

#### All You Need to Know About Forwarding on the Catalyst 8000 platforms BRKENT-2653

David Roten Technical Marketing Engineer / Technical Leader





#### Cisco Webex App

#### **Questions?**

Use Cisco Webex App to chat with the speaker after the session

#### How

- **1** Find this session in the Cisco Live 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 until February 24, 2023.





Agenda

- What are the forwarding mechanisms?
- Where is each mechanism used & why?
- Specific considerations for each?
- How to observe resource utilization for each mechanism?
- Conclusion

## Hardware architectures

cisco ile

#### IOS XE routing hardware architectures

All IOS XE routing platforms use a multi-core and multi-threaded data plane architecture.

BRKENT-2653

5

© 2023 Cisco and/or its affiliates. All rights reserved. Cisco Public

224 cores 896 threads	<b>3.0</b>	Catalyst 8500 platforms	ESP100-X ESP200-X
64 cores 256 threads	QFP 2.0	ASR1002-HX ASR1001-HX	ESP100 ESP200
4+ cores & threads	x86	Catalyst 8500L-8S4X Catalyst 8300 platforms Catalyst 8200 platforms Catalyst 8200L platforms	Catalyst 8000v ISR1100-4G / 6G ISR4300
4+ cores & threads	ARM	ISR1000 platforms	ISR4400

#### Basic data plane function implementation

- Receiving traffic
- Distributing traffic
  - Load based distribution
  - Non-strict flow based distribution
  - Strict flow based distribution
- Crypto processing
- Forwarding action
- Queuing and scheduling

#### Baseline platforms comparison

Autonomous mode	ARM LBD ISR1161X	x86 Non-Strict FBD C8300-2N2S-4X2T	x86 Strict FBD C8500L-8S4X	QFP 3.0 C8500-12X4QC
Traffic acceptance (Rx)	Shared Rx thread	2 shared cores	2 dedicated cores	Dual priority with <b>hardware assist</b>
Traffic distribution (Rx)	Shared Rx thread 1 stage	2 shared cores 1 stage	2 dedicated cores 2 stage	Hardware distributor
Crypto processing	shared thread	2 shared cores	1 encrypt core 1 decrypt core	16 engines dedicated hardware
Forwarding (PPE)	2 threads	10 hyperthreads	12 hyperthreads	896 threads
Queuing (TM)	1 shared thread	2 shared cores	2 dedicated cores	Hardware

#### Baseline platform comparison



This data should not be used for official performance guidance.

cisco live!

#### SDWAN Baseline platform comparison



This data should not be used for official performance guidance.



## Concepts

۲

cisco live!

## Lexicon (1/2)

• **Rx** (receive):

Function that receives packet from an ingress interface and either [ first pass hashes to PP or PP-Rx function | or distributes packet to PP based on load ]

• **PP-Rx** (packet processor – receive):

Function that does second pass hashing of traffic to PP function. Specific to FBD.

• **PP** (packet processor):

Function that does network processing of protocol traffic (NAT, MPLS, Firewall, ACLs, classification, etc)

• **COFF** (crypto):

Function that encrypts and decrypts traffic

#### • **Tx** (transmit):

Function that schedules and transmits a packet on an interface

### Lexicon (2/2)

- Stateful flow features
  - NAT, ZBFW
  - NBAR, Flexible Netflow, UTD
- Cores and threads
  - Non-hyperthreaded x86 systems: a single thread per CPU core
  - Hyperthreaded x86 systems: two threads associated with a given CPU core. Not all cores in a hyperthreaded system must use hyperthreading. Two threads per core **does not** double compute performance. Typically, two hyper threads on a single core will have 20-30% more compute power as a single non-hyperthreaded core.



