



You make **possible**



Cisco Catalyst 3850 and 3650 Switching Architecture

Dimitar Hristov – Technical Marketing Engineer
BRKARC-3438

Cisco *live!*
June 9-13, 2019 • San Diego, CA

#CLUS



Dimitar Hristov
dhristov@cisco.com



I am a Technical Marketing Engineer (TME), and I have been with Cisco for 12 years.

As TME within Enterprise team, I work primary on increasing and innovating new Capabilities for Catalyst 3850 / 3650 and 9000 family.

I have a strong background in providing and troubleshooting customer solutions.
@ CCIE#20505 R&S and Data Center

Is Your Network Ready for Digitization of Tomorrow?

IP Display/DMS



Printer



IP Camera



LED Lights / IoT



AP



PC/Laptop



IP Phone



Does the platform support new PoE devices efficiently?


Does the platform make it easy to provision and scale?

Does the platform support enough Programmability?

Does the platform ensure secure network access?

Does the platform let you adapt to new connectivity requirements?

Enterprise Trends Driving Digital Transformation



Mobility

Fabric Enabled Wireless
Multigigabit



IoT

CoAP / IoT Device profiling
SD Bonjour
Perpetual PoE
AVB



Security

256bit MACsec
Trustworthy Systems
Group based policy
Full Netflow



Programmability
Cloud

DevOps Toolkit
NETCONF/RESTCONF
Yang Models
Streaming telemetry

“The goal of this session is to give you an in depth view of Catalyst 3850/3650 so you can understand its strength as well as its limitations ...”

Agenda

- Introduction to Catalyst 3850/3650 platform
- Platform and UADP ASIC architecture
- Packet walks
- High Availability using data and power stacks
- Resource scale
- Software innovations
- Conclusion

Cisco Webex Teams


Questions?

Use Cisco Webex Teams 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 Webex Teams or go directly to the team space
- 4 Enter messages/questions in the team space

Webex Teams will be moderated by the speaker until June 16, 2019.




cs.co/cicolivebot# BRKARC-3438

Introduction to Catalyst 3850 and 3650



You make networking **possible**

Cisco MultiGigabit – Now IEEE 802.3bz



Cisco MultiGigabit enables various use cases



Digital Imaging

Cisco *live!*



Server in a Branch



High Definition Cameras

#CLUS



Uplink Extension



802.11ac wave 2 APs

The Catalyst 3K Family

Catalyst 3850
Stackwise-480,
Stackpower
Data/PoE/PoE+/UPoE
FRU Uplinks

Catalyst 3650
Stackwise-160,
-
Data/PoE/PoE+/
Fixed Uplinks

Catalyst 3850 SFP
Stackwise-480,
Stackpower
12 and 24 Port Versions
FRU Uplinks

Catalyst 3850 mGig
Stackwise-480,
Stackpower
24 and 48 Port Versions
Stacks with any Catalyst 3850

Catalyst 3850 SFP+
Stackwise-480,
Stackpower
12, 24 and 48* Port Versions
Enabling 10G Aggregation

Catalyst 3650 Mini
Stackwise-160,
Data/PoE/PoE+/
Fixed Uplinks
Stacks with any 3650

Catalyst 3650 mGig
Stackwise-160,
Data/PoE/PoE+/UPOE
Fixed Uplinks
Stacks with any 3650



Jan
2013

Oct
2013

May
2014

Jan
2015



Jun/Aug
2015

Jan
2016

Oct
2016

Built on Cisco's Innovative "UADP" ASIC

One Switch – Multiple Deployment scenarios



Catalyst 3850 Copper

Copper SKUs Data and PoE/UPoE Switches

480G Stacking Capacity

Catalyst 3650 Copper

Copper SKUs Data and PoE/UPoE Switches

160G Stacking Capacity

Catalyst 3850 Fiber SFP

Fiber SKUs SFP Versions

Catalyst 3850 Fiber SFP+

Fiber SKUs SFP+ Versions

Enterprise Class Access Layer

Smaller Core & Aggregation Option

Based on a Common ASIC and Software

Catalyst 3850/3650 Components



You make security **possible**

Uplink Network Module Options on Catalyst 3850



C3850-NM-4-1G

4x1Gig


SFP



C3850-NM-2-10G

2x1Gig+2x10Gig

SFP/SFP+




C3850-NM-4-10G

4x10Gig

SFP/SFP+


48 Ports or 12+ SFP+



C3850-NM-2-40G

2x40Gig

QSFP



C3850-NM-8-10G

8x10Gig

SFP/SFP+

For MultiGigabit and SFP+ Versions only

Flexibility & Investment Protection

Fixed Uplink Options on Catalyst 3650



4x1Gig

SFP



2x1Gig+2x10Gig

SFP/SFP+



4x10Gig

SFP/SFP+



2x40Gig

QSFP



8x10Gig

SFP/SFP+

For MultiGigabit Versions only

Make Uplink Decision at the time of Purchase

48 Port SFP+ Version – 750 WAC PS

4x40G Fixed Uplinks



Line Rate – 640G Switching Capacity

Stack-Cables and Components

Catalyst 3850



3 lengths of cable, 0.5 1 and 3 Meters

Catalyst 3650



1 ring in 3650 vs 3 rings in 3850

Catalyst 3850/3650 models

Similarities & Differences



For Your Reference


Features / Scale	Catalyst 3850	Catalyst 3850 SFP	Catalyst 3850 MultiGigabit	Catalyst 3850 SFP+ (12,24 Ports)	Catalyst 3850 SFP+ (48Port)	Catalyst 3650	Catalyst 3650 Mini	Catalyst 3650 MultiGigabit
ASIC	UADP 1.0	UADP 1.0	UADP 1.1	UADP 1.1	UADP 1.1	UADP 1.0	UADP 1.1	UADP 1.1
Stacking BW	480G / 9	480G / 9	480G / 9	480G / 9	NA	160G / 9	160G / 9	160G / 9
Stack Power	Stackpower, XPS	Stackpower, XPS	Stackpower, XPS	Stackpower, XPS	No	No	No	No
Uplinks	Modular Uplinks	Modular Uplinks	Modular Uplinks	Modular Uplinks	Fixed Uplinks	Fixed uplinks	Fixed Uplinks	Fixed Uplinks
Wireless	100 AP max	100 AP max	100 AP max	100 AP max	100 AP max	50AP max	50AP max	100 AP max
Stacking Module	Built-in	Built-in	Built-in	Built-in	N/A	Optional	Optional	Optional
Memory/Flash	4GB /2GB	4GB /2GB	4GB /4GB	4GB /4GB	8GB /8GB	4GB /2GB	4GB /2GB	4GB / 4GB
Power	Dual (FRU)	Dual (FRU)	Dual (FRU)	Dual (FRU)	Dual (FRU)	Dual (FRU)	Single (Fixed), RPS 2300	Dual (FRU)
MACSEC	128 bit	128 bit	256 bit	256 bit	256bit	128 bit	256 bit	256 bit
HA	SSO	SSO	SSO	SSO, Stackwise Virtual	SSO, Stackwise Virtual	SSO	SSO	SSO


Looking Inside the Switch



You make the power of data **possible**


Catalyst 3850 / 3650: Under the Covers...



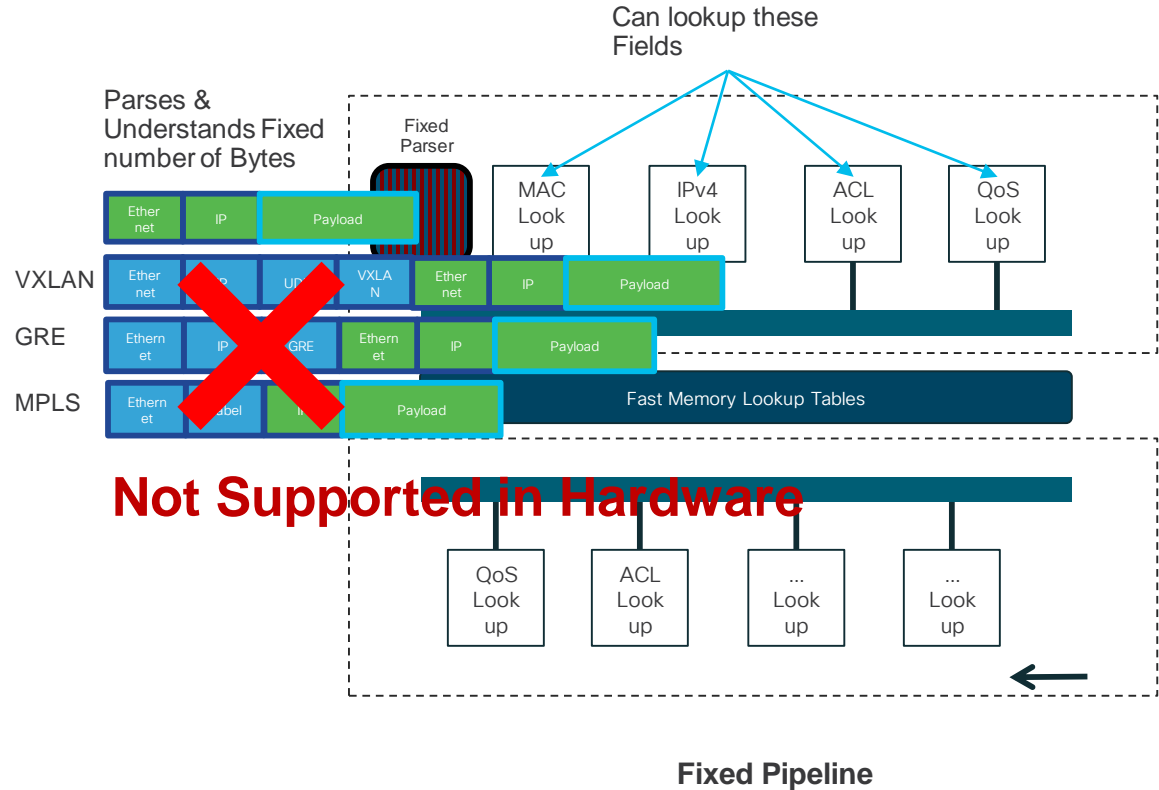


ASICs are a Pillar of Cisco Innovation...

Traditional ASIC Pipeline



Traditional ASIC



Cisco Custom ASIC Innovation



In 2013 Cisco Introduced a new family of switches called Catalyst 3850



Based on UADP ASIC
(Unified Access Data Plane)

Future proofed for the technologies of tomorrow

Some of the Key Capabilities of UADP ASIC



Flex Parser (256 Bytes)
&
Programmable Pipeline
(15 Ingress and 7
Egress)



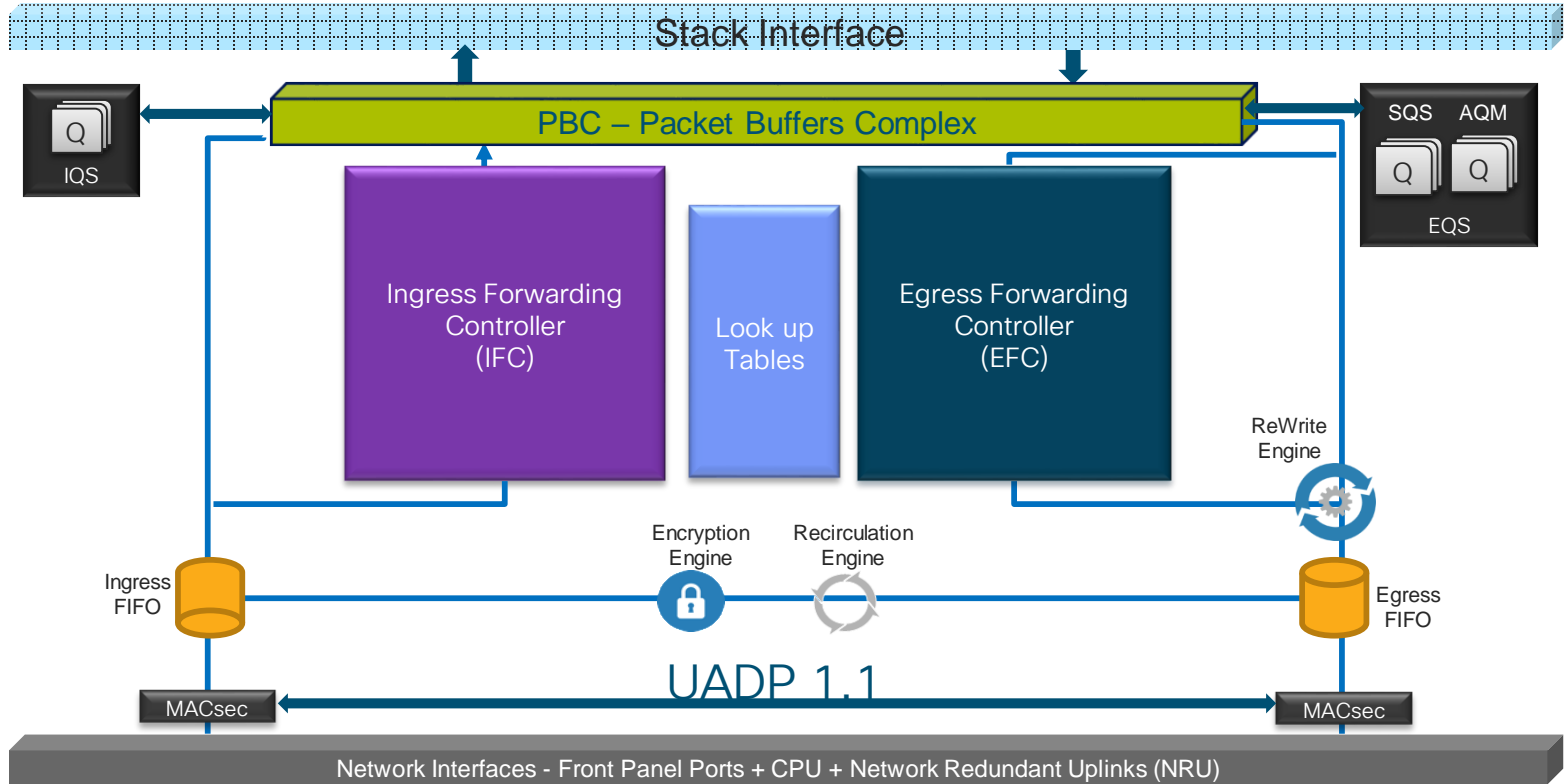
Recirculation
Capability
(upto 16 times)



Micro Engines
(Fragmentation
Encryption
e.g, MACsec 256
bit)

No Compromise on Performance

UADP 1.0 and 1.1 ASIC Block Diagram





LAN MASEC Support per Platform

	MACSEC	Cat 3650 *		Cat 3850*	
		SW	License	SW	License
Switch to Switch	128 Bits SAP	3.7.0E	IPB +	3.7.0E	IPB +
	128 Bits MKA	3.7.1E	IPB +	3.7.1E	IPB +
	256 Bits MKA	16.3.3 +	IPB +	16.3.3 +	IPB +
	Over L2 Hops	Roadmap		Roadmap	
Host to Switch	128 Bits MKA	3.7.1E	IPB +	3.7.1E	IPB +
	256 Bits MKA	16.3.3 +	IPB +	16.3.3 +	IPB +

* Only 3650 Minis and mGig support MACSEC 256
 * Only 3850 mGig and 3850 48 SFP+ support MACSEC 256

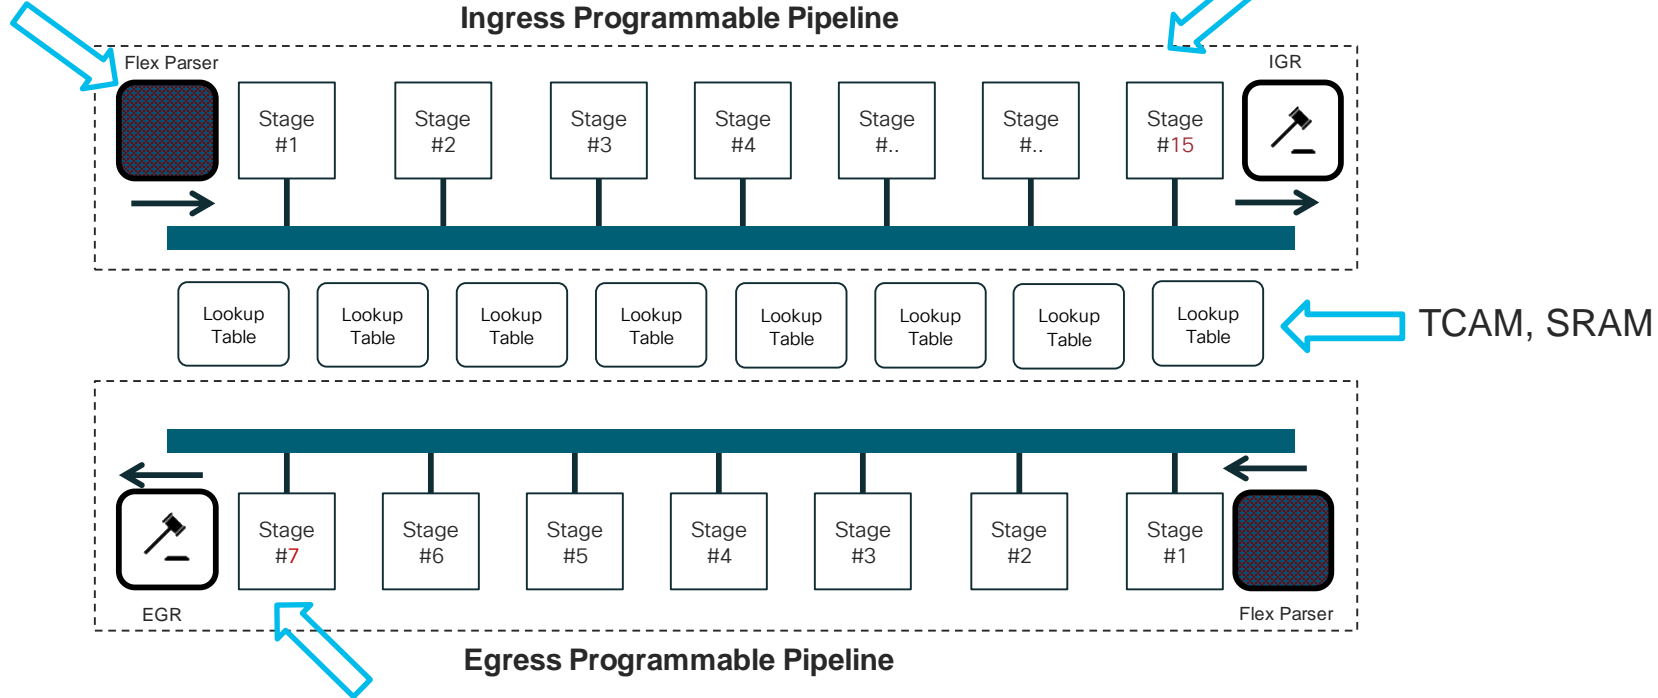
IPB = IP Base

	MACSEC	Cat 3650 *	Cat 3850 *
Switch to Switch	128 Bits SAP	Line Rate	Line Rate
	128 Bits MKA	Line Rate	Line Rate
	256 Bits MKA	Line Rate	Line Rate
Host to Switch	128 Bits MKA	Line Rate	Line Rate
	256 Bits MKA	HW ready	HW ready

Lets take a look at the Programmable Pipelines

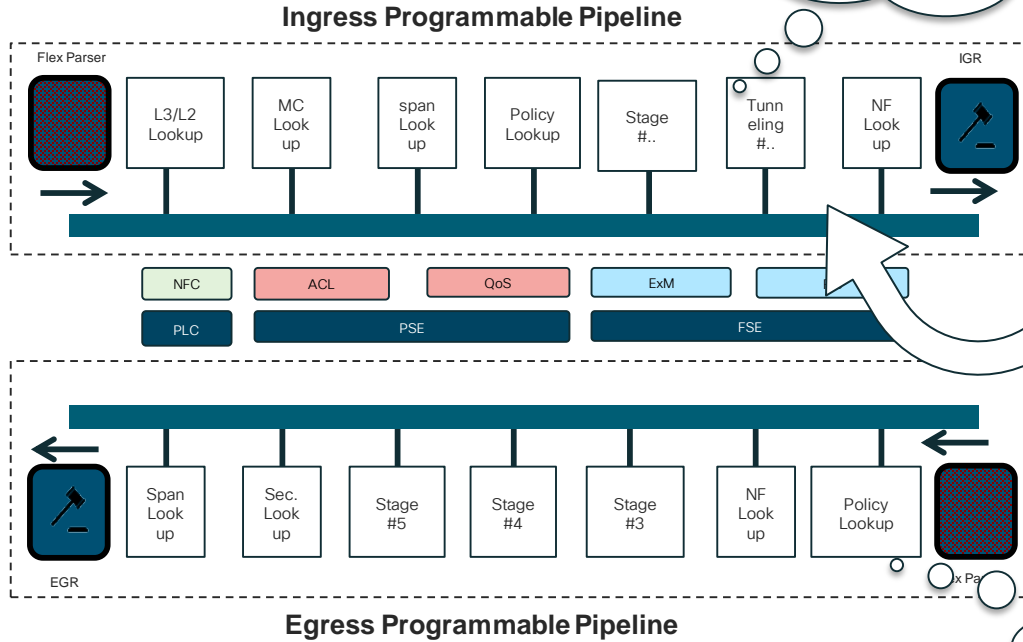
256 B Header

15 Ingress Stages



Micro Code

Programmed to understand
VXLAN



OSPF MPLS ACLs
Software Features

```
if result == east  
    value = east  
else  
    value =  
    }  
return value;
```

ASIC

Programmed to understand
MPLS

UADP Variants



You make multi-cloud **possible**


UADP 1.0




1G/10G
Ethernet



24K
Netflow Records



240 Gbps
Stacking Capacity



6MB
Packet Buffer



56 Gbps
Bandwidth



First Generation of UADP
ASIC



Catalyst
3850 Copper



Catalyst
3650



Catalyst SFP
Fiber

First Programmable ASIC

UADP 1.1



Dual Core
Running @ 500MHz



1G/10G/40G
Ethernet




256 Bit
MACsec
Encryption



24K x2
Netflow Records



1588
IEEE



240 Gbps
Stacking Capacity



6MB x2
Packet Buffer




160 Gbps
Bandwidth

Stackwise-Virtual AVB GRE
VXLAN AVB MPLS GRE
DNS-AS GRE
PTP NBAR Wireless
40G 1588 NBAR SGT 40G
AVC AVB GRE AVB AVB
FnF AVB GRE AVB AVB
SGT AVB DNS-AS 1588 SGT
AVC ERSPAN VXLAN GRE FnF GRE
FnF 40G FnF GRE
1588 ERSPAN NBAR NBAR
ERSPAN ERSPAN NBAR PTP DNS-AS
SGT PTP DNS-AS
Stackwise-Virtual SGT
Stackwise-Virtual FnF GRE AVB
MPLS ERSPAN
ERSPAN SGT PTP PTP



