3140
- 8x10/100/1000BaseT & 6x1GE, 6x10GE, 6x25GE SFP FTW

cisco Live!

Secure Firewall 3100 Series

Architecture

FTD 7.1

ASA 9.17

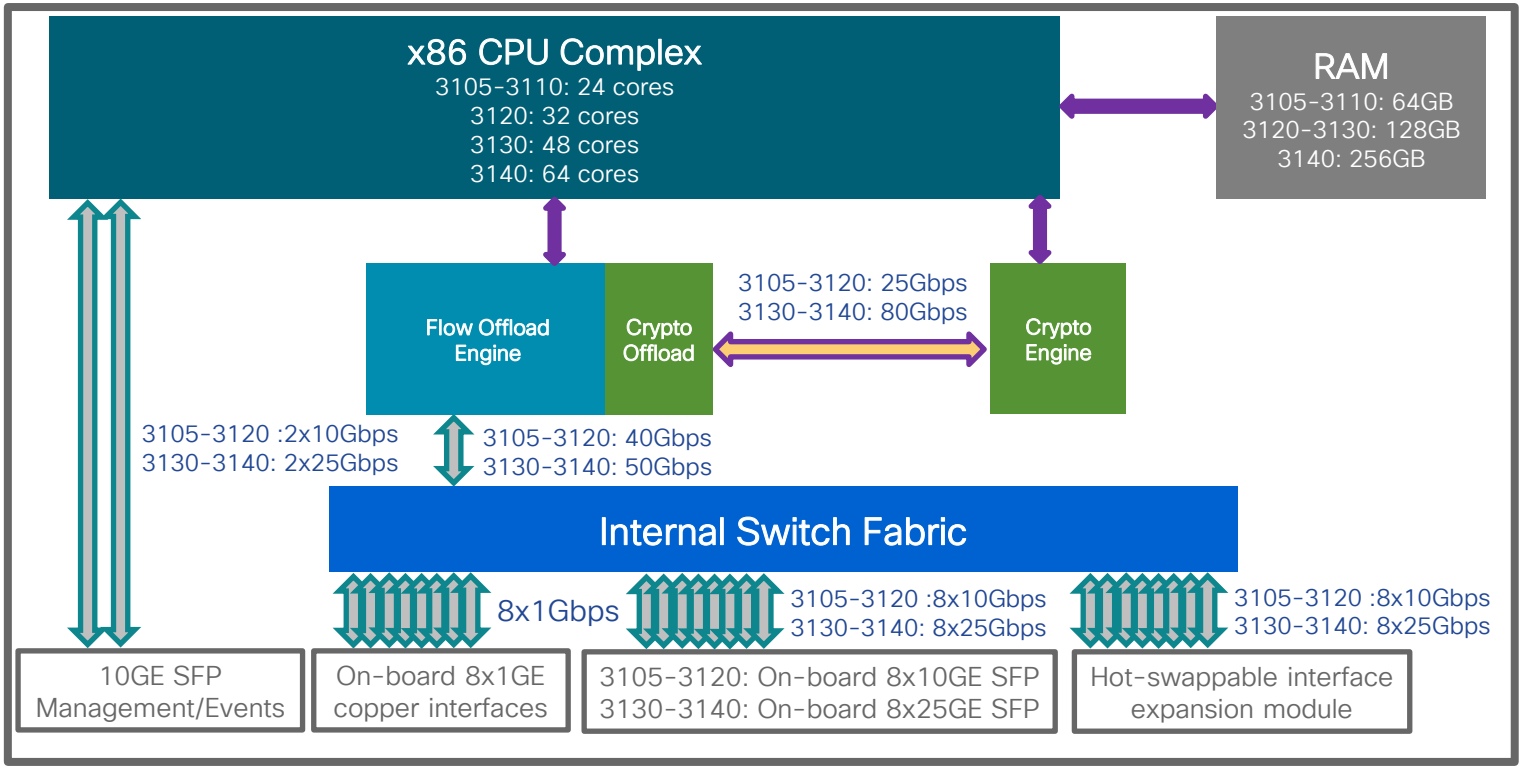

System Bus

Ethernet

Chip-to-Chip Link

Secure Firewall 1200 Series Compact

FTD 7.6

ASA 9.22

- 3 models – 1210CE, 1210CP, 1220CX
  - Network/Security SoC with 8 ARM cores design
  - 16GB of RAM
  - 480GB of NVMe storage
  - Fixed 8x1GE:
    - 1210CP – 4 ports with UPoE+ support (120W total, max of 90W per port)
    - 1220CX – plus 2x 1/10G SFP+
- Multiple SoC-embedded accelerators
  - encryption/decryption
  - traffic processing
- Up to 2.6Gbps (450B) or up to 9Gbps (1024B) for NGFW traffic profiles (~10x over 1010, ~3x over 11xx)
- Up to 10Gbps for IPsec VPN, and up to 1.5Gbps for TLS 1.2/1.3

Secure Firewall 1200 Series Compact

Overview

FTD 7.6

ASA 9.22

Appliance-mode Security Platform for FTD or ASA Application

- Desktop form factor (1210, 1220)
- Fully integrated System-on-a Chip (SoC) with embedded networking/security acceleration
- Active/standby HA support (no clustering, no multi-instance)
- Optional rack mounting kit
- Quiet blower for active cooling
- External brick-style AC power adapter

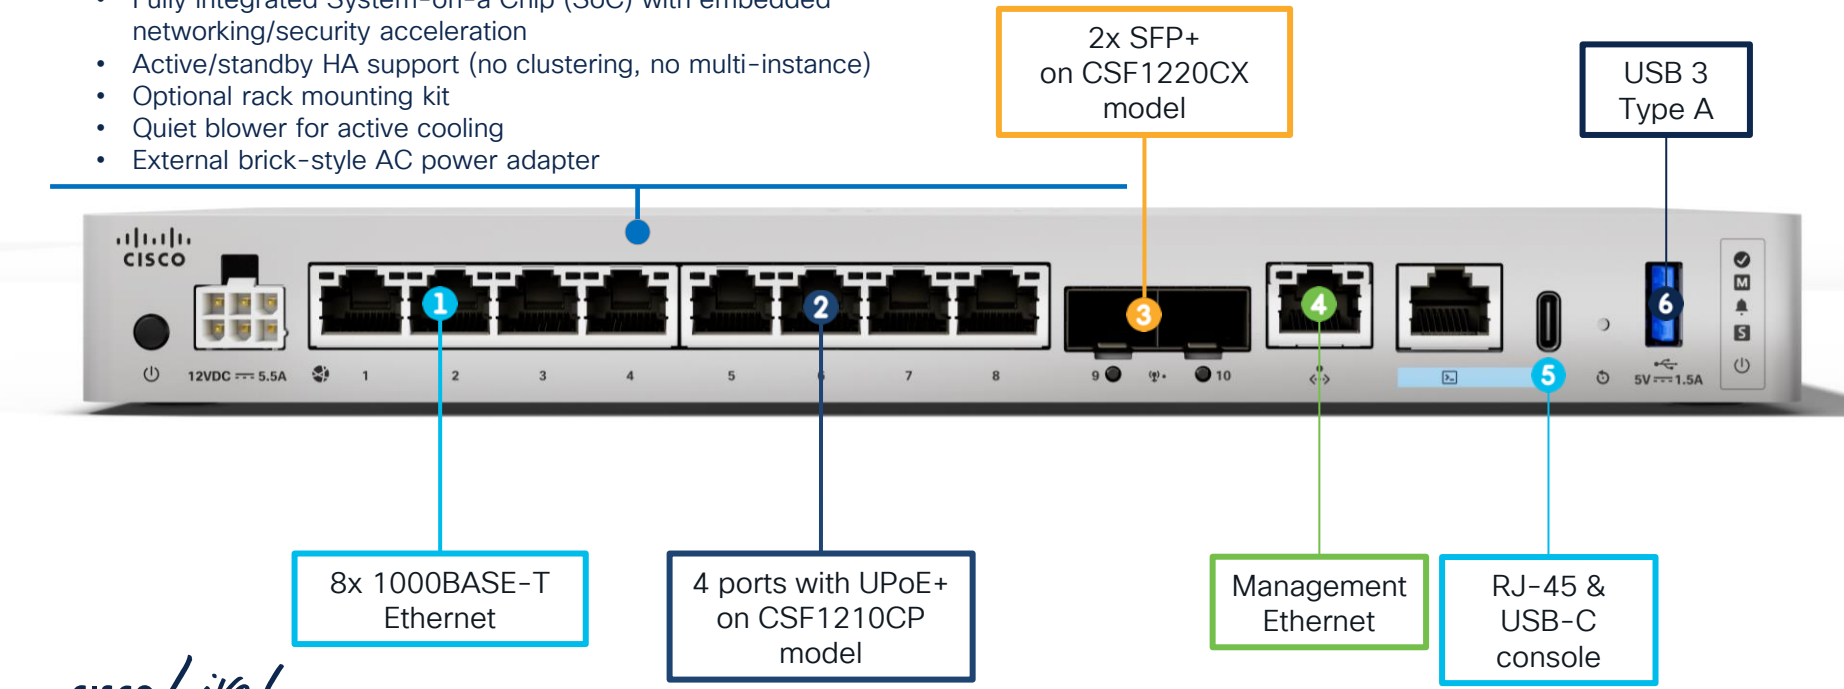

Secure Firewall 1200 Series Compact

FTD 7.6

ASA 9.22

Architecture

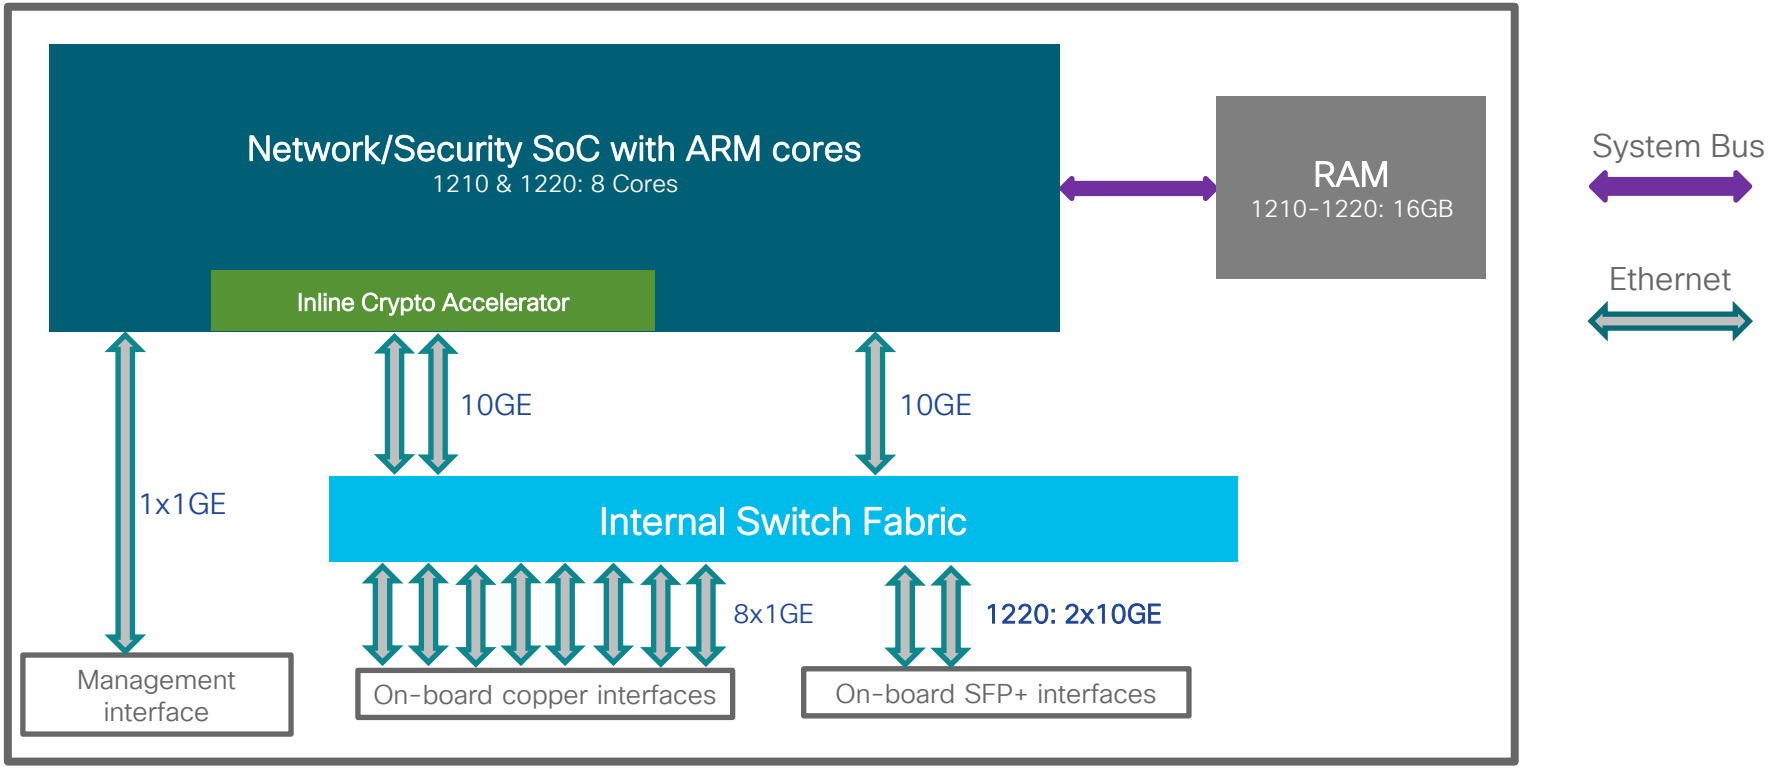

Secure Firewall 1200 Series Compact

FTD 7.6

Key Metrics

	1210CE/CP	1220CX
FTD AVC+IPS HTTP 1024B average packet size	6 Gbps	9 Gbps
IPsec VPN 1024B TCP w/FastPath	5 Gbps	10 Gbps
TLS 50% decrypt	1 Gbps	1.5 Gbps
Concurrent sessions with AVC	200k	300k
New connections per second	35k	50k
Maximum VPN peers	200	300
Maximum VRFs	5	10

Secure Firewall 1200 Series Compact

ASA 9.22

Key Metrics

	1210CE/CP	1220CX
ASA UDP 1500B average packet size	6.5 Gbps	15 Gbps
ASA multiprotocol HTTP, SMTP, FTP, IMAPv4, BitTorrent, DNS mix	6 Gbps	12 Gbps
IPsec 450B site to site, AES-256	5.5 Gbps	12 Gbps
Concurrent sessions full stateful tracking and inspection	200k	300k
New connections per second	175k	250k
Maximum VPN peers	200	300

Secure Firewall 1200 Series

FTD 7.7

ASA 9.23

- 3 models – 1230, 1240 and 1250
  - Network/Security SoC with 12-16 ARM cores design
  - 16-32GB of DDR5 RAM
  - 960GB of NVMe storage
  - Fixed 8x1GE (1230 & 1240) and 8x1/2.5GE (1250)
  - Fixed 4x SFP+ (1/10G)
- Multiple SoC-embedded accelerators
  - encryption/decryption
  - traffic processing
- Up to 12Gbps (450B) or up to 18Gbps (1024B) for NGFW traffic profiles
- Up to 22 Gbps for IPsec VPN, and up to 4 Gbps for TLS 1.2/1.3

Secure Firewall 1200 Series

Overview

FTD 7.7

ASA 9.23

Copper Data Interfaces
- 1230-1240: 8x1000BaseT
- 1250: 8x1/2.5GBaseT

SFP Data Interfaces
- 1230 and 1240: 4x1GE/10GE SFP+
- 1250: 4x1GE/10GE SFP+

1RU

Management
- 10/100/1000BaseT Ethernet
- RJ-45 and USB-C console
- USB-A for external flash

Appliance-Mode Security Platform for FTD or ASA Application
- Rack-Mount (1230, 1240, and 1250)
- Fully integrated System-on-a Chip (SoC) with embedded networking/security acceleration
- Active/standby HA support (no clustering, no multi-instance)

Secure Firewall 1200 Series

Architecture

FTD 7.7

ASA 9.23

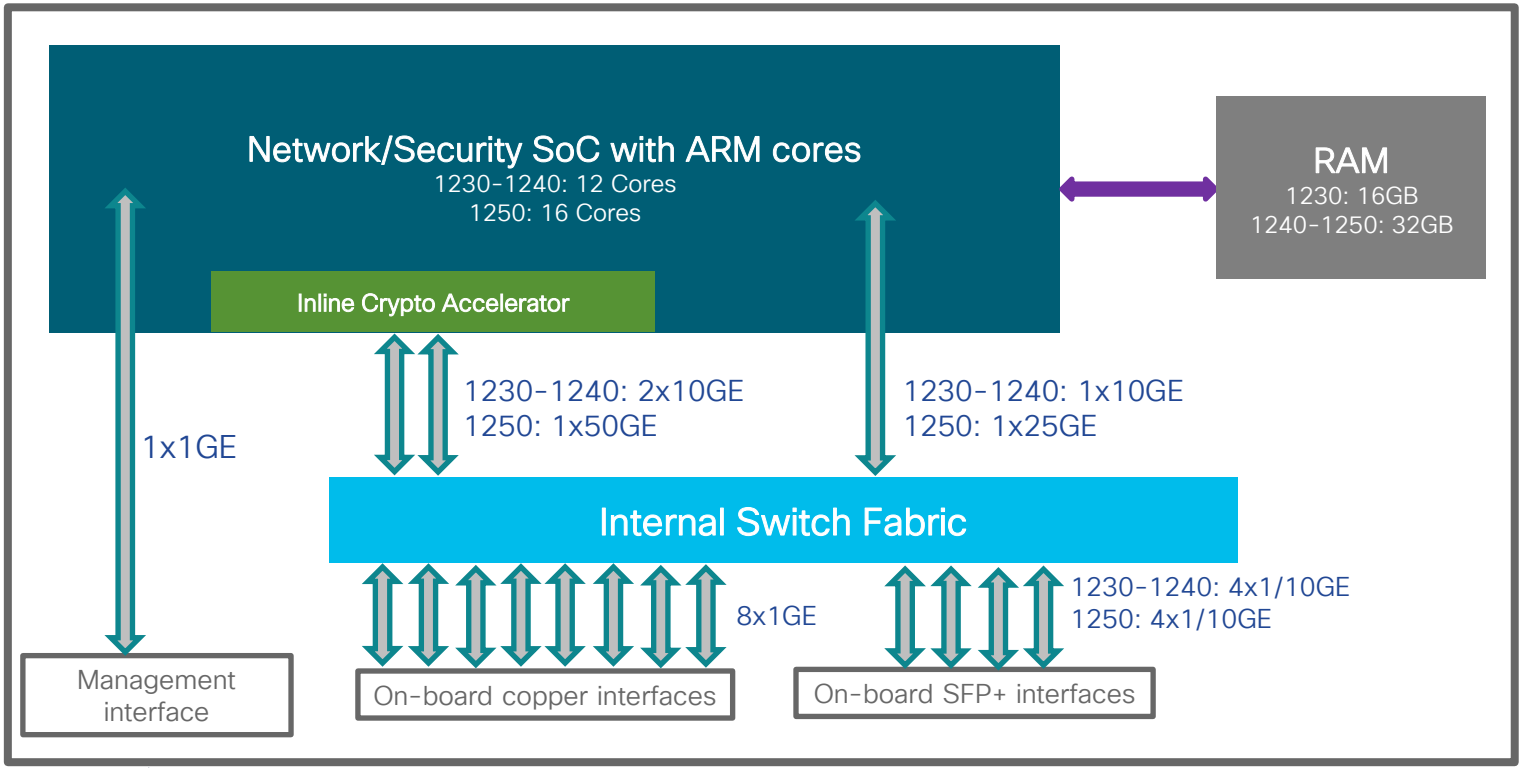

Secure Firewall 1200 Series

FTD 7.7

Key Metrics

	1230	1240	1250
FTD AVC+IPS HTTP 1024B average packet size	9 Gbps	12 Gbps	18 Gbps
IPsec VPN 1024B TCP w/FastPath	13 Gbps	18 Gbps	22 Gbps
TLS 50% decrypt	2.5 Gbps	3.1 Gbps	4.1 Gbps
Concurrent sessions with AVC	0.4M	0.6M	1M
New connections per second	50k	80k	100k
Maximum VPN peers	500	1000	1500
Maximum VRFs	5	5	10

All performance estimates are subject to change in public release.

Secure Firewall 1200 Series

ASA 9.23

Key Metrics

	1230	1240	1250
ASA UDP 1500B average packet size	20+ Gbps	20+ Gbps	20+ Gbps
ASA multiprotocol Mix of HTTP, SMTP, FTP, IMAPv4, BitTorrent, and DNS	20+ Gbps	20+ Gbps	20+ Gbps
IPsec 450B site to site, AES-256	13 Gbps	18 Gbps	22 Gbps
Concurrent sessions full stateful tracking and inspection	0.4M	0.6M	1M
New connections per second	350k	450k	550k
Maximum VPN peers	500	1000	1500

All performance estimates are subject to change in public release.