#### What are forwarding mechanisms?

- IOS XE routing architecture is capable of using a multi-core and multi-threaded data plane –
- Multiple metrics involved:
  - Hardware architecture in a given platform
  - Targeted performance for a given product
  - Targeted price point for a given product
  - Economy of scale in build components
  - Optimizations in software development



#### Routing platform consistencies

- IOS XE platforms use the same software for dataplane and control plane features.
- Features can be delivered across platforms using the same source software providing consistency and compatibility across platforms.
- $\boldsymbol{\cdot}$  Distribution of traffic amongst the dataplane threads varies.
- Various platforms offer differing amount of compute power.
- The forwarding architecture:
  - does not affect feature availability
  - **does** affect the forwarding capacity and potential performance advantages as well as pinch points



cisco live!

• Packet Ordering and Ordered Mutual Exclusion (ordered access to data)

If packets from a given flow are handled by different threads, the cores holding on to the 2 through n packets might have to wait as the system deals with all the complications of the first packet to hand stateful processing of the flow.

- Stateful features have to allocate and initialize structures when encountering a new flow. This makes the processing time for the first packet longer than for subsequent packets.
- Crypto acceleration hardware differs across platforms. Some crypto hardware requires walking through the packet feeding its data into the hardware and getting a result back. On these platforms, larger packets take longer to process and if a flow has packets varying in size, this waiting situation can occur.



#### Memory contention

If many threads are handling a flow, they may all need access to the same memory to update statistics, do flow rate calculations, get and update stateful information about the flow, etc.

This overlaps somewhat with the issues of ordered mutual exclusion.

#### • L1 and L2 x86 / ARM caching

Caching is not optimized if different threads need to pass information through the L1 and L2 caches instead of keeping it local

#### Packet ordering issues



BRKENT-2653





#### L1/L2 cache efficiency

- QFP based platforms have an optimized design that makes L1/L2 caching less of a concern as various threads process the same flows
- General purpose CPUs have L1/L2 caches that are dedicated to each thread/core. The caches are **not** shared amongst the threads/cores. L3 caches may be shared depending on the implementation



## Forwarding Mechanisms

cisco ive

#### Forwarding architectures

#### Load based (LBD)

Packets are be handled by any available data plane thread.

Generally used on platforms with more threads than interfaces.

#### Strict Flow based (S-FBD)

Packets are **strictly** distributed based on flow hashing.

Available data plane threads **will not** assist busy data plane threads.

#### **Non-strict Flow based (NS-FBD)**

Packets can be handled by any available data plane thread.

Efforts are made to keep packets from a given flow on a given thread but no guarantees.



#### Forwarding architectures – IOS XE 17.9

The three forwarding architectures available today are and there uses



### Load based architecture for x86 and QFP



## $\overbrace{\rightarrow}\overset{}{\searrow}$

#### Load Based Distribution

- Packets are distributed to threads strictly based on availability by a hardware distributor (QFP) or Rx thread (x86/ARM based systems)
- Efforts (*but not guarantees*) are made at ingress to keep traffic for a given ingress port on the same thread.
  - Challenging if there are a son software platforms mall number of active ports/queues
  - Challenging with a larger number of threads available vs ingress ports
- Load based distribution allows parallelism of processing for a single flow
  - Strict FBD does not. Non-strict FBD runs between the two.

#### Load Based Distribution - characteristics

- A packet from any given flow can land across multiple threads
- State for a flow must be readily available to any core at any time
- **QFP platforms provide hardware assist** for this functionality making it easier and simple to implement without significant performance hit
- x86 and ARM platforms do not have hardware assist The following are valid in x86 and ARM systems but may not cause significant determent.
  - Packet ordering and ordered access to data challenge
  - Memory contention for tracking flow state challenge
  - X86/ARM L1 and L2 cache implications challenge

#### QFP based load distribution



- QFP has Dispatcher Hardware which finds an available thread to process packets
- Dedicated resource so no other contention
- Dispatcher has excess capacity so no "Rx" bottleneck concerns in QFP platforms

cisco ile

#### x86 based load distribution



- x86 and ARM platforms use a generic thread for the Rx function. This thread is responsible to queue packets for distribution to PPEs.
- Since this is not a dedicated resource in x86 and ARM platforms, there is the potential that the Rx thread can be a bottleneck.