Catalyst 3850
Multigigabit



Catalyst 3850
SFP+



Catalyst 3650
Mini



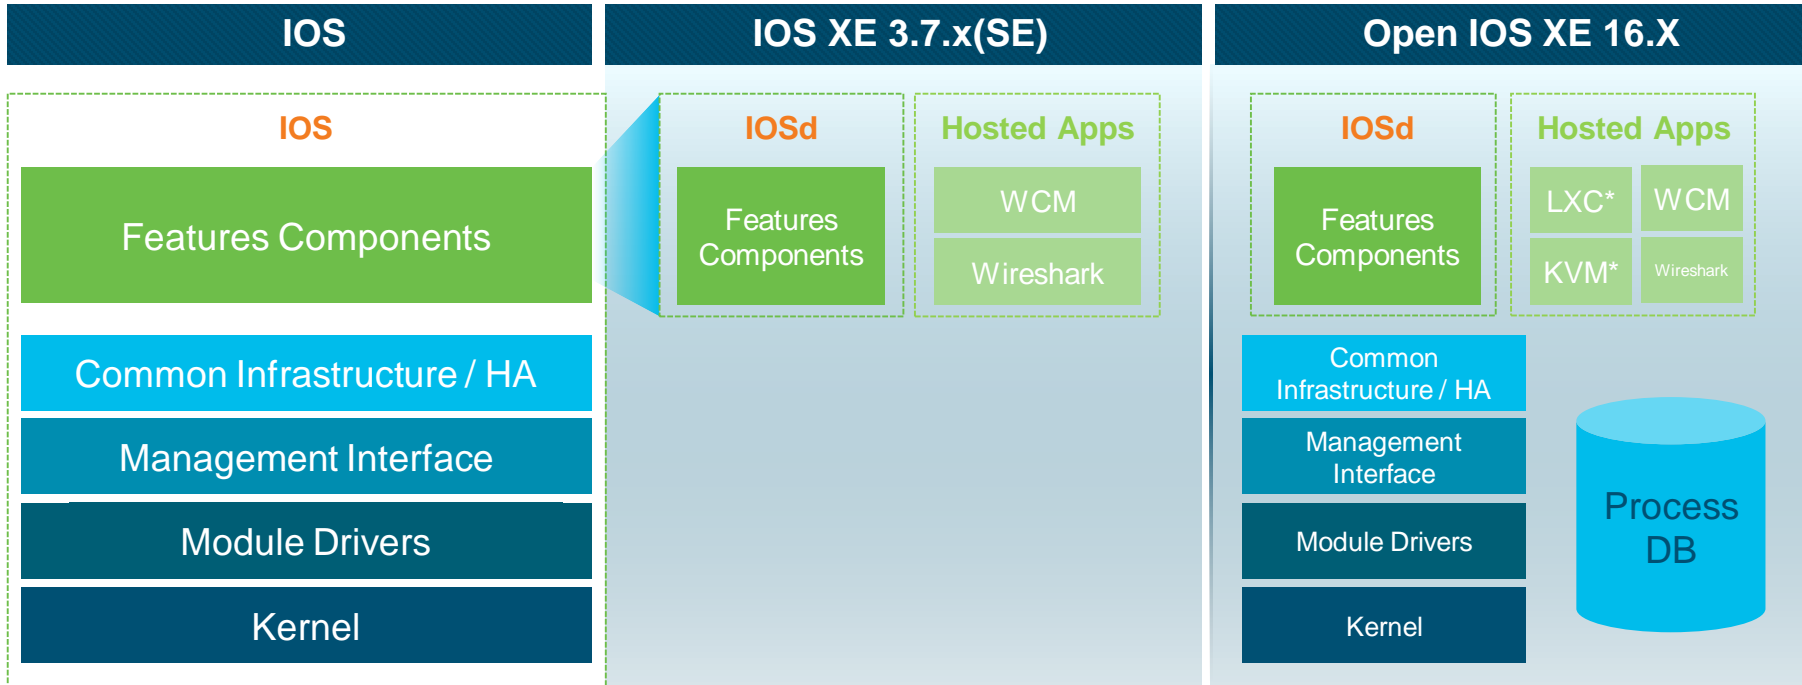
Catalyst 3650
Multigigabit

Enhanced Version of UADP
1.0 ASIC


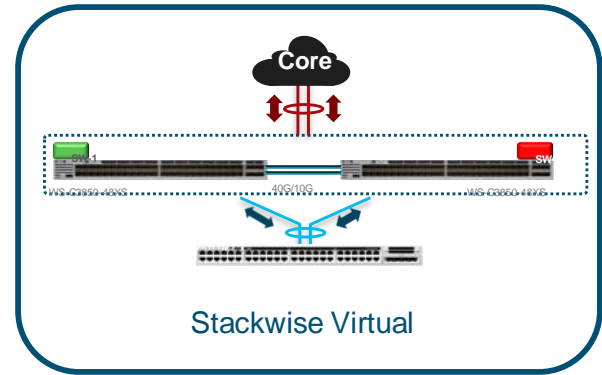

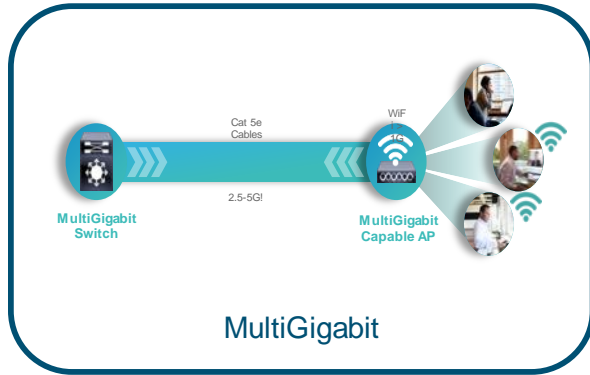
Enhanced Power & Security Capability

IOS XE Evolution

Same Look & Feel, More Powerful Architecture



UADP ASIC & Open IOS-XE Enables...



Most Importantly : Software Defined – Access



Identity-based Policy & Segmentation

Decoupled security policy definition from VLAN and IP Address to enable rapid policy updates



Single Network Fabric

Automation across wired and wireless for optimized traffic flows, and workflow-based management provide consistency at scale



Insights & Telemetry


Analytics and insights into User and application behavior for proactive issue identification and resolution

Industry's first policy-based automation from the Edge to the Cloud

Catalyst 3850 / 3650 - Simplify network operations open-standards APIs

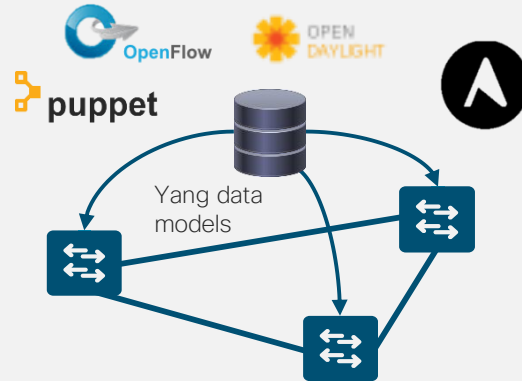


Day 0 Device onboarding



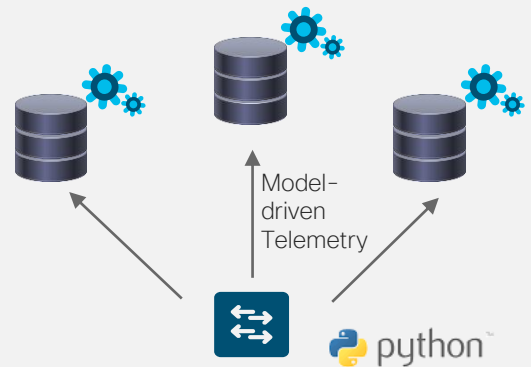
Automatic device bring-up using industry-standard interface

Day 1 Model-driven configuration mgmt



Open-source tool support for configuration and provisioning

Day N Monitoring, analytics, SW management



Model-driven telemetry for monitoring and analytics; On-Box python 2.7

Plug in to operational in minutes

Catalyst 3650 with Cisco DNAC demo

Cisco DNA Center

DESIGN

POLICY

PROVISION

ASSURANCE

PLATFORM



[View Overall Health](#)

Network Snapshot

Sites

As of May 30, 2019 5:16 PM

1

DNS Servers : 1
NTP Servers : 0

[Add Sites](#)

Network Devices

As of May 30, 2019 5:16 PM

7

Unclaimed : 0
Unprovisioned : 7
Unreachable : 0

[Find New Devices](#)

Application Policies

As of May 30, 2019 5:16 PM

1

Successful Deploys: 1
Errored Deploys: 0

[Add New Policy](#)

Network Profiles

As of May 30, 2019 5:16 PM

0

[Manage Profiles](#)

Images

As of May 30, 2019 5:16 PM

3

Untagged Images : 1
Unverified Images : 0

[Import Images/SMUs](#)

DNA Licensed Devices

As of May 30, 2019 5:16 pm

5

Switches : 5
Routers : 0
Access Points : 0

[Manage Licenses](#)

Old car with a huge value



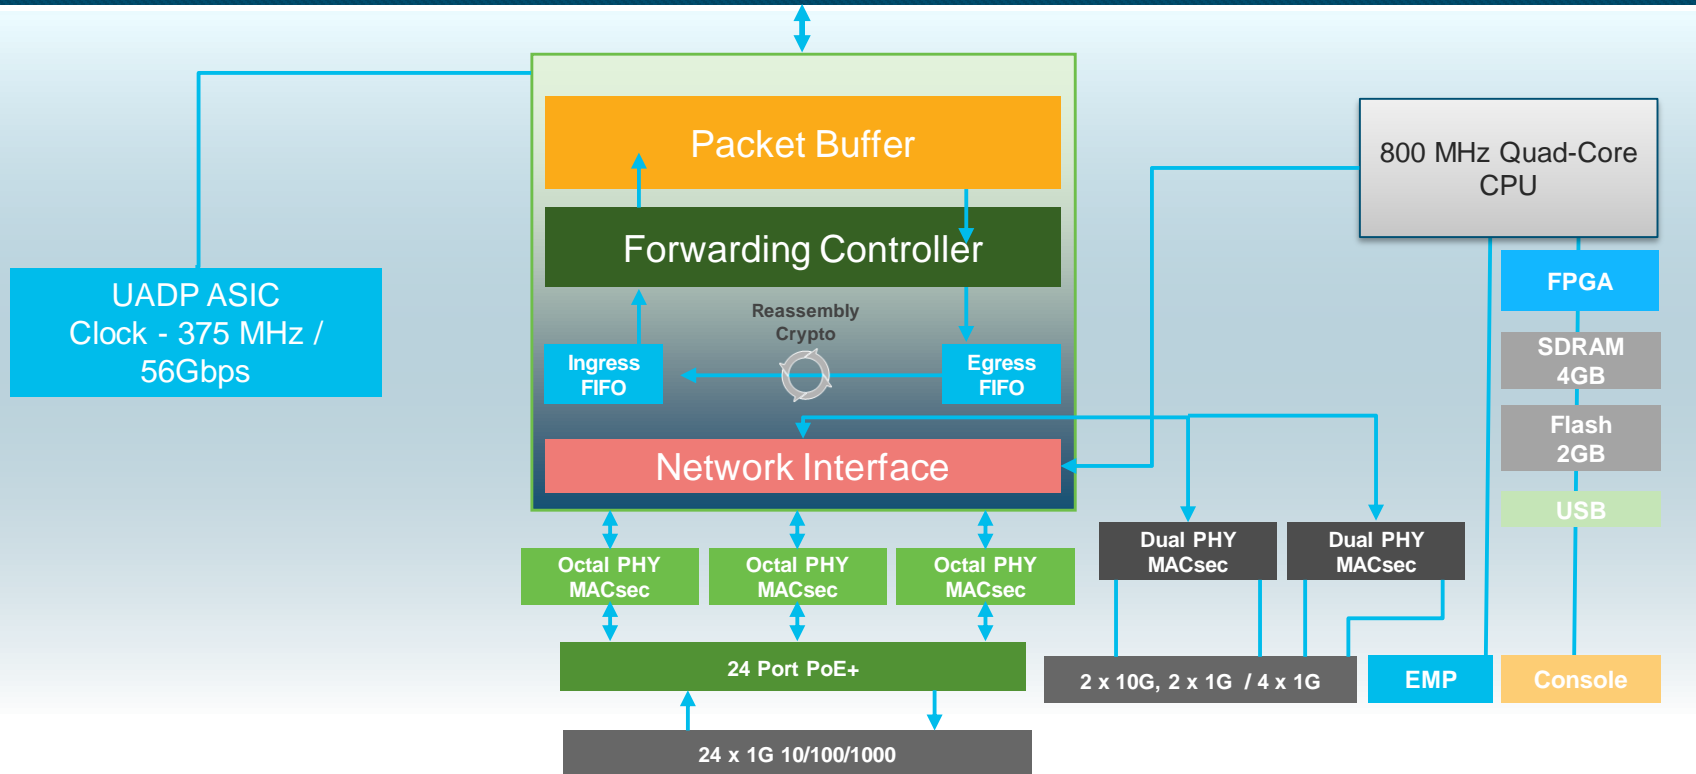
Platform and ASIC Architecture, Packet Walk