Secure Firewall 9300 Series

- 1 chassis, choice of three Service Modules
  - central Supervisor with switching fabric – 2x40GE towards each Service Module, 5x40GE towards Network Module bays
  - 8xSFP/SFP+ ports built-in plus one SFP management port
  - two Network Module bays – choice of 1/10/40/100GE interfaces & FTW
  - each Service Module can run either ASA or FTD – support for mixed mode operation
  - AC/DC redundant PS (3000W)
- Advanced NPU and VPN crypto hardware on each Service Module
- Clustering support on all models – up to 16x
- up to 64 Gbps for FW+AVC+IPS with 1024 bytes average packet size per Service Module
- up to 51 Gbps for IPsec with 1024 bytes average packet size with release 7.2 per Service Module

Secure Firewall 9300 Series Overview

Supervisor
- Application deployment and orchestration
- Network attachment and traffic distribution
- Clustering base layer for ASA or FTD

Network Modules
- 10GE, 40GE, 100GE
- Hardware bypass for inline NGIPS

3RU

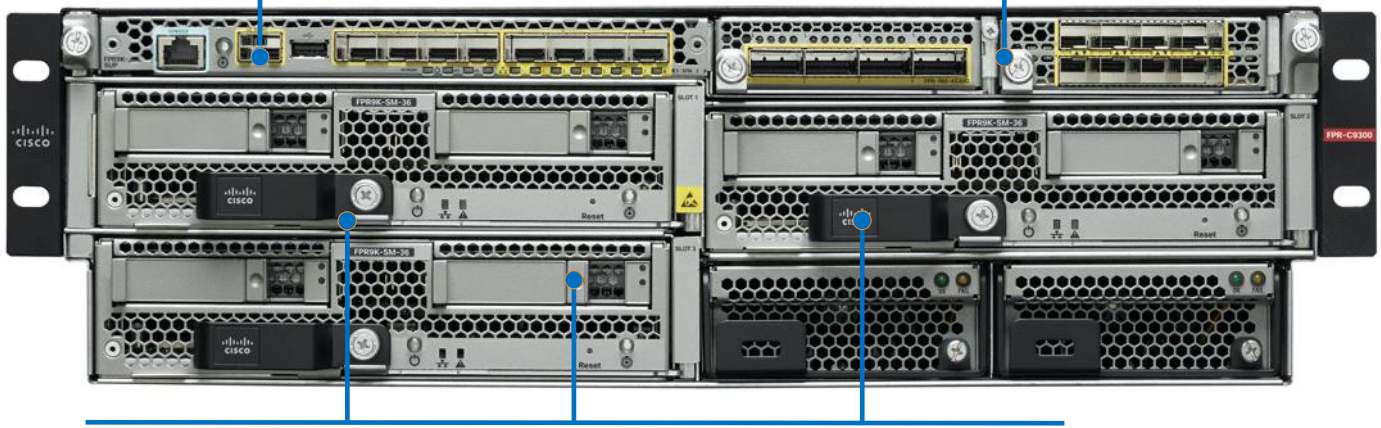

Security Modules
- Embedded Smart NIC and crypto hardware
- Cisco (ASA, FTD) and third-party (Radware DDoS) applications
- Standalone or clustered within and across chassis

Secure Firewall 9300 Series

Supervisor Module

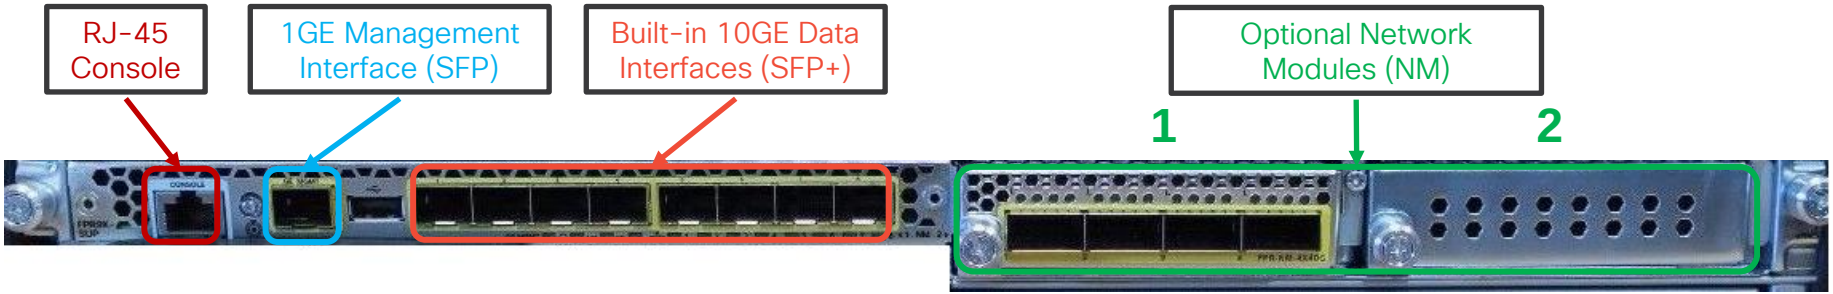

- Network interface allocation and security module connectivity
  - LACP or Static (in FXOS 2.4.1) Port-Channel creation with up to 16 member ports
  - Up to 500 VLAN subinterfaces for Container instances in FXOS 2.4.1
- Application image storage, deployment, provisioning, and service chaining
- Clustering infrastructure for supported applications
- Smart Licensing and NTP for entire chassis

Secure Firewall 9300 Series

Supervisor Architecture

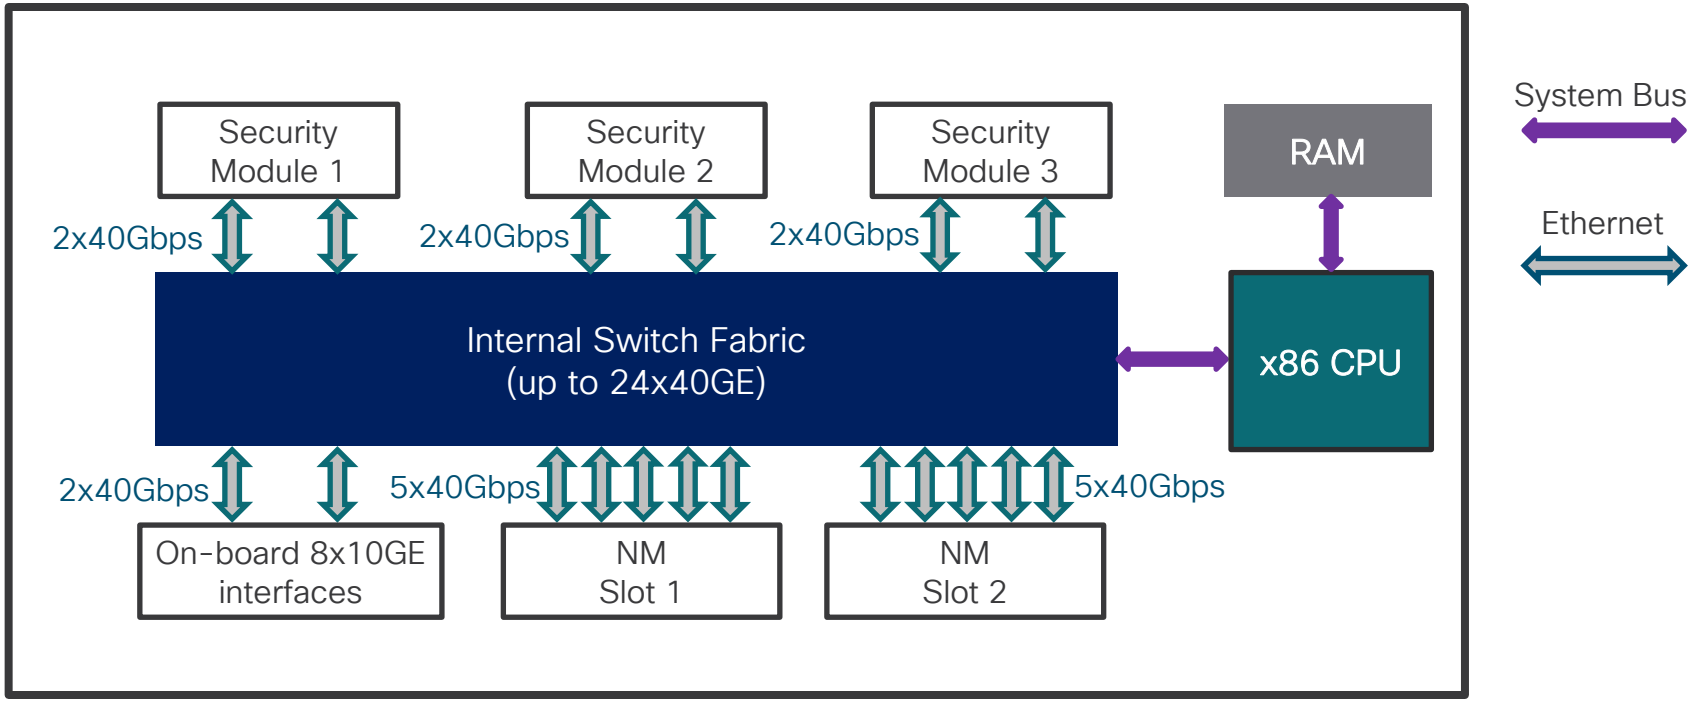

Secure Firewall 9300 Series

Security Module Architecture

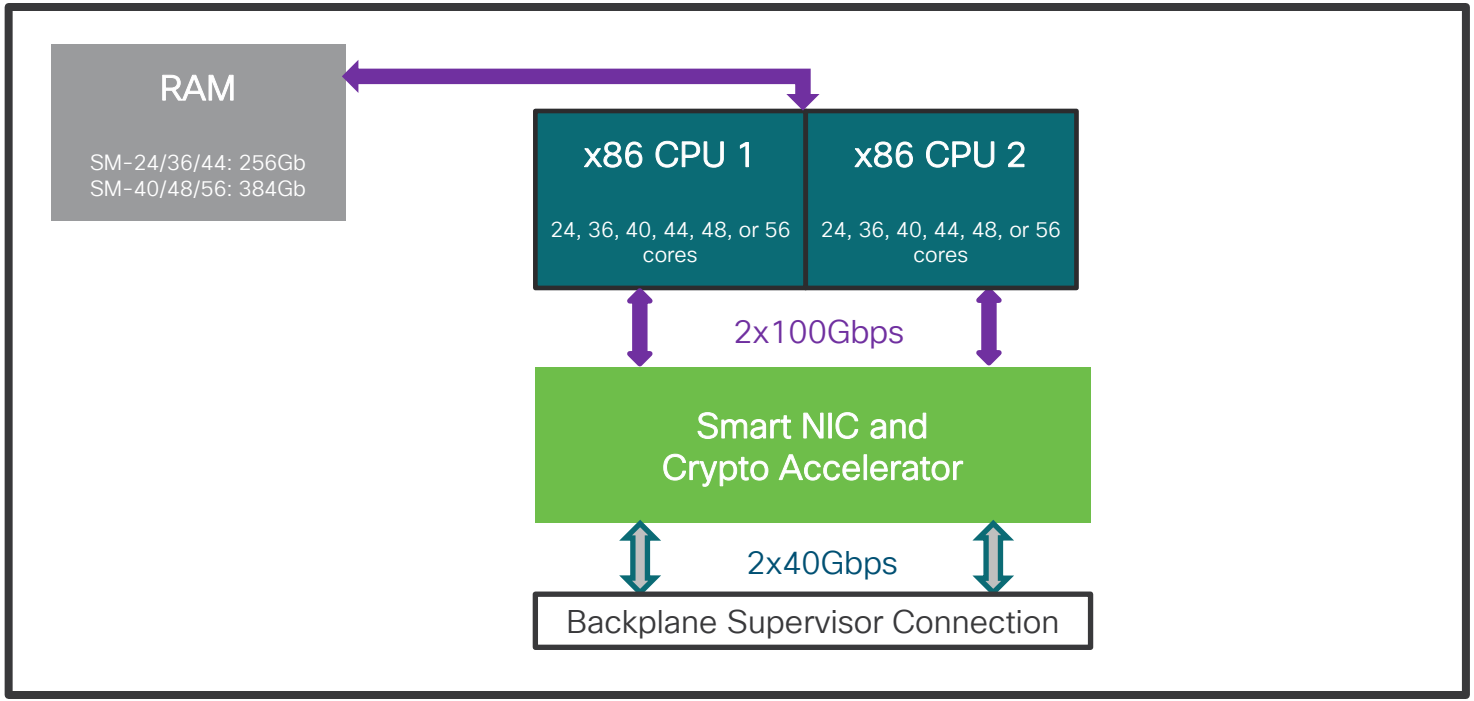

Secure Firewall 9300 Series

Security Modules

- Built-in hardware Smart NIC and Crypto Accelerator
- SM-40, SM-48, and SM-56
  - Dual 1.6TB SSD in RAID1 by default
  - Higher performance on cryptographic operations
- Previous generation SM-24, SM-36, and SM-44
  - Dual 800GB SSD in RAID1 by default
  - SM-24 is NEBS Level 3 Certified
- Mixed standalone modules supported in FXOS 2.6.1
  - Mixed modules supported with FTD multi-instance clustering in FXOS 2.8.1

Secure Firewall 4100 Series

- 4 models, 4112/4115/4125/4145
  - 12-44 CPU physical cores
  - 8xSFP/SFP+ built-in
  - two Network Module bays
  - AC/DC redundant PS (1100W AC/950W DC)
- Advanced NPU and VPN crypto hardware
- Clustering support on all models, 16x
- 53 Gbps for FW+AVC+IPS with 1024 bytes average packet size
- 24 Gbps for IPsec with 1024 bytes average packet size with release 7.2

Secure Firewall 4100 Series Overview

Built-in Supervisor and Security Module
- Same hardware and software architecture as 9300
- Fixed configurations (4110-4150)

Solid State Drives
- Independent operation (no RAID)
- Default slot 1 provides 200-800GB of total storage
- Slot 2 adds 400GB of AMP storage

1RU

Network Modules
- 10GE and 40GE interchangeable with 9300
- Partially overlapping fail-to-wire options

Secure Firewall 4100 Series Architecture

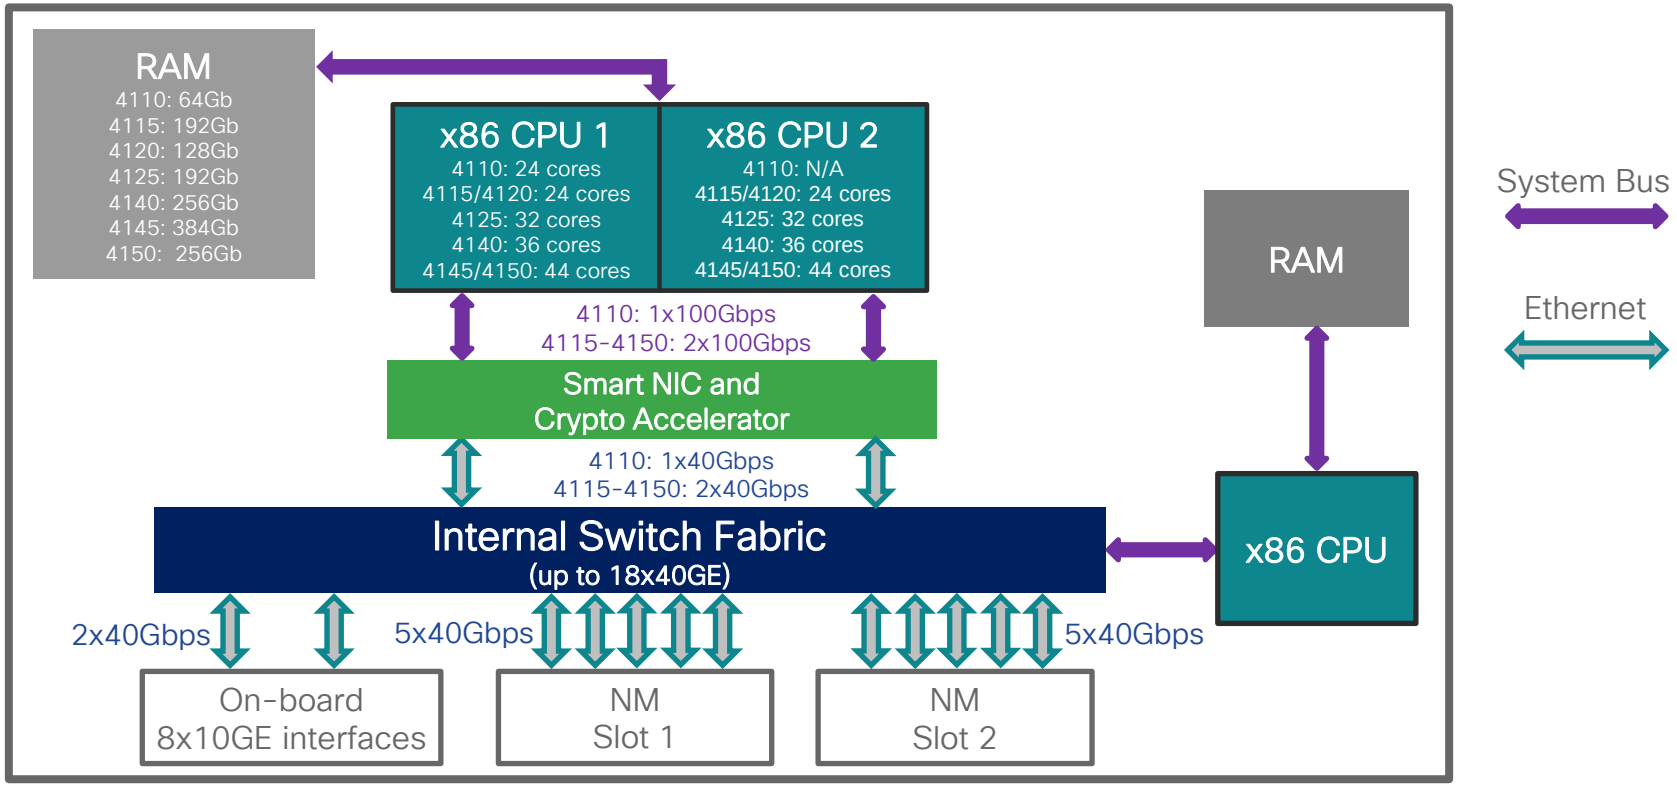

Secure Firewall 4100/9300 Series

Smart NIC and Crypto

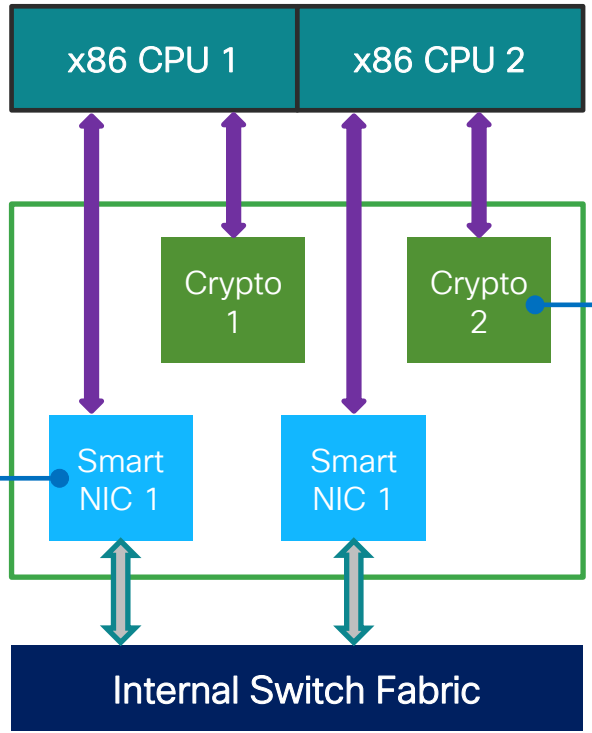

Crypto Accelerator

- Single on 4110, dual elsewhere
- Configurable core bias to IPsec/TLS on 4110, 4120, 4140, 4150 and 9300 SM-24, SM-36, SM-44; shared elsewhere
- IPsec S2S and RAVPN
- TLS/DTLS RAVPN
- TLS inspection assistance

Cisco Programmable NIC

- Single on 4110, dual elsewhere
- 40Gbps connectivity each
- Packet Matching and Rewrite
- Tracks 2M flows for Flow Offload

FXOS 2.3.1

System Bus

Ethernet

Secure Firewall 2100 Series

Last day of sales coming on May 2025

- 4 models (2110, 2120, 2130, 2140)
  - 4-16 cores
  - 12x1G TX
  - 4x SFP (2110/20) or 4x SFP+ (2130/40)
  - 16-64GB of RAM
  - one 200GB SSD disk with one optional for redundancy
  - 250-400W AC (2110-2140)
    350W DC (2130-2140) power supply
- Advanced x86 processing with multi-core NPU
- 2.5Gbps to 10Gbps for FW+AVC+IPS with 1024 bytes average packet size
- 365Mbps to 1.4Gbps for TLS decryption performance
- 950Mbps to 3.5Gbps for IPsec with 1024 bytes average packet size

Secure Firewall 2100 Series Overview

Integrated Security Platform for FTD or ASA Application

- Lightweight virtual Supervisor module
- Embedded x86 and NPU with Hardware Crypto Acceleration
- Fixed configurations (2110, 2120, 2130, 2140)
- Dual redundant power supplies on 2130 and 2140 only

SFP/SFP+ Data Interfaces

- 4x1GE on 2110 and 2120
- 4x10GE on 2130 and 2140

1RU

Copper Data Interfaces

- 12x1GE Ethernet

Network Module

- 2130 and 2140 only
- Same 8x10GE SFP module as on 4100/9300

Secure Firewall 2100 Series Architecture

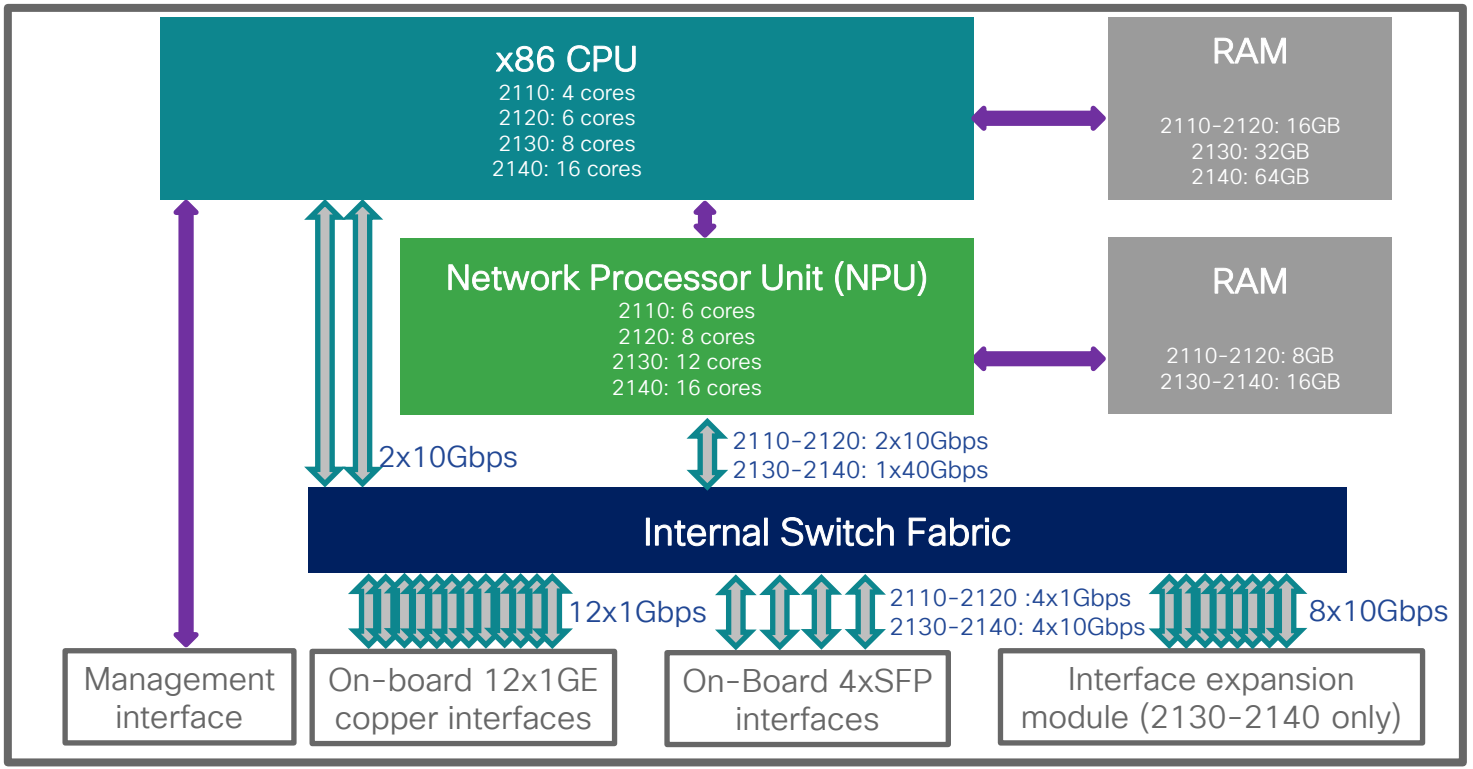

Secure Firewall 1010/1010E

- 1 model – 1010/1010E
  - 4 physical cores
  - 8x1G TX, 2 ports (7/8) with PoE IEEE 802.3at on 1010
  - 8GB of RAM
  - one 200GB SSD disk
  - AC 115W (1010 for PoE) or 55W (1010E has no PoE support)
- x86 with hardware assisted cryptographic processing (QAT) for IPsec & TLS
- 0.85Gbps for FW+AVC+IPS with 1024 bytes average packet size
- 195Mbps for TLS decryption performance
- 400Mbps for IPsec with 1024 bytes average packet size

Secure Firewall 1100 Series

- 3 models – 1120, 1140 & 1150
  - 12-16 physical cores
  - 8x1G TX
  - 4x SFP (1120/1140) or 2x SFP + 2x SFP+ (1150)
  - 16-32GB of RAM
  - one 200GB SSD disk
  - AC 100W (1120/1140/1150) power supply
- x86 with hardware assisted cryptographic processing (QAT) for IPsec & TLS
- 2.3Gbps to 5Gbps for FW+AVC+IPS with 1024 bytes average packet size
- 850Mbps to 1.4Gbps for TLS decryption performance
- 1.2Gbps to 2.4Gbps for IPsec with 1024 bytes average packet size

Secure Firewall 1010/E Overview

Desktop

Integrated Security Appliance with ASA or FTD

- Embedded x86 CPU with QuickAssist Crypto Acceleration
- Fixed non-modular configuration

Copper Data Interfaces

- 8x1GE Ethernet
- Built-in Layer 2 switch
- Power over Ethernet (PoE) on ports 7 and 8

Secure Firewall 1100 Series Overview

Integrated Security Appliance with ASA or FTD
- Embedded x86 CPU with QuickAssist Crypto Acceleration
- Fixed non-modular configurations (1120, 1140, 1150)