# Non-strict flow based architecture





## Non-Strict Flow Based Distribution - characteristics

- Packets from a flow may, but not always, land on different threads
- State for a stateful flow must be readily available to any core at any time *just in case*
- x86 platforms are **do not have hardware assist** The following are valid in x86 systems but may not cause significant detrement.
  - Packet ordering and ordered mutual exclusion in software challenge
  - Memory contention for memory tracking flow state challenge
  - x86 lower L1 and L2 cache implications challenge

#### Non-Strict Flow Based Distribution

- Packets are classified on a tuple from the **outermost** encapsulation by Rx thread
- Fragments use 3-tuple classification
- If that initially targeted thread is busy (elephant flow scenario / uneven flow hashing), the packet can be handled by a PPE thread that has free cycles
- In heavy load situations with few flows, the behavior is very similar to a load based distribution scenario
- In the case we are the destination of an IPSEC tunnel, hashing is done based on the interior address after decryption

#### Non-Strict Flow Based Distribution



- x86 platforms use a generic thread for the Rx function. This thread is responsible to find an available data plane thread to process the packet.
- On non strict flow based platforms, this requires the additional step of hashing the outermost header of the packet (higher cost than the load based scenario)

Since this is not a dedicated resource in x86 platforms, there is the potential that the Rx threads can be a bottleneck (CEF / racetrack).

BRKENT-2653

# Strict flow based architecture



cisco life!



### Strict Flow Based Distribution - characteristics

- Packets from a flow always land on the same PP thread
- State for a stateful flow must be readily available to only a single thread

- x86 platforms are optimized
  - Packet ordering and ordered access handled by hardware optimized
  - Less memory contention for memory tracking flow state optimized
  - More efficient x86 L1 and L2 cache usage optimized
- Elephant flows are a concern as flows are strictly mapped to a thread with no offload for high-utilization

#### Strict Flow Based Distribution

- Packets are classified on a tuple on the outermost encapsulation by Rx thread in the first pass by the Rx function
- Fragments use a 3-tuple classification
- Packets are assigned to a PPE-Rx thread based on the initial hashing by the Rx thread
- PPE-Rx function will check the flow database and based on that potentially redirect the packet to an alternate core
- There is no offloading at PPE-Rx if a thread is 100% utilized because there is no sharing of stateful flow information between threads.
- In the case we are the destination of an IPSEC or GRE tunnel, hashing is done based on the interior address after decryption
#### Strict Flow Based Distribution



- x86 platforms use a generic thread for the Rx function. This thread is responsible to find an available data plane thread to process the packet.
- On strict flow based platforms, this requires the additional step of hashing the header of the packet (higher cost than the load based scenario)
- Rx thread inspect the outermost header and PPE-Rx thread may also inspect inner header.
- Since Rx compute is not a dedicated resource in x86 platforms, there is the potential that the Rx thread can be a bottleneck.



#### The tuple classifications

	VRF ID	Src IP v4/v6	Dest IP v4/v6	Protocol	Src port	Dest port	SPI	GRE key	Dst MAC	Src MAC
Fragment	$\bigcirc$	$\bigotimes$	$\bigcirc$	$\bigcirc$						
UDP / TCP	$\bigcirc$	$\bigcirc$	$\bigotimes$	$\bigotimes$	$\bigotimes$	$\bigcirc$				
ESP	$\bigcirc$	$\bigcirc$	$\bigotimes$	$\bigotimes$			$\bigcirc$			
GRE	$\bigcirc$	$\bigcirc$	$\bigcirc$	$\bigcirc$				$\bigcirc$		
L2									$\bigcirc$	$\bigcirc$
Default	$\bigotimes$	$\bigotimes$	$\bigcirc$	$\bigotimes$						

MPLS will generally be inspected for recognized payload encapsulation and MPLS tags.

cisco

#### Placement guidance for S-FBD platforms

- Strict flow-based distribution platforms are best suited for network locations where there is a diversity of flows.
  - High flow count generic routing
  - High flow count internet gateway
  - High flow count crypto action
  - High flow count SDWAN edge router
- The theme here is **High flow count**!