You make networking **possible**

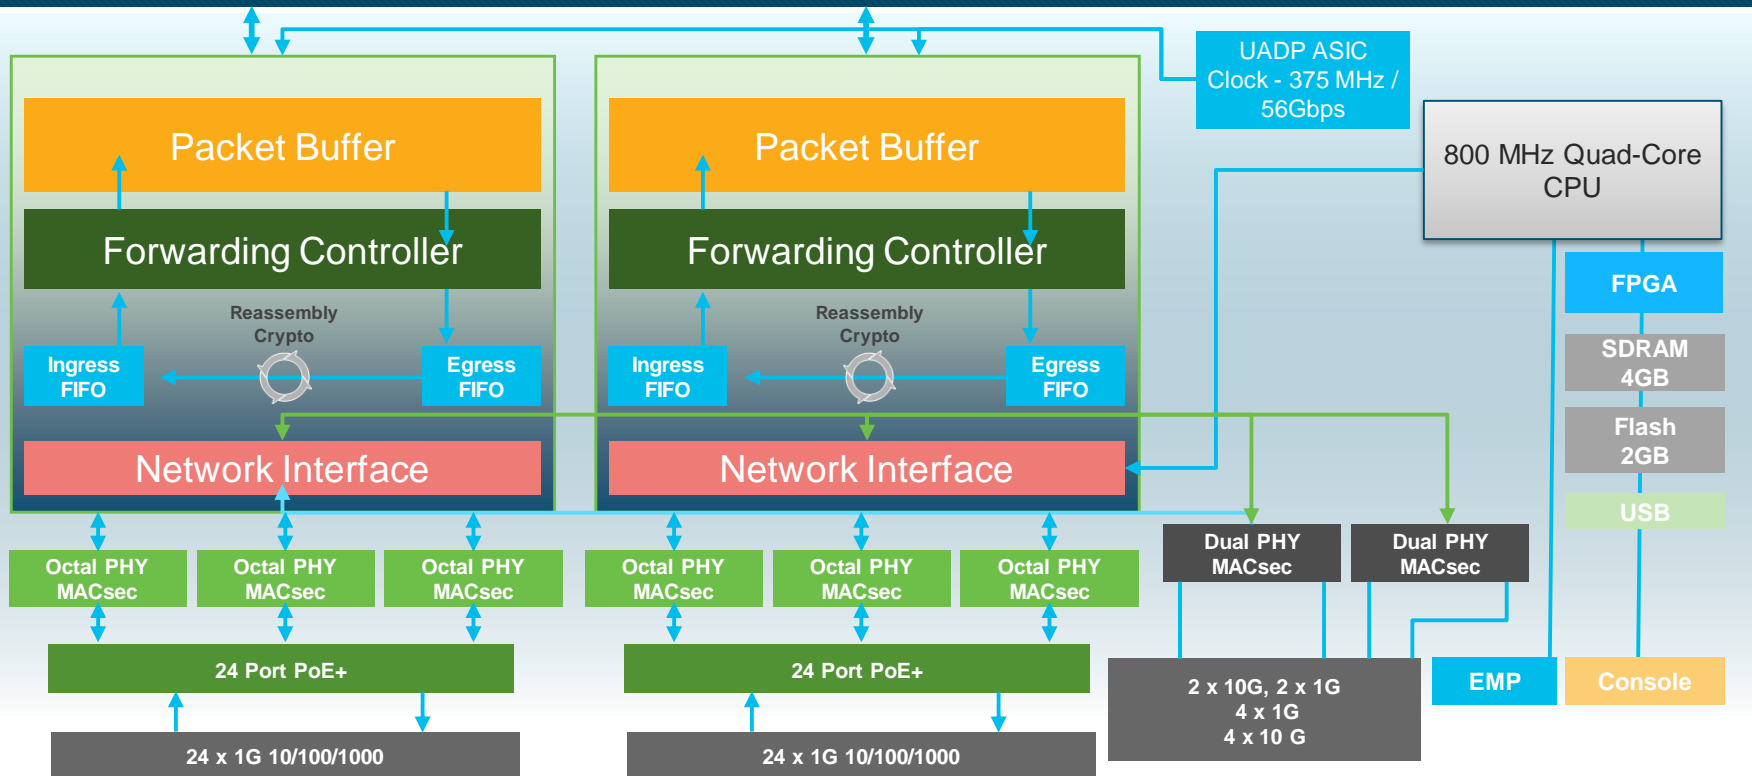
Catalyst 3850/3650–24 Port Layout

480 / 160 Gbps STACK INTERFACE



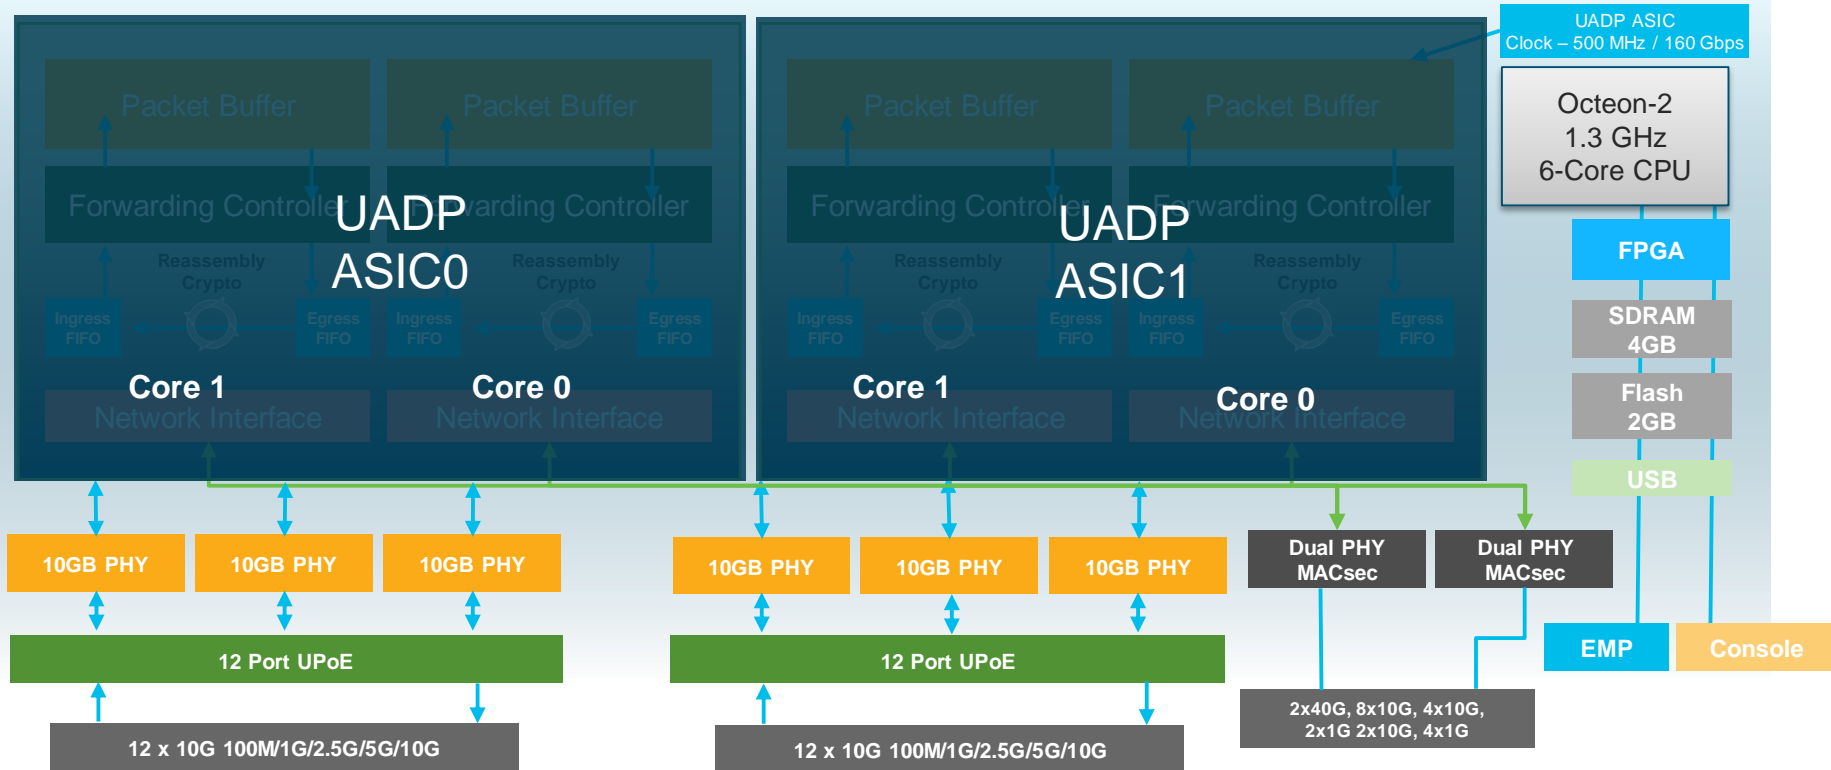
Catalyst 3850/3650-48 Port Layout

480 / 160 Gbps STACK INTERFACE



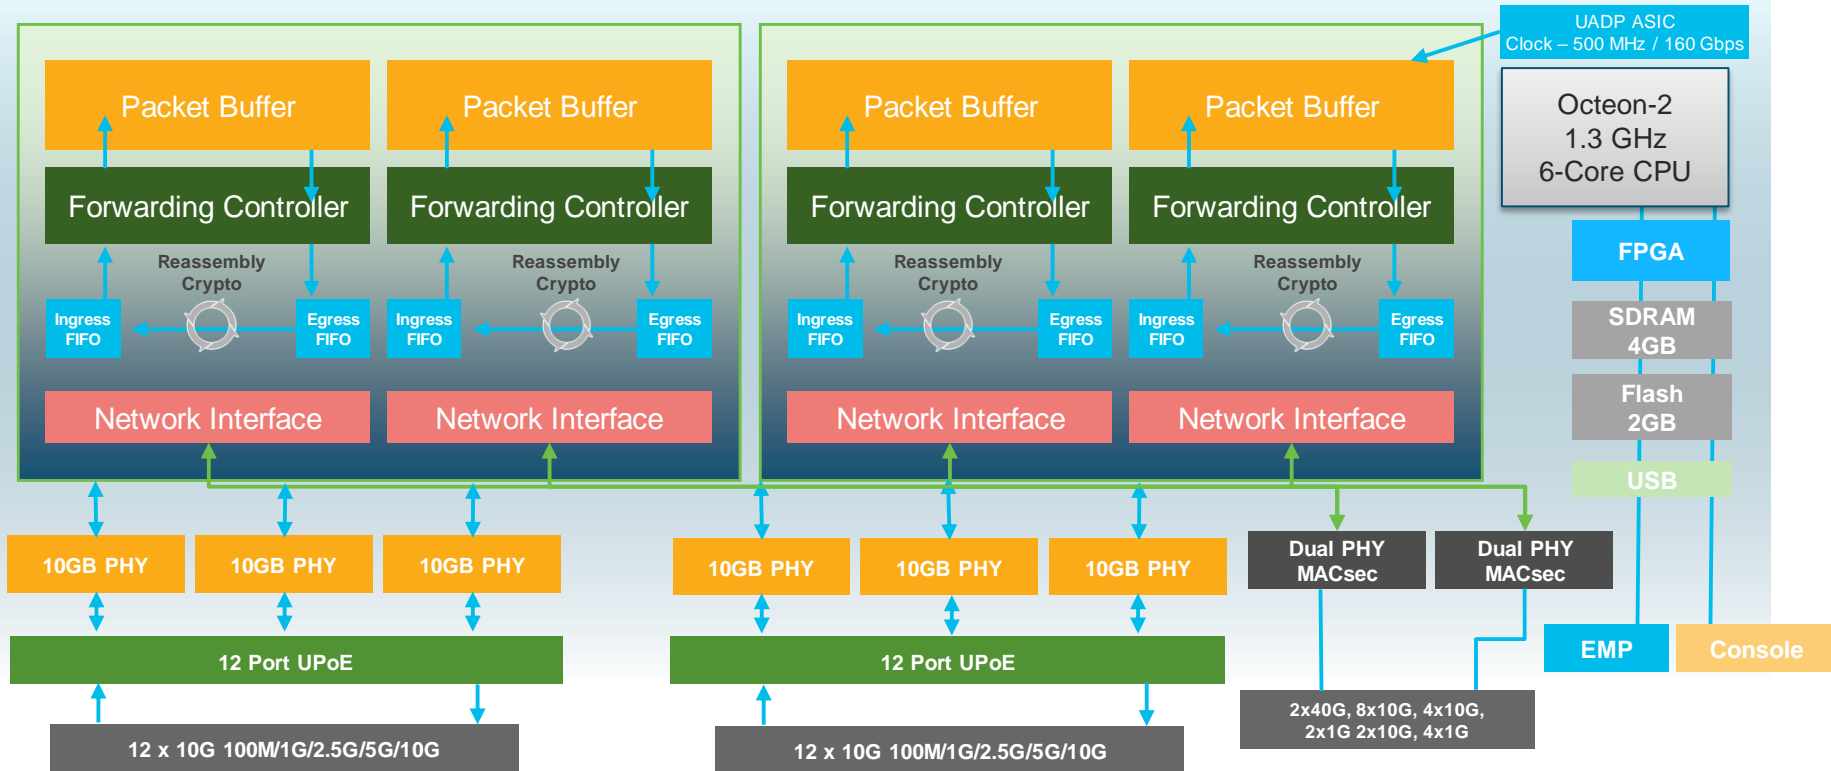
Catalyst 3850 MultiGigabit – 24 Port Layout

480 Gbps STACK INTERFACE



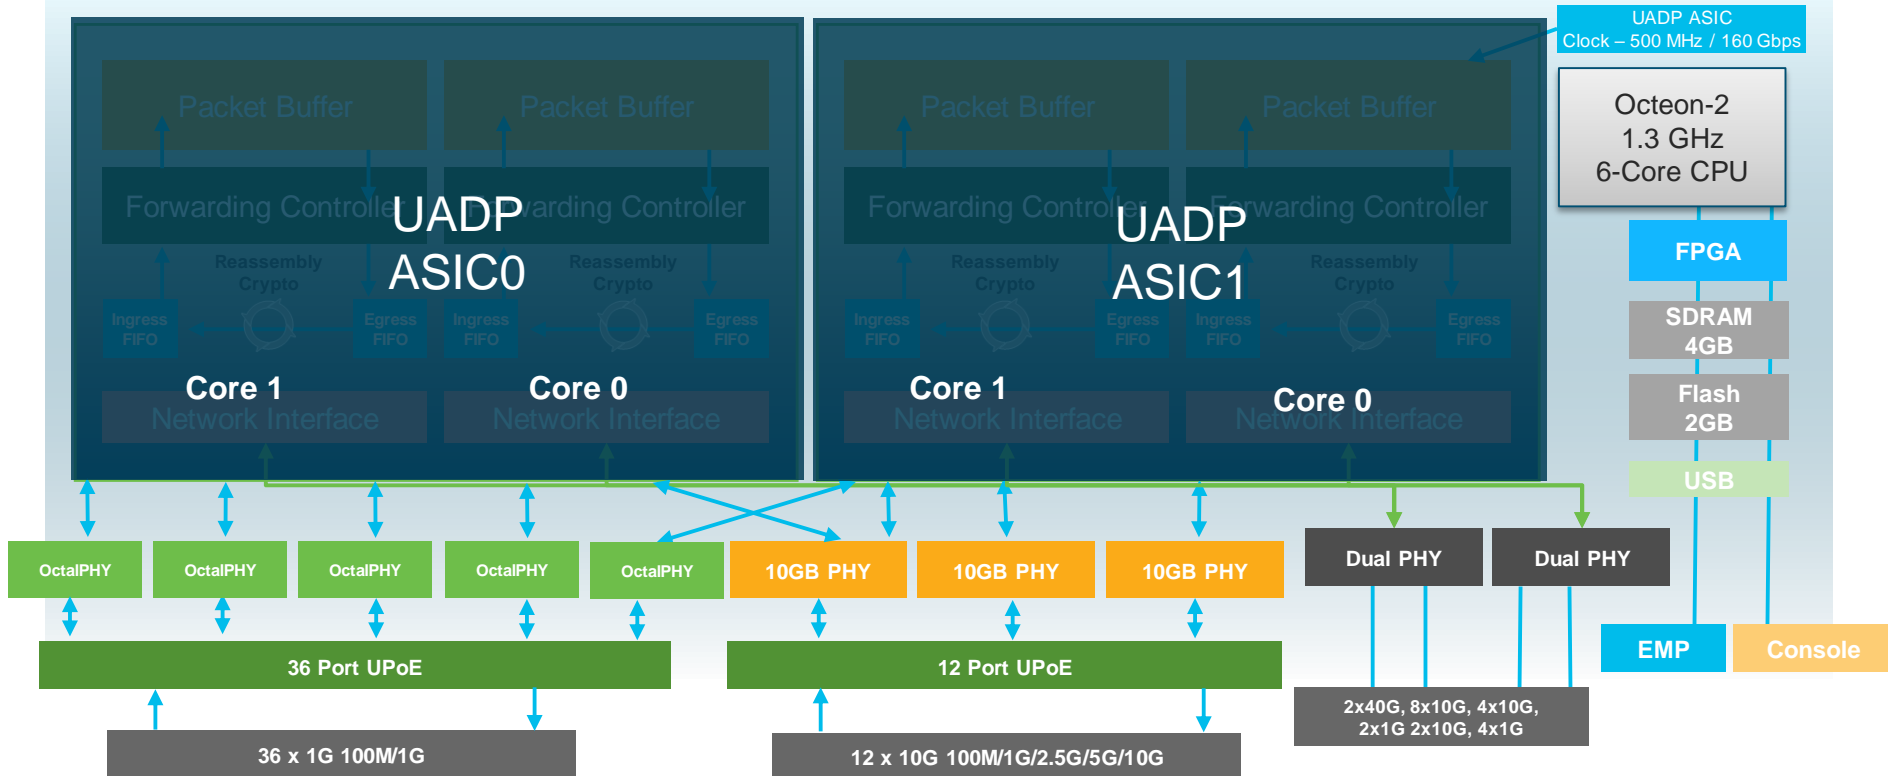
Catalyst 3850 MultiGigabit – 24 Port Layout

480 Gbps STACK INTERFACE



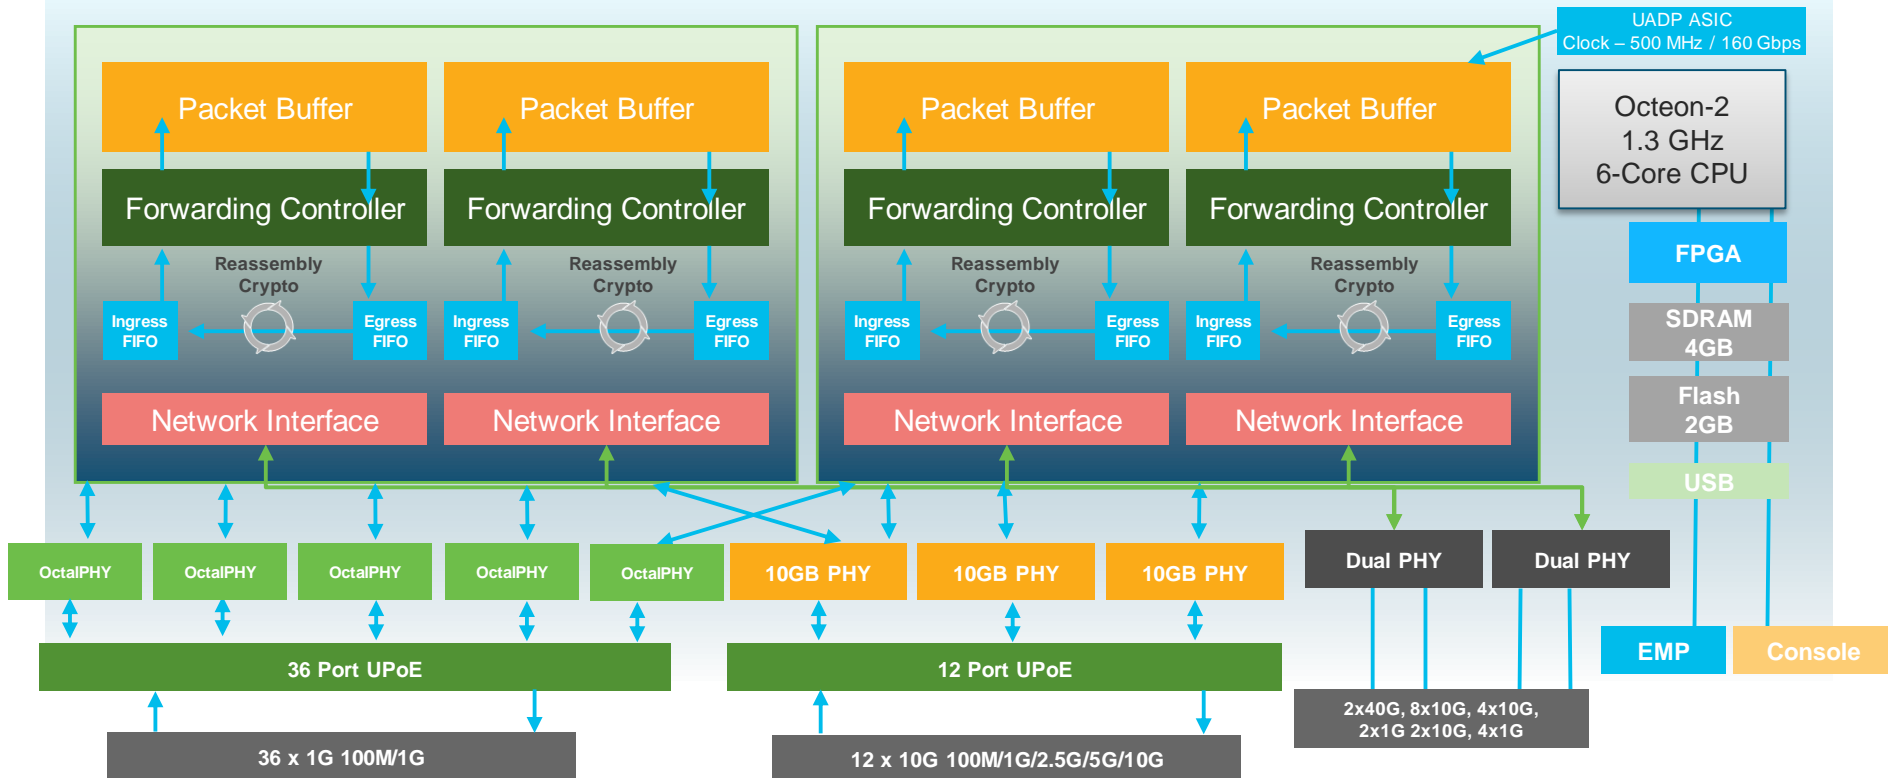
Catalyst 3850/3650 MultiGigabit-48 Port Layout

480 / 160 Gbps STACK INTERFACE



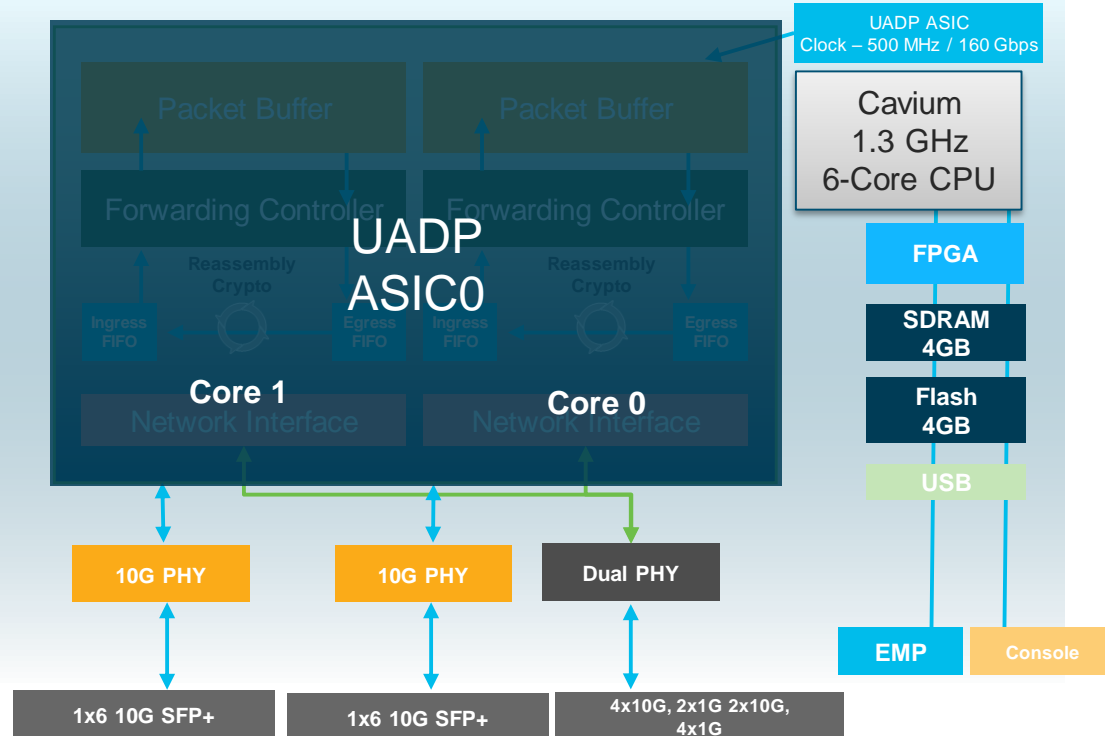
Catalyst 3850/3650 MultiGigabit–48 Port Layout