SFP Data Interfaces
- 4x1GE on 1120 and 1140
- 2x1GE, 2x10GE on 1150

1RU

Copper Data Interfaces
- 8x1GE Ethernet

Field Replaceable SSD

Secure Firewall 1100 Series Architecture

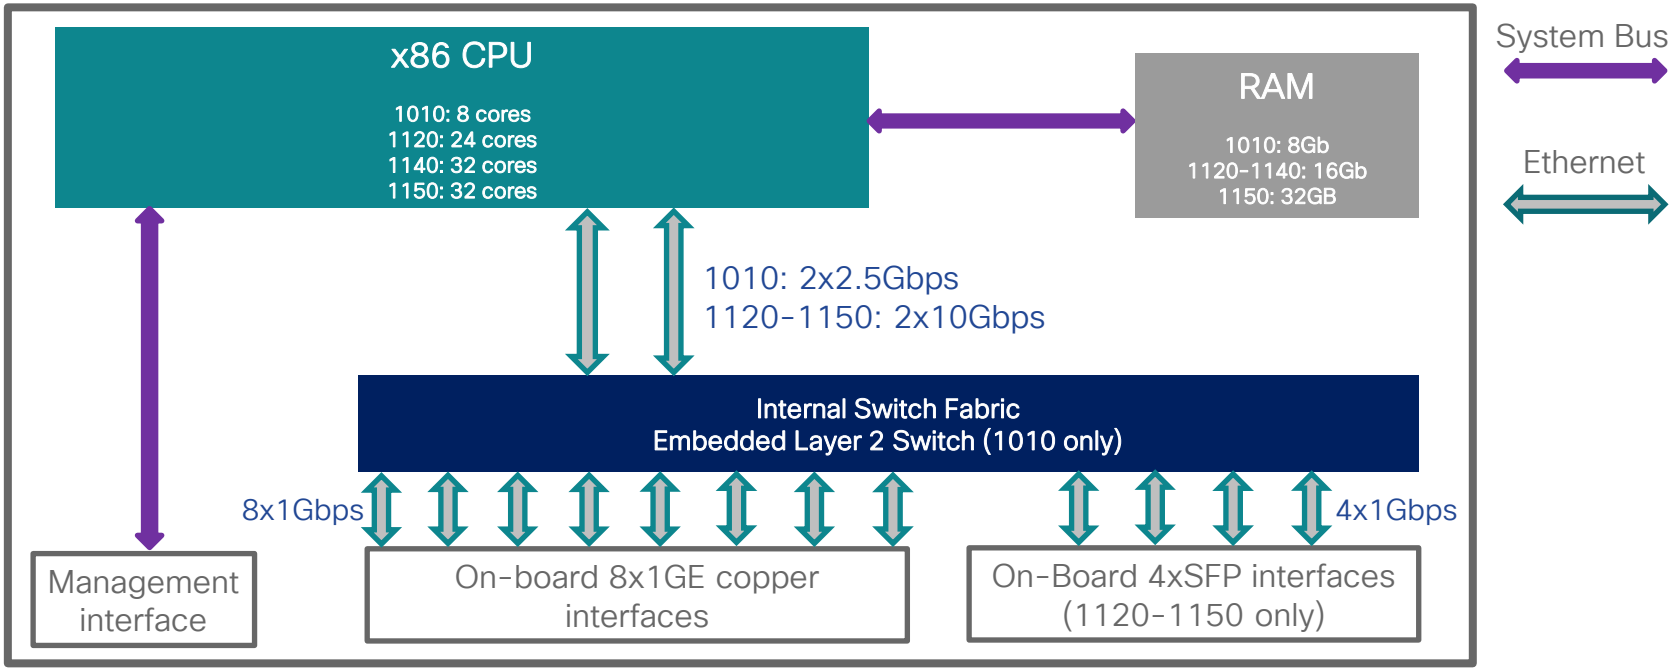

Secure Firewall ISA 3000 Series

- 2 models
  - Intel 4-core Atom CPU, I-Temp compliant
  - 4x 10/100/1000TX or 2x10/100/1000TX & 2xSFP; dedicated 10/100/1000 Management Port
  - 8GB of RAM, 16GB of flash memory + mSATA 64GB with 1GB removable SD flash card
  - Dual internal DC power supplies
- Built for harsh environments and temperature ranges (-40F to 158F; -40C to 70C)
- Hardened for vibration, shock, surge, and electrical noise immunity
- Broad OT protocol coverage (universal to all Snort 3 based sensors): BACnet, CIP, COSEM, COTP, DNP3, GOOSE, GSE, ECP, FDC, Honeywell CS/NIF Server & Esperion DSA Server monitor, IEC 60870-5-104, IEC 61850 MMS, Modbus, Omron FINS, OPC-UA, Q.931, Siemens S7, SRC, TPKT – plus all (3000+) OpenAppID applications
- Can run either ASA or FTD code

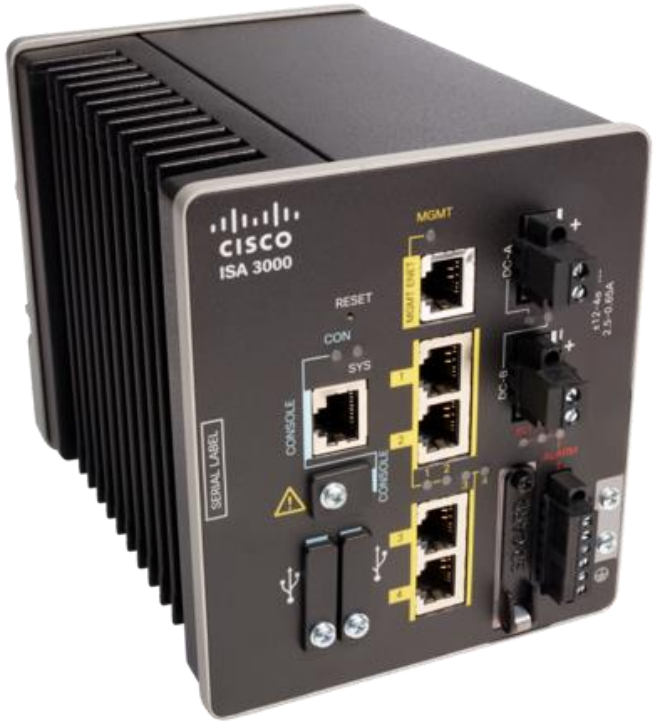

Secure Firewall FMC 1700/2700/4700

- 3 models – 1700/2700/4700
  - 1x AMD CPU (8-24 cores)
  - 2x10G NIC for connectivity (Intel X710)
  - 2x10/25G (Intel E810XXVDA2) additional ports in 4700
  - 32-128GB of RAM
  - 2.4TB-120TB of HDD space
  - 240GB SSD recovery disk
- 50 (1700), 300 (2700) and 1000 (4700) sensors supported
- 30, 60, 400M IPS events supported
- 5/12/30k FPS flow rate
- 50, 150, 600k network hosts

Firewall Management Center Appliances Scale

			FMC 1700 FMC 1600	FMC 2700 FMC 2600	FMC 4600	FMC 4700
FMCv2	FMCv10	FMCv25		FMCv300		
HA and lab deployments	small networks	small networks	medium networks	medium networks	big Enterprise/SPs	big Enterprise/SPs
2 FTDs	10 FTDs	25 FTDs	50 FTDs	300 FTDs	750 FTDs	1000 FTDs
Maximum number of FTD sensors supported						
10 million	10 million	10 million	30 million	60 million	300 million	400 million
Maximum number of IPS events						
< 5,000	< 5,000	< 5,000	5,000	12,000	20,000	30,000
Maximum event rate (EPS)						

Secure Firewall Network Modules

2100/4100/9300 and 3100/4200 portfolio

3100 network modules		SW release
FPR3K-XNM-8X10G	8x 1/10G SFP+	7.1
FPR3K-XNM-8X25G	8 port 1/10/25G SFP+	7.1 (3130/40)
FPR3K-XNM-4X40G	4x 40G QSFP+ (breakout supported to 4x10G)	7.2 (3130/40)
FPR3K-XNM-8X1GF	8x 1GE TX FTW	7.3
FPR3K-XNM-6X1SXF	6x 1GE SX FTW	7.2.3/7.3.1
FPR3K-XNM-6X10SRF/LRF	6x10G FTW	7.2.3/7.3.1
FPR3K-XNM-6X25SRF/LRF	6x25G FTW	7.2.3/7.3.1
FPR3K-XNM-2X100G	**3130/3140 only:** 2x100G QSFP/QSFP28 (40/100G + breakout to 4x10G or 4x25G supported)	7.4.1

4200 network modules		SW release
FPR4K-XNM-8X1GF	8x 1G FTW	7.4.0
FPR4K-XNM-6X10SRF/LRF	6x10G FTW (SR or LR)	
FPR4K-XNM-6X25SRF/LRF	6x 25G FTW (SR or LR)	
FPR4K-XNM-8X10G	8x 1/10G SFP/SFP+	
FPR4K-XNM-8X25G	8x 1/10/25G SFP/SFP+	
FPR4K-XNM-4X40G	4x 40G QSFP+ (supports 4x10G)	
FPR4K-XNM-2X100G	2x100G QSFP/QSFP28 (supports 4x10/25G or 40G)	
FPR4K-XNM-4X200G	4x200G QSFP+ (supports 40/100G)	
FPR4K-XNM-2X400G	2x400G (supports 4x10, 4x25, 200G*)	7.6 (7.7*)

All FTW modules have built–in optics, and it's fixed.
Same-kind OIR is supported.

Secure Firewall Network Modules

2100/4100/9300 and 3100/4200 portfolio

2100 network modules	
FPR2K-NM-8X10G	8 port SFP+
FPR2K-NM-8X1G	8 port SFP
FPR2K-NM-6X1SX-F	6 port 1G SX Fiber FTW
FPR2K-NM-6X10SR-F	6 port 10G SR FTW
FPR2K-NM-6X10LR-F	6 port 10G LR FTW
FPR2K-NM-8X1G-F	8 port 1G Copper FTW

Last day of sales: May 2025

4100 network modules		SW release
FPR4K-NM-8X1G-F	8x1GE FTW	
FPR4K-NM-6X1SX-F	6x 1GE SX FTW	
FPR4K-NM-6X10SR/LR-F	6x 10G FTW (SR or LR)	
FPR4K-NM-8X10G	8x 1/10G SFP+	
FPR4K-NM-2X40G-F	2x 40G FTW	
FPR4K-NM-4X40G	4x 40G QSFP+	
FPR4K-NM-2X100G	2x 100G QSFP/QSFP28	7.3.1 (4112/15/ 4125/45)

All FTW modules have built–in optics, and it's fixed.
Same-kind OIR is supported.

Secure Firewall Network Modules

2100/4100/9300 and 3100/4200 portfolio

9300 network modules		SW release
FPR9K-NM-8X10G	8x 10G SFP+	every release
FPR9K-NM-6X10SR-F/LR-F	6x 10G FTW Does not support hot-swapping.	FXOS 2.0.1
FPR9K-NM-4X40G	4x 40G QSFP+	every release
FPR9K-NM-2X40G-F	2x 40G FTW Does not support hot-swapping.	FXOS 2.0.1
FPR9K-DNM-2X100G	2x 100G QSFP28 (doube-wide) Does not support hot-swapping.	FXOS 1.1.4
FPR9K-NM-2X100G	2x 100G QSFP28	FXOS 2.4.1
FPR9K-NM-4X100G	4x 100G QSFP28	FXOS 2.4.1

All FTW modules have built–in optics, and it's fixed.
Same-kind OIR is supported.

Secure Firewall Network Modules

Fail-to-Wire network module internals

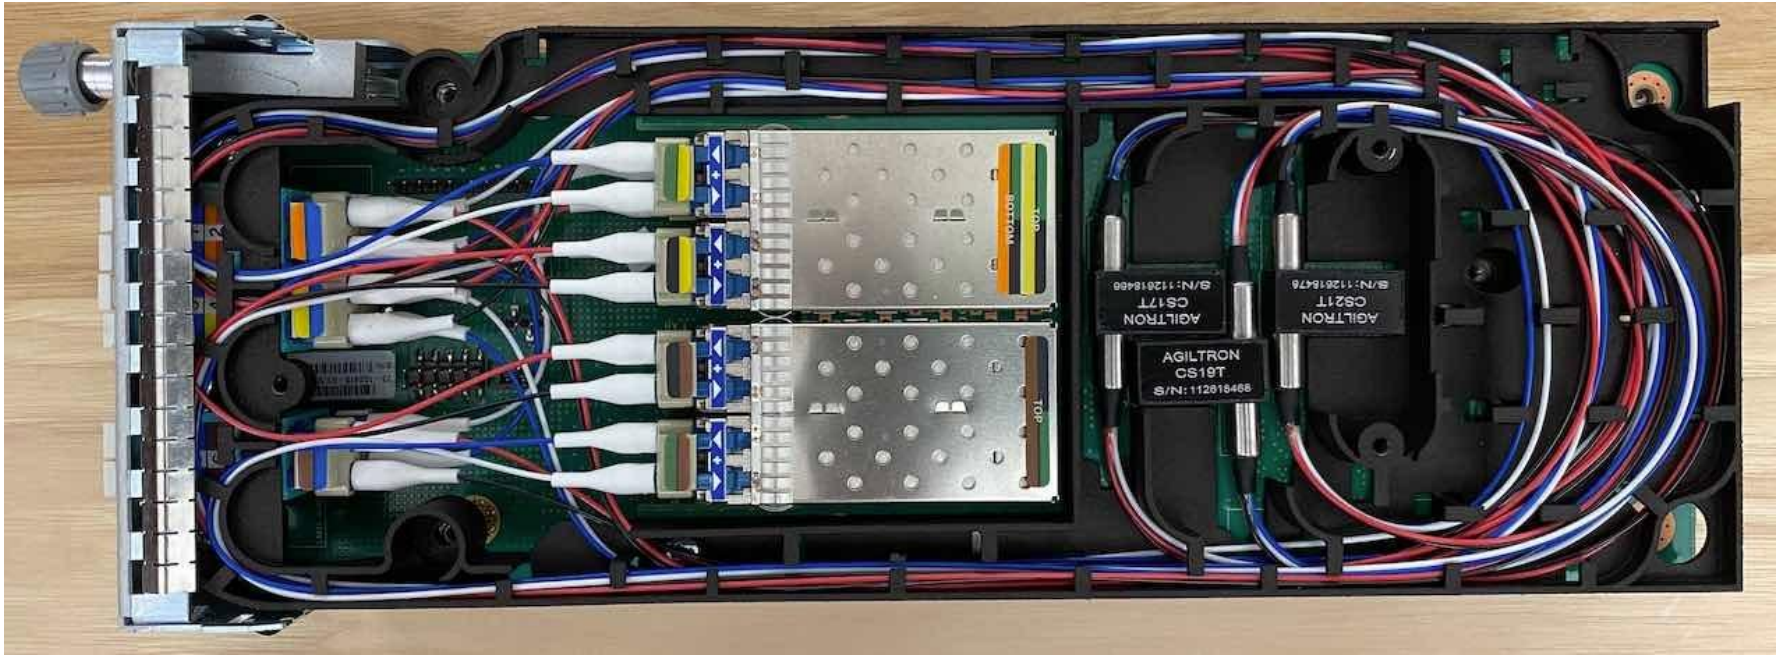

Last Day of Support (LDoS)

Please plan migration to 1200, 3100 and 4200 series

2020

Oct 31, 2020
- FP8250
- FP8260
- FP8270
- FP8290

2022

Aug 31, 2022
- ASA 5512
- ASA 5515
- ASA 5505

Dec 31, 2022
- FP7010
- FP7020
- FP7030
- FP8020
- FP8030
- FP8040

2023

May 31, 2023
- ASA 5585

Sep 30, 2023
- ASA 5506W

2024

Jun 30, 2024
- FP7050
- FP7110
- FP7115
- FP7120
- FP7125
- FP8350
- FP8360
- FP8370
- FP8390

2025

August 31, 2025
- 4120
- 4140
- 4150
- 9300 SM-24
- 9300 SM-36
- 9300 SM-44

Sep 30, 2025
- ASA 5525
- ASA 5545
- ASA 5555

We're here!

2026

Aug 31, 2026
- ASA 5506
- ASA 5508
- ASA 5516

Throughput Considerations

Third-Party Security Reference Evaluations

FORRESTER WAVE LEADER 2024

Secure Firewall
Leader in enterprise Firewall

Secure Workload
Leader in Microsegmentation

Secure Firewall
Cybersecurity Excellence Award

Secure Firewall
Global InfoSec Award

NetSecOPEN

University of New Hampshire InterOperability Laboratory

Secure Firewall
Best inspected throughput

Secure Firewall
2024 Best Next Gen Firewall

Multicloud Defense
Finalist

How would you test your firewall?

Methodology? Tools?

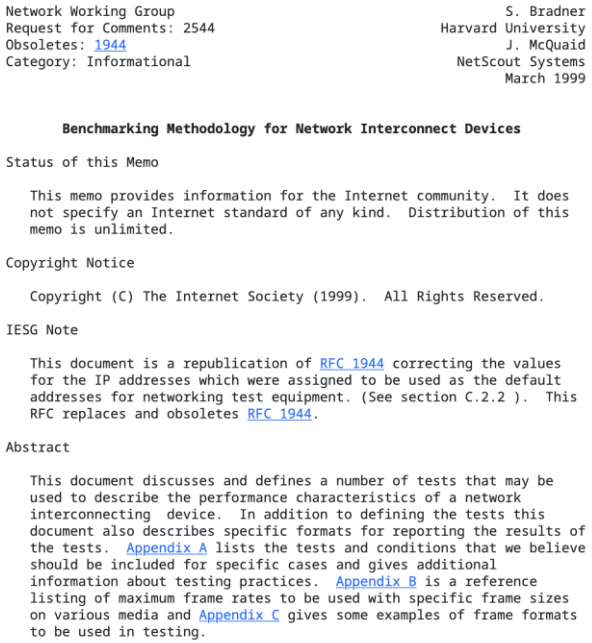

Network Working Group
Request for Comments: 2544
Obsoletes: 1944
Category: Informational

S. Bradner
Harvard University
J. McQuaid
NetScout Systems
March 1999

Benchmarking Methodology for Network Interconnect Devices

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (1999). All Rights Reserved.

IESG Note

This document is a republication of RFC 1944 correcting the values for the IP addresses which were assigned to be used as the default addresses for networking test equipment. (See section C.2.2). This RFC replaces and obsoletes RFC 1944.

Abstract

This document discusses and defines a number of tests that may be used to describe the performance characteristics of a network interconnecting device. In addition to defining the tests this document also describes specific formats for reporting the results of the tests. Appendix A lists the tests and conditions that we believe should be included for specific cases and gives additional information about testing practices. Appendix B is a reference listing of maximum frame rates to be used with specific frame sizes on various media and Appendix C gives some examples of frame formats to be used in testing.

https://datatracker.ietf.org/doc/html/rfc2544

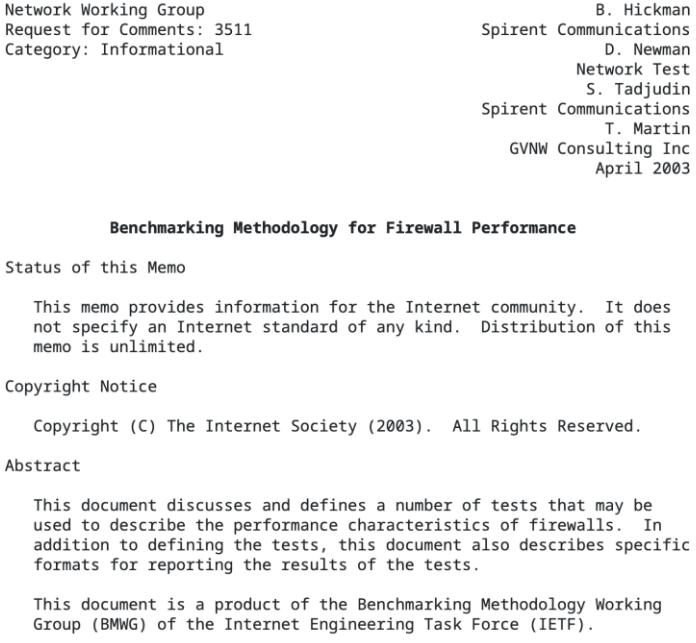

Network Working Group
Request for Comments: 3511
Category: Informational

B. Hickman
Spirent Communications
D. Newman
Network Test
S. Tadjudin
Spirent Communications
T. Martin
GVNW Consulting Inc
April 2003

Benchmarking Methodology for Firewall Performance

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2003). All Rights Reserved.

Abstract

This document discusses and defines a number of tests that may be used to describe the performance characteristics of firewalls. In addition to defining the tests, this document also describes specific formats for reporting the results of the tests.

This document is a product of the Benchmarking Methodology Working Group (BMWG) of the Internet Engineering Task Force (IETF).

https://datatracker.ietf.org/doc/html/rfc3511

How would you test your firewall?

Methodology? Tools?

Change between iPerf 2.0, iPerf 3.0 and iPerf 3.1

- **iPerf2 features currently supported by iPerf3 :**
  - TCP and UDP tests
  - Set port (-p)
  - Setting TCP options: No delay, MSS, etc.
  - Setting UDP bandwidth (-b)
  - Setting socket buffer size (-w)
  - Reporting intervals (-i)
  - Setting the iPerf buffer (-l)
  - Bind to specific interfaces (-B)
  - IPv6 tests (-6)
  - Number of bytes to transmit (-n)
  - Length of test (-t)
  - Parallel streams (-P)
  - Setting DSCP/TOS bit vectors (-S)
  - Change number output format (-f)
- **New Features in iPerf 3.0 :**
  - Dynamic server (client/server parameter exchange) – Most server options from iPerf2 can now be dynamically set by the client
  - Client/server results exchange
  - A iPerf3 server accepts a single client simultaneously (multiple clients simultaneously for iPerf2)
  - iPerf API (libiperf) – Provides an easy way to use, customize and extend iPerf functionality
  - -R, Reverse test mode – Server sends, client receives
  - -O, --omit N : omit the first n seconds (to ignore TCP slowstart)
  - -b, --bandwidth n[KM] for TCP (only UDP for IPERF 2): Set target bandwidth to n bits/sec (default 1 Mbit/sec for UDP, unlimited for TCP).
  - -V, --verbose : more detailed output than before
  - -J, --json : output in JSON format
  - -Z, --zerocopy : use a 'zero copy' sendfile() method of sending data. This uses much less CPU.
  - -T, --title str : prefix every output line with this string
  - -F, --file name : xmit/recv the specified file
  - -A, --affinity n/n,m : set CPU affinity (cores are numbered from 0 - Linux and FreeBSD only)
  - -k, --blockcount #[KMG] : number of blocks (packets) to transmit (instead of -t or -n)
  - -4, --version4 : only use IPv4
  - -6, --version6 : only use IPv6
  - -L, --flowlabel : set IPv6 flow label (Linux only)
  - -C, --linux-congestion : set congestion control algorithm (Linux and FreeBSD only) (-Z in iPerf2)
  - -d, --debug : emit debugging output. Primarily (perhaps exclusively) of use to developers.
  - -s, --server : iPerf2 can handle multiple client requests. iPerf3 will only allow one iperf connection at a time.
- **New Features in iPerf 3.1 :**
  - -I, --pidfile file write a file with the process ID, most useful when running as a daemon.
  - --cport : Specify the client-side port.
  - --sctp use SCTP rather than TCP (Linux, FreeBSD and Solaris).
  - --udp-counters-64bit : Support very long-running UDP tests, which could cause a counter to overflow
  - --logfile file : send output to a log file.

How would you test your firewall?

Methodology? Tools?

Traffic Patterns Used/Referenced in Tests

450B HTTP Test (11KB Object)

This test measures throughput with a lot of clients and servers that use a transactional HTTP pr
client who downloads a relatively small object (11KB). Due to the TCP protocol overhead, the av
While most real-world deployments would rarely experience such a traffic pattern, this measurer
baseline with a lot of room to grow.

1024B HTTP Test (256KB Object)

This test is very similar to the 450B HTTP one, but it uses a larger and more realistic object size
overhead, the average frame size is around 1024 bytes. This represents typical production conc
metric to leverage when choosing a firewall appliance.

1500B UDP

This test uses a transactional UDP profile with 1500-byte frames. Due to the stateless nature of
Many vendors use this profile to measure maximum firewall performance, but it is only practical
world conditions.

TLS

This test follows the 1024B HTTP test conditions with 50% of sessions encapsulated into TLS (I
Client TLS sessions use AES256-SHA cipher with 2048-bit RSA keys, and the server is assume
decryption. These test results can be linearly extrapolated for other percentages of TLS traffic; f
twice as high with 25% of HTTPS connections in the overall traffic mix.