Per internal flow performance will always be limited by single thread forwarding capacity.

Encapsulated traffic as a transit router will be seen as a single flow between tunnel endpoints.

High  $\rightarrow$  1000s, 10000s, 10000s Maximum supported flow count is 2 million. A bidirectional flow will consume two entries.

## Core distribution

.

cisco ive!

#### Default in autonomous mo

DONE

#### Core dist – C8300–1N1S–4T2X – DP heavy



cisco live!

#### Default in SD-WAN mode

DONE

#### Core dist – C8300–1N1S–4T2X – SP heavy



Default in autonomous mode



Default in SD-WAN mode



cisco liver



cisco / ille

#### C8500L x86 thread allocation (XE 17.9)

C8500L-8S4X**#show platform software cpu alloc** CPU alloc information:



Slow control plane cpu alloc: Template used: CLI-data\_plane\_heavy C8500L-8S4X**#show platform software cpu alloc** CPU alloc information:

Control plane cpu alloc: 0-1,12-13 Data plane cpu alloc: 6-11,18-21 Service plane cpu alloc: 2-5,14-17

Slow control plane cpu alloc: Template used: **CLI-service\_plane\_heavy** 



# Packet flow and analysis

cisco ive

#### Ingress packet flow for Strict FB



#### Testbed config for following data

- C8500L-8S4X running IOS XE 17.8.01a
- Data-plane heavy mode
- Using TenGig 0/1/2 and TenGig0/1/3
- Single IP address configured on main-interfaces
- Injected traffic is 500 byte packets, 20% line rate on each port
- Randomized UDP source and destination ports
  - 0x00FF mask for UDP ports, 256x256x2 = 131072 flows total

#### Step 1a – Ingress processing

C8500L-8S4X#show plat hardware qfp active datapath infrastructure bindings Port Instance Bindings:

ID	Port	IOS Port	WRKR12	WRKR13
1	rcl0	rclC	Rx	Tx
2	ipc	ipc	: Tx	Rx
3	vxe_punti	vxe_puntif	Tx	Rx
4	fpe0	GigabitEthernet0/0/0	Tx	Rx
5	fpe1	GigabitEthernet0/0/1	. Tx	Rx
6	fpe2	GigabitEthernet0/0/2	Rx	Tx
7	fpe3	GigabitEthernet0/0/3	Tx	Rx
8	fpe4	GigabitEthernet0/0/4	Rx	Tx
9	fpe5	GigabitEthernet0/0/5	Tx	Rx
10	fpe6	GigabitEthernet0/0/6	Rx	Tx
11	fpe7	GigabitEthernet0/0/7	Tx	Rx
12	fpe8	TenGigabitEthernet0/1/0	Rx	Tx
13	fpe9	TenGigabitEthernet0/1/1	. Tx	Rx
14	fpe10	TenGigabitEthernet0/1/2	Rx	Tx
15	fpe11	TenGigabitEthernet0/1/3	Tx	Rx

Tx / Rx mapping has changed across releases. Be sure to check your platform and release.