480 / 160 Gbps STACK INTERFACE



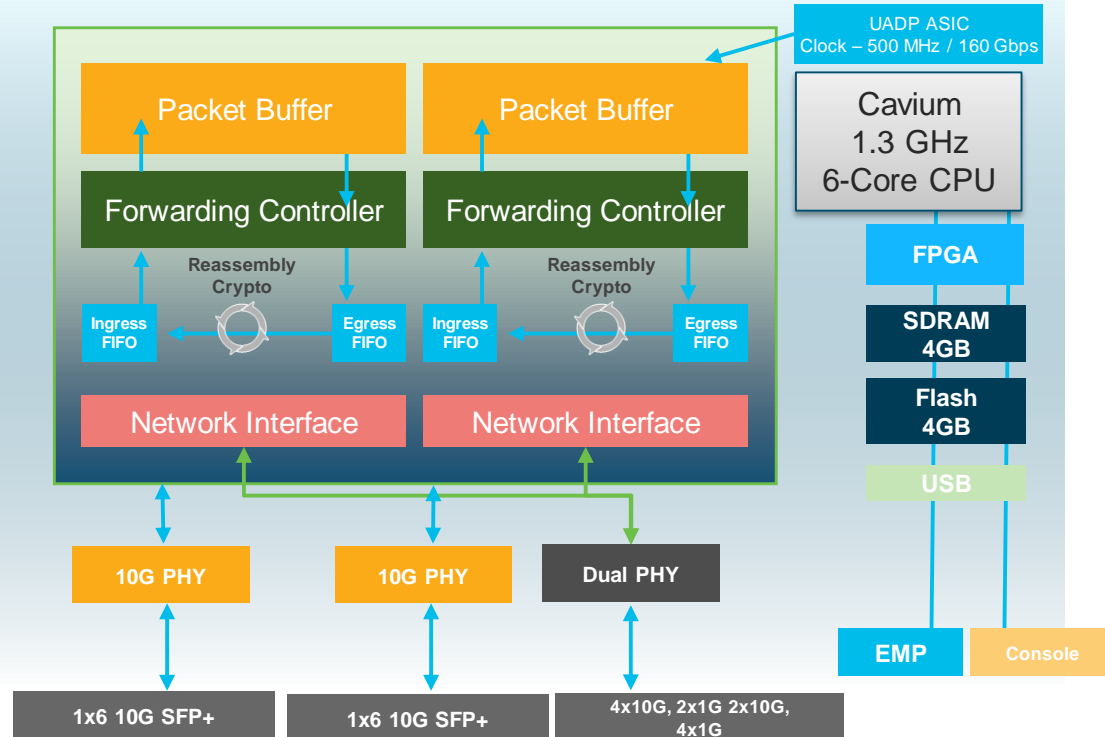
Catalyst 3850 12XS Port Layout

480 Gbps STACK INTERFACE



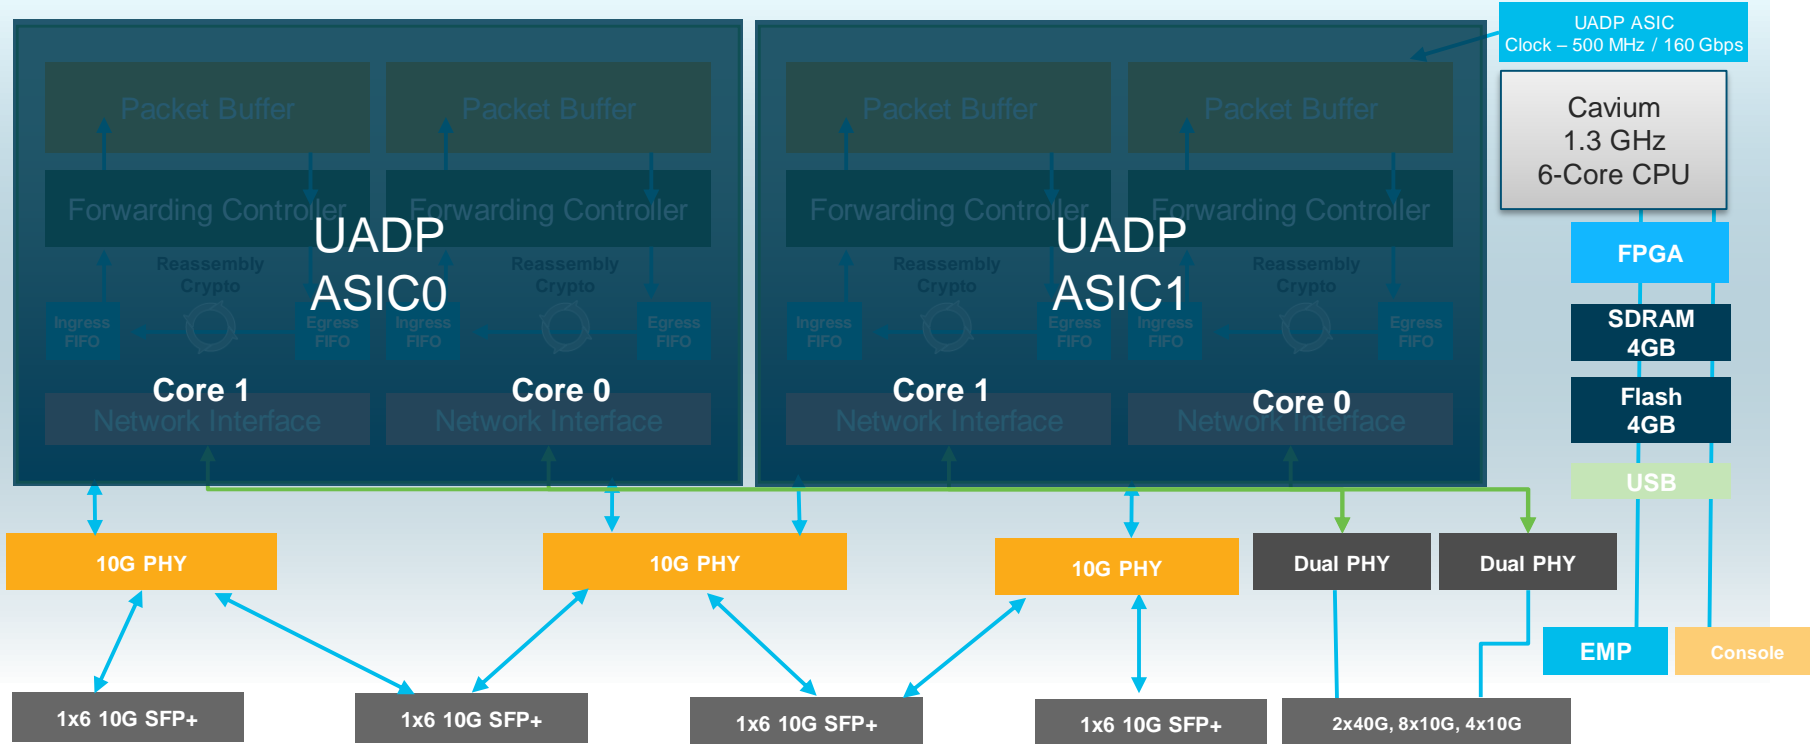
Catalyst 3850 12XS Port Layout

480 Gbps STACK INTERFACE



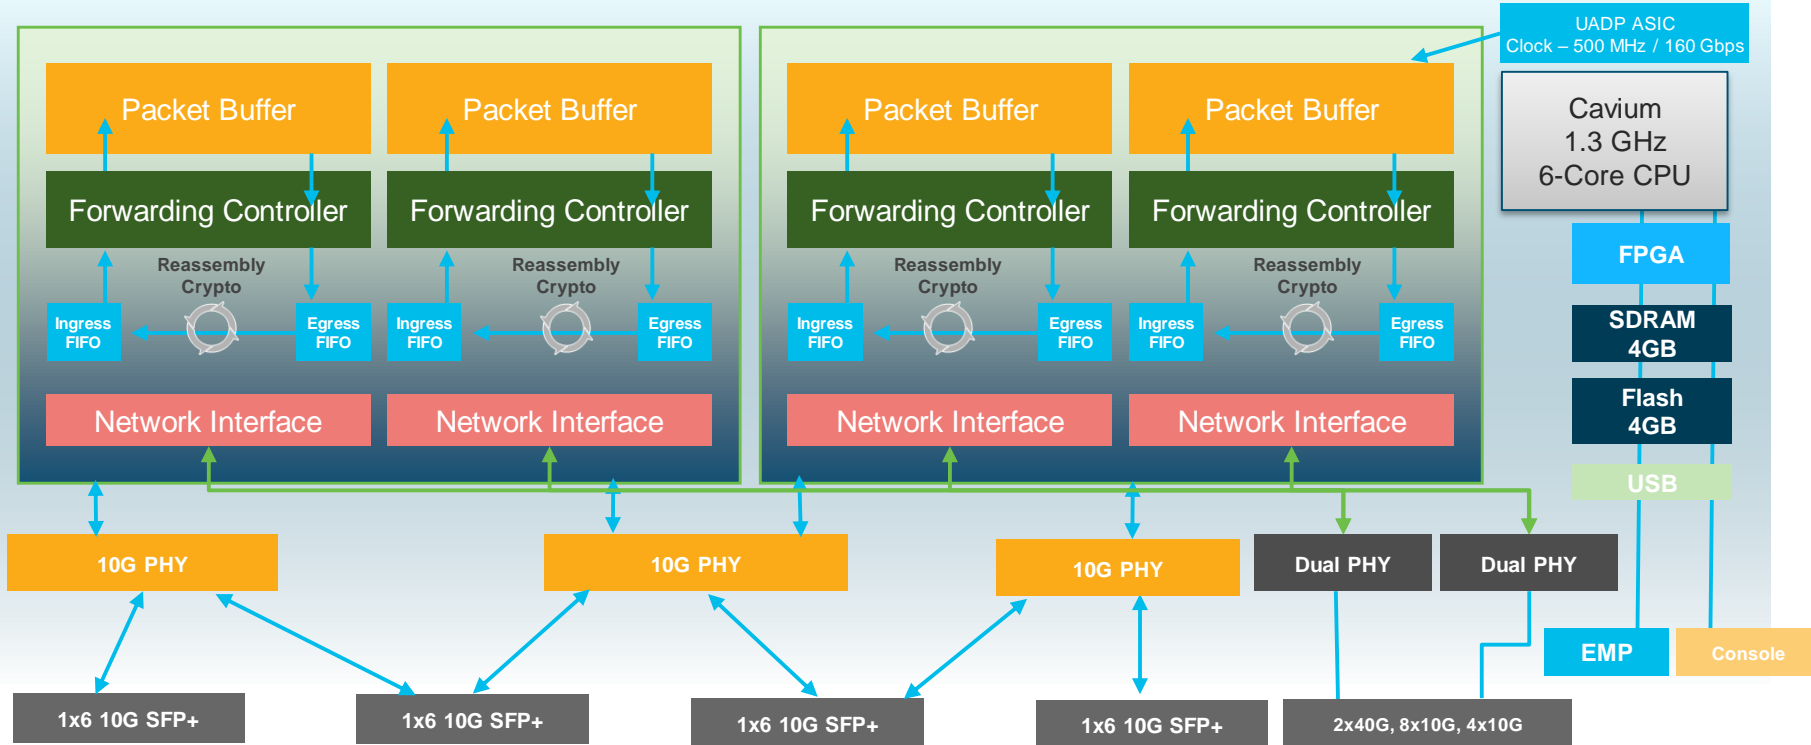
Catalyst 3850 24XS Port Layout

480 Gbps STACK INTERFACE



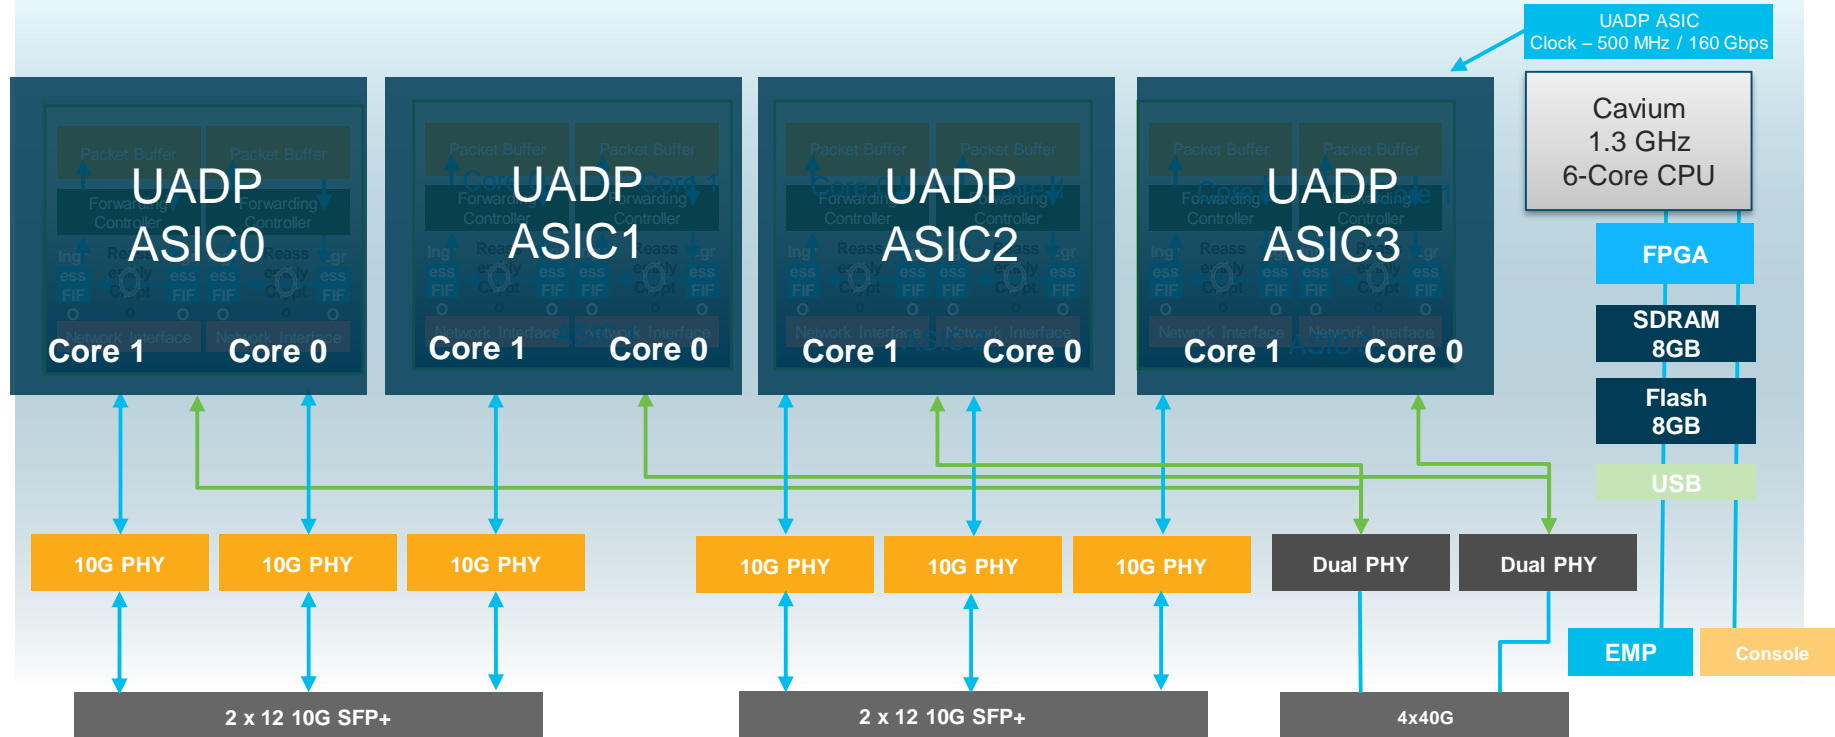
Catalyst 3850 24XS Port Layout

480 Gbps STACK INTERFACE



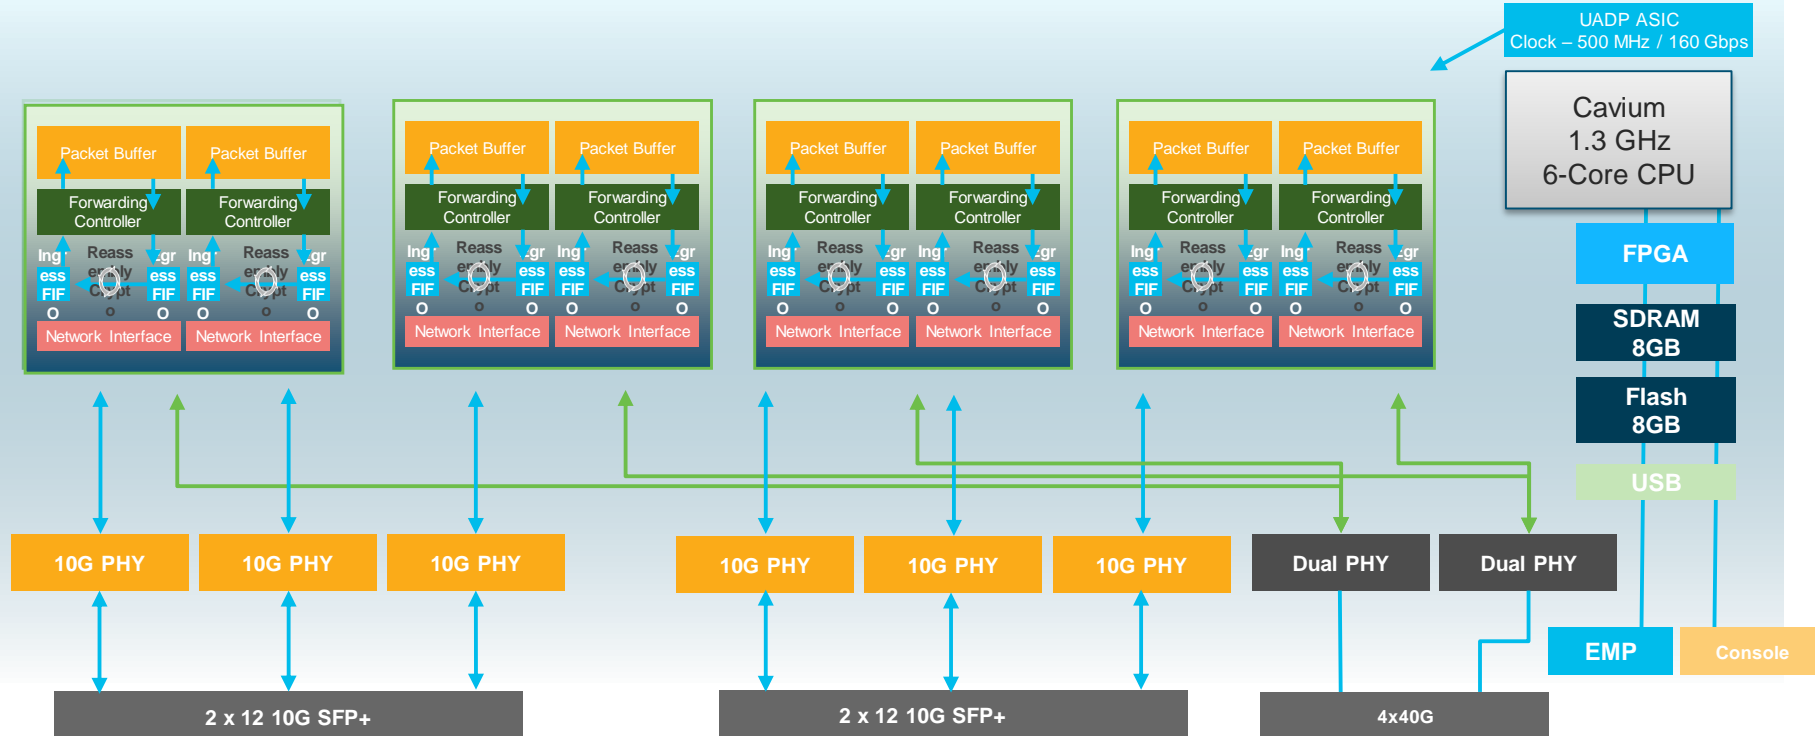
Catalyst 3850 48XS Port Layout

480 Gbps STACK INTERFACE



Catalyst 3850 48XS Port Layout

480 Gbps STACK INTERFACE





Number of ASICs in different versions of Switches

Product Version	UADP Version	Number of ASIC / Total number Cores	Clock Speed	Total Bandwidth Available
24 Port 3850/3650	1.0	1 / 1	375 MHz	56 G
48 Port 3850/3650	1.0	2 / 1	375 MHz	112 G
12/24 Port 3850 SFP	1.0	1 / 1	375 MHz	56 G
12 Port 10G 3850	1.1	1 / 1	500 MHz	160 G
24/48 Port mGig 3850s	1.1	2 / 4	500 MHz	320 G
24 Port mGig 3650	1.1	2 / 4	500 MHz	160 G
48 Port mGig 3650	1.1	2 / 4	500 MHz	320 G
24 Port 10G 3850	1.1	2 / 4	500 MHz	320 G
48 Port 10G 3850	1.1	4 / 8	500 MHz	640 G

Port Mappings – Platform Level Command

```
Cat3850# show platform software fed switch active ifm mappings
```

Interface	IF_ID	Inst	Asic	Core	Port	SubPort	Mac	Cntx	LPN	GPN	Type	Active
TenGigabitEthernet2/0/1	0x8	3	3	0	7	0	7	0	1	97	NIF	Y
TenGigabitEthernet2/0/2	0x9	3	3	0	6	0	6	0	2	98	NIF	Y
TenGigabitEthernet2/0/3	0xa	3	3	0	5	0	5	0	3	99	NIF	Y
TenGigabitEthernet2/0/4	0xb	3	3	0	4	0	4	0	4	100	NIF	Y
TenGigabitEthernet2/0/5	0xc	2	2	0	4	0	4	0	5	101	NIF	Y
TenGigabitEthernet2/0/6	0xd	2	2	0	5	0	5	0	6	102	NIF	Y
TenGigabitEthernet2/0/7	0xe	2	2	0	6	0	6	0	7	103	NIF	Y
TenGigabitEthernet2/0/8	0xf	2	2	0	7	0	7	0	8	104	NIF	Y
TenGigabitEthernet2/0/9	0x10	3	3	0	0	0	0	0	9	105	NIF	Y
TenGigabitEthernet2/0/10	0x11	3	3	0	1	0	1	0	10	106	NIF	Y
TenGigabitEthernet2/0/11	0x12	3	3	0	2	0	2	0	11	107	NIF	Y

```
.....  
<snip>
```