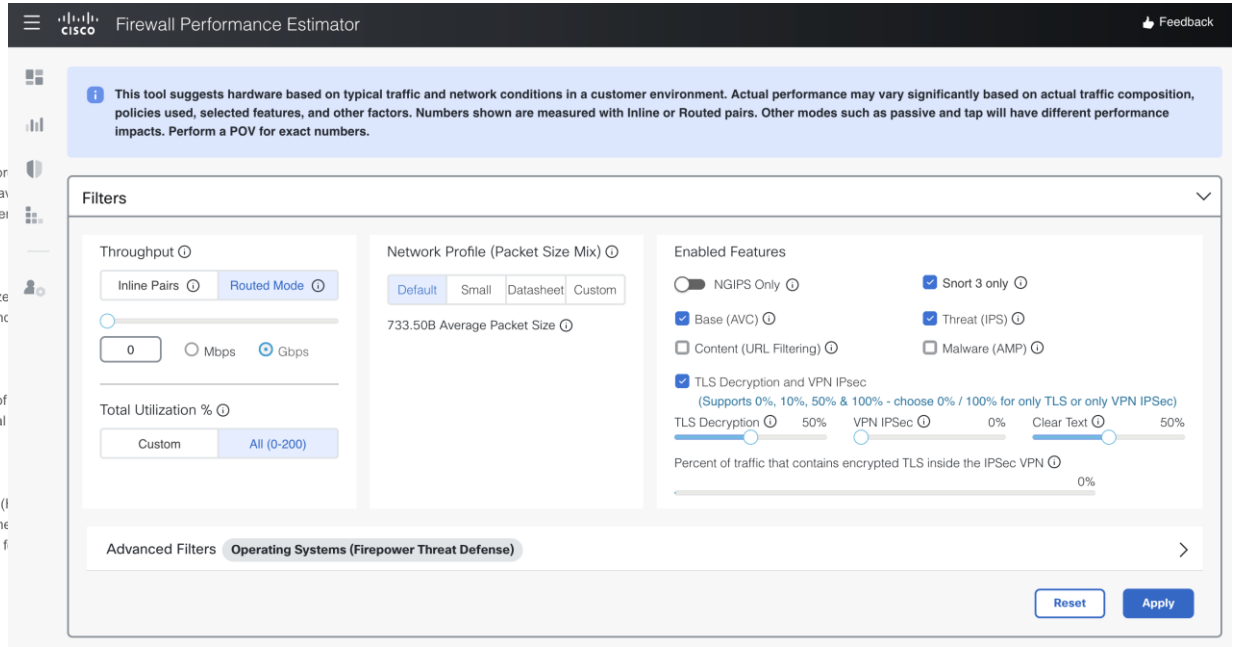

https://techzone.cisco.com/t5/FirePOWER-Threat-Defense/Testing-methodology-used-for-Cisco-Secure-Firewall-Threat/ta-p/1968099

Cisco Partners have access to: https://ngfwpe.cisco.com

How would you test your firewall?

Methodology? Tools?

Internet Engineering Task Force (IETF) — B. Balarajah
Request for Comments: 9411
Obsoletes: 3511 — C. Rossenhoevel, EANTC AG
Category: Informational — B. Monkman, NetSecOPEN
Published: March 2023
ISSN: 2070-1721

Benchmarking Methodology for Network Security Device Performance

Abstract

This document provides benchmarking terminology and methodology for next-generation network security devices, including next-generation firewalls (NGFWs) and next-generation intrusion prevention systems (NGIPSs). The main areas covered in this document are test terminology, test configuration parameters, and benchmarking methodology for NGFWs and NGIPSs. (It is assumed that readers have a working knowledge of these devices and the security functionality they contain.) This document aims to improve the applicability, reproducibility, and transparency of benchmarks and to align the test methodology with today's increasingly complex layer 7 security-centric network application use cases. As a result, this document makes RFC 3511 obsolete.

https://datatracker.ietf.org/doc/html/rfc9411

How would you test your firewall?

Methodology? Tools?

Application Traffic Mix Performance[1]

Key Performance Indicator	Healthcare traffic mix	Education traffic mix
Inspected Throughput	3,589 Mbit/s	3,164 Mbit/s
Application Transactions per second	15,030	17,691

Table 2: Results summary for application mix traffic test

HTTP Traffic Performance

Key Performance Indicator	Values
Connections Per Second (CPS)	42,366 CPS @ 1 KByte and 13,889 CPS @ 64 KByte object sizes
Inspected Throughput	11,254 Mbit/s @ 256 KByte and 922 Mbit/s @ 1 KByte object sizes
Transactions Per Second (TPS)	80,018 TPS @ 1 KByte and 5,241 TPS @ 256 KByte object sizes
Time to First Byte (TTFB)	1.53 ms average TTFB @ 1 KByte and 1.51 ms average TTFB @ 64 KByte object sizes[2]
Time to Last Byte (TTLB)	0.75 ms average TTLB @ 1 KByte and 1.63 ms average TTLB @ 64 KByte object sizes[2]
Concurrent connection	1,999,872 average concurrent connection

Table 3: Results summary for HTTP tests

HTTPS Traffic Performance

Key Performance Indicator	Values
Connections Per Second (CPS)	6,922 CPS @ 1 KByte and 4,927 CPS @ 64 KByte object sizes
Inspected Throughput	4,545 Mbit/s @ 256 KByte and 549 Mbit/s @ 1 KByte object sizes
Transactions Per Second (TPS)	38,352 TPS @ 1 KByte and 2,076 TPS @ 256 KByte object sizes
Time to First Byte (TTFB)	3.02 ms average TTFB @ 1 KByte and 3.01 ms average TTFB @ 64 KByte object sizes[2]
Time to Last Byte (TTLB)	1.01 ms average TTLB @ 1 KByte and 2.29 ms average TTLB @ 64 KByte object sizes[2]
Concurrent connection	149,040 average concurrent connection

Table 4: Results summary for HTTPS tests

https://www.netsecopen.org/_files/ugd/150f3f_c9447032940f4cff96855327329eb013.pdf

Generalized architecture view

Cisco Firewall Threat Defense Architecture

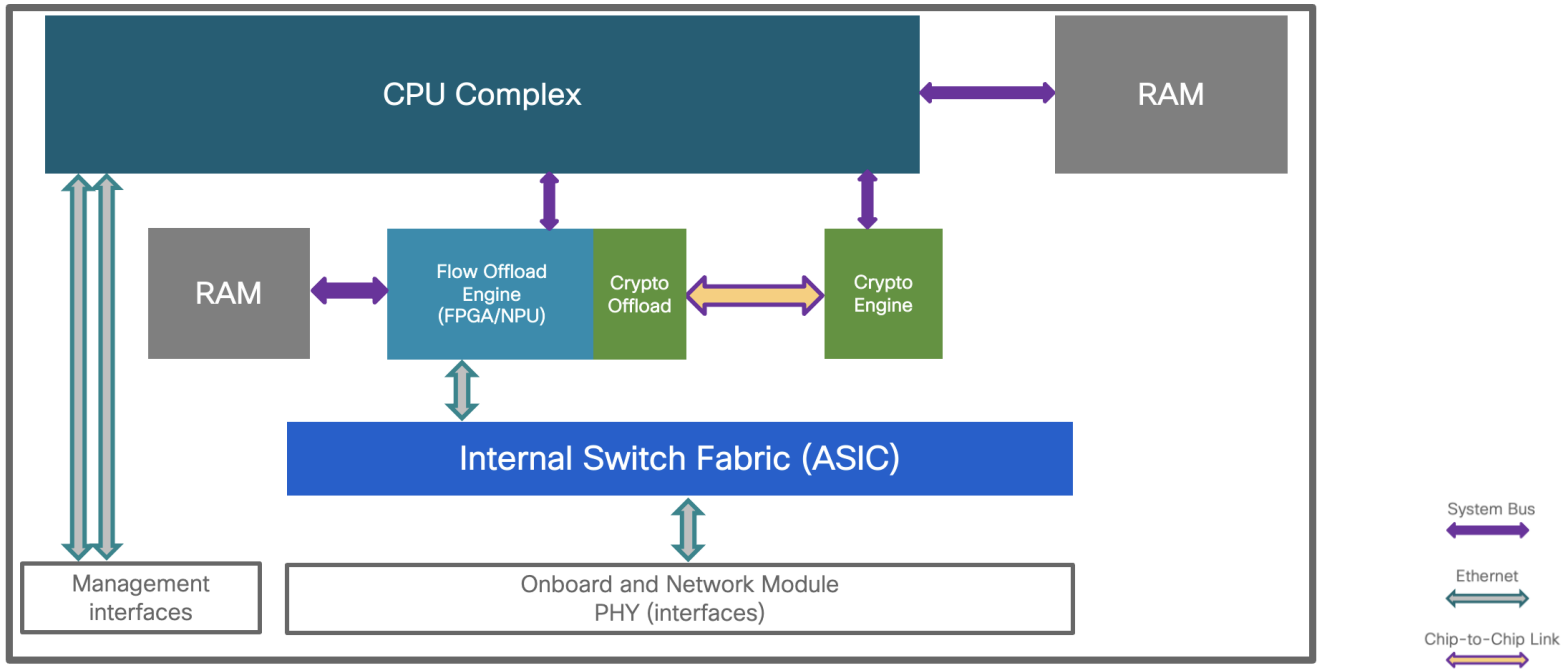

Generalized architecture view

Critical flow components

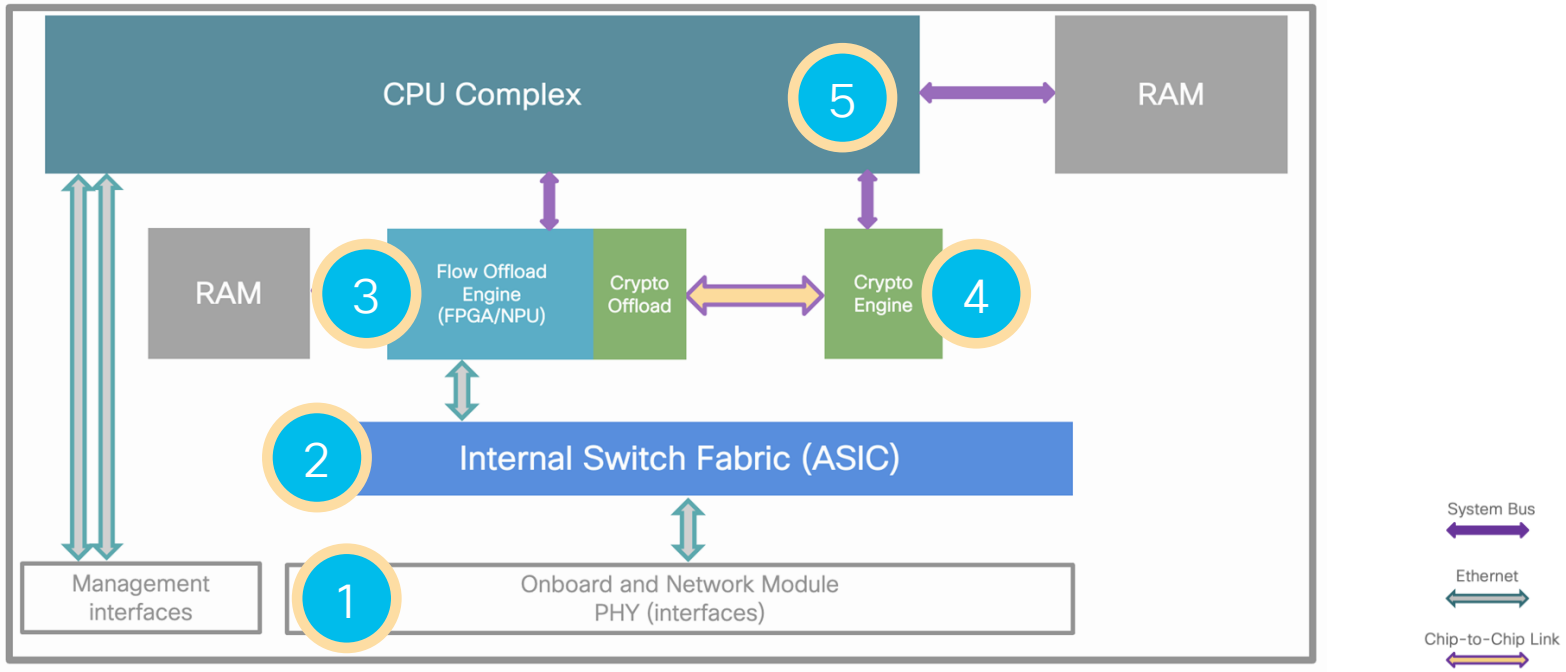

Why the architecture matters?

Traditional design – overall processing flow

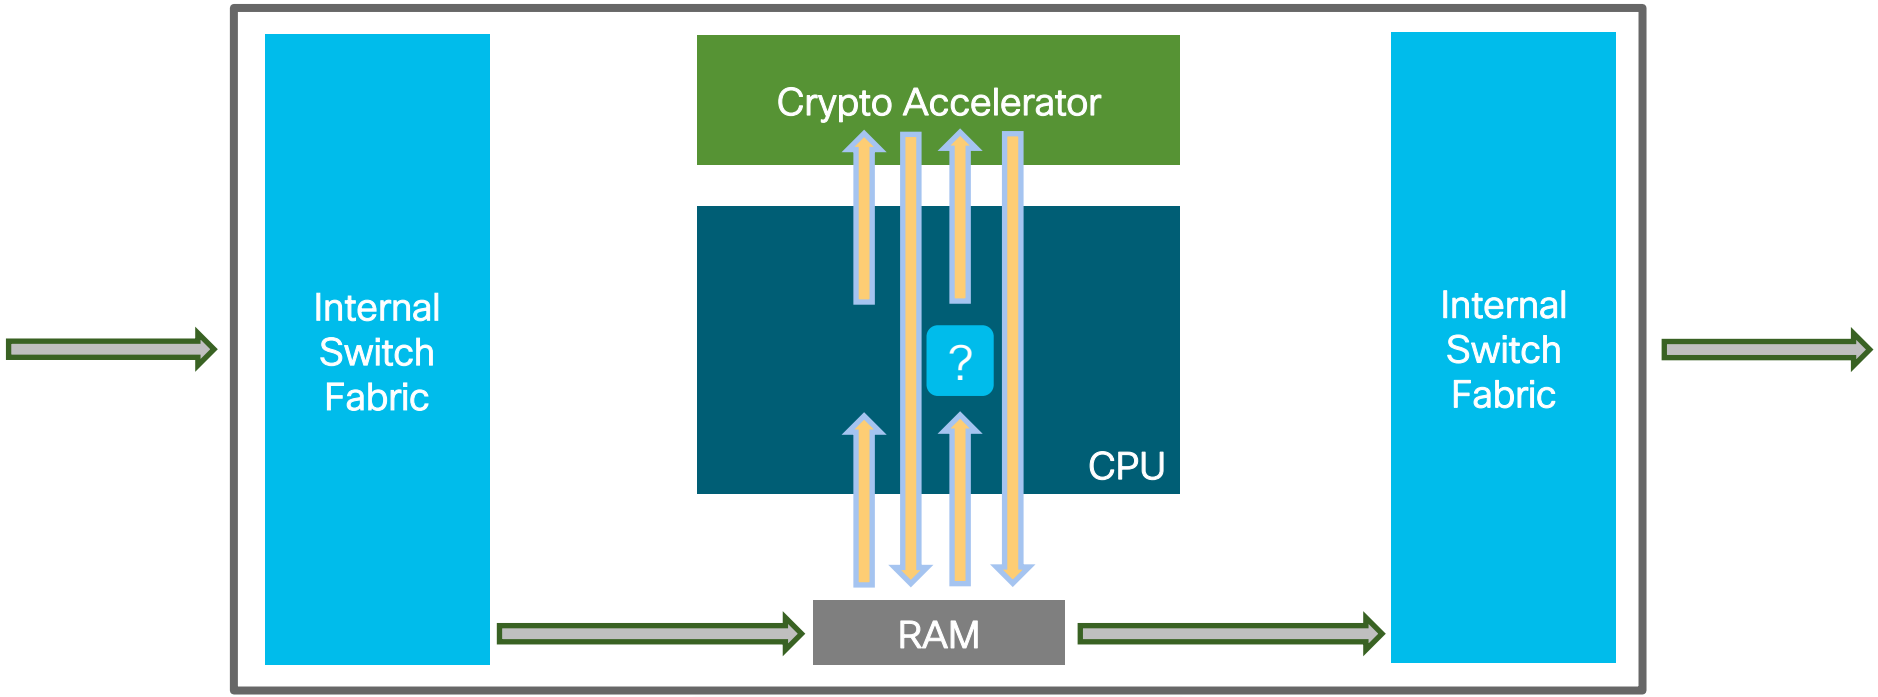

Why the architecture matters?

New Cisco design – inline processing with hardware offload

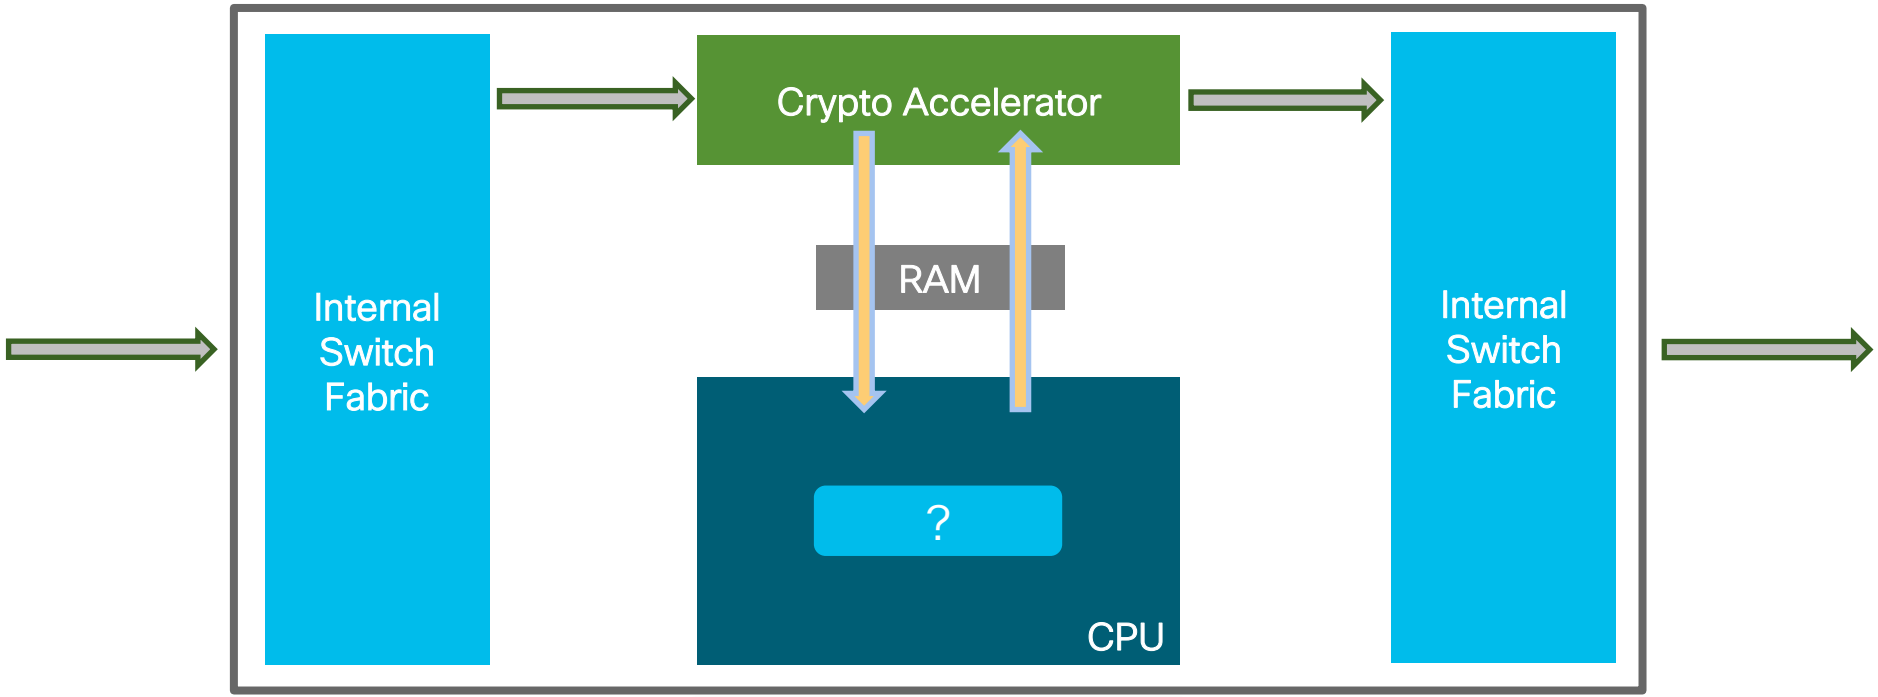

Ingress to device

Heating the planet part
AKA looking for threats

Egress from device

Configurable CPU Core Allocation

FTD 7.3+

- FTD had a static CPU core allocation between Data Plane and Snort

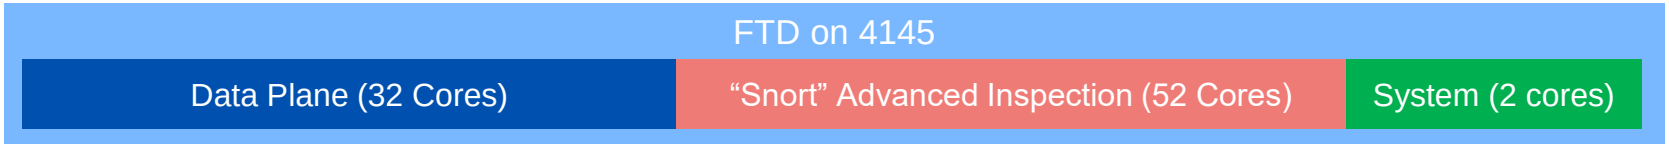

- Tailor FTD to a specific use case with a configurable allocation
  - Select from a few templates in FTD 7.3; dynamic in the future
  - VPN headend or basic stateful firewall would use more Data Plane cores
  - Heavy IPS and file inspection would bias toward more “Snort” cores
- 7.4.1 brings support for 3100 & 4200
  - support already on FTDv, 4100, 9300

Configurable CPU Core Allocation

FTD 7.3+

- FTD had a static CPU core allocation between Data Plane and Snort

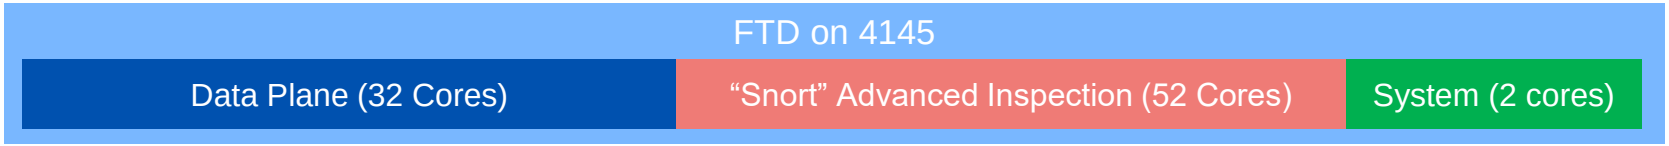

Name	Core allocation
Default	Normal for balanced FTD system
VPN heavy with prefilter	90% cores for data plane, 10% for Snort
VPN heavy	60% cores for data plane, 40% for Snort
IPS heavy	30% cores for data plane, 70% for Snort

Single-Flow Performance Considerations

- A single stateful flow must be processed by one processor core at a time
  - Trying to share a complex data structure leads to race conditions
  - Stateless parallel processing leads to out-of-order packets
- No magic trick to single-flow throughput
  - Deploy more powerful CPU cores
  - Reduce the amount of security inspection
- Pay performance price for real security
  - ...or deploy a router or a switch instead