Rx = ingress processing Tx = TM = queuing and scheduling

#### Step 1b – Ingress processing

#### C8500L-8S4X#show platform hardware qfp active datapath infrastructure sw-cio Credits Usage:

ID	Port	: Wght	Global	wrkr0	WRKR1		WRKR10	WRKR11	WRKR12	WRKR13	WRKR14	WRKR15	Total
1	rcl(	) 1:	6048	0	0		0	0	96	0	0	0	6144
1	rcl(	128:	6048	0	0		0	0	96	0	0	0	6144
2	ipo	: 1:	0	0	0		0	0	0	0	0	0	0
3	vxe punti	. 1:	466	0	0		0	0	0	46	0	0	512
4	fpe	) 1:	1952	0	0		0	0	0	96	0	0	2048
4	fpe	2:	1952	0	0		0	0	0	96	0	0	2048
 1 /	<b>Em a 1</b> (	1.	1050	0	0		0	0	0.6	0	0	0	2040
14	Ipel(		1952	0	0	•••	0	0	96	0	0	0	2048
14	There	Ζ.	1952	0	0		0	0	90	0	0	0	2040
 Core	Utilizati	on over	preced	ing <b>315</b>	4.3457	sec	conds						
	ID:	0	1	10	11		12	13	14	15			
% PPI	E-RX: 4	1.29	4.54	4.46	4.44		0.00	0.00	0.00	0.00	Hashing	/ distribut	on
	% PP: 15	5.26 1	6.45	16.39	16.31		0.00	0.00	0.00	0.00	Feature	processing	q
	% RX: (	0.00	0.00	0.00	0.00		5.19	6.03	0.00	0.00	Rx proc	essing	
										0 00	Troffic N	Assessor (T	
	% TM: (	0.00	0.00	0.00	0.00		31.17	35.17	0.00	0.00		nanager (T	x)
olo (	% TM: ( COFF: (	).00).00	0.00 0.00	0.00	0.00		0.00	35.17	0.00	0.00	Crypto (	offload	x)

#### Step 1c – Ingress processing

From the perspective the process running on the specific thread  $\rightarrow$  detailed visibility.

C8500L-8S4X#show platform hardware qfp active datapath infrastructure sw-cio Credits Usage:

ID	Port	Wght	Global	wrkr0	WRKR1	WRKR10	WRKR	11	WRKR12	WRKR13	WRKR14	WRKR15	Total
1	rcl0	1:	6048	0	0	C		0	96	0	0	0	6144
1	rcl0	100.	6040	<u> </u>	<u> </u>	C		0	96	0	0	0	6144
2	ipc	If the	credit v:	alues fo	r fnell-	fneX annr	nach	0	0	0	0	0	0
3	vxe_punti		n thorout					0	0	46	0	0	512
4	fpe0	0 the	n there a	are issu	es with i	ngress pa	скег	0	0	96	0	0	2048
4	fpe0	rate t	being too	) high fo	or the pla	atform.		0	0	96	0	0	2048
14	fpe10	Limiti	ng the m	nax pac	kets fron	n a given		0	0	0	96	0	2048
14	fpe10	sourc	source in the fashion protects the system.							0	96	0	2048
… Core	Utilizatio	lt doe sourc	e not all e to ove	ow one rwhelm	elephan the syst	t flow fron tem and d	n one eny						
	ID:	servio	ce to all	other flo	DWS.			3	14	15			
% PPE	-RX: 4.	29 4	4.34	4.40	4.44	0.00	0.0	0	0.00	0.00	Rx proc	esssing	
9	PP: 15.	26 1	6.45	16.39	16.31	0.00	0.0	0	0.00	0.00	Feature	processing	)
9	RX: 0.	00	00.0	0.00	0.00	5.19	6.0	3	0.00	0.00	Rx proc	essing	
9	TM: 0.	00	00.0	0.00	0.00	31.17	35.1	7	0.00	0.00	Traffic N	Nanager (T	×)
% C	OFF: 0.	00	00.0	0.00	0.00	0.00	0.0	0	0.00	0.00	Crypto	offload	
% I	DLE: 80.	45 79	9.01	79.15	79.26	63.63	58.5	1	99.74	99.76	Idle		