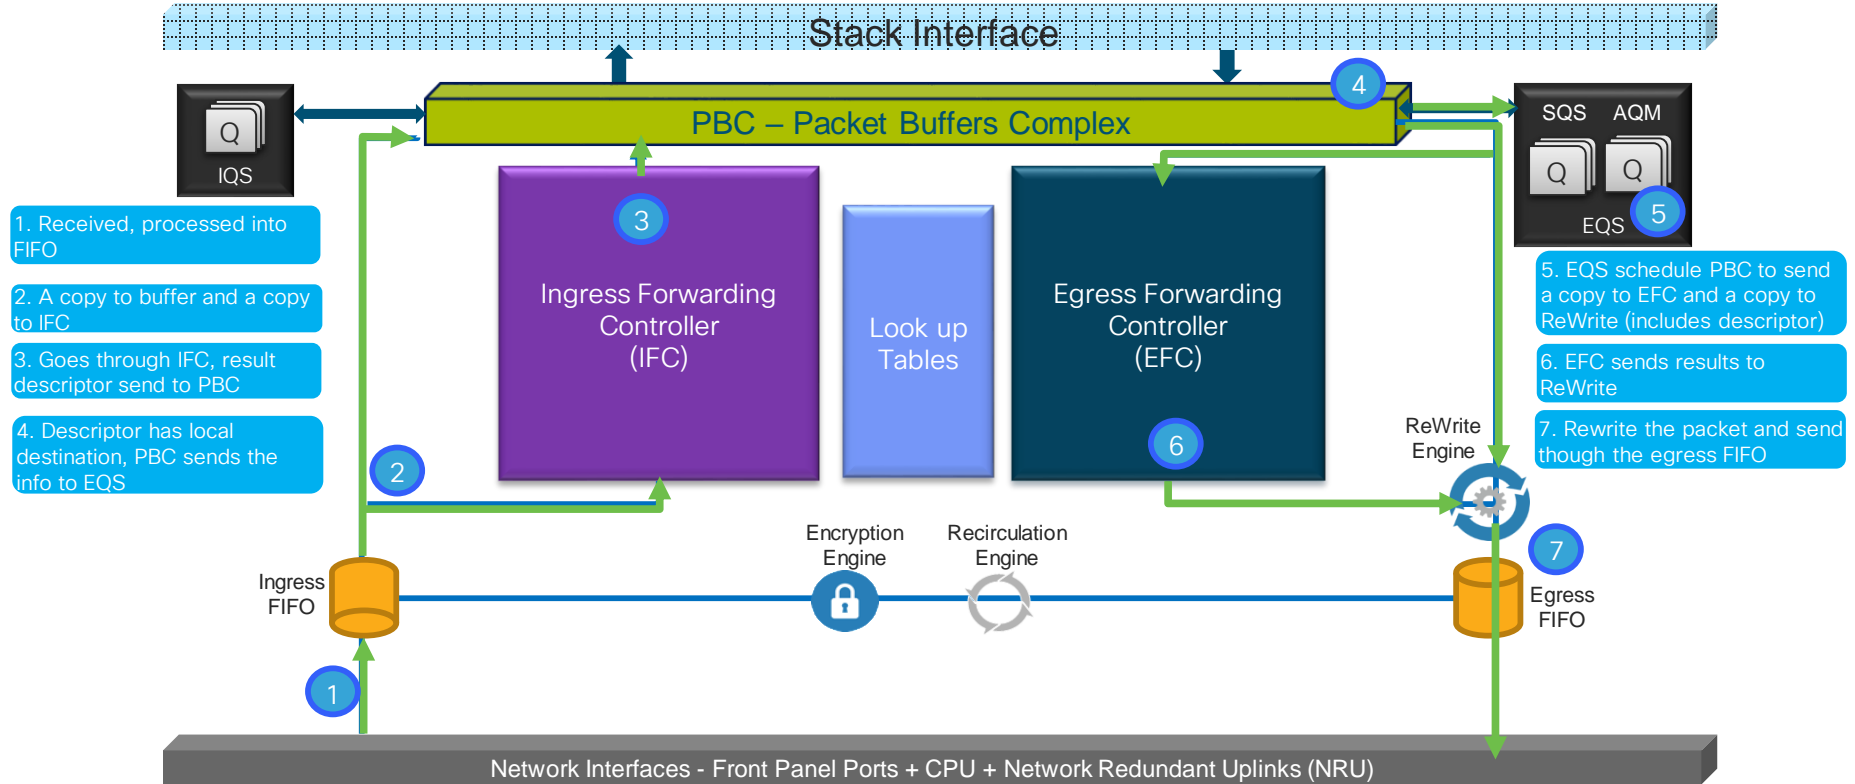
```
.....  
Cat3850#
```

Packet Walks

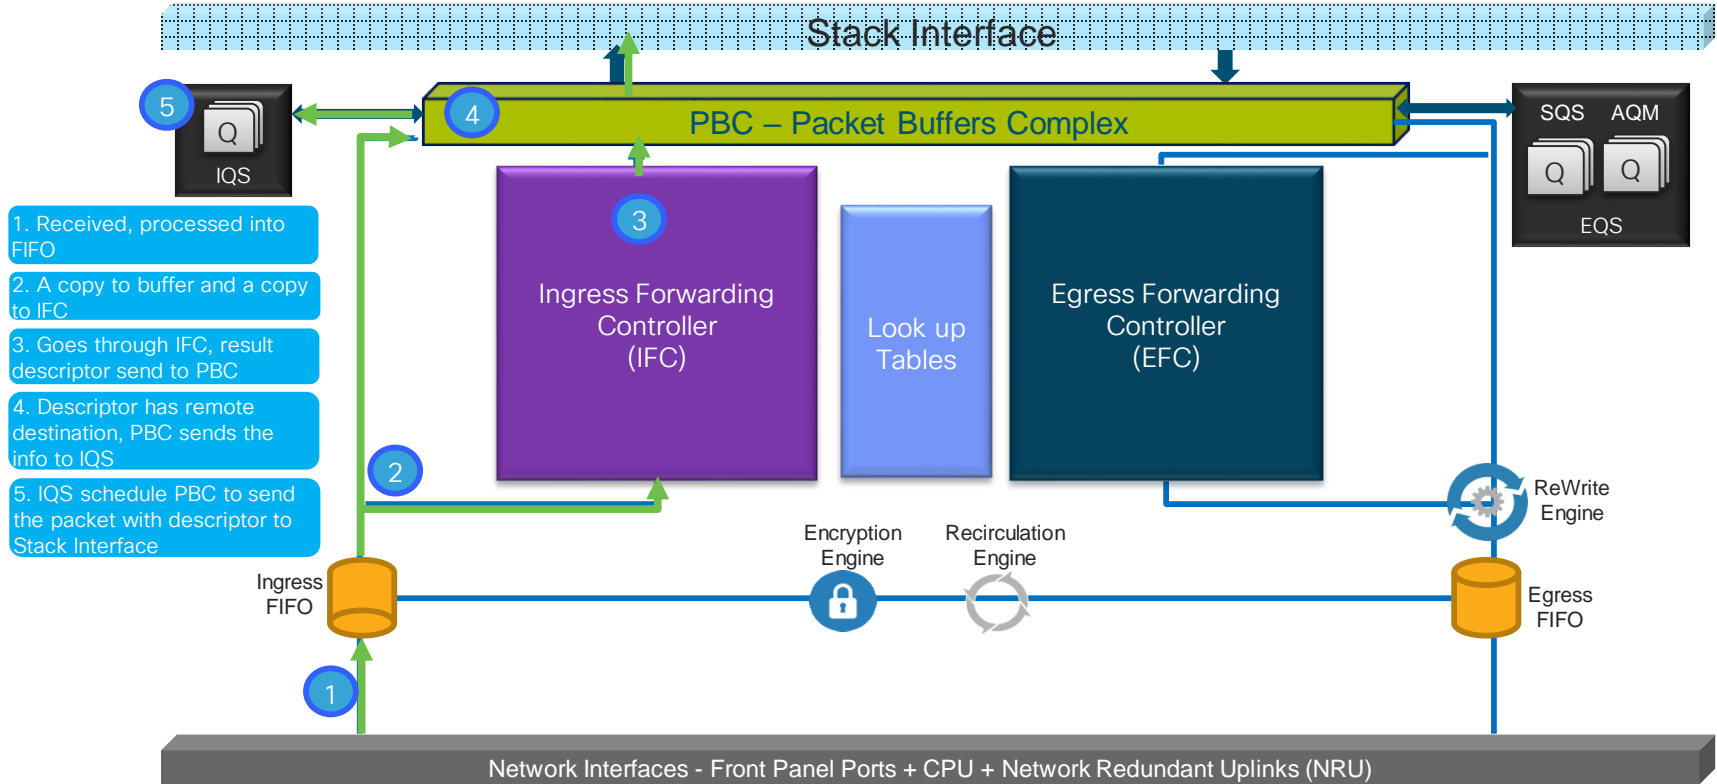


You make security **possible**

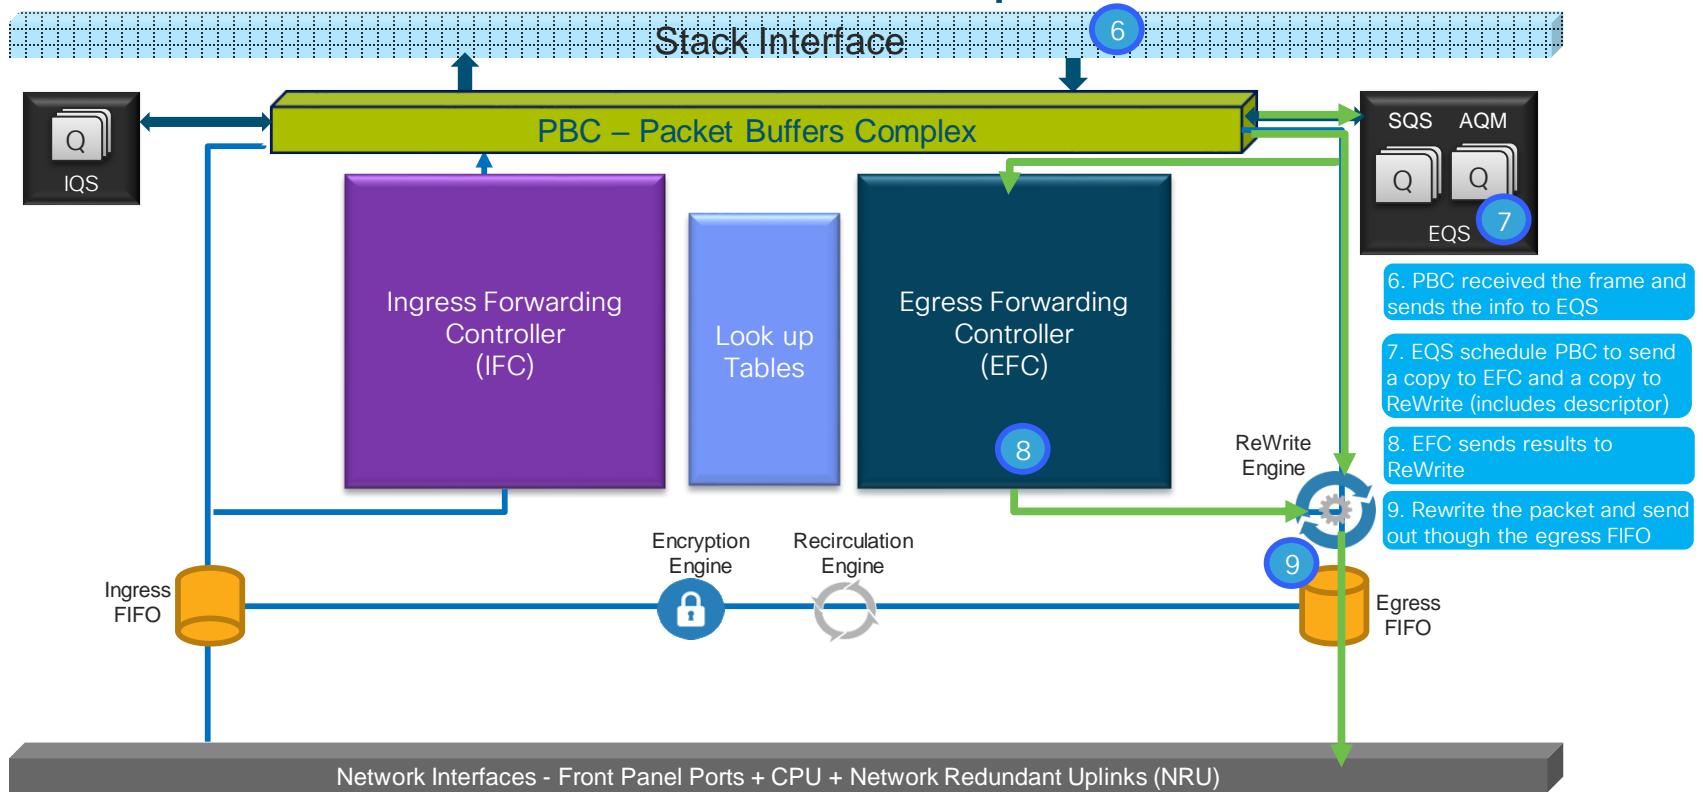
Unicast – within ASIC



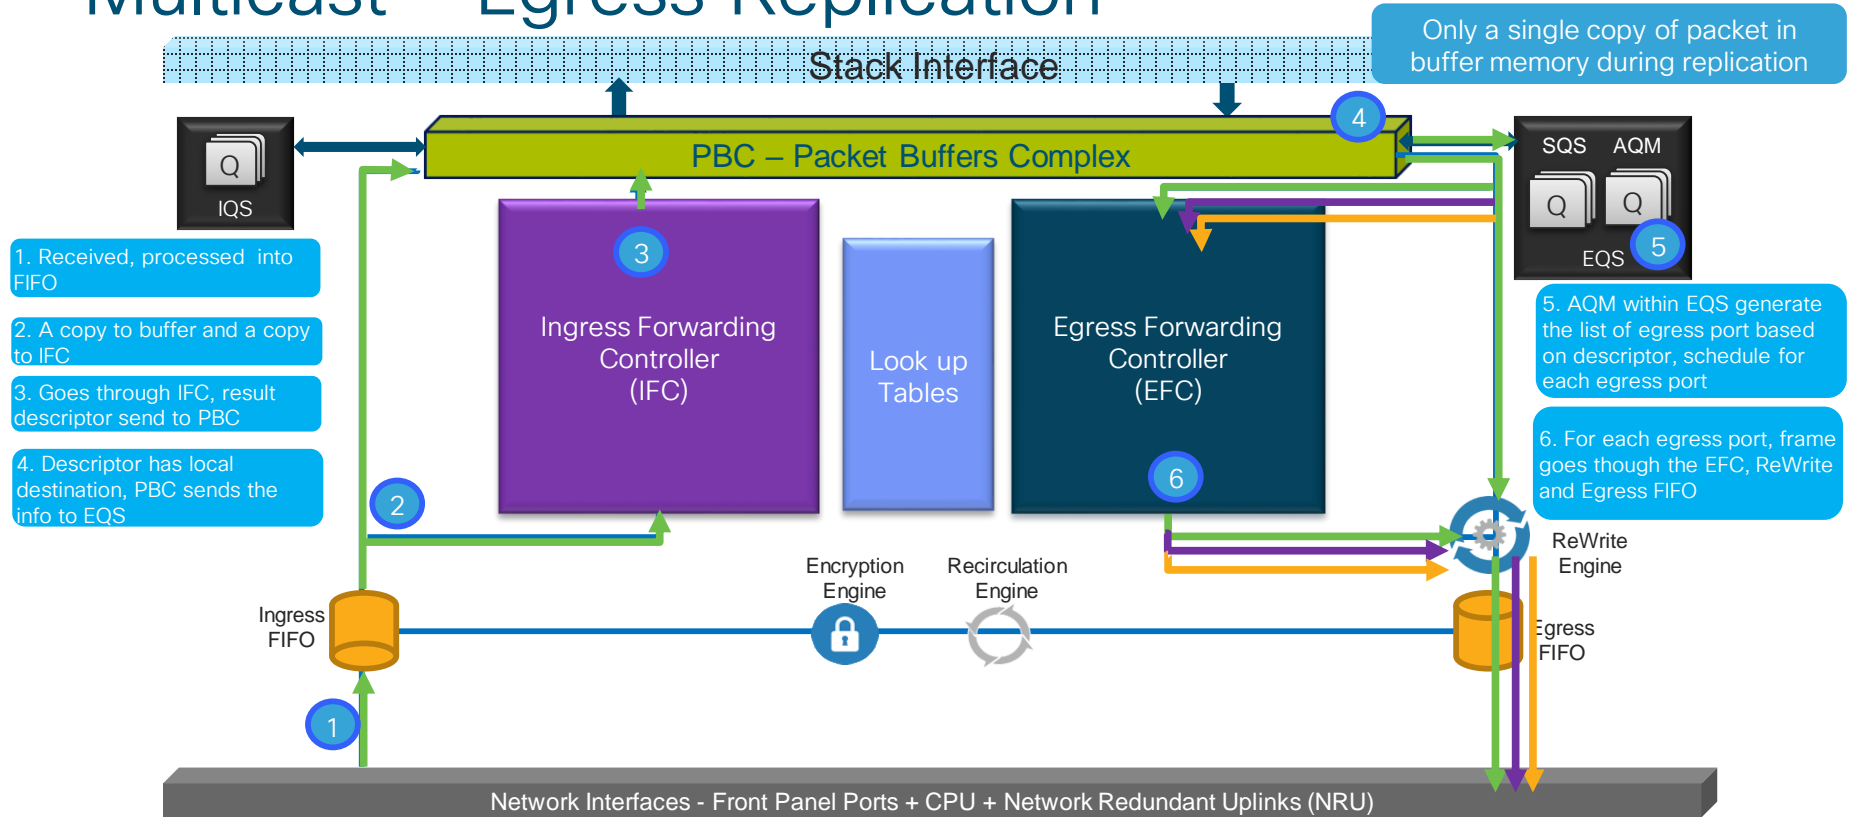
Unicast – Across ASICs on Input





Unicast – Across ASICs on Output



Multicast – Egress Replication





Catalyst 3850 / 3650 - UADP Performance - 375 MHz **Clock Speed**



Future Proofed for 802.11ac and beyond

Catalyst 3850 / 3650 – UADP Performance – 500 MHz Clock Speed



Future Proofed for 802.11ac and beyond


Stacking Architecture



You make the power of data **possible**

The Stack Ring

480 Gbps capacity




- 6 rings in total
- 3 rings go East
- 3 rings go West
- Each ring is 40Gbps
- 240Gbps uni-direction
- Spatial Reuse= 480Gbps

Assuming 4 x 24-port 3850 Switches


Unicast Packet Path on the Stack Ring

- Packet segmented into 256 bytes
- Packet travels half the ring for unicast traffic
- Segments reordered at destination stack port
- Destination strips the packet off the stack ring




Stack Ring Spatial Reuse

- Credit based system on the Stack Ring
- Multiple stack ports grab the ring that is free and they have credits on to transmit




Stack Ring Spatial Reuse

- Unicast flows use only part of the Ring
- Increases the stack ring bandwidth to 480Gbps




Multicast Packet Path on Stack Ring

- One copy of the source packet is placed on the rings
- Interested Stack Ports grab the segments when they see them
- Packet segments travel the whole ring back to source
- The source strips these segments off the ring (Source Stripping)
- Results in efficient replication of multicast traffic for multiple Stack Port receivers




Assuming
4 x 24-port
3850 Switches

Resiliency – StackWise-160




- Modular Stacking (Optional)
 - New stack adapters
 - New connectors and copper cables
- Stack Bandwidth
 - 80 Gbps uni-directional
 - 160 Gbps with spatial reuse
- Stateful Switch Over (SSO)
 - Faster Convergence (vs 3750-X)
 - Active-Standby model
 - Improved Central synchronization on Active Switch for Wired+Wireless
- Tunnel SSO ensures AP, MA-MC connectivity during failover

How many Can I stack together?



Enforced by Software

Up to 9



Limited in ASIC

Did you learn all parts?




Image Source: newatlas.com

High Availability – Data and Power Stacks



You make networking **possible**

Catalyst 3850 Stack vs. Catalyst 6500

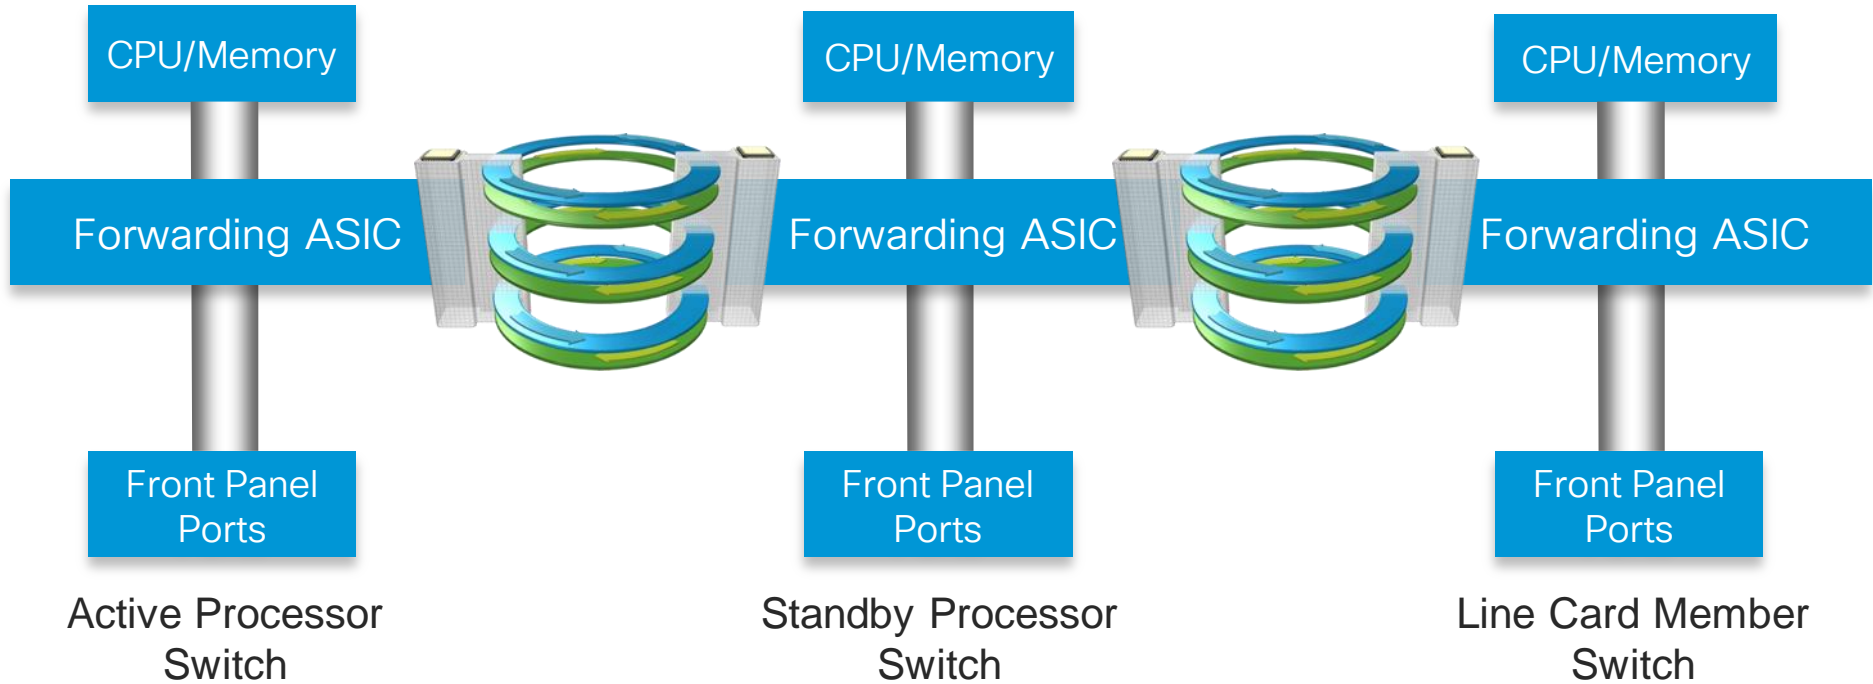
- Active and Standby Members run IOSd, etc.
- Synchronize information
- Active controls Data plane programming for all members
- Member switches act as Line cards—connected via the Stack Cable



- Active and Standby Supervisors
- Run IOS on Supervisors
- Synchronize information
- Active programs all DFCs
- DFCs run a subset of IOS for LCs




Catalyst StackWise Architecture



Centralized Control Plane – Scalable Distributed Data Plane

Stack Discovery

- Stack Interfaces brought online
- Infra and LC Domains boot in parallel
- Stack Discovery Protocol discovers Stack topology – broadcast, followed by neighborcast
- In full ring, discovery exits after all members are found.
- In half ring, system waits for 2mins
- Active Election begins after Discovery exits



```
Stack port 1 cable is connected and the link is up
```

```
Stack port 2 cable is connected and the link is up
```

```
Waiting for 120 seconds for other switches to boot
```

```
%IOSXE-1-PLATFORM: process stack-mgr: %STACKMGR-1-DISC_START: Switch 3 is starting stack discovery.
```

```
##All switches in the stack have been discovered
```


```
Switch number is 3
```

```
%IOSXE-1-PLATFORM: process stack-mgr: %STACKMGR-1-DISC_DONE: Switch 3 has finished stack discovery.
```

```
%IOSXE-1-PLATFORM: process stack-mgr: %STACKMGR-1-SWITCH_ADDED: Switch 3 has been added to the stack.
```


Stack Active Election


- 1) The stack (or switch) whose member has the higher user configurable **priority 1–15**
- 2) The switch or stack whose member has the **lowest MAC address**



```
%IOSXE-1-PLATFORM: process stack-mgr: %STACKMGR-1-ACTIVE_ELECTED: Switch 3 has been elected ACTIVE.
```

Stack Initialization

- Active starts RP Domain (IOSd, WCM, etc) locally
- Programs hardware on all LC Domains
- Traffic resumes once hardware is programmed
- Starts 2min Timer to elect Standby in parallel
- Active elects Standby
- Standby starts RP Domain locally
- Starts Bulk Sync with Active RP
- Standby reaches “Standby Hot”



```
%STACKMGR-1-STANDBY_ELECTED: 3 stack-mgr: Switch 2  
has been elected STANDBY.
```

```
Switch#show switch  
Switch/Stack Mac Address : 2037.0652.a580 - Local Mac Address  
Mac persistency wait time: Indefinite
```

Switch#	Role	Mac Address	Priority	H/W Version	Current State
1	Member	2037.0653.ca80	5	P6A	Ready
2	Standby	2037.0653.db00	10	P6A	HA sync in progress
*3	Active	2037.0652.a580	15	V01	Ready

HA Best Practices & Recommendations

- Power up the first Switch that you want to make it as Active
- Configure Priority of the switch (1-15) – 1 by default – the higher the better
- Power up the second member that you want to make as Standby & then power up rest of the members
- To add a member to an existing stack plug in the stack cable first, then power up the switch
- Avoid stack Merge & Stack split if possible

```
Catalyst3850#switch 1 priority 15
```



```
Catalyst3850#switch 2 priority 14
```



```
Catalyst3850#switch 3 priority 13
```




```
Catalyst3850#switch 4 priority 12
```