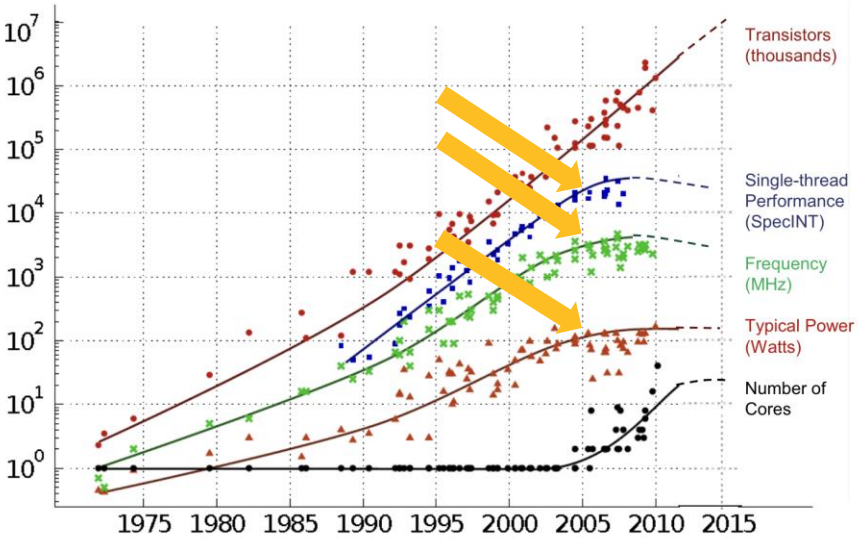

Source:
https://science.osti.gov/-/media/ascr/ascac/pdf/reports/2013/SC12_Harrod.pdf
https://www.lanl.gov/conferences/salishan/salishan2011/3moore.pdf

Managing Single-Flow Throughput

- Roughly estimated as overall throughput divided by Snort cores
  - 145Gbps of 1024-byte AVC+IPS on 4245 / 63 Snort cores = ~2.3Gbps
  - 65Gbps of 1024-byte AVC+IPS on 4215 / 15 Snort cores = ~4.3Gbps
  - Egress Optimization improves throughput by up to 20% in FTD 6.4 NGIPS mode, and in some VPN scenarios with 7.0
  - Reducing impact on all flows from few Superflows is more important
- "What does your security policy tell you to do?"
  - NGFW performance capacity must not dictate your security policy
  - Flow Offload vs Snort 3 Elephant Flow Offload (7.2+) or Intelligent Application Bypass (IAB) (pre 7.2)

Elephant Flow Detection

Per-flow tracking replaces Intelligent Application Bypass (IAB)

FTD 7.2

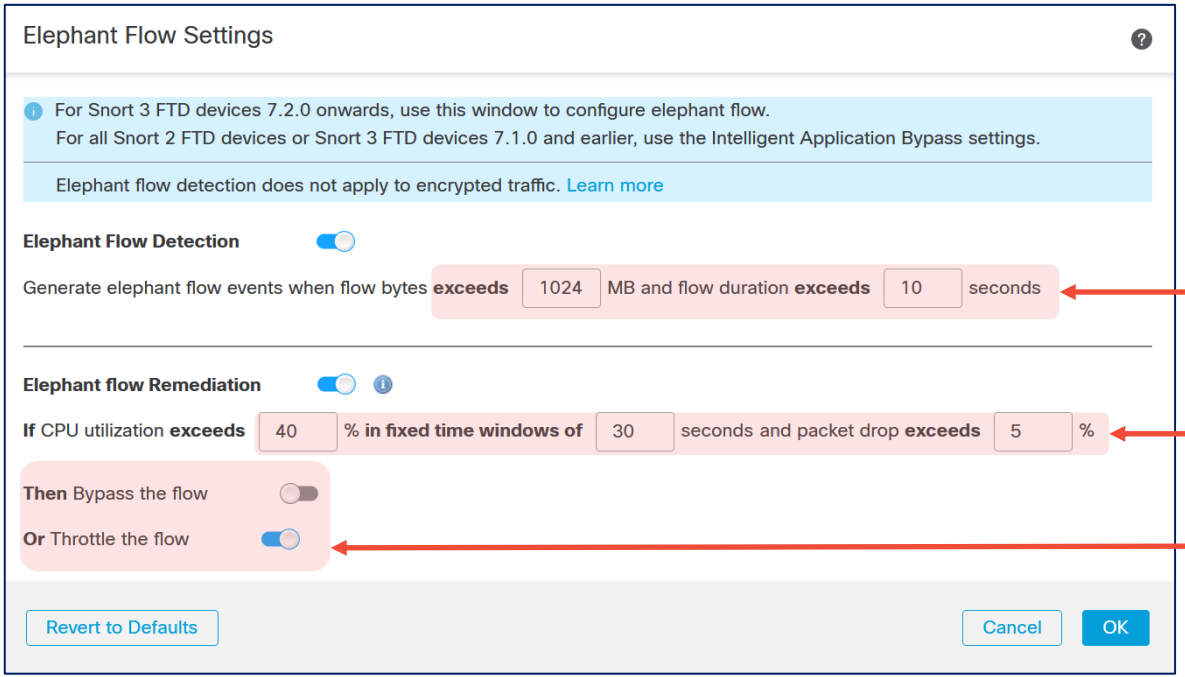

Throughput threshold to qualify as an Elephant Flow

Optional flow-specific CPU resource consumption and packet drop thresholds for remediation.

Optional flow remediation actions.

Flow Offload Operation

Full Inspection

- Dynamically program Offload engine after flow establishment
- Ability to switch between Offload and full inspection on the fly

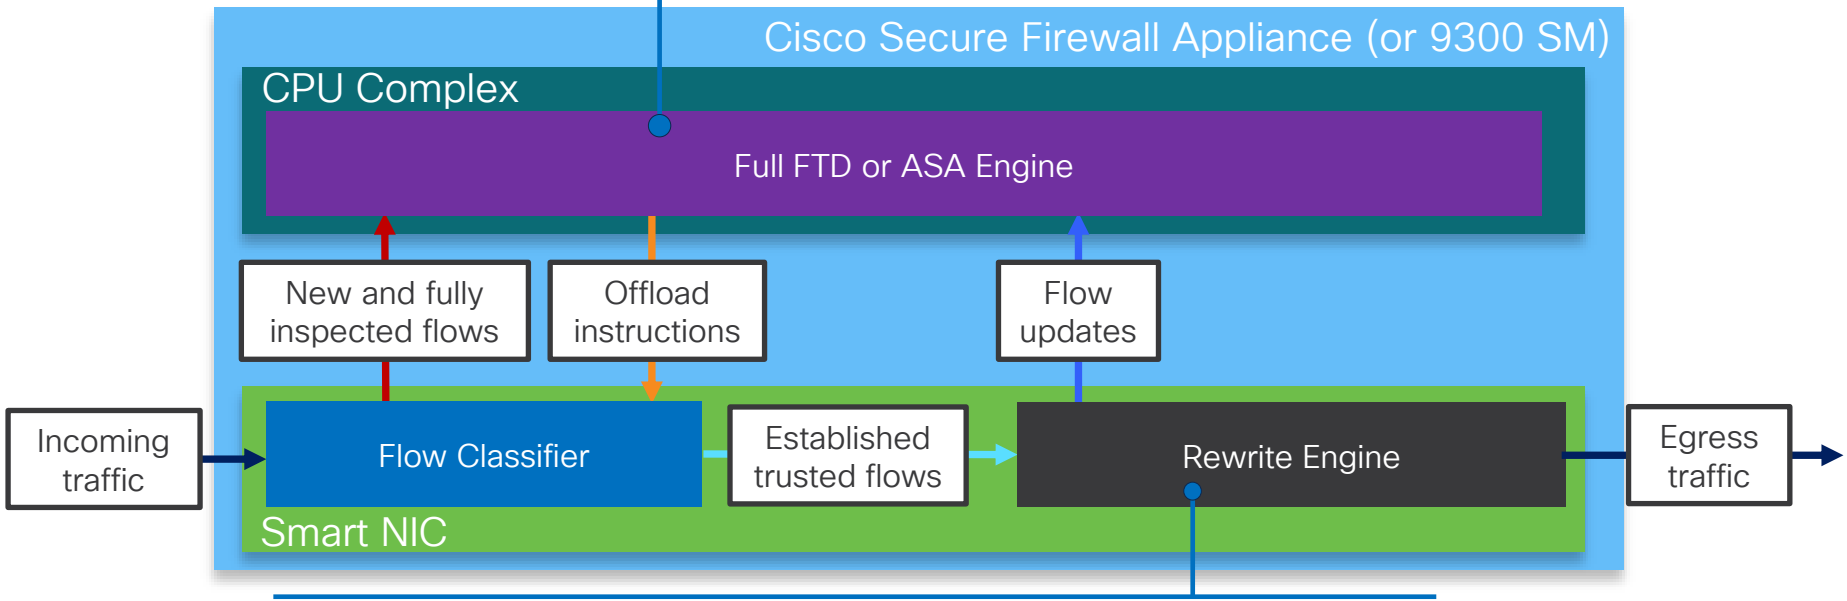

Flow Offload

- Limited state tracking, NAT/PAT, TCP Seq Randomization, <5µs for 64B UDP traffic

Dynamic Flow Offload for 3100 & 4200

FTD 7.7

Supported for IPv4 flows with Snort 3

- Snort may mark flow as trusted in following use cases:
  - AC Policy with Action set to Trust
  - Elephant Flow Offload or Intelligent Application Bypass (IAB) Policy match to Trust
  - File Policy with Detection Action
  - IPS Policy that leads to Trust
- Much higher scale than in 4100/9300
- Much more effective hash algorithm as well (>50%)

Scale out encryption in clustering

FTD 7.7

Enabling Security Gateway use cases for Mobile Core Protection

- IPsec Cluster Offload
  - IPsec is fully accelerated (offloaded to data plane – dedicated cryptographic hardware) by distributed cluster members
- Distributed Control Plane for IKE & IPsec across Cluster
  - Enabling processing of IKE and IPsec traffic on the node that becomes flow owner rather than centralizing control plane only on cluster control unit (mode available so far only on 9300)
- Cluster Hardware Redirect
  - Offload traffic redirected using CCL (Cluster Control Link) with hardware (directly via FPGA) without involving CPU

Virtual Firewall on Data Processing Unit (DPU)

Future

- Network Interface Controller (NIC) with a DPU in a server or switch
  - Inline hardware acceleration for broad packet processing functionality
  - Perfect opportunity to accelerate and scale firewall in hybrid data centers

ASAv/FTDv software and Multicloud Defense is deployed on x86 CPU in generic private and public cloud environments.

If a DPU is present, additional ARM software components program inline acceleration of flow processing, IPsec and (D)TLS encryption, and other capabilities.

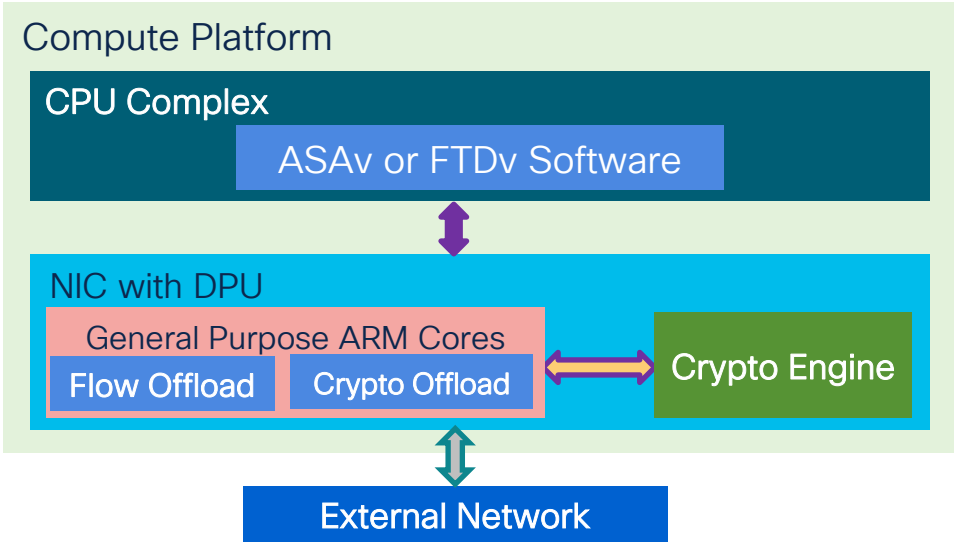

Scale
Considerations

”What’s maximum size of policy I can use?”

ACE = Access Control Entry, ACP = Access Control Policy

- Starting from 7.2, FTD by default uses OGS on greenfield deployments
  - OGS = Optimized Group Search
  - OGS allows for higher scale for policies and connections per second, at the expense of per-packet performance
- With 7.6, OGS implementation was upgraded, to handle more corner cases, execute with higher scale and provide hit counters (and timestamps) also on folded entries
  - this was further improved on 7.7 with new corner cases we’ve found
- While FMC will warn you before deploying rulesets close to those limits, please use following slide as guidance only and consult your Partner or Cisco Security Specialist before deploying policies

Maximum supported policy sizes for FTD

As of release 7.6

Appliance model	Maximum tested FTD ACEs	UI Rule Count (assuming 1 rule expands to 50 ACEs)	UI Rule Count (assuming 1 rule expands to 100 ACEs)
1010/1010E	10,000	200	100
1120	90,000	1,800	900
1140	110,000	2,200	1,100
1150	185,000	3,700	1,850
1200C	50,000	1,000	500
2110	60,000	200	100
2120	100,000	1,800	900
2130	250,000	2,200	1,100
2140	500,000	3,700	1,850

Maximum supported policy sizes for FTD

As of release 7.6

Appliance model	Maximum tested FTD ACEs	UI Rule Count (assuming 1 rule expands to 50 ACEs)	UI Rule Count (assuming 1 rule expands to 100 ACEs)
3105	2,750,000	55,000	27,500
3110	2,750,000	55,000	27,500
3120	3,000,000	60,000	30,000
3130	3,500,000	70,000	35,000
3140	4,000,000	80,000	40,000
4112	2,000,000	40,000	20,000
4115	4,000,000	80,000	40,000
4125	5,000,000	100,000	50,000
4145	8,000,000	160,000	80,000

Maximum supported policy sizes for FTD

As of release 7.6

Appliance model	Maximum tested FTD ACEs	UI Rule Count (assuming 1 rule expands to 50 ACEs)	UI Rule Count (assuming 1 rule expands to 100 ACEs)
4215	6,000,000	120,000	60,000
4225	8,000,000	160,000	80,000
4245	10,000,000	200,000	100,000
9300 w/SM-40	6,000,000	120,000	60,000
9300 w/SM-48	8,500,000	170,000	85,000
9300 w/SM-56	9,500,000	190,000	95,000

Designing for High Availability

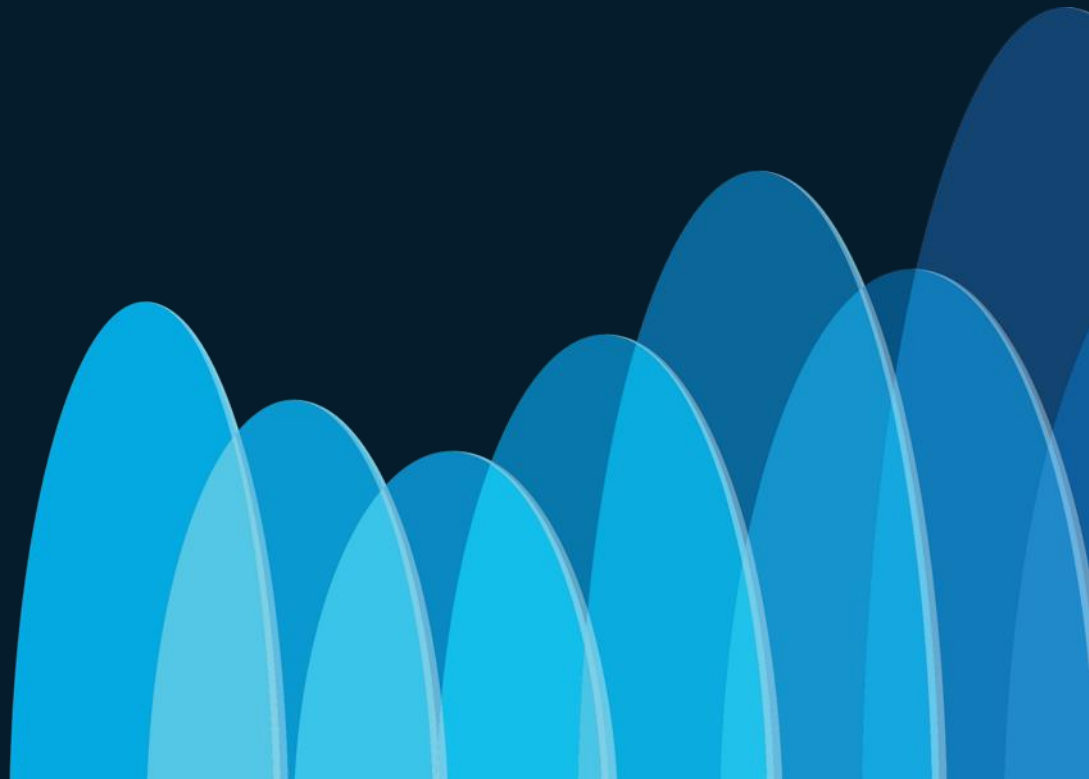

cisco Live!

How to achieve high scale & redundancy?

That's a philosophical question

- HA or Clustering
- HA = Active/Standby (Active/Active for ASA with multi-context)
- Clustering = true horizontal scaling: with every device added you add capacity to handle traffic and scale to do so
- Clustering howtos for:
  - 3100/4200 FTD: https://www.cisco.com/c/en/us/td/docs/security/secure-firewall/management-center/cluster/ftd-cluster-sec-fw.html
  - 3100/4200 ASA: https://www.cisco.com/c/en/us/td/docs/security/asa/special/cluster-sec-fw/secure-firewall-cluster.html
  - 4100/9300 FTD: https://www.cisco.com/c/en/us/td/docs/security/firepower/fxos/clustering/ftd-4100-9300-cluster.html
  - 4100/9300 ASA: https://www.cisco.com/c/en/us/td/docs/security/firepower/fxos/clustering/asa-cluster-solution.html




FTD High Availability and Clustering

- FTD inherits failover and clustering infrastructure from ASA
  - Replicates full NGFW/NGIPS configuration and opaque flow state
  - Supports all NGFW/NGIPS interface modes
  - Interface and Snort instance (at least 50%) health monitoring
  - Zero-Downtime upgrades for most applications
- Ensures full stateful flow symmetry in both NGIPS and NGFW modes

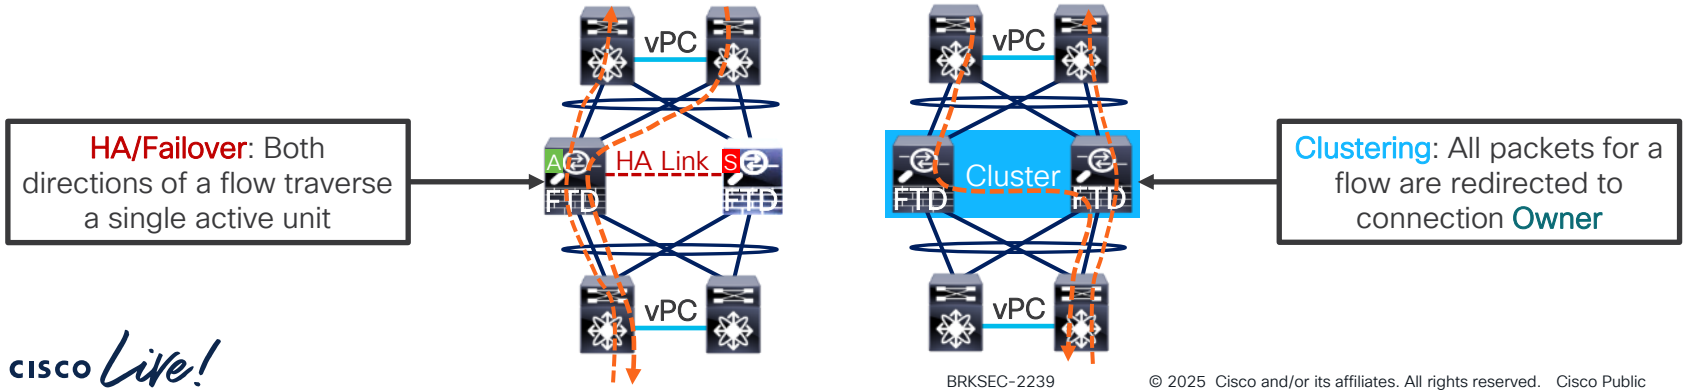

Firewalling with Redundancy

Standard High Availability – “Active/Standby” concept

Minimal impact on switchover

FTD

ASA

Active unit – control & data plane

Standby unit – control & data plane

Active unit – control & data plane

Standby unit – control & data plane

Failover event
Some form of failure detected or manual switchover

Firewalling with Redundancy

All Active Mode – “Clustering” concept

No impact on cluster node loss, join or upgrade*

FTD

ASA

Clustering – example for 3140

Active unit – control & data plane	45Gbps, 6M conn 300k cps
Active unit – control & data plane	72Gbps, 12M conn 300k cps
Active unit – control & data plane	108Gbps, 18M conn 450k cps
Active unit – control & data plane	144Gbps, 24M conn 600k cps
Keep getting more active units	Each unit adds scale and performance
Active unit – control & data plane	576Gbps, 96M conn 784k cps

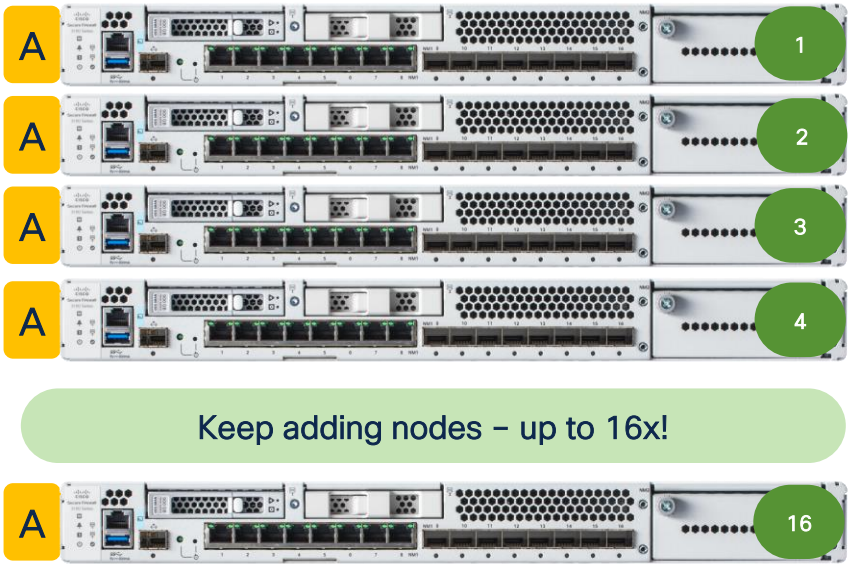

example for NGFW 1024B profile

* for non-centralized features and protocols

New TCP Flow with FTD Inter-Chassis Clustering

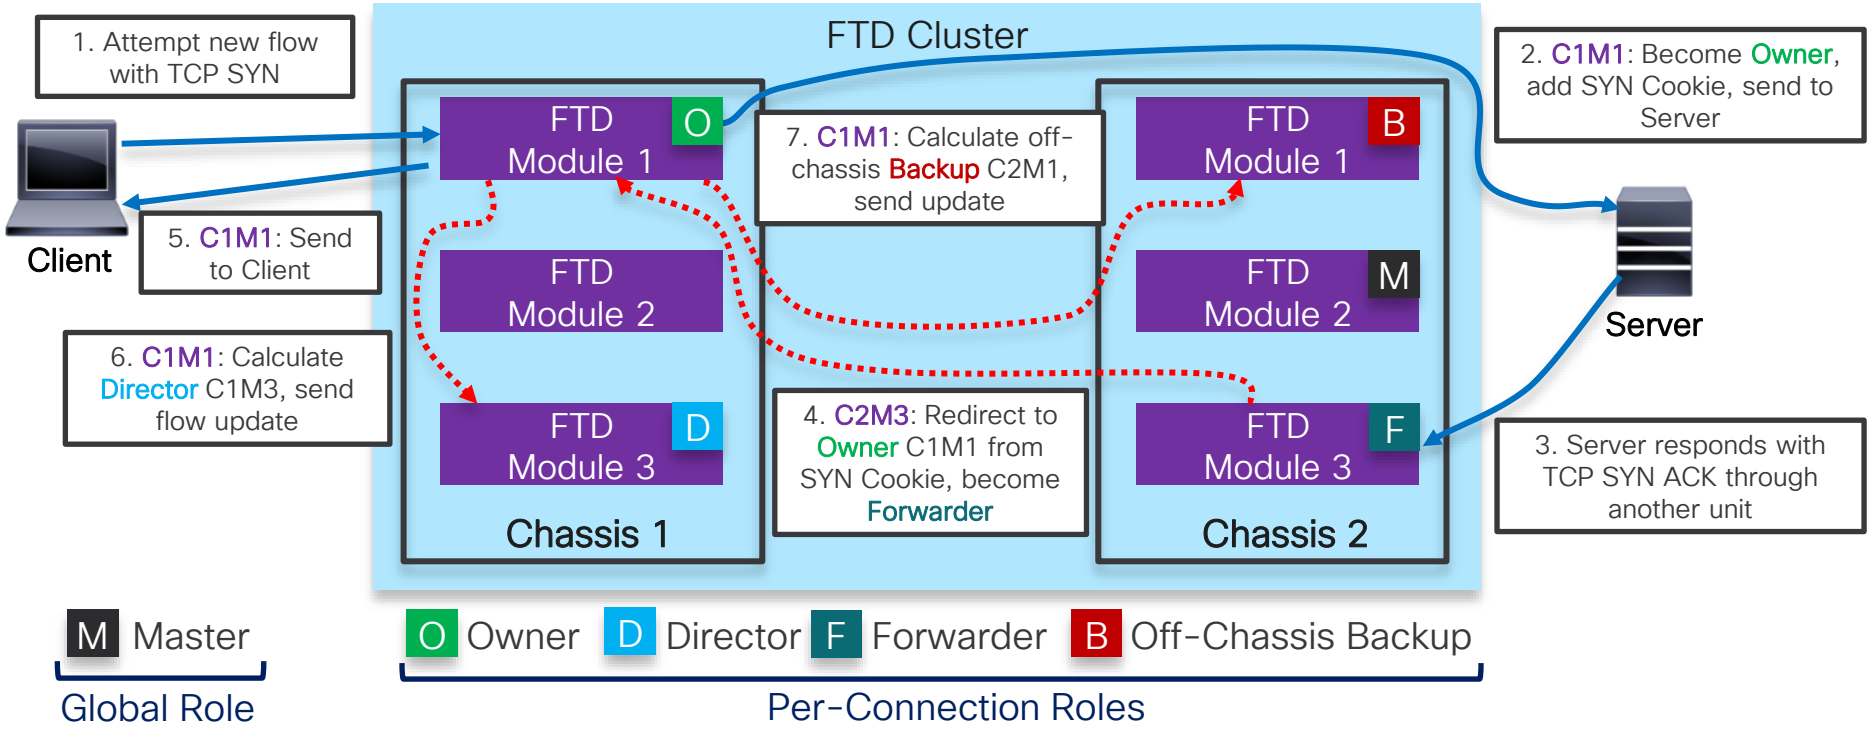

Secure Firewall Clustering sizing

There are three major factors in calculating cluster performance and scale (1/3)