BRKENT-2653

#### Step 1d – Ingress processing

#### C8500L-8S4X#show platform hardware qfp active datapath infrastructure sw-cio



#### Step 1e - Tricky information

#### C8500L-8S4X#show platform resources r0 cpu

CPU utilization for five seconds: 71%, one minute: 71%, five minutes: 71% Core 0: CPU utilization for five seconds: 1%, one minute: 2%, five minutes: 28 Core 1: CPU utilization for five seconds: 1%, one minute: 2%, five minutes: 2% Core 2: CPU utilization for five seconds: 100%, one minute: 100%, five minutes: 100% Core 3: CPU utilization for five seconds: 100%, one minute: 100%, five minutes: 100% Core 4: CPU utilization for five seconds: 100%, one minute: 100%, five minutes: 100% Core 5: CPU utilization for inutes: 100% This output is from the perspective Core 6: CPU utilization for inutes: 100% of underlying Linux OS. More Core 7: CPU utilization for inutes: 100% Core 8: CPU utilization for detailed visibility is available from the Core 9: CPU utilization for inutes: 100% show platform commands. Core 10: CPU utilization fo nutes: 11% Core 11: CPU utilization fo nutes: 11% Core 12: CPU utilization fo 28 nutes: Core 13: CPU utilization fo Linux can not see the difference in nutes: 28 Core 14: CPU utilization fo minutes: 100% polling with active or idle results. Core 15: CPU utilization fo minutes: 100% Core 16: CPU utilization for five seconds: 100%, one minute: 100%, five minutes: 100% Core 17: CPU utilization for five seconds: 100%, one minute: 100%, five minutes: 100% Core 18: CPU utilization for five seconds: 100%, one minute: 100%, five minutes: 100% Core 19: CPU utilization for five seconds: 100%, one minute: 100%, five minutes: 100%

#### Don't always trust the penguin!

#### Steps 3 and 6 - Crypto

From the perspective the process running on the specific thread  $\rightarrow$  accurate.

C8500L-8S4X#show platform hardware qfp active datapath infrastructure sw-cio Credits Usage:

ID	Port	Wght	Global	wrkr0	WRKR1	V	VRKR10	WRKR11	WRKR12	WRKR13	WRKR14	WRKR15	Total
1	rcl0	1:	6048	0	0		0	0	96	0	0	0	6144
1	rcl0	128:						0	96	0	0	0	6144
2	ipc	1:	No	crypto	in this	SE	etun so		0	0	0	0	0
3	vxe_punti	1:				, 00		0	0	46	0	0	512
4	fpe0	1:	the	cryptc	) threa	d s	tay idi	e. <sub>0</sub>	0	96	0	0	2048
4	fpe0	2:		÷				0	0	96	0	0	2048
14	fpe10	1:	1952	0	0		0	0	0	0	96	0	2048
14	fpe10	2:	1952	0	0		0	0	Q	0	96	0	2048
Core	Utilizatic	on over	preced	ing 60.	34 seco	nds							
	ID:	0	1	10	11		12	13	14	15			
% PPE	-RX: 4.	29	4.54	4.46	4.44		0.00	0.00	0.00	0.00	Rx proc	esssing	
90	PP: 15.	26 1	6.45	16.39	16.31		0.00	0.00	0.00	0.00	Feature	processing	g
90	RX: 0.	00	0.00	0.00	0.00		5.19	6.03	0.00	0.00	Rx proc	essing	
90	TM: 0.	00	0.00	0.00	0.00	3	31.17	35.17	0.00	0.00	Traffic N	lanager (T	x)
% C	OFF: 0.	00	0.00	0.00	0.00		0.00	0.00	0.00	0.00	Crypto o	offload	
% I	DLE: 80.	45 7	9.01	79.15	79.26	6	53.63	58.51	99.74	99.76	Idle		

BRKENT-2653

#### Step 4,5 - Data plane processing

From the perspective the process running on the specific thread  $\rightarrow$  high detail

C8500L-8S4X#show platform hardware qfp active datapath infrastructure sw-cio Credits Usage:



BRKENT-2653

#### Step 5 - Data plane processing

C8500L-8S4X show platform hardware qfp active fbd-flowdb balance distribution





Ideal distribution of traffic for a C8500L system. Maximum system performance in this scenario.