Stack Split – Double Failure – Scenario #1

- Stack speed is reduced by half because of the half ring
- The top side of the split remains stable, Active initiating Clean up for the members data
- Lower Side of the Stack reboots - Active election is held on the lower side of the split
- Active starts RP domain locally and programs local hardware as well as that of the member
- Active elects Standby (after 2 min timer), and signals Standby to start its RP Domain
- Active and Standby perform Bulk Sync as part of HA – where lower member is Standby-Hot




Stack Merge – Two Active members in one Stack

- Full ring is restored
- Stack Discovery runs to build Stack topology with broadcast and neighbor cast packets
- HA detects there are two Active switches (1 and 2) in Stack
- Whole Stack reboots
- Stack initialization happens as before
- Configuration of the Active elected is downloaded on all members




Stack Member Addition

- Stack discovery initiated and completed
- Plug in the member, completing full ring
- Power up the member
- Stack Discovery process runs and completes immediately after discovery happens
- Active detects the new addition, and programs the hardware of the member
- Active is not pre-empted by powering on another member even if it was High Priority



Stack Member Deletion

- Stack discovery initiated and completed
- Active detects member removal – and Clean up process is initiated
- Clean-up involves removing TCAM entries referencing removed member, MAC addresses, CDP tables – more like all ports on the member are shutdown
- Configuration is moved to Pre-Provisioned state



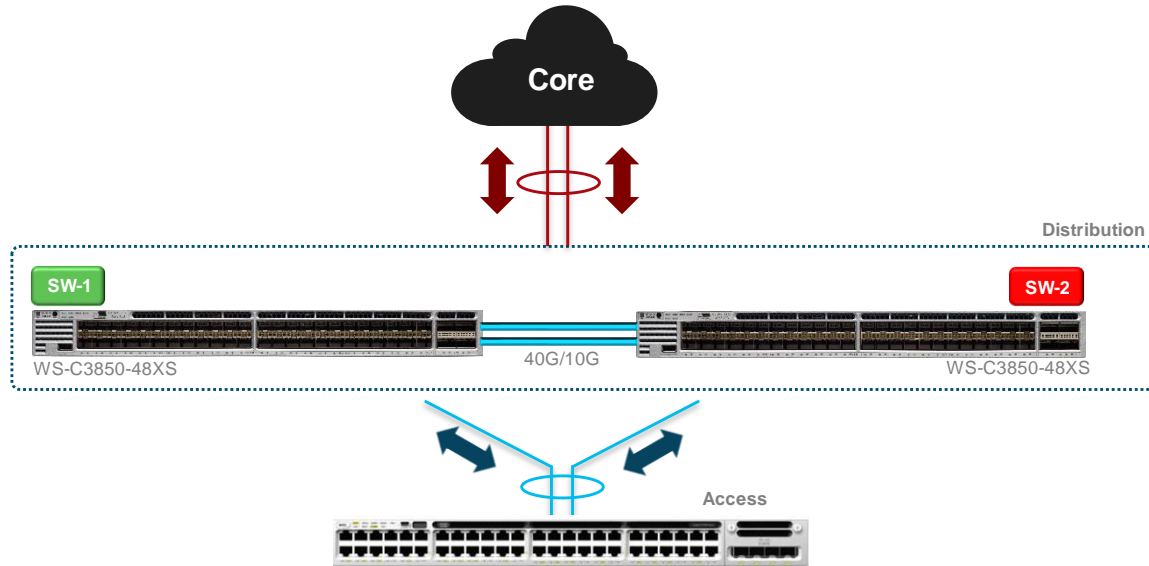
Stackwise Virtual



You make security **possible**

StackWise Virtual Architecture

- Cisco StackWise Virtual extends proven back-panel technology over front-panel network ports
- Unified control and management plane architecture
- Fully distributed forwarding architecture
- Multi-Chassis EtherChannels (MEC)
- StackWise Virtual is supported on 3850-12XS / 24XS / 48XS
- Similar to Cat 6500 VSS



Stack Power



You make the power of data **possible**

StackPower - Overview


“Zero-Footprint” RPS deployment



- Provides RPS functionality with Zero RPS Footprint
- Pay-as-you-grow architecture – similar to the Data Stack
- 1+N Redundancy with Inline Power
- Up to 4 Switches in a StackPower Ring
- Multiple StackPower Possible within one Data Stack
- Up to 9 Switches in a star topology with XPS

Power Budget Modes


Power Sharing Mode



2530W - 30W

- The Default Mode – Pools Power from All PS
- Total POE Budget = Sum of All PS – 30W

Redundant Mode




2530W - 1100W - 30W

- User Configurable – Reserves the Largest PS
- Total POE Budget = Sum of All PS – Largest PS – 30W

How StackPower Works?

Power Sharing Mode – Default Mode





```
WS-C3850-24U# show stack-power
```

Power Stack Name	Stack Mode	Stack Topolgy	Total Pwr(W)	Rsvd Pwr(W)	Alloc Pwr(W)	Unused Pwr(W)	Num SW	Num PS
-----	-----	-----	-----	-----	-----	-----	---	---
MDF	SP-PS	Ring	5115	55	1010	4050	4	5

How StackPower Works?

Redundant Mode



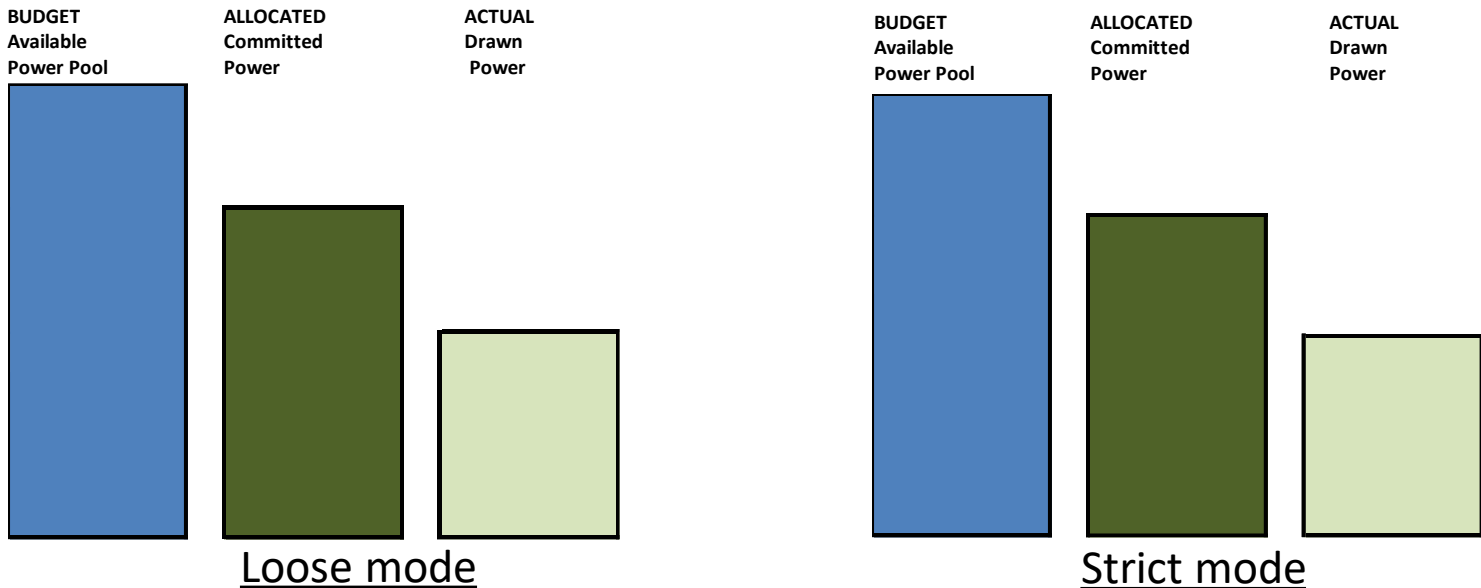
```
3850(config)# stack-power stack MDF
3850(config-stackpower)# mode redundant
```

WS-C3850-24U# show stack-power detail

Power Stack Name	stack Mode	Stack Topolgy	Total Pwr(W)	Rsvd Pwr(W)	Alloc Pwr(W)	Unused Pwr(W)	Num SW	Num PS
MDF	SP-R	Ring	5115	1135	1010	2970	4	5

Enforcement Modes

Strict & Loose Modes Control The Behavior of Load Shed



- Loose mode allows for a negative power budget
- Strict mode sheds load as soon as the power budget goes below the Allocated power level

Power Supplies

Catalyst 3850



350WAC

440WDC

715WAC

1100WAC

Same as 3750-X—
Interchangeable New PIDs

Catalyst 3650



250WAC

640WAC

640WDC

1025WAC

Wider Than 3850/3750-X PSs
Different Watts Capacity

MultiGigabit SKUs
Same PS as 3850s



715WAC


1100WAC

System Power Reserved

- 30~60W StackPower Reserve
- System Power is Reserved based on different PIDs

Catalyst 3850/3650 Version	System Power Reserved	Catalyst3850/3650 Version	System Power Reserved
24 Port 3850 Copper (Data/PoE/PoE+/UPoE)	200	48 Port 3850 Fiber SFP+	280
48 Port 3850 Copper (Data/PoE/PoE+/UPoE)	280	24 Port 3850 mGig	520
12 Port 3850 Fiber SFP	200	48 Port 3850 mGig	470
24 Port 3850 Fiber SFP	200	24 Port 3650	200
12 Port 3850 Fiber SFP+	300	48 Port 3650	280
24 Port 3850 Fiber SFP+	410		

LOW BATTERY



NEED CAFFEINE

Image Source: [pinterest.com](https://www.pinterest.com)

BRK.



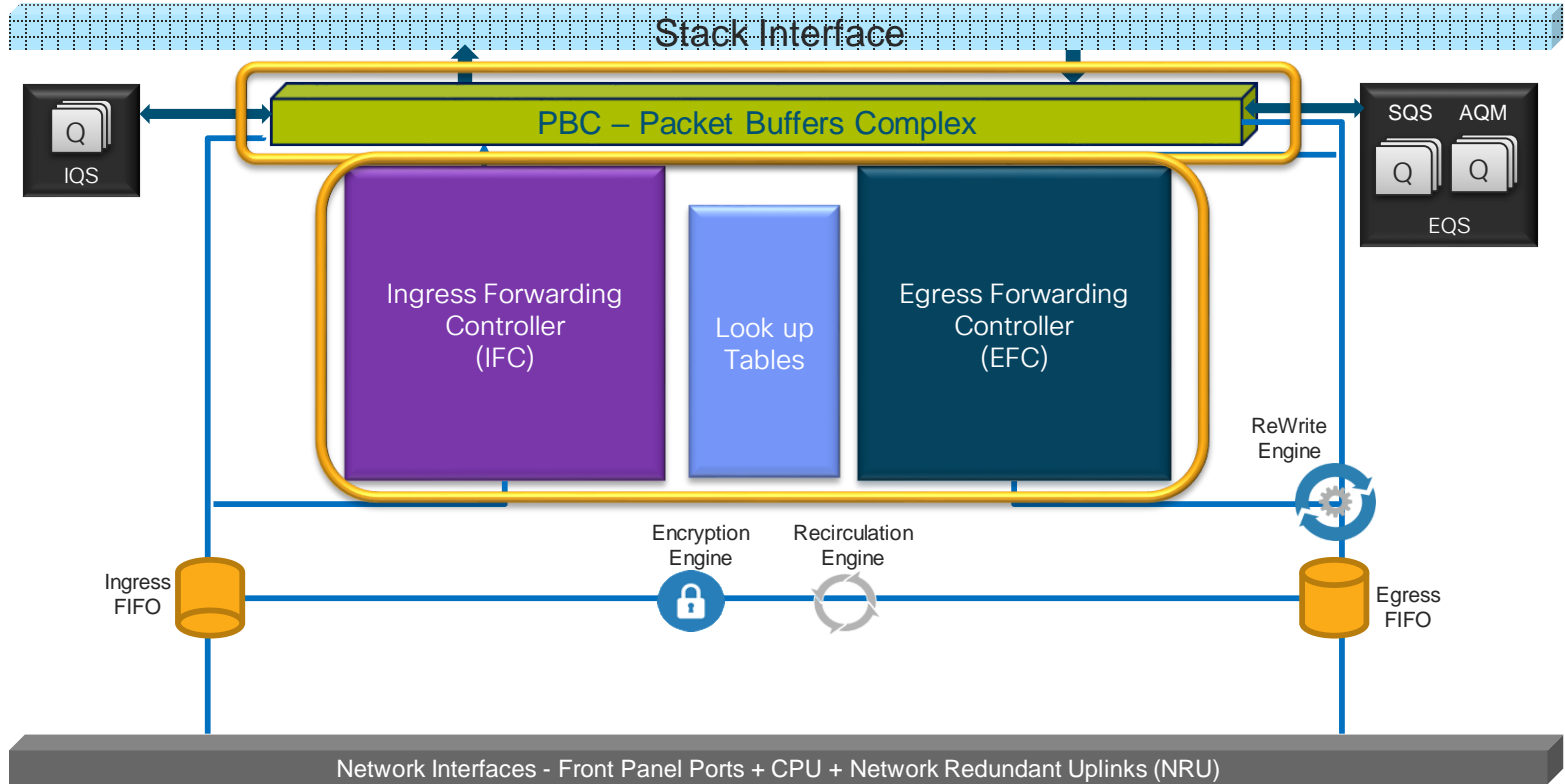
Image Source: [teecentury.com](https://www.teecentury.com)

Scale



You make networking **possible**

UADP ASIC Block Diagram



TCAM and SRAM Resource Usage

```
Cat3850# show platform hardware fed switch active fwd-asic resource tcam utilization
```


```
CAM Utilization for ASIC [0]
```

```
Table
```

	Max Values	Used Values
Unicast MAC addresses	32768/512	18/22
L3 Multicast entries	4096/512	0/7
L2 Multicast entries	4096/512	0/9
Directly or indirectly connected routes	16384/7168	2/19
QoS Access Control Entries	2560	620
Security Access Control Entries	3072	150
Netflow ACES	768	15
Policy Based Routing ACES	1024	9
Flow SPAN ACES	512	5
Output Flow SPAN ACES	512	8
Control Plane Entries	512	230
Tunnels	256	17
Lisp Instance Mapping Entries	256	3
Input Security Associations	256	4
Output Security Associations and Policies	256	5
SGT_DGT	4096/512	0/1
CLIENT_LE	4096/256	0/0
INPUT_GROUP_LE	6144	0
OUTPUT_GROUP_LE	6144	0
MACsec SPD	256	2

Catalyst 3850 – TCAM & ACL Scale

ACL Resources	
IPv4 Entries	3000 Entries
IPv6 Entries	Half the IPv4
One type of IPv4 ACL (RACL, PACL, VACL, GACL*)	1500 Entries
L4OPs/Label	8 L4OPs
Ingress VCUs	196
Egress VCUs	92

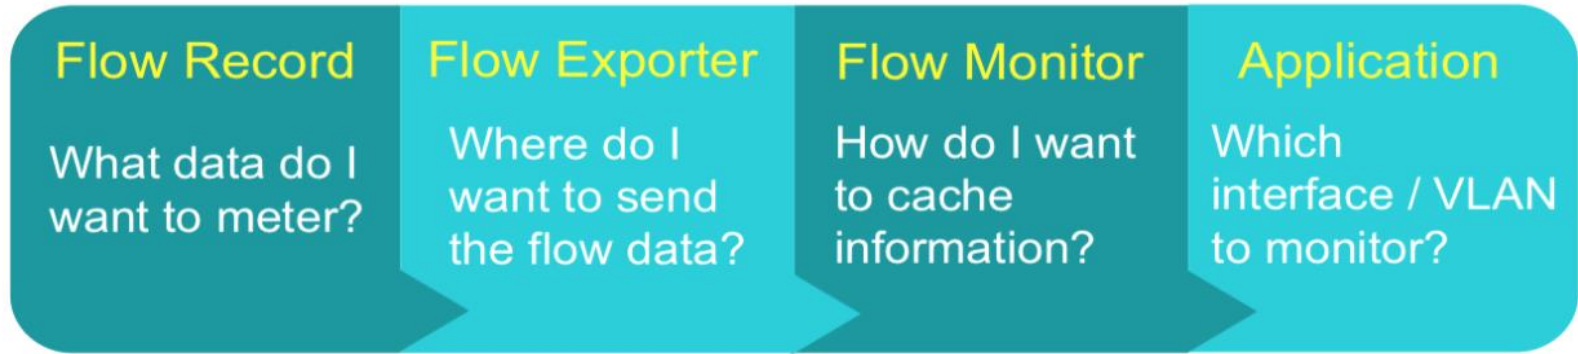


Regions are flexible BUT cannot span across TAQs

- GACL (Group Client ACL) – Any dot1x client attached features like dACL, QoS, Filter ID, Per User ACLs are in GACL region
- Order of Processing : GACL → PACL → VACL → RACL
- TAQ – ACL TQD (TCAM Quads)
- VCU = Value Comparison Unit

Catalyst 3850 – Netflow Scale

Configuring FNF involves 4 major steps:



FnF fully integrated in the ASIC – NO performance impact

Ingress & Egress FnF supported on all ports, VLANs & SSIDs

System Scalability: **24K flows / ASIC. 1.0 and 2 x 24k / ASIC 1.1**

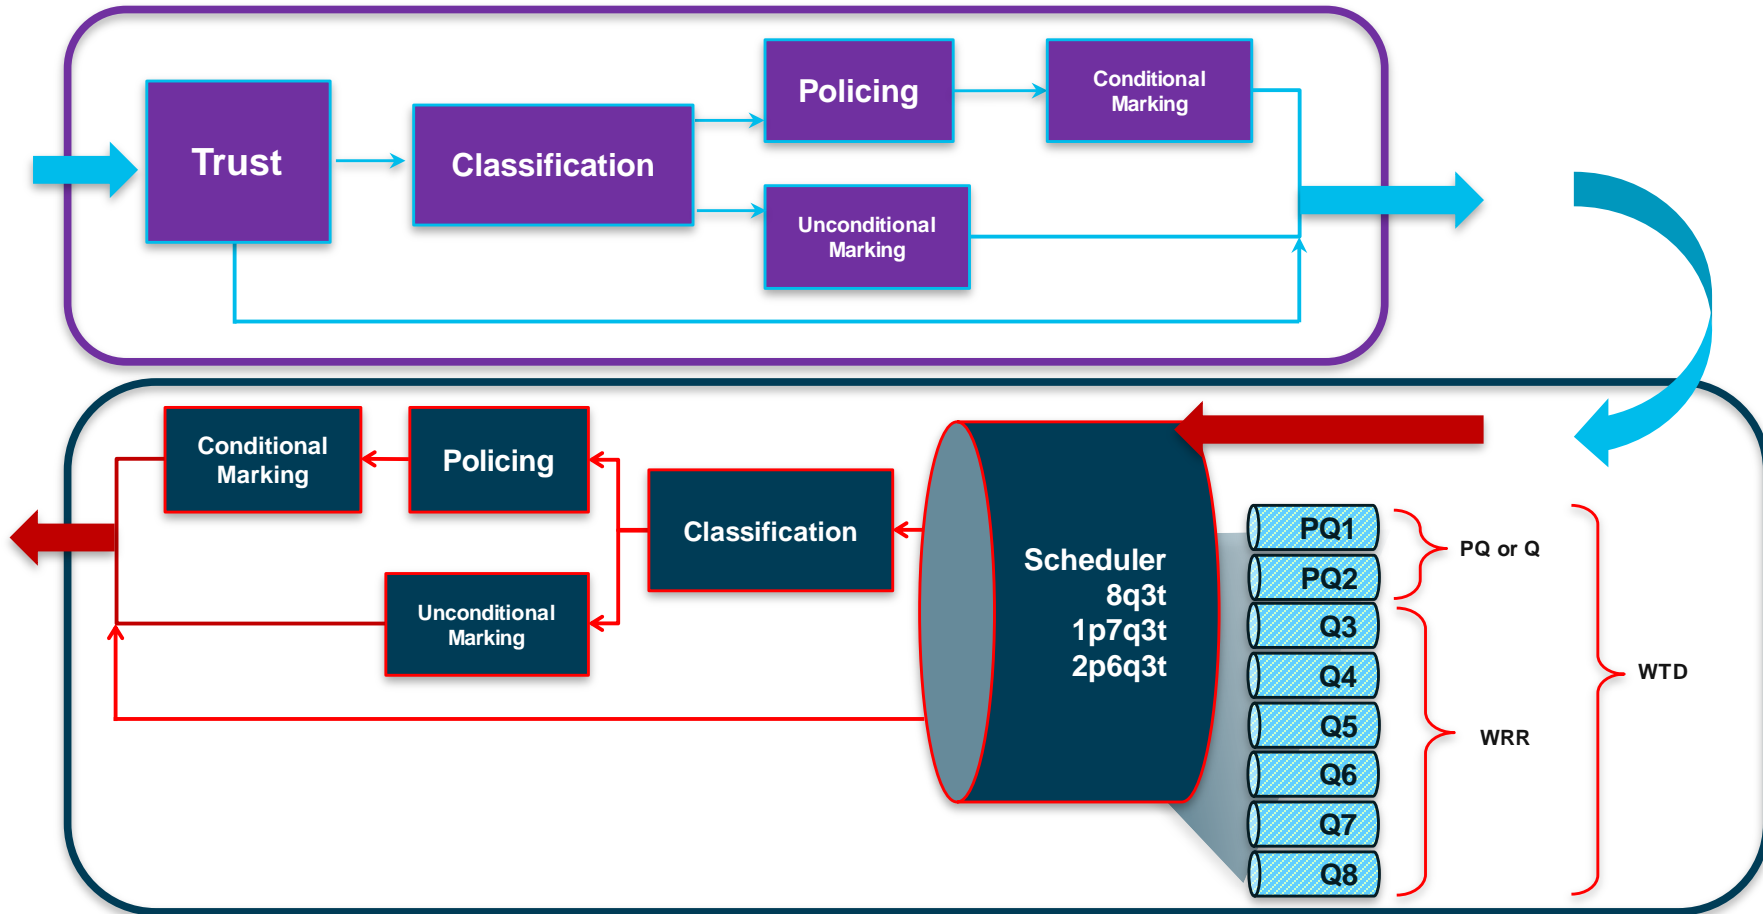
IPv4 & IPv6 capable

QoS



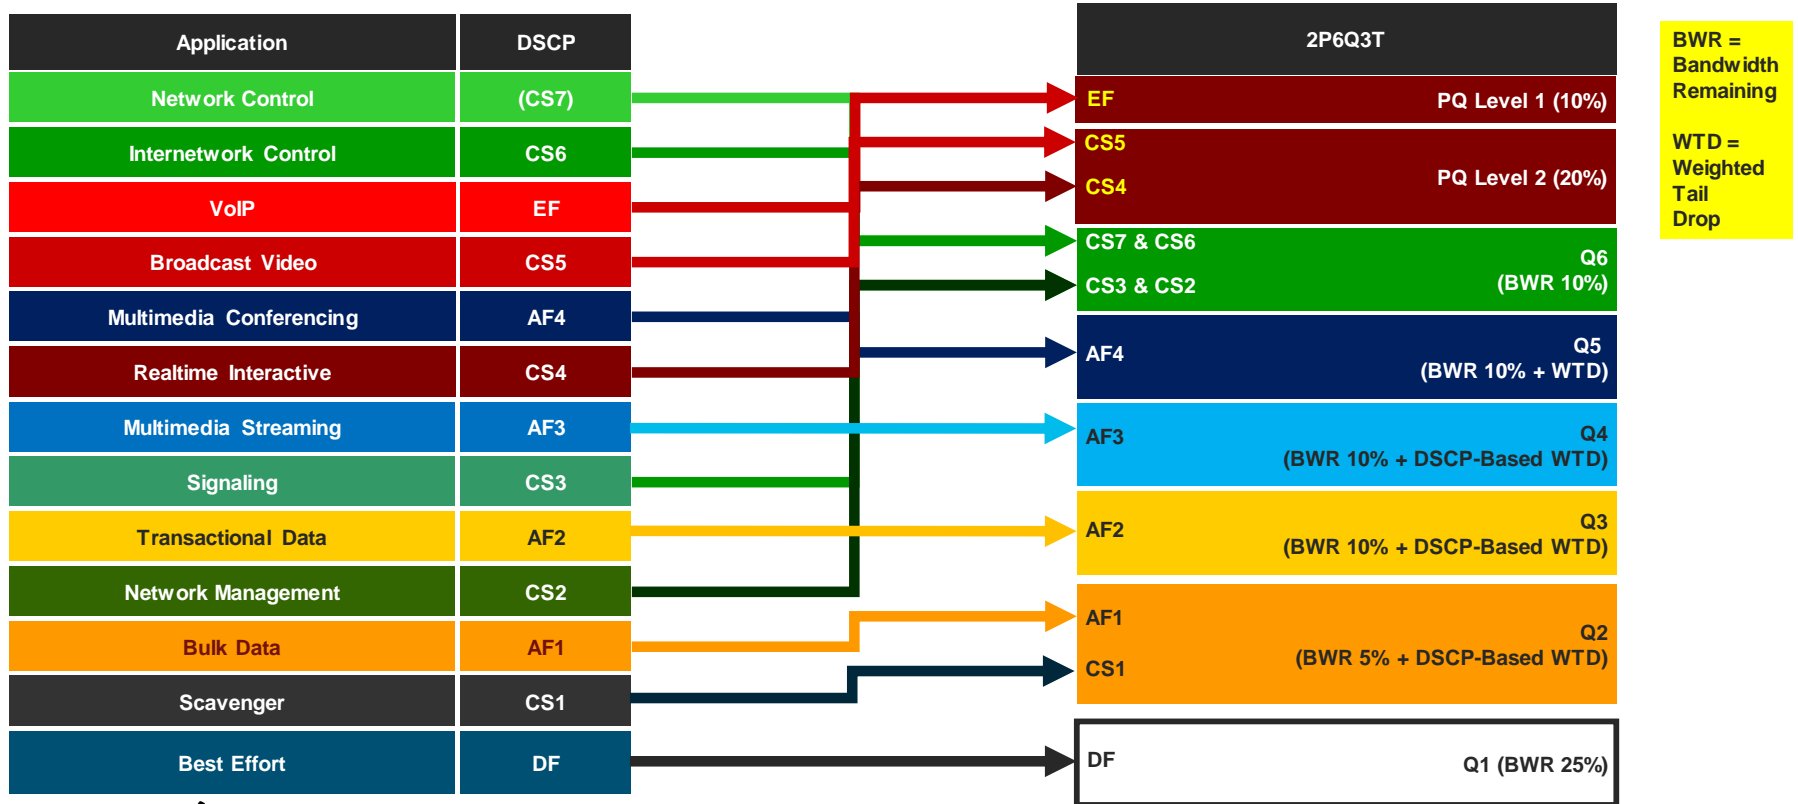
You make the power of data **possible**

Catalyst 3850/3650 - QoS Fundamental Actions

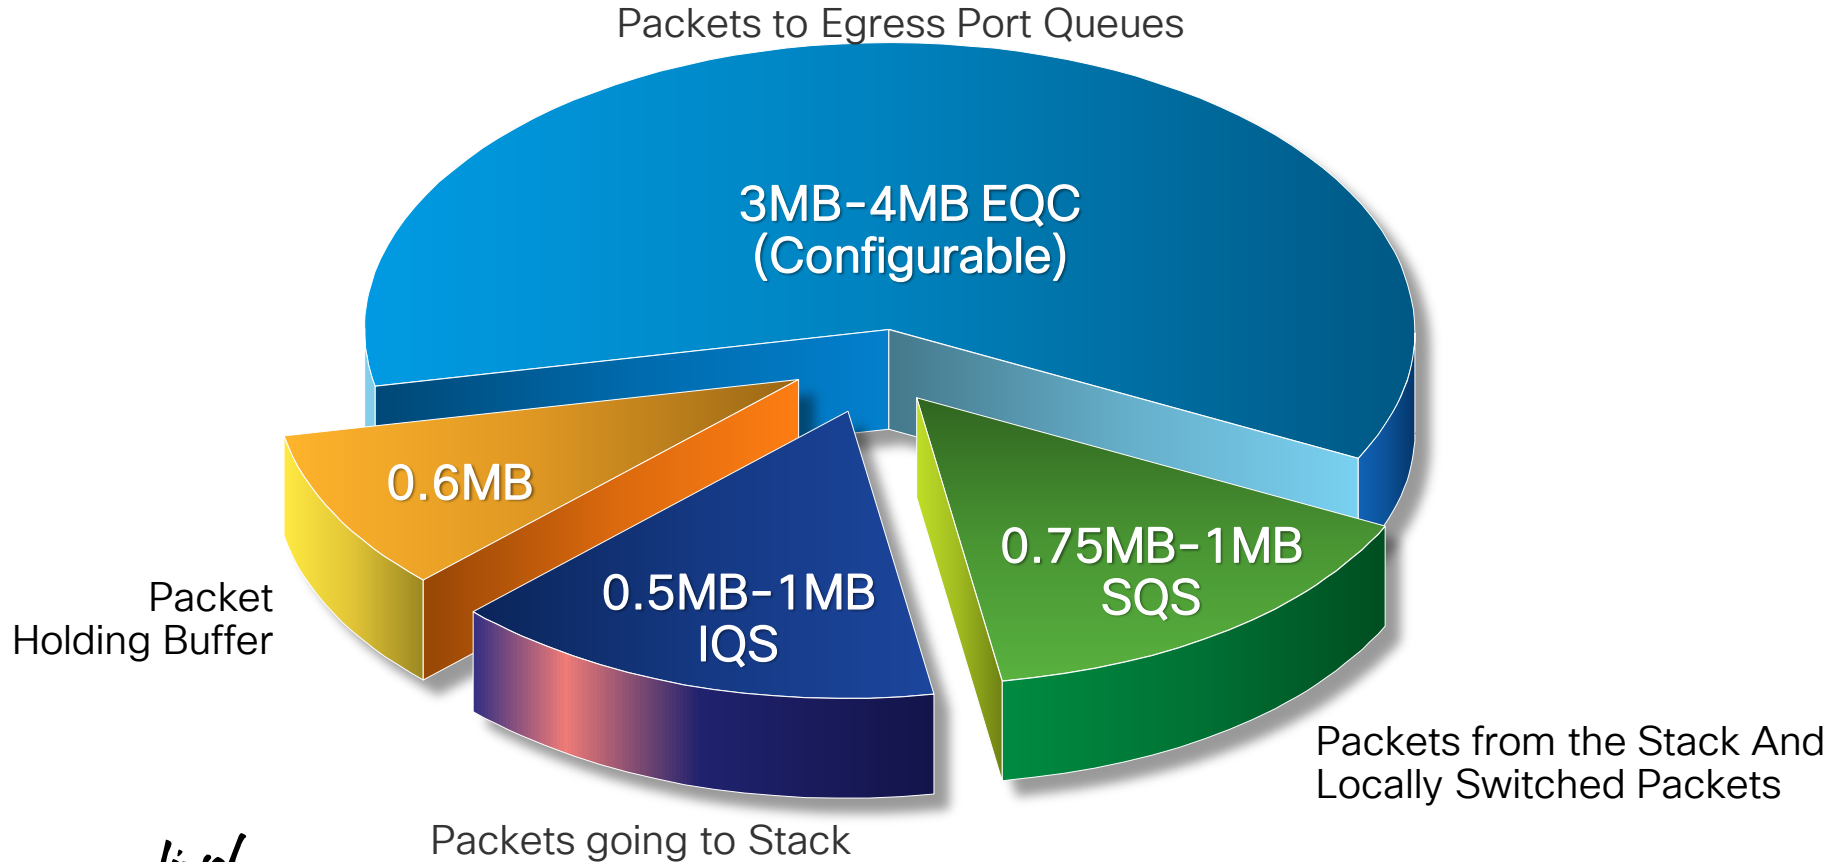


Catalyst 3650/3850 Campus QoS Design

Wired Port Egress Queuing (2P6Q3T with WTD) Model

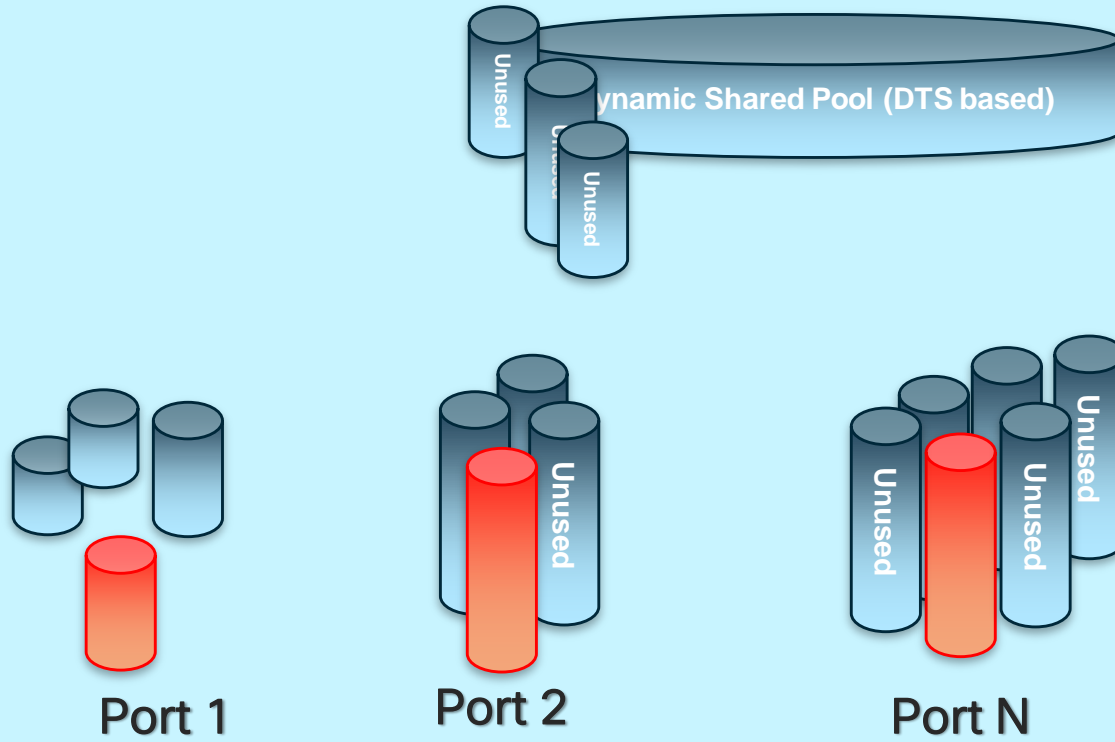


6MB Packet Buffer



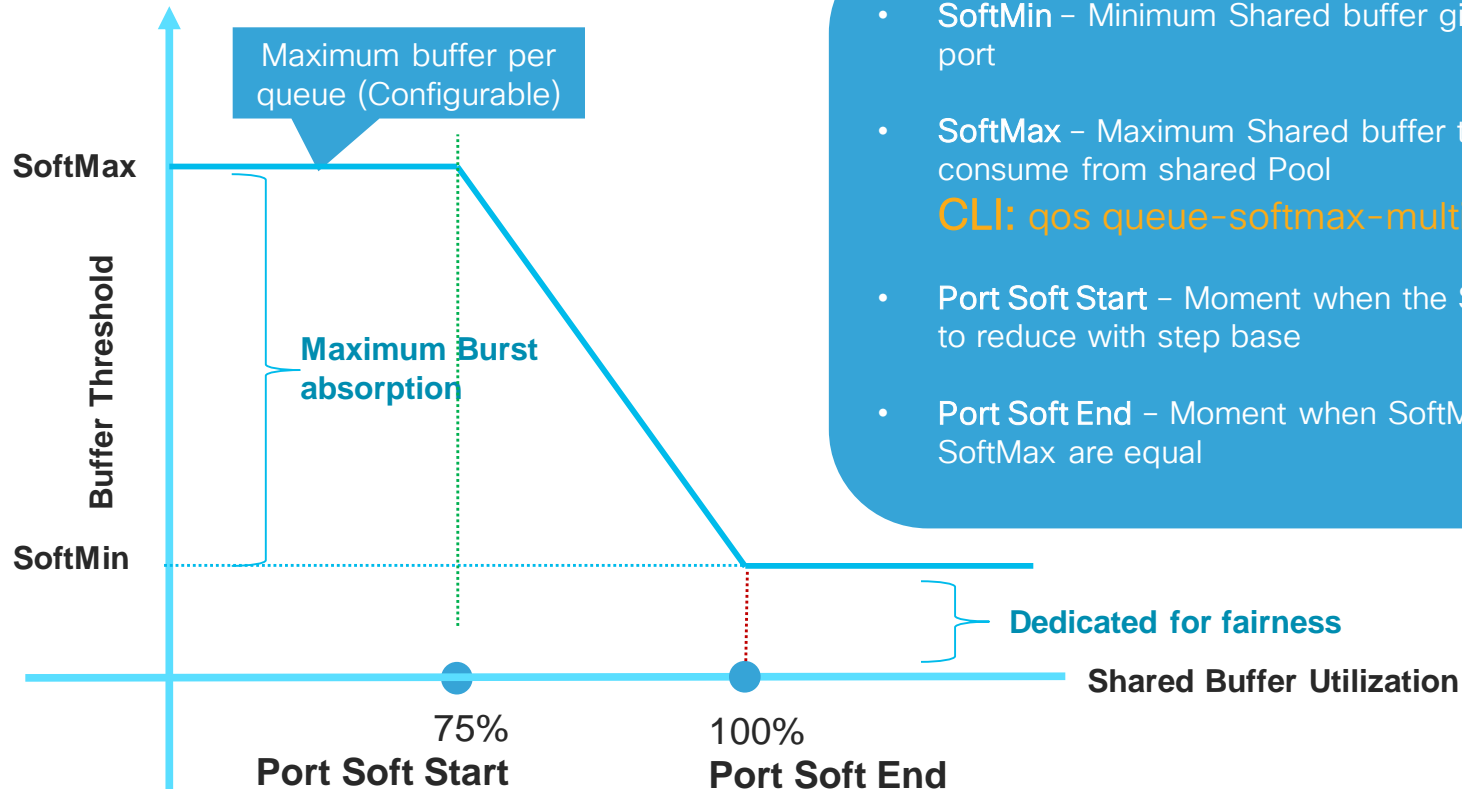
Dynamic Threshold Scale (DTS)

Switch



- Shared buffer is good for burst absorption.
- Dedicated buffer is good for predicted performance for each port.
- Buffer management is flexible: Dedicated plus shared.
- Configurable dedicated threshold per port/queue
- Configurable global maximum shared threshold
- Automatically adjusted depends on the available shared pool

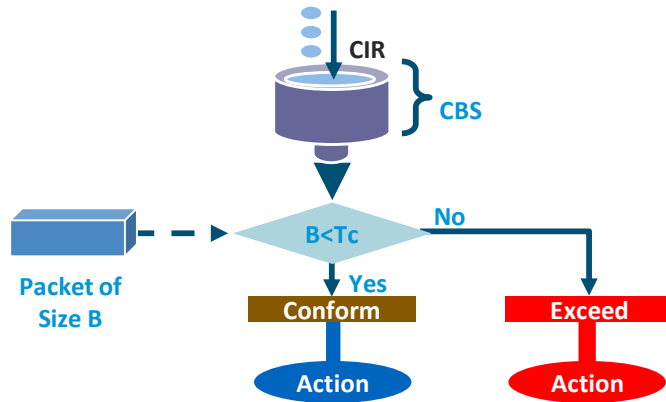
DTS - Dynamic Fair Buffer Sharing



- **SoftMin** - Minimum Shared buffer given to the port
- **SoftMax** - Maximum Shared buffer the port can consume from shared Pool
CLI: `qos queue-softmax-multiplier 1200`
- **Port Soft Start** - Moment when the Softmax start to reduce with step base
- **Port Soft End** - Moment when SoftMin and SoftMax are equal

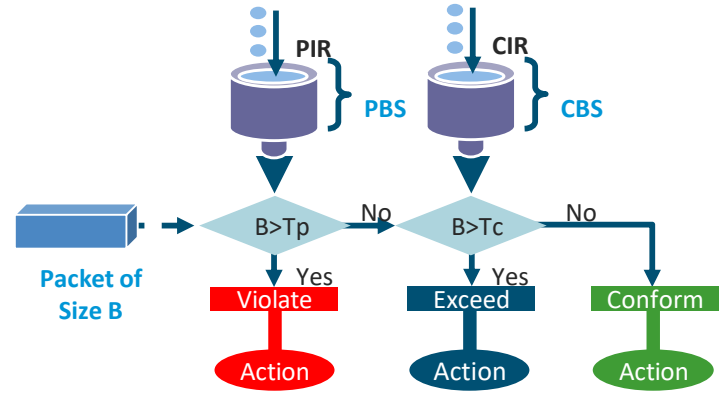
Policing

1 Rate 2 Color



```
police cir 100000000 bc 3125000  
conform-action set-dscp-transmit af41  
exceed-action drop
```

2 Rate 3 Color



```
police cir percent 10 pir percent 50  
conform-action transmit  
exceed-action set-dscp-transmit dscp table MARKDOWN  
violate-action drop
```

Catalyst 3850 – QoS Scale

QoS Scale Numbers	
Class-maps (Ingress)	1024
Class-maps (egress)	512
Table-maps (ingress)	16
Table-maps (egress)	16
Aggregate Policers	2000
Wired Queues/port	8 queues
Wireless Queues/port	4 queues
Buffer per ASIC Core	6 MB



Key Differences – Catalyst 3850/3650 vs 3750-X/E



Catalyst 3850

- MQC
- Trust by Default
- 8 Queues 1P7Q3T OR 2P6Q3T
- 6 MB Buffer per ASIC or ASIC Core
- HQoS – 2 Level Hierarchical Policy



Catalyst 3750-X

- MLS QoS
- Untrust by Default
- 4 Queues Per Port – 1P3Q3T
- 2 MB Buffer per ASIC
- No HQoS

Reference: <https://www.cisco.com/c/en/us/support/docs/switches/catalyst-3850-series-switches/118629-technote-qos-00.html>

Catalyst 3850 – SD-Access Scale

Fabric Constructs	Maximum Supported
Fabric Edge Relevant Scale	
Virtual Networks	64
Groups (SGTs)	4096 / 255 DGT
SGACLs (Security ACEs)	1350
Fabric Border Relevant Scale	
Virtual Networks	64
Groups (SGTs)	4096
SGACLs (Security ACEs)	1350
Fabric Control Plane Entries	4096
IPv4 Routes	8K
IPv4 Host Entries	16K

Need to Copy yourself to do more?