- **Throughput**
  - for L2 assume **80%** of combined maximum throughput of all members
  - for modern switches that can do L2 etherchannel load-balancing using L2/L3/L4 information even when just forwarding L2 frames, and for L3 routing deployments this factor can go up to **100%**
  - **example for FTD:** cluster of 4x 3140 has NGFW 1024B profile maximum throughput of 144Gbps (4x 45Gbps * 0,8)
  - **example for ASA:** cluster of 4x 3140 has ASA multiprotocol profile maximum throughput of 137.6Gbps (4x 43Gbps * 0,8)

Note:
Theoretical maximum for NGFW 1024B profile for:
- 16x 3140 – 0.57Tbps
- 16x 4245 – 1.79Tbps

Secure Firewall Clustering sizing

There are three major factors in calculating cluster performance and scale (2/3)

- **Connections per second**
  - due to additional tasks associated with the flow creation process, assume nodes can do up to **50%** of their rated connections per second
  - **example for FTD:** cluster of 4x 3140 has maximum of 600k cps (4x 300k * 0,5)
  - **example for ASA:** cluster of 4x 3140 has maximum of 2.2M cps (4x 1.1M * 0,5)

Note:
Theoretical maximum for FTD:
- 16x 3140 – 2.4M cps
- 16x 4245 – 6.4M cps

Secure Firewall Clustering sizing

There are three major factors in calculating cluster performance and scale (3/3)

- **Maximum connections**
  - as cluster members maintain additional stub connection, assume maximum number of sessions at a level of 60% of combined scale
  - example for FTD: cluster of 4x 3140 can hold up to 24M of connections (4x 10M * 0,6)
  - example for ASA: cluster of 4x 3140 can hold up to 24M of connections (4x 10M * 0,6)

Note:
Theoretical maximum for FTD:
- 16x 3140 – 96M cps
- 16x 4245 – 576M cps

How to achieve high scale & redundancy?

Advanced setup – geo-redundant cluster, with traffic localization

- North-South insertion with LISP inspection and owner reassignment
- East-West insertion for first hop redundancy with VM mobility
- Underlying fabric can be anything transporting Ethernet with RTT up to 20ms
  - ideally – dark fiber
  - also tested – VPLS, VPWS, EVPN

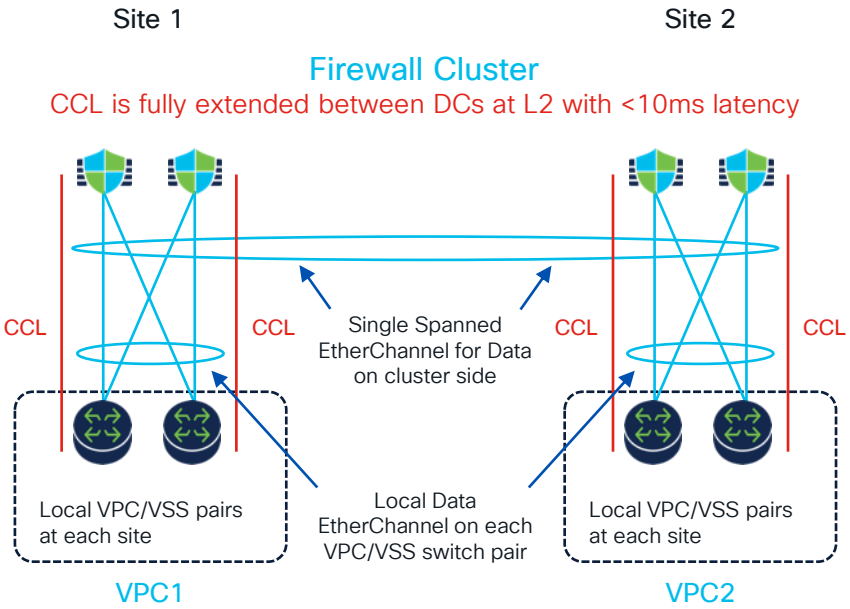

Clustering for Virtual Firewalls

FTD 7.1+

ASA 9.17+

- Clustering combines multiple firewalls into one logical device
  - Seamless scalability up to 16 FTD units with no traffic disruption
  - Stateful handling of asymmetric traffic and failure recovery
  - Single point of management and unified reporting
- Better elasticity and failure handling in hybrid cloud with clustering
  - aws Google Cloud Platform Azure KVM vmware
  - Individual data interface IP addresses instead of a single Port-channel
  - VxLAN-based Cluster Control Link for unicast control plane
  - No source NAT requirement for handling traffic asymmetry
  - Existing flow re-hosting on failure in supported environments

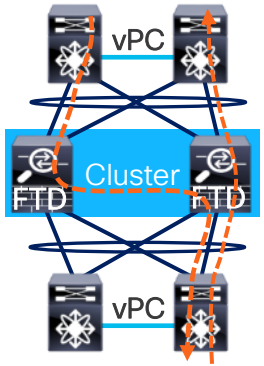

Cluster Health Dashboard

FMC 7.3

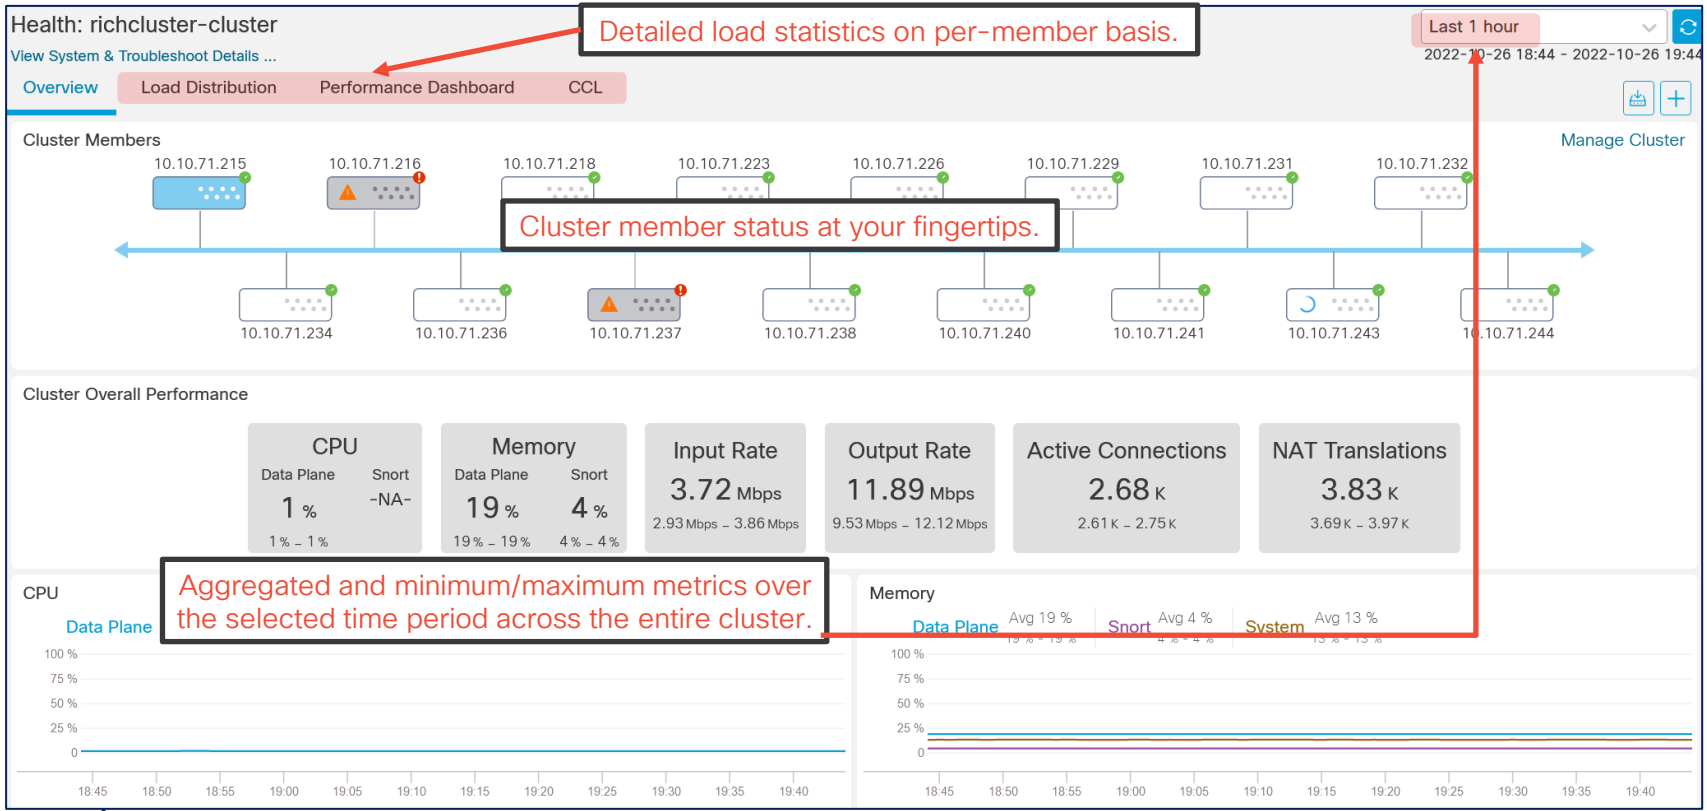

Cluster Enhancements

Layer 3 insertion at the edge

FTD 7.6

ASA 9.22

Individual Interface Mode

- Layer 3
- Load-balancing via routing: PBR, ITD, static ECMP or ECMP with dynamic routing
- Routed mode
- FTDv & 3100/4200

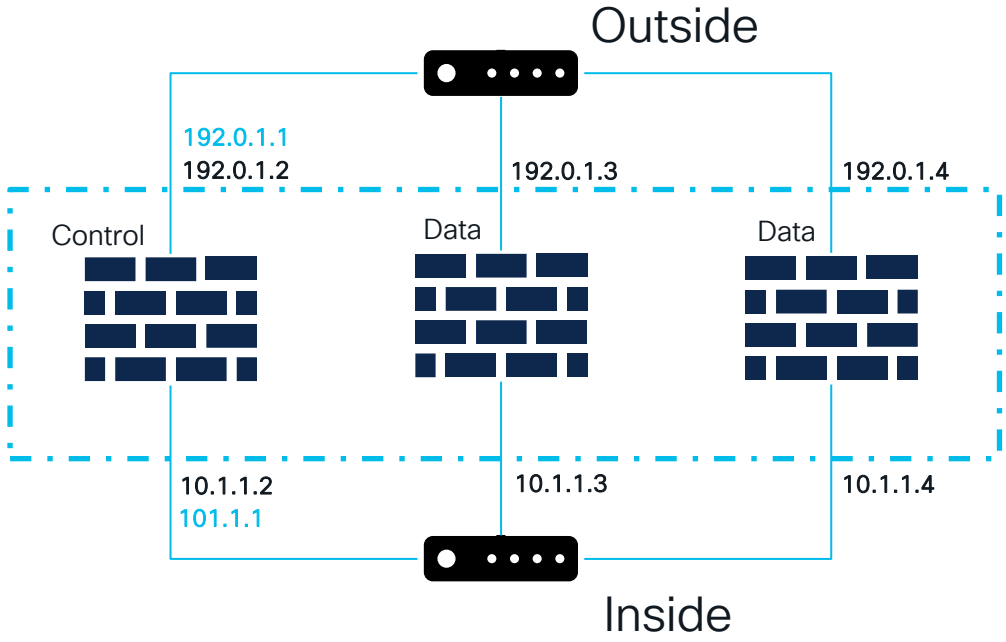

Cluster Enhancements

FTD 7.6

ASA 9.22

Fully routed mode for FTDv, 3100 and 4200

- On legacy ASA hardware, both spanned and routed clustering modes were supported
- Since then, we supported only spanned as that was initially most popular for Enterprise/DC high scale deployments
- With routed mode gaining more and more popularity (ECMP/UCMP), we're bringing routed/individual mode back
- Each unit runs its own as independent routing instance
- Feature supported with multi-context mode (ASA), but not (yet) on Multi-Instance as clustering support is coming soon

Cluster Enhancements

Fully routed mode for FTDv, 3100 and 4200

FTD 7.6

ASA 9.22

Appliance model	Spanned Mode Cluster	Individual Mode Cluster
Layer used for ingress/egress traffic	L2	L3
Data Interface	Grouped to form a single spanned EtherChannel across all nodes	Each data interface has its own IP address received from cluster pool
Data Traffic Load Balancing	Handled by EtherChannel (upstream and downstream switches)	Uses ECMP/UCMP or PBR for load balancing (upstream and downstream routers)
Routing Modes	Routed or Transparent mode	Routed mode only

Designing for Multi-Tenancy

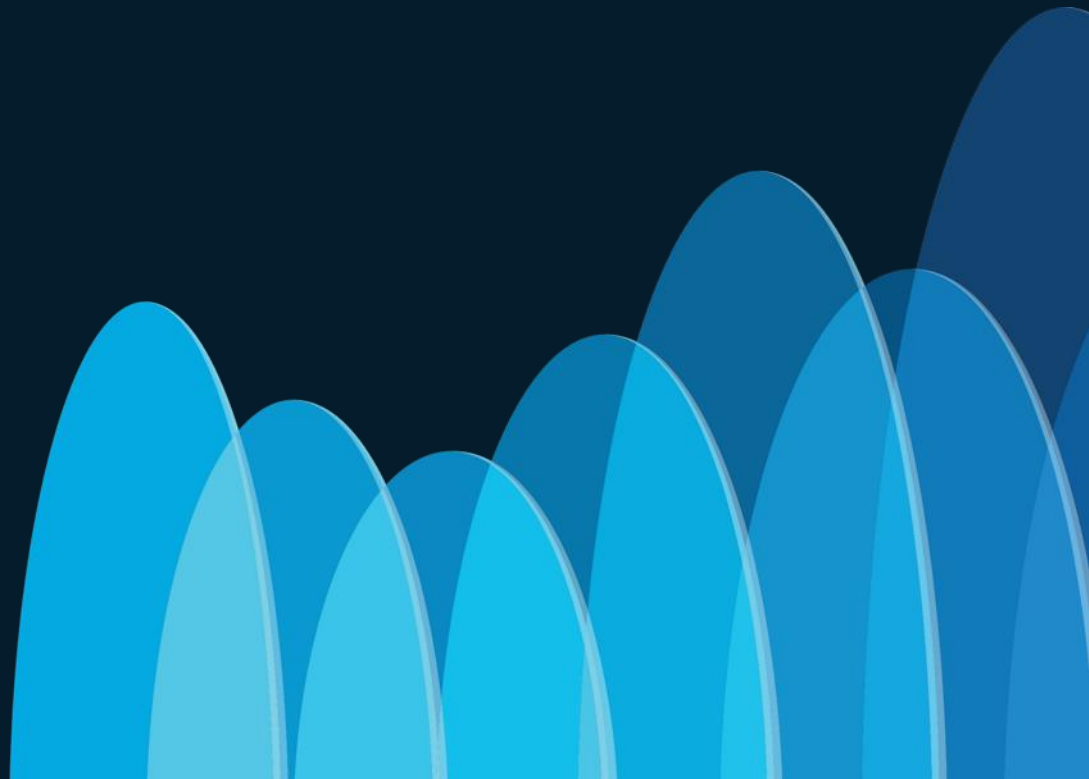

cisco Live!

Multi-tenancy at scale

Granular RBAC, separation using domains, VRFs and Multi-Instance

- Users see only devices assigned within their domain (up to 1024)
- FMC RBAC provides granular separation of duties between operators
- Multi-Instance and VRFs can be mixed in the same environment

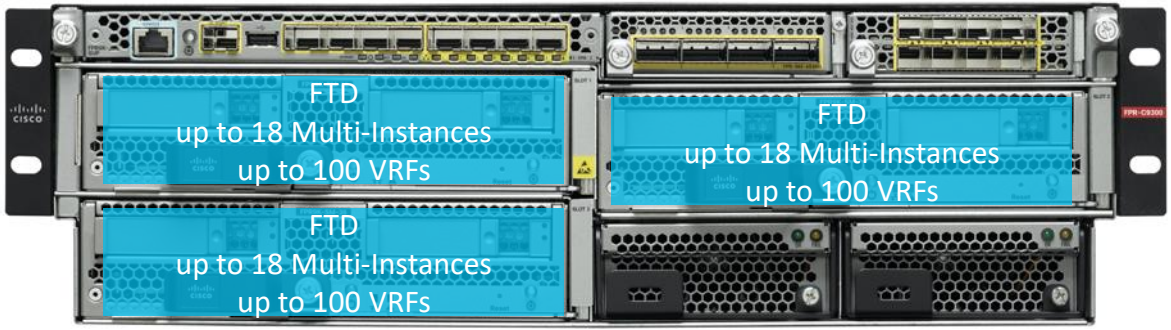

9300 service chaining – ASA + FTD

Unique capability for chassis with multiple Service Modules

- Example configuration:
  - SM-40 for ASA RA VPN duties
    up to 20k tunnels, and up to 15Gbps DTLS throughput
    with 450 byte packets
  - SM-56 for FTD NGFW/NGIPS duties
    up to: 64Gbps of NGFW (IPS+AVC) throughput,
    35M connections, 490K CPS, 12Gbps TLS inspection
    (50% of overall traffic)

Decrypted traffic from AnyConnect sessions terminated at ASA moves to inspection by NGFW/NGIPS, on the way back is again encrypted by ASA and sent to remote endpoint

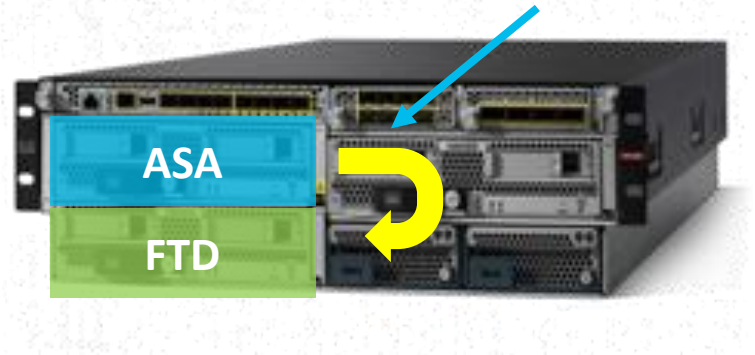

Incoming AnyConnect users – full RA VPN feature set on ASA

Incoming traffic to NGFW/NGIPS protected services in DMZ

Outgoing traffic from NGFW/NGIPS protected users & AnyConnect users (if working with centralized internet access)

Available from FXOS 2.6(1), ASA 9.12(1) and FTD 6.4.0:
https://www.cisco.com/c/en/us/td/docs/security/firepower/fxos/fxos261/release/notes/fxos261_rn.html#id_113895

Virtual Routing and Forwarding (VRF) Lite

- Starting from FTD 6.6, interfaces can be in different Routing Domains
  - Overlapping IP address support between user and Global VRF
  - Traffic forwarding between different VRF with static routes and NAT

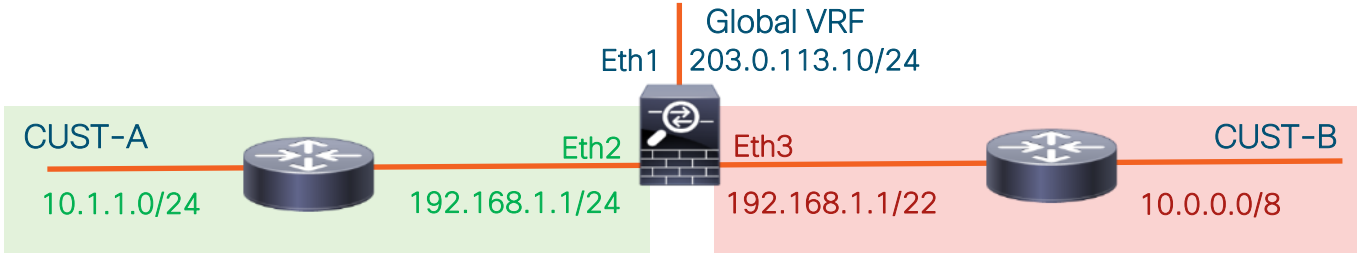

- Existing single security policy across all VRFs, no per-VRF rules
  - Connection events are enriched with VRF ID for usability
- Can be combined with FTD multi-instance

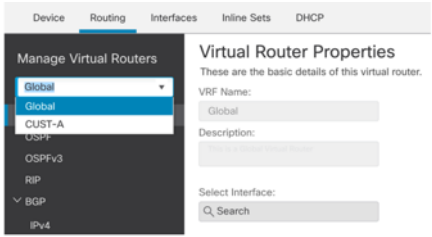

Multi-tenancy at scale

“How to achieve massive scale” (for Fun & Profit)

Interface	Logical Name	Type	Security Zones	Virtual Router
Diagnostic0/0	diagnostic	Physical		Global
GigabitEthernet0/0		Physical		
GigabitEthernet0/0.100	T10_GI0_INSIDE	SubInterface	T10_INSIDE	T10
GigabitEthernet0/0.101	T11_GI0_INSIDE	SubInterface	T11_INSIDE	T11
GigabitEthernet0/1		Physical		
GigabitEthernet0/1.200	T10_GI1_OUTSIDE	SubInterface	T10_OUTSIDE	T10
GigabitEthernet0/1.201	T11_GI1_OUTSIDE	SubInterface	T11_OUTSIDE	T11
GigabitEthernet0/2	Passive	Physical		
GigabitEthernet0/3		Physical		

Multi-tenancy at scale

“How to achieve massive scale” (for Fun & Profit)