Even distribution amongst all data plane threads.

All of these flows sum up to 131072 which was the 256x256x2 based on the injected traffic.

This command is not available on NS-FBD or FBD systems.

#### Step 7a – Egress processing

C8500L-8S4X#show plat hardware qfp active datapath infrastructure bindings Port Instance Bindings:

ID	Port	IOS Port	WRKR12	WRKR13
1	rcl0	rcl0	Rx	Tx
2	ipc	ipc	Tx	Rx
3	vxe_punti	vxe_puntif	Tx	Rx
4	fpe0	GigabitEthernet0/0/0	Тх	Rx
5	fpe1	GigabitEthernet0/0/1	Tx	Rx
6	fpe2	GigabitEthernet0/0/2	Rx	Тх
7	fpe3	GigabitEthernet0/0/3	Tx	Rx
8	fpe4	GigabitEthernet0/0/4	Rx	Тх
9	fpe5	GigabitEthernet0/0/5	Tx	Rx
10	fpe6	GigabitEthernet0/0/6	Rx	Тх
11	fpe7	GigabitEthernet0/0/7	Tx	Rx
12	fpe8	TenGigabitEthernet0/1/0	Rx	Тх
13	fpe9	TenGigabitEthernet0/1/1	Тх	Rx
14	fpe10	TenGigabitEthernet0/1/2	Rx	Тх
15	fpe11	TenGigabitEthernet0/1/3	Tx	Rx

cisco il

#### Step 7b – egress processing

From the perspective the process running on the specific thread  $\rightarrow$  more visibility.

#### ${\tt C8500L-8S4X} {\#} {\tt show} {\tt platform} {\tt hardware} {\tt qfp} {\tt active} {\tt datapath} {\tt infrastructure} {\tt sw-cio}$

Credits	: Usage:											
0100100	, couge.	Low values	for the	e %ID	LE for th	ne Rx/Tx						
ID	Port	cores does	not ab	osolut	ely indica	nte no moi	re KR11	WRKR12	WRKR13	WRKR14	WRKR15	Total
1	rcl0	processing	power	r for ir	ngress pa	ackets.	0	96	0	0	0	6144
1	rcl0	%IDLE can	be zer	o. wh	ile there	is still	0	96	0	0	0	6144
2	ipc	additional c	apacity	v avai	lable. If t	here is	0	0	0	0	0	0
3 vx	ke_punti	some idle th	apaon	oro is	clearly a	dditional	0	0	46	0	0	512
4	fpe0	some fuic ti				uantional	0	0	96	0	0	2048
4	fpe0	packet per	second	u cap	acity.		0	0	96	0	0	2048
14	fpe10	Packets car	Packets can be bundled for release which						0	96	0	2048
14	fpe10	can increas	e effic	ciency	but still s	show the	0	0	0	96	0	2048
		same level	of utiliz	zation	l.							
Core Ut	cilizatio	on over bred	еатну	J JIJ4	1.3437 50	econas		1				
				1 0		10	10	1.4	1 -			
0 222 2	.D:			IU			13	14	15			
% PPE-F	RX: 4.	.29 4.54	4	.46	4.44	0.00	0.00	0.00	0.00	RX proc	esssing .	
% F	PP: 15.	.26 16.45	16	5.39	16.31	0.00	0.00	0.00	0.00	Feature	processing	g
% F	XX: 0.	.00 0.00	0	0.00	0.00	5.19	6.03	0.00	0.00	Rx proc	essing	
% I	:M: 0.	.00 0.00	0	0.00	0.00	31.17	35.17	0.00	0.00	Traffic N	<i>l</i> anager	
% COF	FF: 0.	.00 0.00	0	0.00	0.00	0.00	0.00	0.00	0.00	Crypto o	offload	
% IDI	LE: 80.	45 79.01	79	.15	79.26	63.63	58.51	99.74	99.76	ldle		