Interface	Logical Name	Type	Sec
Diagnostic0/0	diagnostic	Physical	
GigabitEthernet0/0		Physical	
GigabitEthernet0/0.100	T10_GI0_INSIDE	SubInterface	T10
GigabitEthernet0/0.101	T11_GI0_INSIDE	SubInterface	T11
GigabitEthernet0/1		Physical	
GigabitEthernet0/1.200	T10_GI1_OUTSIDE	SubInterface	T10
GigabitEthernet0/1.201	T11_GI1_OUTSIDE	SubInterface	T11
GigabitEthernet0/2	Passive	Physical	
GigabitEthernet0/3		Physical	

Virtual Router	Interfaces
Global	diagnostic
T10	T10_GI1_OUTSIDE, T10_GI0_INSIDE
T11	T11_GI1_OUTSIDE, T11_GI0_INSIDE

Multi-tenancy at scale

“How to achieve massive scale” (for Fun & Profit)

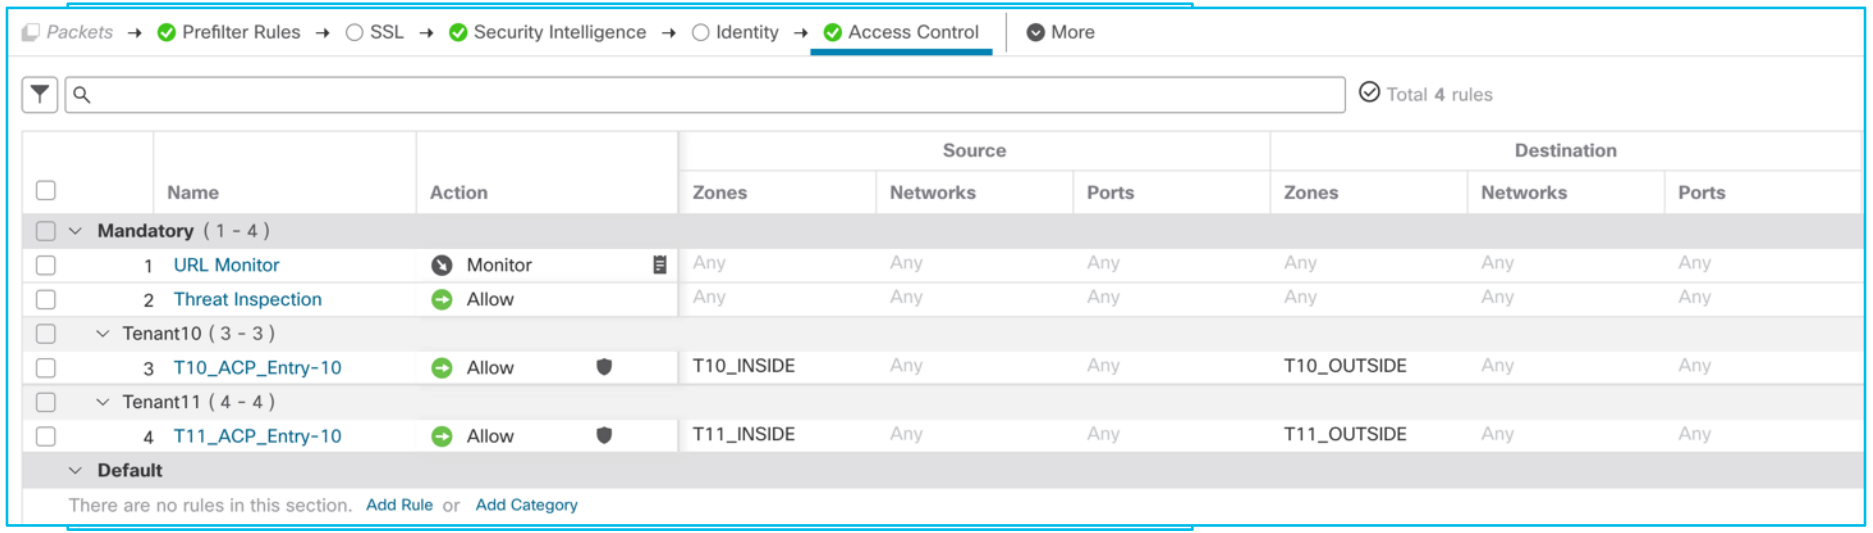

VRF Scalability as for FTD 7.7

Current generation platforms

Platform	VRF Count	Platform	VRF Count	Platform	VRF Count
1010/1120	5	**2110**	10	**4112**	60
1140	10	**2120**	20	**4115**	80
1150	10	**2130**	30	**4125/45**	100
		2140	40		
1210CE/CP	5				
1220CX	10			**4215/25/45**	100
		3105	10		
		3110	15	**9300 SM-44/48/56**	100
1230	10	**3120**	25		
1240	10	**3130**	50	**FTDv**	30
1250	15	**3140**	100	**ISA 3000**	10

NEW 7.7

VRF Scalability as of last FTD version supported

Previous generation platforms

Platform	VRF Count	Platform	VRF Count
ASA5508-X	10	9300 SM-24	100
ASA5516-X	10	9300 SM-36	100
ASA5525-X	10	9300 SM-40	100
ASA5545-X	20		
ASA5555-X	20		
4110	60		
4120	80		
4140	100		
4150	100		

Multi-Instance Capability Summary

Supported on 3100, 4100, 4200 and 9300

- Instantiate multiple logical devices on a single module or appliance
  - FTD application in 6.3 for 4100 and 9300
  - FTD application in 7.6 for 4200 and 7.4.1 for 3100
  - Leverage Docker infrastructure and container packaging
- Allows tenant management separation, independent instance upgrade and resource protection

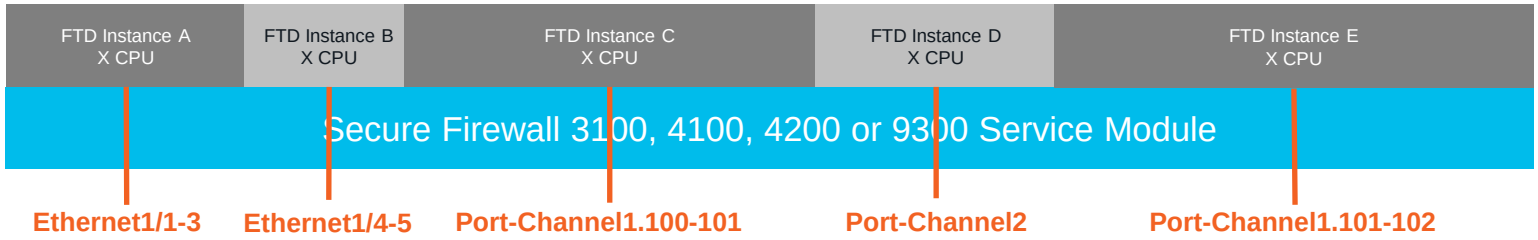

Multi-Instance Mode

Full migration and configuration support in FMC for 3100 and 4200

FTD 7.6

FMC 7.6

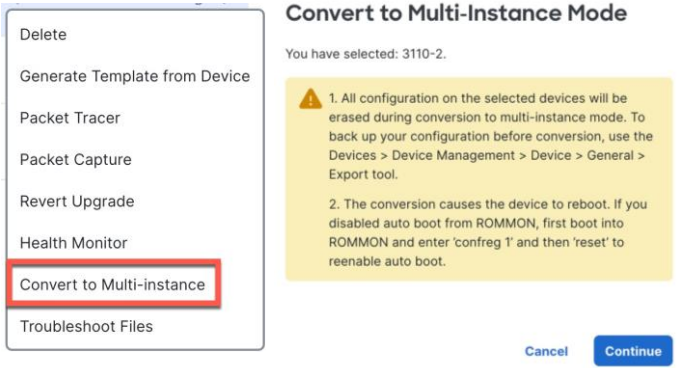

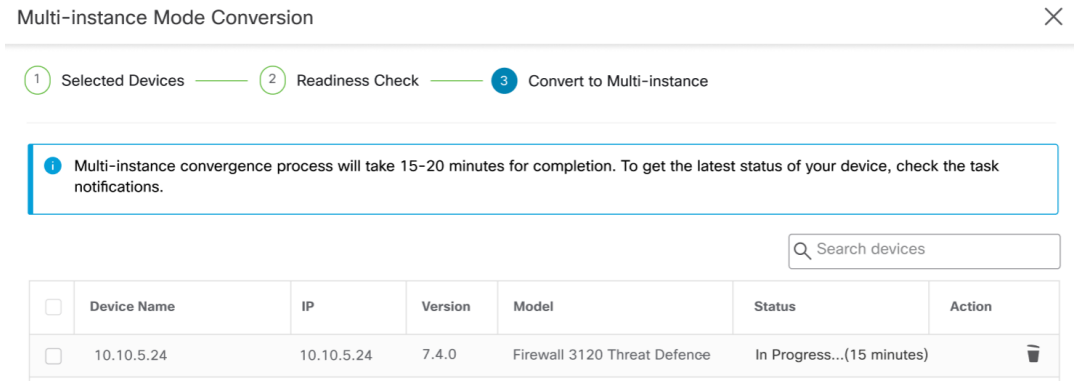

Multi-Instance

Scale Summary 1/3

Appliance model	Initial FTD support	Management Solution	Maximum number of instances
Virtual FTD (FTDv)	-	-	-
1010/11xx	-	-	-
1200C/1230/40/50	-	-	-
3105	-	-	-
3110	7.4.1	FMC	3
3120	7.4.1	FMC	5
3130	7.4.1	FMC	7
3140	7.4.1	FMC	10

Reference:
https://www.cisco.com/c/en/us/td/docs/security/secure-firewall/threat-defense/use-case/multi-instance-sec-fw/multi-instance-sec-fw.html

Multi-Instance

Scale Summary 2/3

Appliance model	Initial FTD support	Management Solution	Maximum number of instances
4110	6.3.0	FMC & FXOS	3
4120	6.3.0	FMC & FXOS	3
4140	6.3.0	FMC & FXOS	7
4150	6.3.0	FMC & FXOS	7
4112	6.6.0 / 2.8.1	FMC & FXOS	3
4115	6.4.0 / 2.6.1	FMC & FXOS	7
4125	6.4.0 / 2.6.1	FMC & FXOS	10
4145	6.4.0 / 2.6.1	FMC & FXOS	14

Reference:
https://www.cisco.com/c/en/us/td/docs/security/firepower/fxos/multi-instance/multi-instance_solution.html

Multi-Instance

Scale Summary 3/3

Appliance model	Initial FTD support	Management Solution	Maximum number of instances
4215	7.6.0	FMC	10
4225	7.6.0	FMC	15
4245	7.6.0	FMC	34
9300 SM-24	6.3.0	FMC & FXOS	7
9300 SM-36	6.3.0	FMC & FXOS	11
9300 SM-44	6.3.0	FMC & FXOS	14
9300 SM-40	6.4.0 / 2.6.1	FMC & FXOS	13
9300 SM-48	6.4.0 / 2.6.1	FMC & FXOS	15
9300 SM-56	6.4.0 / 2.6.1	FMC & FXOS	18

Network Interfaces

Multiple modes for Secure Firewall appliances

- Physical, EtherChannel, and VLAN subinterfaces are an option
  - FXOS supports up to 500 total VLAN subinterfaces since FXOS 2.4.1
  - FTD can also create VLAN subinterfaces on physical and EtherChannel interfaces
  - Each instance can have a combination of different interface types

Data (Dedicated)

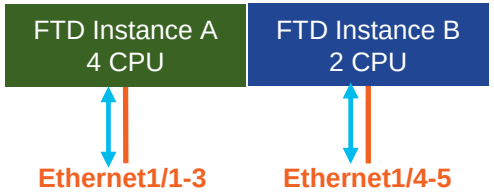

Supported Modes: Routed, Transparent, Inline, Inline-tap, Passive, HA
Supported Traffic: unicast, broadcast, multicast

Data-Sharing (Shared)

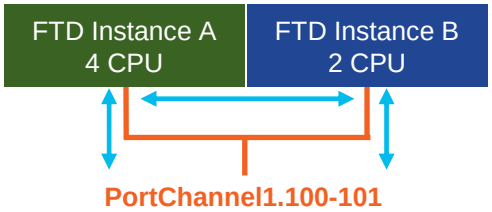

Supported Modes: Routed (no BVI members), HA
Supported Traffic: unicast, broadcast, multicast

Mgmt/Firewall-Eventing

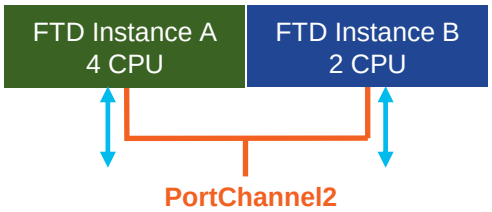

Supported Modes: Management, Eventing
Supported Traffic: unicast, broadcast, multicast

Designing for Internet Edge

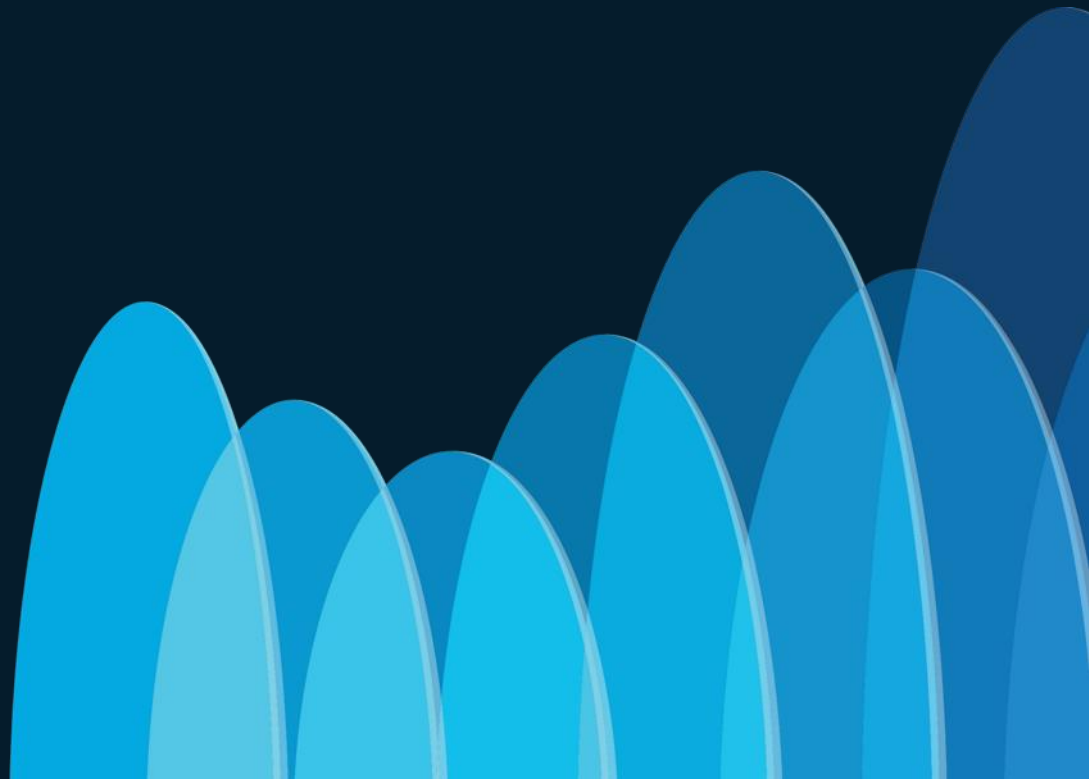

cisco Live!

Routing on Cisco Firewall at the edge

- Multiple use cases
  - Redundant/optimal internet access
  - SDWAN scenarios
  - Internal network routing architecture
- Both ASA and FTD support all major routing protocols:
  - RIP, OSPFv2, OSPFv3, IS-IS, EIGRP and BGP
  - PIM-SM for multicast routing (with IGMPv1/v2)

How we test our FTD appliances?

Appliance model	Maximum # of BGP routes tested	Maximum # of BGP neighbors
1010/1100	5k / 10k	5
1200C	50k	100
1230/1240/1250	50k	100
3100	100k	500 (w/BFD)
4100	200k	500 (w/BFD)
4200	200k	500 (w/BFD)
9300	200k	500 (w/BFD)

How we test our FTD appliances?

Appliance model	Maximum # of BGP routes tested	Maximum # of BGP neighbors
5505	5k	2
5512	20k	20
5525	15k	60
5545	15k	100
5555	15k	100
5508	10k	10
5516	10k	10
ASA 5585 SSP-10	20k	200
ASA 5585 SSP-60	100k	500

Internet access scenario - BGP

Topology and major assumptions

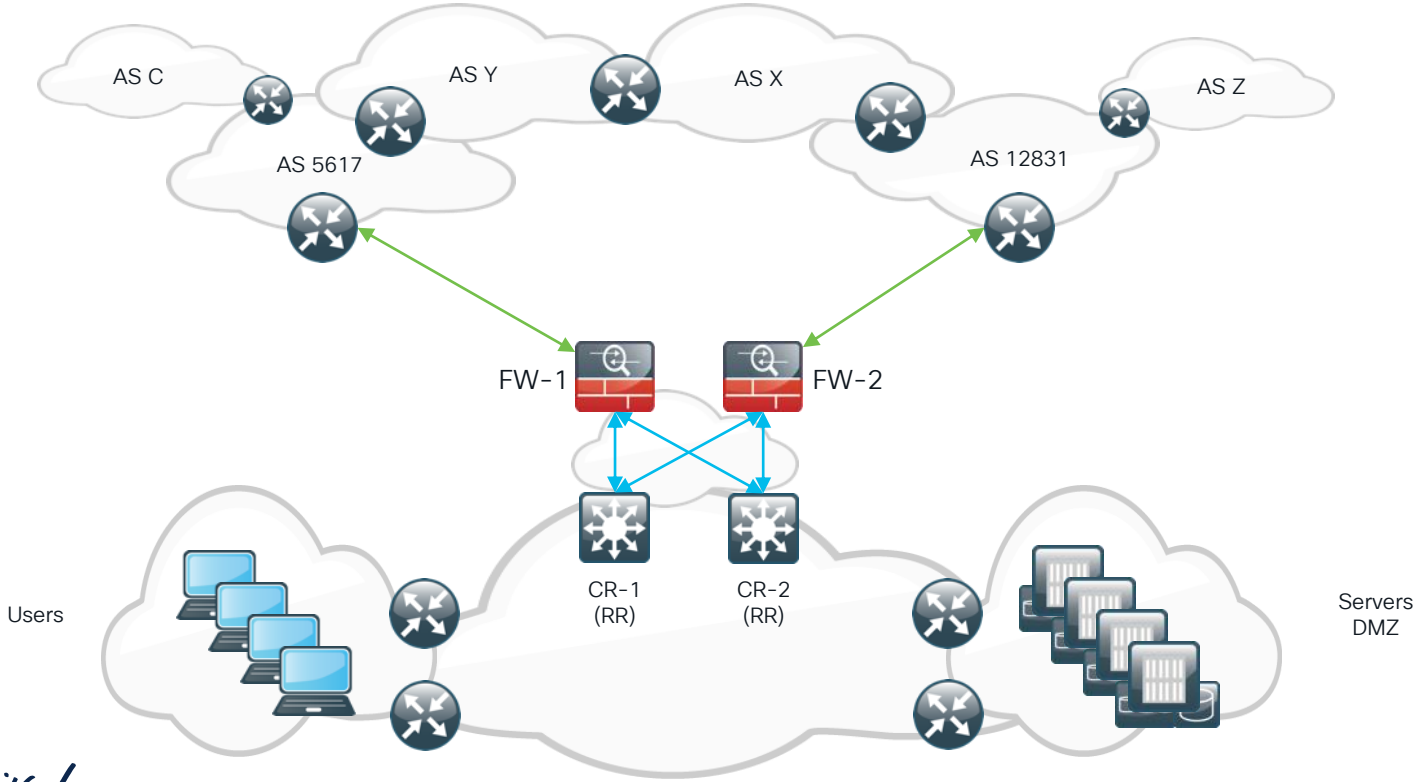

Internet access scenario - eBGP

Option 1: full BGP routes

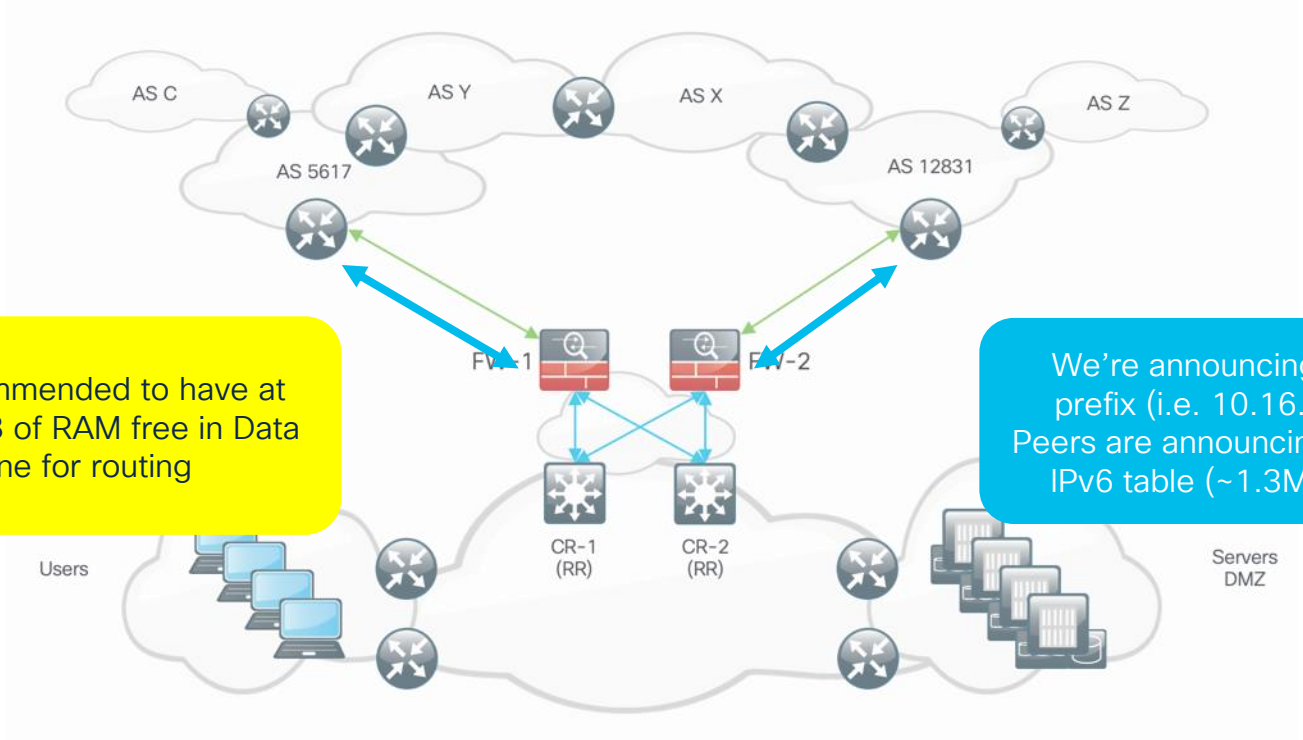

It is recommended to have at least ~1GB of RAM free in Data Plane for routing

We're announcing our own prefix (i.e. 10.16.24.0/24). Peers are announcing full IPv4 & IPv6 table (~1.3M prefixes)

Internet access scenario – eBGP

Option 1: full BGP routes

```
> sh bgp ipv4 unicast summary
BGP router identifier 169.254.10.254, local AS number 65055
BGP table version is 984072, main routing table version 984072
983198 network entries using 196639600 bytes of memory
983198 path entries using 78655840 bytes of memory
155154/155133 BGP path/bestpath attribute entries using 32272032 bytes of memory
173187 BGP AS-PATH entries using 9067894 bytes of memory
15389 BGP community entries using 1229164 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 317884530 total bytes of memory
BGP activity 3584448/2388995 prefixes, 3584909/2389459 paths, scan interval 60 secs

Neighbor        V         AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
85.232.240.179  4      65055 155728  6       984072    0    0 00:03:16  983198

> sh bgp ipv6 unicast summary
BGP router identifier 169.254.10.254, local AS number 65055
BGP table version is 212960, main routing table version 212960
212252 network entries using 50091472 bytes of memory
212252 path entries using 22074208 bytes of memory
54970/54970 BGP path/bestpath attribute entries using 11433760 bytes of memory
173187 BGP AS-PATH entries using 9067894 bytes of memory
15389 BGP community entries using 1229164 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 93896498 total bytes of memory
BGP activity 3584448/2388995 prefixes, 3584909/2389459 paths, scan interval 60 secs

Neighbor        V         AS MsgRcvd MsgSent  TblVer  InQ OutQ Up/Down  State/PfxRcd
2001:1A68:2C:2::179
                4      65055 55611  6       212960    0    0 00:03:20  212204
```

NOTE

~304MB for IPv4
~90MB for IPv6

This is single session. Additional sessions will increment the values by amount needed to store (mostly) additional paths and unique attributes.

"Your mileage will vary" – you'll also need additional 200-300MB at minimum to cover for route churn.

Internet access scenario – eBGP

Option 2: partial BGP routes – limit AS_PATH to 2-3 (neighbor++)

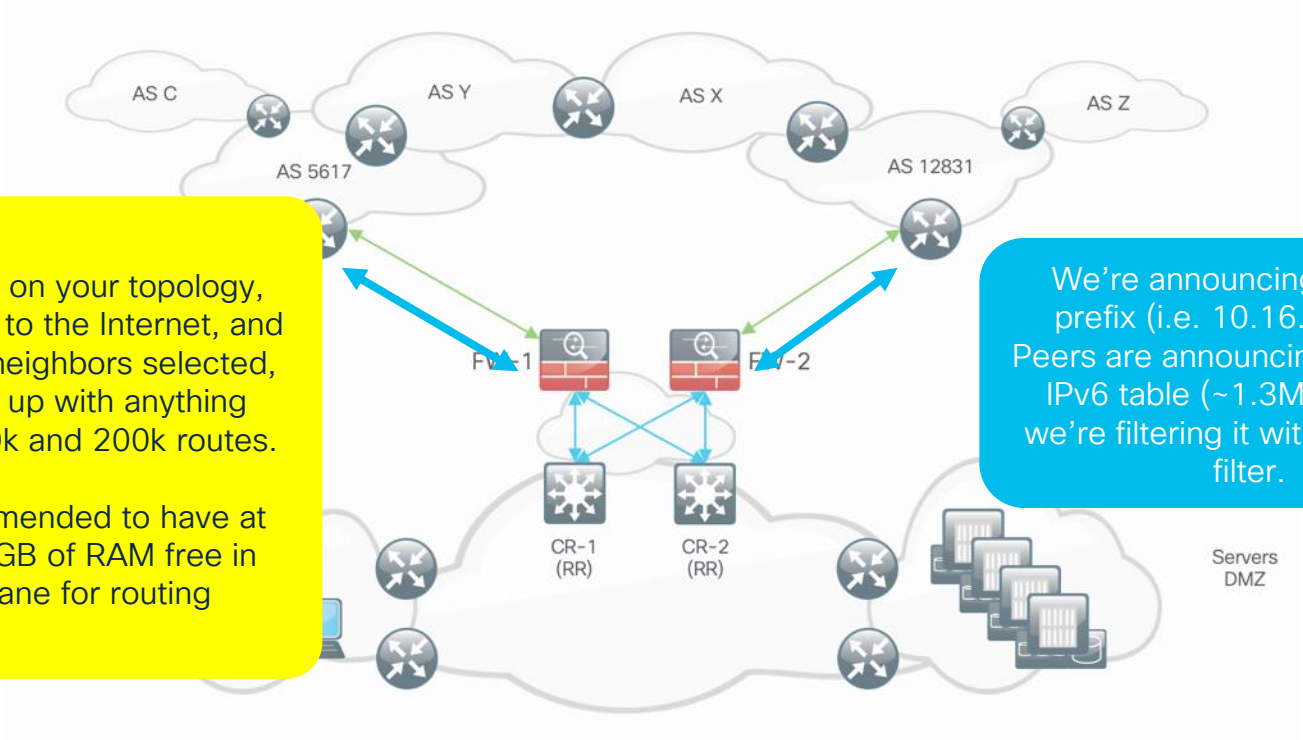

Depending on your topology, connectivity to the Internet, and number of neighbors selected, you'll end up with anything between 30k and 200k routes.

It is recommended to have at least ~0.5GB of RAM free in Data Plane for routing

We're announcing our own prefix (i.e. 10.16.24.0/24). Peers are announcing full IPv4 & IPv6 table (~1.3M prefixes), we're filtering it with AS_PATH filter.

Internet access scenario – eBGP

Option 2: partial BGP routes – limit AS_PATH to 2-3 (neighbor++)

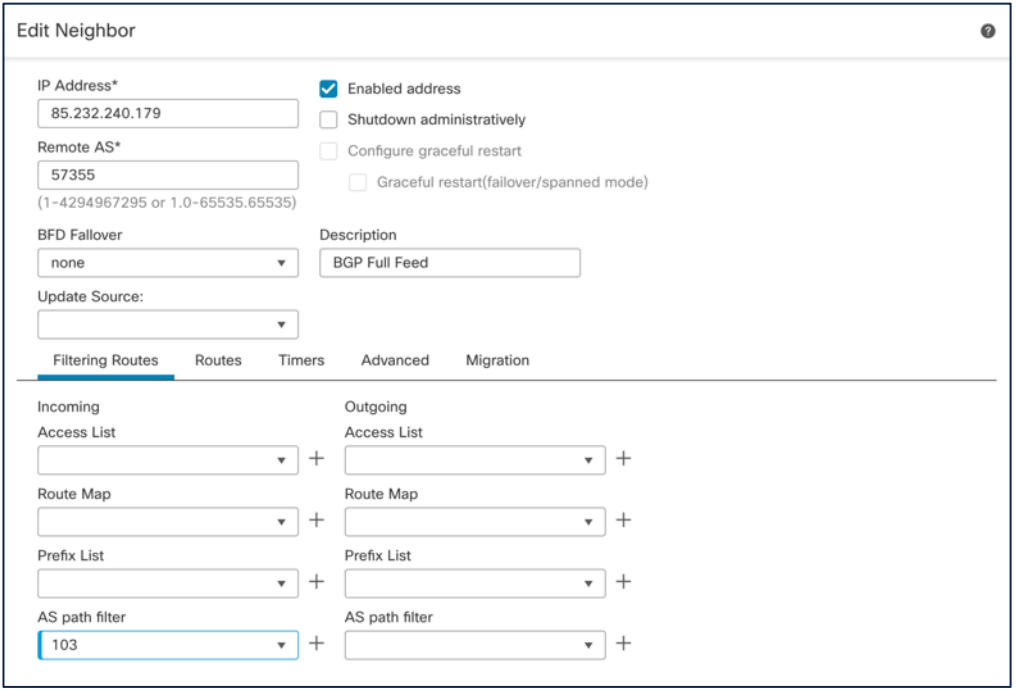

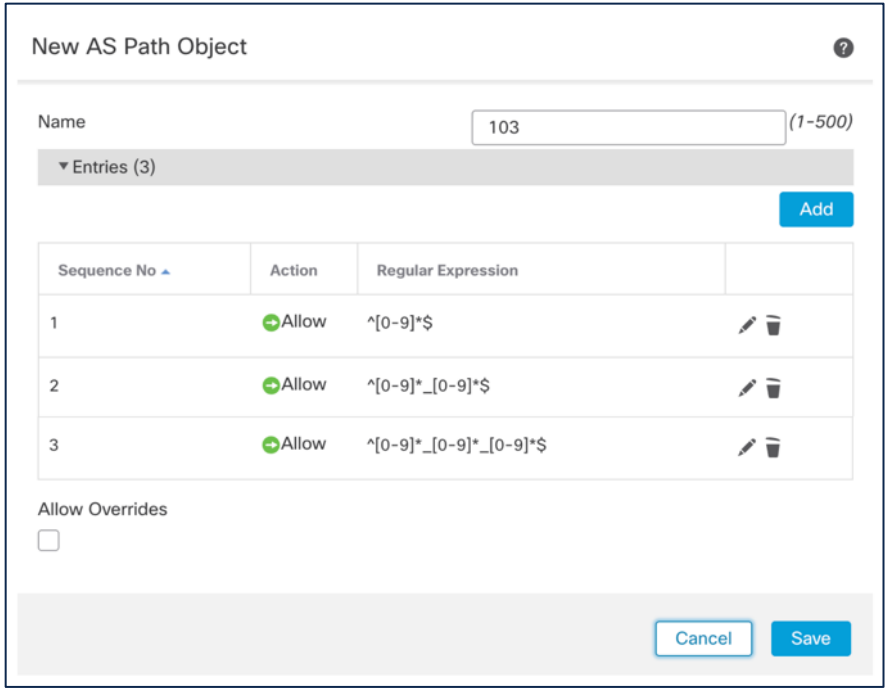

Internet access scenario - eBGP

Option 2: partial BGP routes - limit AS_PATH to 2-3 (neighbor++)

```
> sh bgp ipv4 unicast summary
BGP router identifier 169.254.10.254, local AS number 65055
BGP table version is 984072, main routing table version 984072
176782 network entries using 35356400 bytes of memory
176782 path entries using 14142560 bytes of memory
11834/11740 BGP path/bestpath attribute entries using 2461472 bytes of memory
54002 BGP AS-PATH entries using 3138824 bytes of memory
15389 BGP community entries using 1229164 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
52656 BGP filter-list cache entries using 1684992 bytes of memory
BGP using 56784248 total bytes of memory
BGP activity 96290761/96065182 prefixes, 139438390/139212814 paths, scan interval 60 secs

Neighbor        V         AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
85.232.240.179  4      65055 155449  5       176794    0    0 00:02:08  176782

> sh bgp ipv6 unicast summary
BGP router identifier 169.254.10.254, local AS number 65055
BGP table version is 212960, main routing table version 212960
48794 network entries using 11515384 bytes of memory
48794 path entries using 5074576 bytes of memory
52558/10560 BGP path/bestpath attribute entries using 10932064 bytes of memory
54002 BGP AS-PATH entries using 3138824 bytes of memory
15389 BGP community entries using 1229164 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
52656 BGP filter-list cache entries using 1684992 bytes of memory
BGP using 32345840 total bytes of memory
BGP activity 96290761/96065182 prefixes, 139438390/139212814 paths, scan interval 60 secs

Neighbor        V         AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
2001:1A68:2C:2::179
                4      65055 54441  4       57725    0    0 00:00:17  48794
```

NOTE

~54MB for IPv4
~31MB for IPv6

This is single session. Additional sessions will increment the values by amount needed to store (mostly) additional paths and unique attributes.

"Your mileage will vary" - you'll also need additional 80-120MB at minimum to cover for route churn.

Internet access scenario – eBGP

Option 3: only default routing, BGP used as link keepalive (and for ECMP)

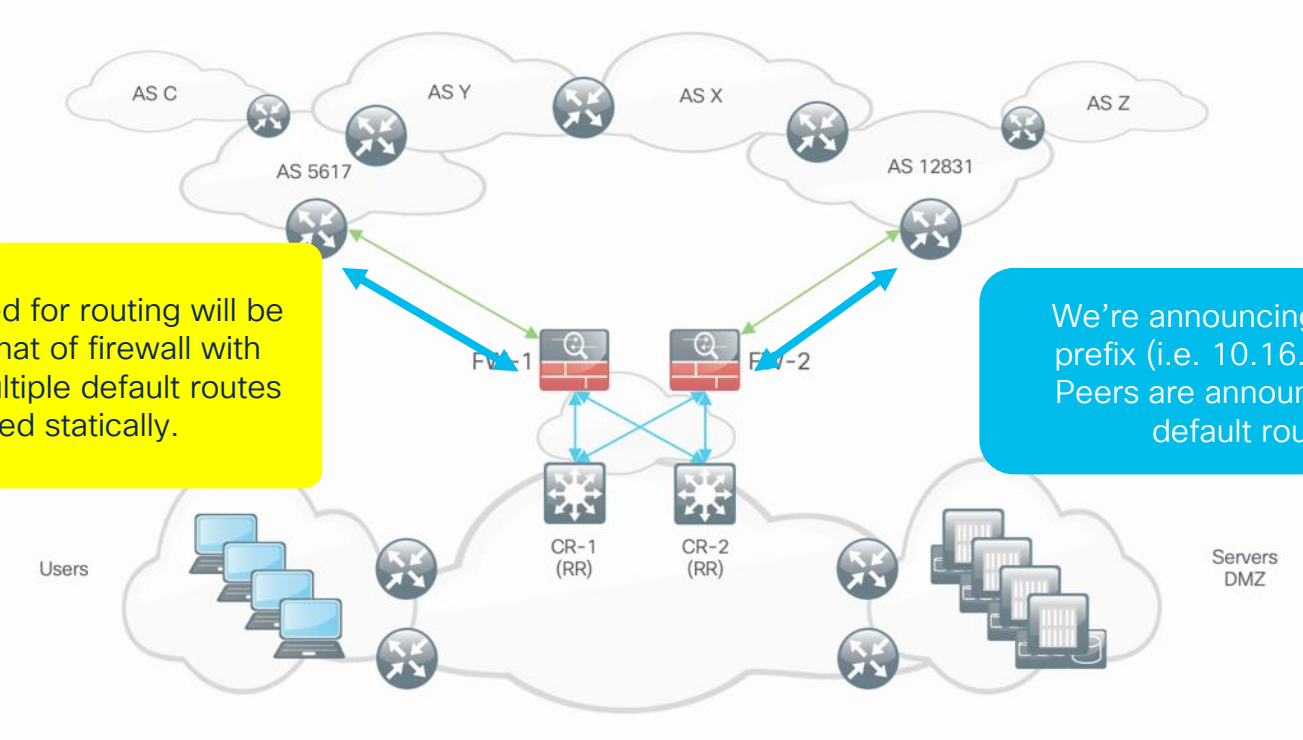

Memory used for routing will be similar to that of firewall with single or multiple default routes defined statically.

We're announcing our own prefix (i.e. 10.16.24.0/24). Peers are announcing only default route.

Internet access scenario – eBGP

Option 3: only default routing, BGP used as link keepalive (and for ECMP)

```
> sh bgp ipv4 unicast summary
BGP router identifier 169.254.10.254, local AS number 65055
BGP table version is 4093684, main routing table version 4093684
1 network entries using 200 bytes of memory
1 path entries using 80 bytes of memory
1/1 BGP path/bestpath attribute entries using 208 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 488 total bytes of memory
BGP activity 4853424/4853422 prefixes, 4861587/4861585 paths, scan interval 60 secs

Neighbor        V         AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
169.254.10.1    4      65055 69     57       4093684    0    0 00:58:40  1

> sh bgp ipv6 unicast summary
BGP router identifier 169.254.10.254, local AS number 65055
BGP table version is 1078776, main routing table version 1078776
1 network entries using 236 bytes of memory
1 path entries using 104 bytes of memory
1/1 BGP path/bestpath attribute entries using 208 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 548 total bytes of memory
BGP activity 4853424/4853422 prefixes, 4861587/4861585 paths, scan interval 60 secs

Neighbor        V         AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
2001:db8:100::1 4      65055 69     57       1078776    0    0 00:58:35  1
```

NOTE

~0.5kB for IPv4
~0.5kB for IPv6

This is single session. Additional sessions will increment the values by amount needed to store (mostly) additional paths and unique attributes.

"Your mileage **will** vary" – but that's least stressing option to choose if it fits your requirements.

Internet access scenario – eBGP

Option 3: only default routing, BGP used as link keepalive (and for ECMP)

```
> sh resource usage
Resource              Current    Peak   Limit    Denied Context
Telnet                      1       1       5         0 System
Conns                       3       6  400000         0 System
Hosts                       6       8     N/A         0 System
Inspects [rate]             0      30     N/A         0 System
Routes                     15 1195471 unlimited         0 System

> sh route bgp
[...]
Gateway of last resort is 169.254.10.1 to network 0.0.0.0

B*      0.0.0.0 0.0.0.0 [200/0] via 169.254.10.1, 00:59:17

> sh ipv6 route bgp
[...]
IPv6 Routing Table - 5 entries
B   ::/0 [200/0]
    via 2001:db8:100::1,
```

Summary

Webex App

Questions?

Use the Webex app to chat with the speaker after the session

How

1. Find this session in the Cisco Events mobile app
2. Click “Join the Discussion”
3. Install the Webex app or go directly to the Webex space
4. Enter messages/questions in the Webex space

Webex spaces will be moderated by the speaker until February 28, 2025.

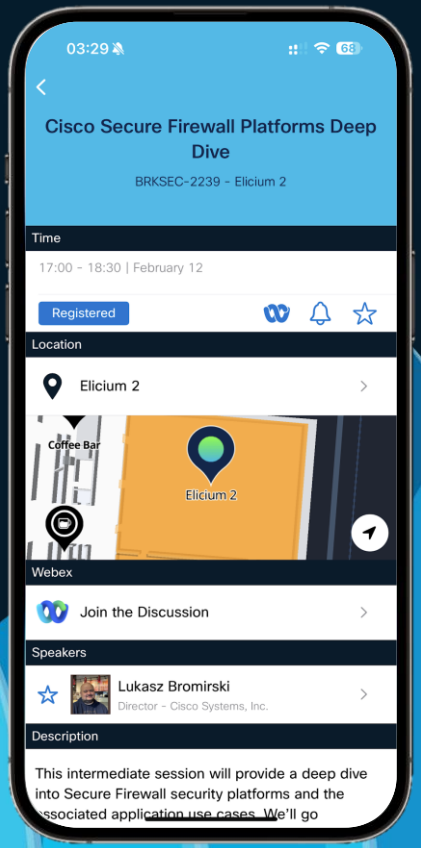

Fill Out Your Session Surveys

Participants who fill out a minimum of 4 session surveys and the overall event survey will get a unique Cisco Live t-shirt.

(from 11:30 on Thursday, while supplies last)

All surveys can be taken in the Cisco Events mobile app or by logging in to the Session Catalog and clicking the ‘Participant Dashboard’

Content Catalog

Continue your education

- Visit the Cisco Showcase for related demos
- Book your one-on-one Meet the Engineer meeting
- Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs
- Visit the On-Demand Library for more sessions at ciscolive.com/on-demand. Sessions from this event will be available from March 3.

Contact me at: lbromirs@cisco.com

Security

Network Security

Learn about a broad range of solution and technologies which will help you better understand how to secure your network. You will find topics such as FTD, VPN, SASE, Meraki Security Policies and Network Analytics.

Arrow of time in this Universe goes one way (at least, it seems so...)

START

Monday, February 10 | 2:00 p.m.
BRKSEC-2708
Cisco SDWAN Use Cases & Best Practices

Tuesday, February 11 | 8:00 a.m.
BRKSEC-2057
Secure Connectivity Anywhere - The Evolution of Cisco Remote Access Technologies

Tuesday, February 11 | 12:00 p.m.
BRKSEC-2236
Keeping Up on Network Security with Cisco Secure Firewall

Wednesday, February 12 | 1:00 p.m.
BRKSEC-3274
TAC and Engineering on Cisco Secure Firewall Threat Detection Performance - Performance Profiling tools, Tuning and Best Practices

Wednesday, February 12 | 5:00 p.m.
BRKSEC-2239
Cisco Secure Firewall Platforms Deep Dive

Thursday, February 13 | 8:30 a.m.
BRKSEC-3320
Pig-in-the-Middle - TLS Decryption and Encrypted Visibility Engine Deep Dive on Cisco Secure Firewall

Thursday, February 13 | 10:45 a.m.
BRKSEC-3935
Think Like a TAC Engineer: Troubleshooting Secure Client Remote Access Issues

Thursday, February 13 | 1:00 p.m.
BRKSEC-2821
Securing Industrial Networks: Strategies and Best Practices

Friday, February 14 | 9:15 a.m.
BRKSEC-3533
Think Like a TAC Engineer: A Guide to Cisco Secure Firewall most Common Pain Points

Friday, February 14 | 11:15 a.m.
FINISH
BRKSEC-2086
Optimizing Security and Agility: Leveraging SD-WAN Capabilities in Cisco Secure Firewall

The bridge to possible

Cisco Security Beta Programs

Attention to Feedback
Beta customer bugs and enhancements receive high visibility & priority

Top notch communication
Private conference calls with product team

Training
Customers receive early training and experience with new features

Customer Support
Our feature experts will be on-hand and responsive to your issues

Sign-Up Now: https://cs.co/security-beta

Thank you

cisco Live!

cisco Live!

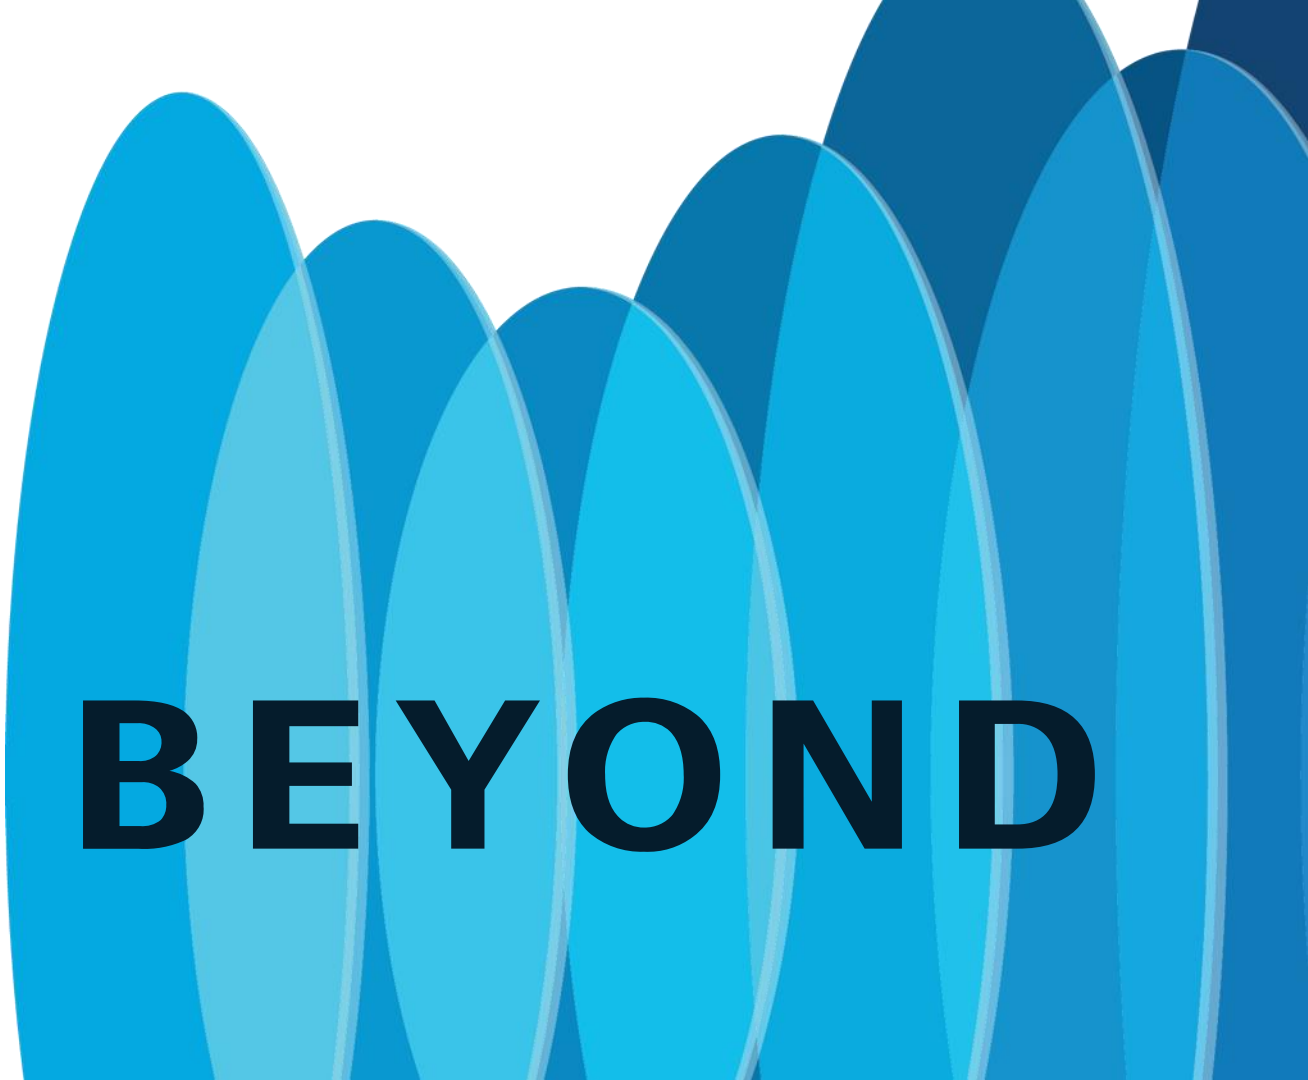

GO BEYOND