BRKENT-2653

#### Summary on troubleshooting x86 performance

• Use "**sw-cio"** output:

To get I/O (Rx/Tx) core utilization guidance To get data plane thread utilization To get crypto thread utilization

• Use "show plat resources r0 cpu":

To get control plane core utilization guidance To get service plane core utilization guidance

• Use "show proc cpu":

To get IOSd thread utilization

#### Network management access to x86 performance

- Use "**sw-cio"** output:
- To get I/O (Rx/Tx) core utilization guidance To get data plane thread utilization ceqfpUtilProcessingLoad 1.3.
  - To get crypto thread utilization ciscoEntityPerformanceMIB

```
Use "show plat resources r0 cpu":
```

- To get control plane core utilization guidance To get service plane core utilization guidance
- Use "show proc cpu":

To get IOSd thread utilization

lance No SNMP/YANG method 1.3.6.1.4.1.9.9.715.1.1.6.1.14 1.3.6.1.4.1.9.9.756

# L2 and L3 forwarding

cisco live!

### Wide range of L3 forwarding in the platform

- All L3 forwarding is via the data plane cores
- There is no L3 forwarding via the control plane threads
- Only traffic destined for the router itself will be punted to the control plane threads

### L3 forwarding options

- IPv4 via CEF
- IPv6 via CEF
- MPLS via CEF
- L3VPN

Full hashing in Strict Flow-based platforms.

Maximum performance with flow diversity.

- Tunnels
  - GRE
  - IPSEC

If router is the tunnel source or destination, full hashing in strict flowbased platforms.

If router is a transit router for the tunnel, elephant flows could be a concorn with small number of tunnels (outermost header bashing)

- concern with small number of tunnels (outermost header hashing).

Maximum performance with multiple tunnels and flow diversity in the tunnel.



## L2 forwarding options

- With L3 underlay:
  - OTV
  - L2VPN
  - L2TP
  - EVPN -
  - VxLAN
- Native L2 forwarding
  - L2 bridging
  - L2 port connect
  - L2 switching on C8300 platforms
    with switch modules

EVPN L3 gateway

Cross-leaf VLAN over L2 VNI

## L2 forwarding options

• With L3 underlay:

• 12VPN

• 12TP

FVPN

VxLAN

Ingress L2 side traffic uses L2 for hashing. Incoming remote OTV traffic hashed against remote OTV overlay IP address.

Ingress L2 side traffic uses L2 for hashing. Incoming remote traffic hashed against remote IP encapsulation source address.

Ingress L3 addressing for hashing of VxLAN encapsulated traffic.

- Native L2 forwarding
  - L2 bridging —
  - L2 port connect

L2 addresses always used for hashing to data plane cores.

L2 switching on C8300 platforms
 with switch modules

Not part of the router dataplane.

## Conclusion

۲

cisco live!

#### Key takeaways

- Feature and functional compatibility across the portfolio
- Understand the forwarding mechanism for the platform before you dive into troubleshooting
  - When you suspect a bottleneck, use the sw-cio output to begin your troubleshooting process (except QFP platforms)
- C8500L requires a diversity of flows for maximum performance
  - Inner packet inspection is supported for GRE, IPSEC, and SDWAN tunnels
- IOS XE x86 platforms have flexibility for core allocation to suit your network needs

#### Complete your Session Survey

- Please complete your session survey after each session. Your feedback is important.
- Complete a minimum of 4 session surveys and the Overall Conference survey (open from Thursday) to receive your Cisco Live t-shirt.



https://www.ciscolive.com/emea/learn/sessions/session-catalog.html



### Continue Your Education

abab.

Visit the Cisco Showcase for related demos.



Book your one-on-one Meet the Engineer meeting.



Attend any of the related sessions at the DevNet, Capture the Flag, and Walk-in Labs zones.



Visit the On-Demand Library for more sessions at <u>ciscolive.com/on-demand</u>.



CISCO The bridge to possible

# Thank you

cisco life!