Phonlamai Photo/Shutterstock.com

Software Innovations



You make networking **possible**

Trustworthy Systems



You make multi-cloud **possible**

Cisco Catalyst 3850 / 3650 Trustworthy Systems

Design/
Develop

Plan/
Order

Source

Make

Quality

Delivery

Service/End
of Life (EOL)

Physical security practices + security technology innovations + logical security processes

Secure boot
Boot sequence
check

**Integrity
verification**
Malware protection

**Runtime
defenses**
64-bit ASLR




**PnP SUDI
support**
Two-way trust

Image signing
Authentic OS

**Hardware
authenticity**
Genuine hardware

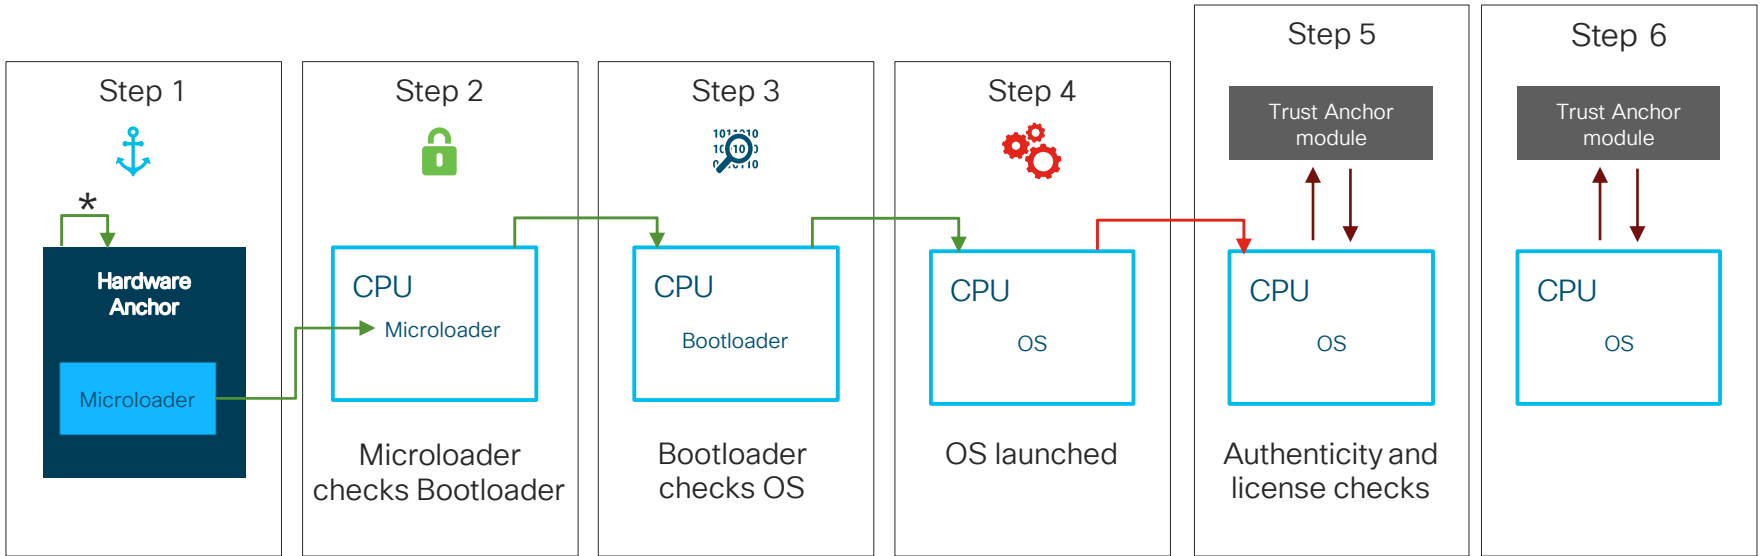
Cisco® trustworthy systems use industry best practices to help ensure full development lifecycle integrity and end-to-end security

REAL WORLD: Converted 3750-X 24 to 48 Ports




Cisco Secure Boot and Trust Anchor Module

Validating the Authenticity of Software Followed by Hardware



* The first instructions that run on a CPU are either stored in immutable hardware so that they cannot be tampered with or are validated by the hardware



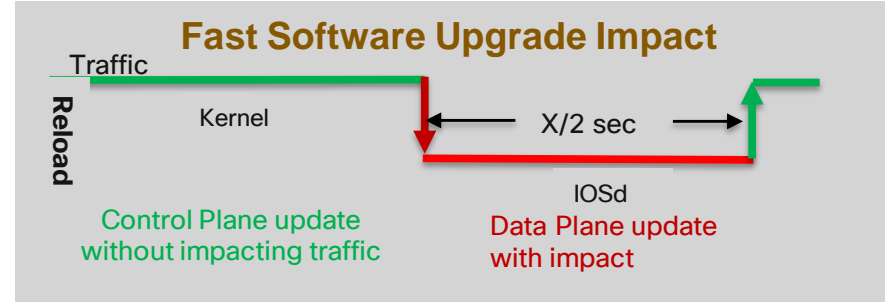
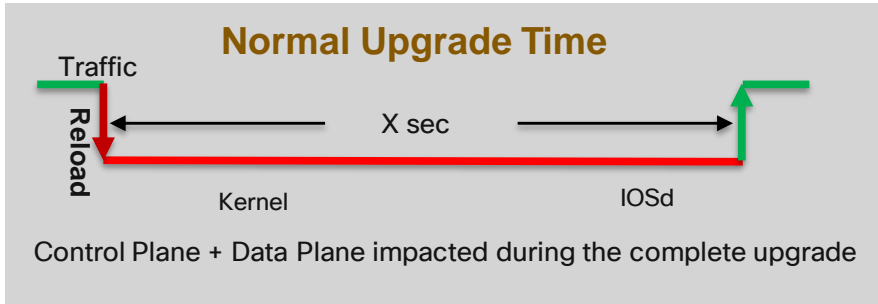
High Availability




You make customer experience **possible**

Fast Software Upgrade (16.8)

Regular Upgrade Vs Fast Software Upgrade Process




#Install add file image activate commit



Traffic is impacted throughout the upgrade cycle



#Install add file image activate Fastreload commit






Traffic impact is reduced by more than half by separating the control and data plane updates



Fast Software Upgrade

Supported and Unsupported Designs

Layer 2 Access Layer Designs– FSU Supported



Unsupported Designs



Software Maintenance Upgrades (SMU)

Software Upgrades are Challenging

Cost


- Expensive Upgrades - Business Loss
- Each device upgrade causes Network outage

Time

- Reduced IT staff slows software roll out
- Physical presence required

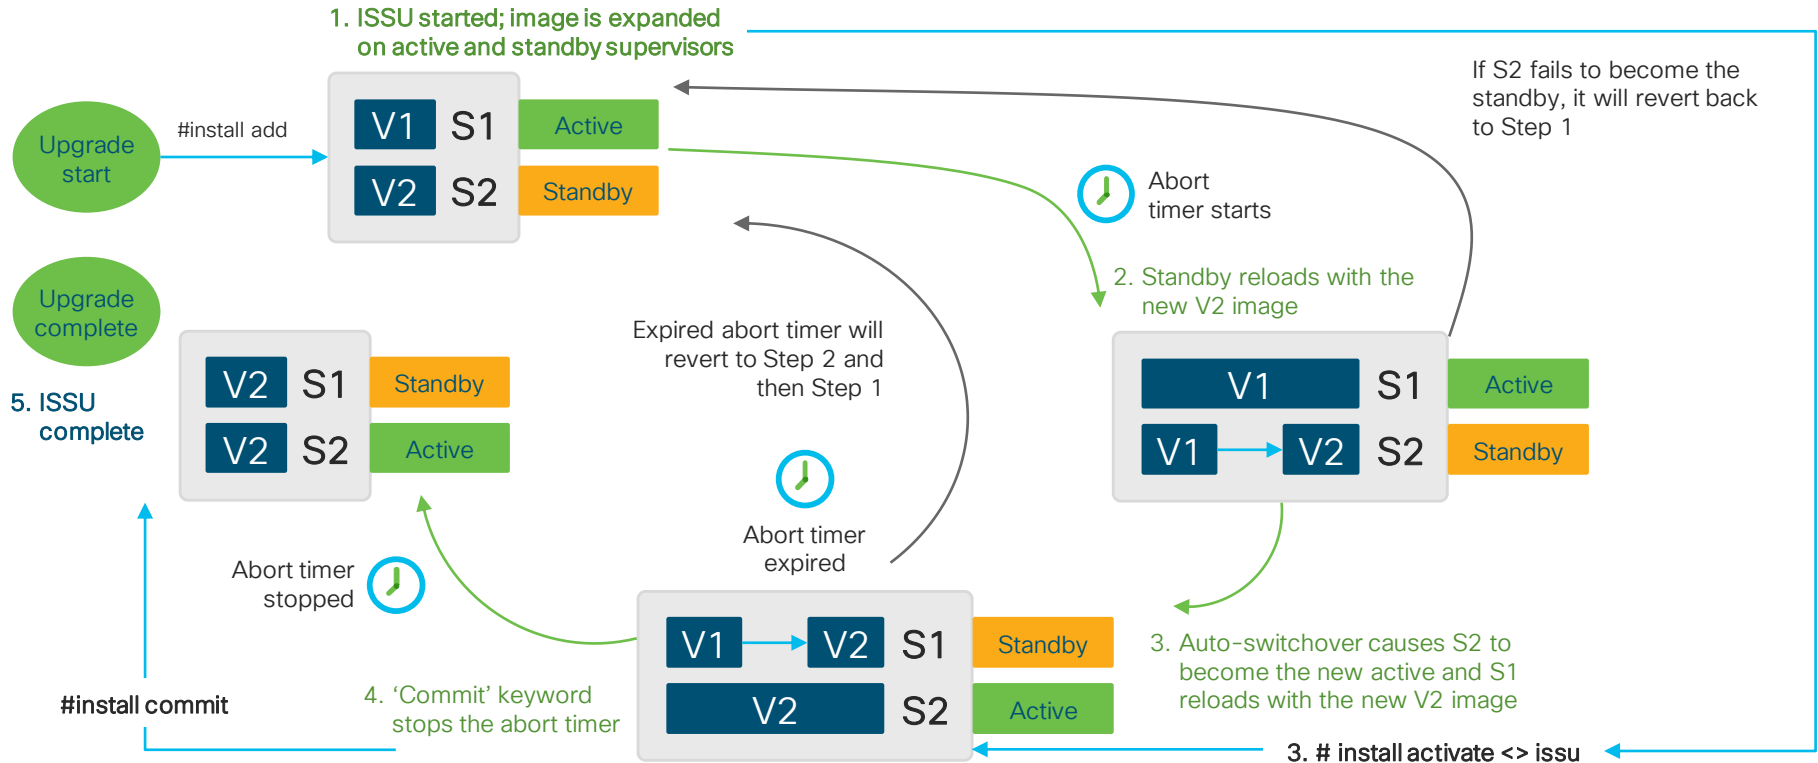
Scope

- New Code requires bug analysis, certification



Install of COLD SMU will require a system reload

Cisco Catalyst 3850 12/24/48XS ISSU workflow



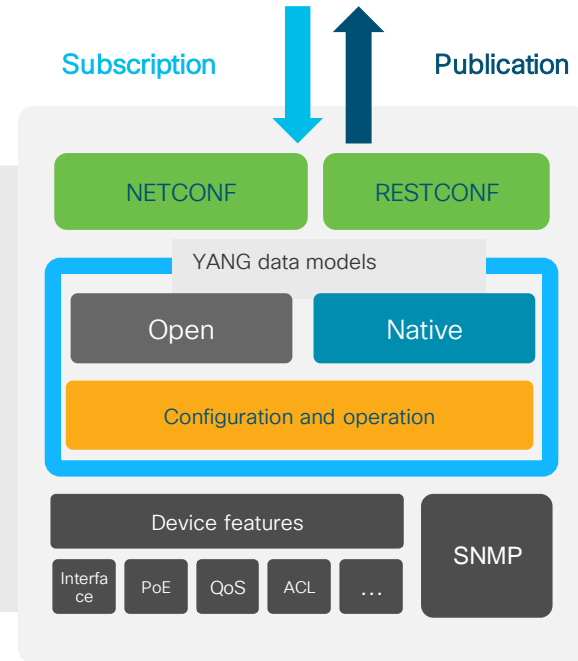
Programmability



You make networking **possible**

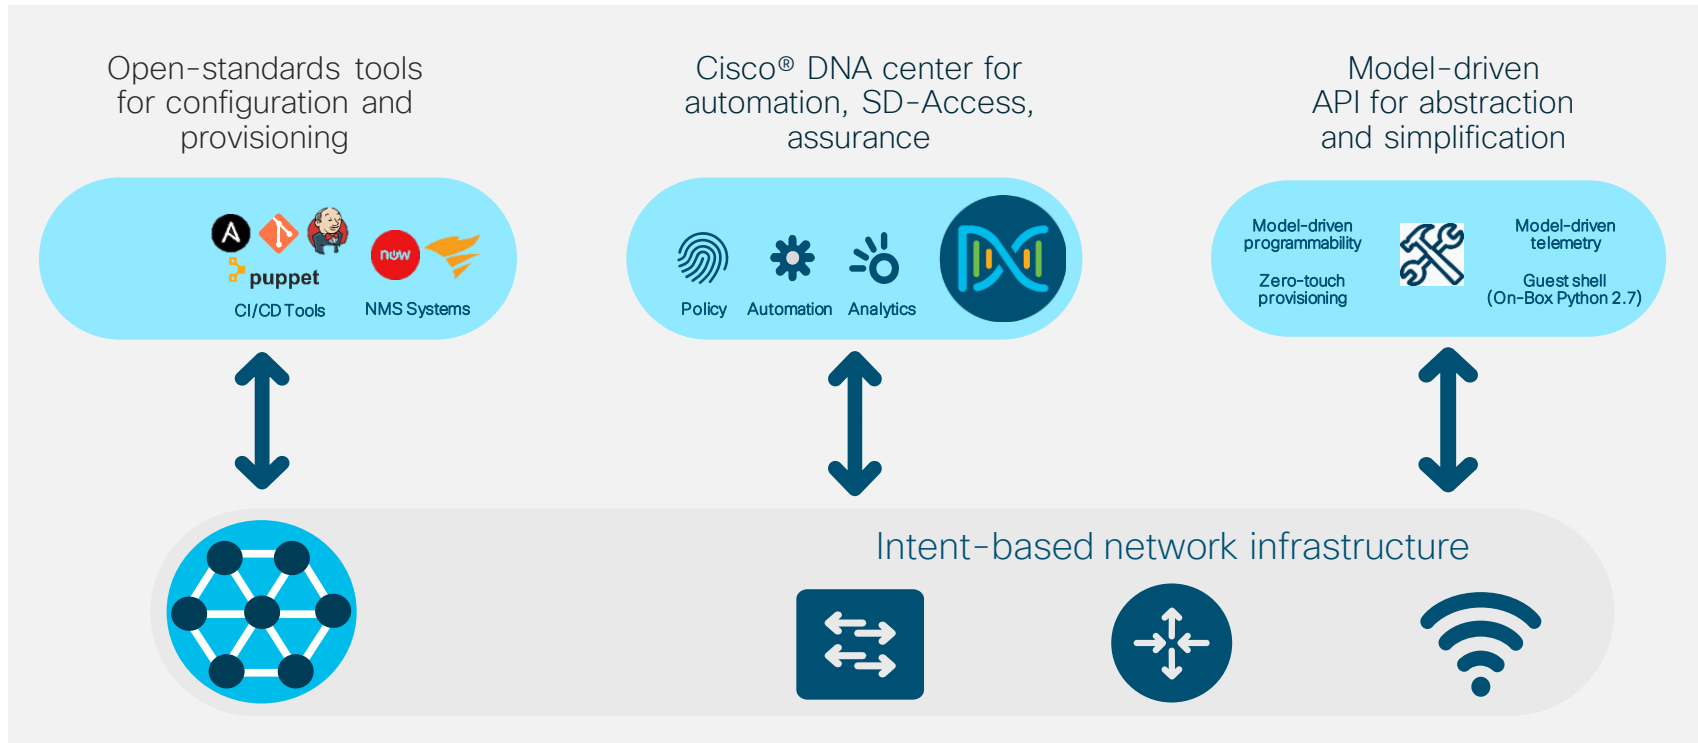
Model-driven telemetry

- Support for any YANG subtree
- Structured data
- XML/JSON encoding
- Periodic or on change
- Reduced CPU load



Export enriched, consistent, and concise data with context from devices for a better user and operator experience

Software Defined – Open, Programmable, Flexible Scaled operations



Network Subscription over NETCONF

A subscription is a contract between a subscription service and a subscriber that specifies the type of data to be published.




Instruction on:

- What data to collect (only Dial In option)?
- How often (≥ 1 sec)



Periodic subscription



Counters / Measures




Periodic subscription RPC

```
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <establish-subscription xmlns="urn:ietf:params:xml:ns:yang:ietf-event-notifications"
  xmlns:yp="urn:ietf:params:xml:ns:yang:ietf-yang-push">
    <stream>yp:yang-push</stream>
    <yp:xpath-filter>/mdt-oper:mdt-oper-data/mdt-subscriptions</yp:xpath-filter>
    <yp:period>1000</yp:period>
  </establish-subscription>
</rpc>
```

- Xpath-filter:
 - XML Xpath filter defining the data object to which you want to subscribe.
- Period:
 - The time period, in centiseconds (100th of a second), between push updates containing the subscribed information

On-change subscription



State / Configuration / Identifiers




On-change subscription RPC

```
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">  
  <establish-subscription xmlns="urn:ietf:params:xml:ns:yang:ietf-event-notifications"  
xmlns:yp="urn:ietf:params:xml:ns:yang:ietf-yang-push">  
    <stream>yp:yang-push</stream>  
    <yp:xpath-filter> /cdp-ios-xe-oper:cdp-neighbour-details/cdp-neighbour-detail</yp:xpath-filter>  
    <yp:dampening-period>0</yp:dampening-period>  
  </establish-subscription>  
</rpc>
```

CDP-neighbor with XPath
and Dampening period

Candidate Config Datastore

- Prior to Polaris 16.8, a NETCONF client would modify the Running config by directly sending an edit-config RPC to the Running in CDB, which will then be synced to the Running config in IOS

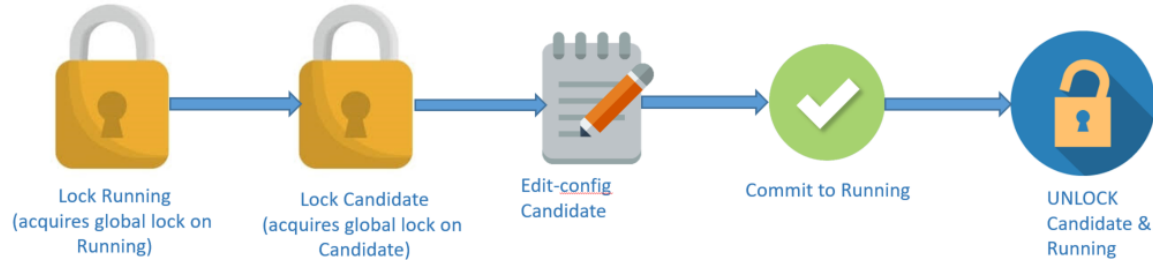


- Polaris 16.9 has introduced Candidate Data store functionality, which let's a user "stage" his/her configuration before pushing it to the device



How does it effect device configuration?

- Configuration comparison between device configuration through “Running” config versus “Candidate” config



Switch(config)#[netconf-yang](#) feature candidate-datastore

netconf-yang and/or restconf is transitioning from running to candidate

netconf-yang and/or restconf will now be restarted, and any sessions in progress will be terminated.

A

Spaghetti code




Image source: medium.com

Audio Video Bridging



You make security **possible**

Audio Video Bridging

Convergence of AV on Ethernet with Cisco Catalyst Switches



Cisco Catalyst family of switches are now Avnu-certified

The **Avnu Alliance** (www.avnu.org) is a community creating an interoperable ecosystem servicing the precise timing and low latency requirements of diverse applications using open standards through certification

Supported Platforms (standalone)



C9300



C9500



C3850



C3650

Select models

See www.cisco.com/go/avb

Open Standards

- IEEE 802.1BA: Audio Video Bridging
- IEEE 802.1AS: Generalized PTP
- IEEE 802.1Qat: Stream Reservation Protocol
- IEEE 802.1Qav: Forwarding and Queuing for Time-Sensitive Streams
- IEEE 1588v2: PTP Support**

Ecosystem Partners



Power over Ethernet



You make the power of data **possible**

Cisco POE Innovations

2-event classification

- Fast power negotiation without LLDP
- Physical layer negotiation < 1s
- **(config-if)#power inline port 2-event**

Perpetual POE

- Uninterrupted POE power during control plane reboot
- **(config-if)#power inline port poe-ha**

Fast POE

- Bypasses IOS control plane boot
- Restores power to PD within 30sec of power restore
- **(config-if)#power inline port poe-ha**

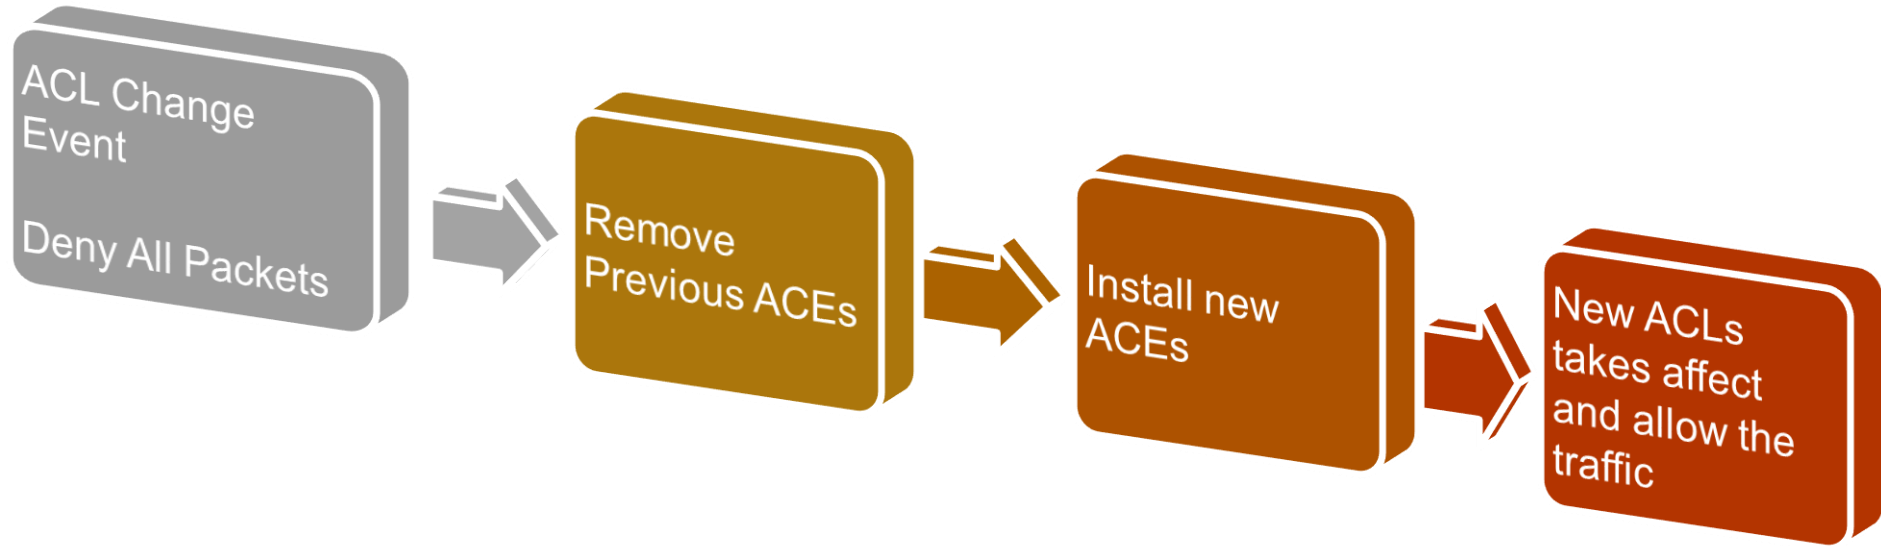
Delivering robust and resilient power infrastructure

ACL Hitless Update



You make multi-cloud **possible**

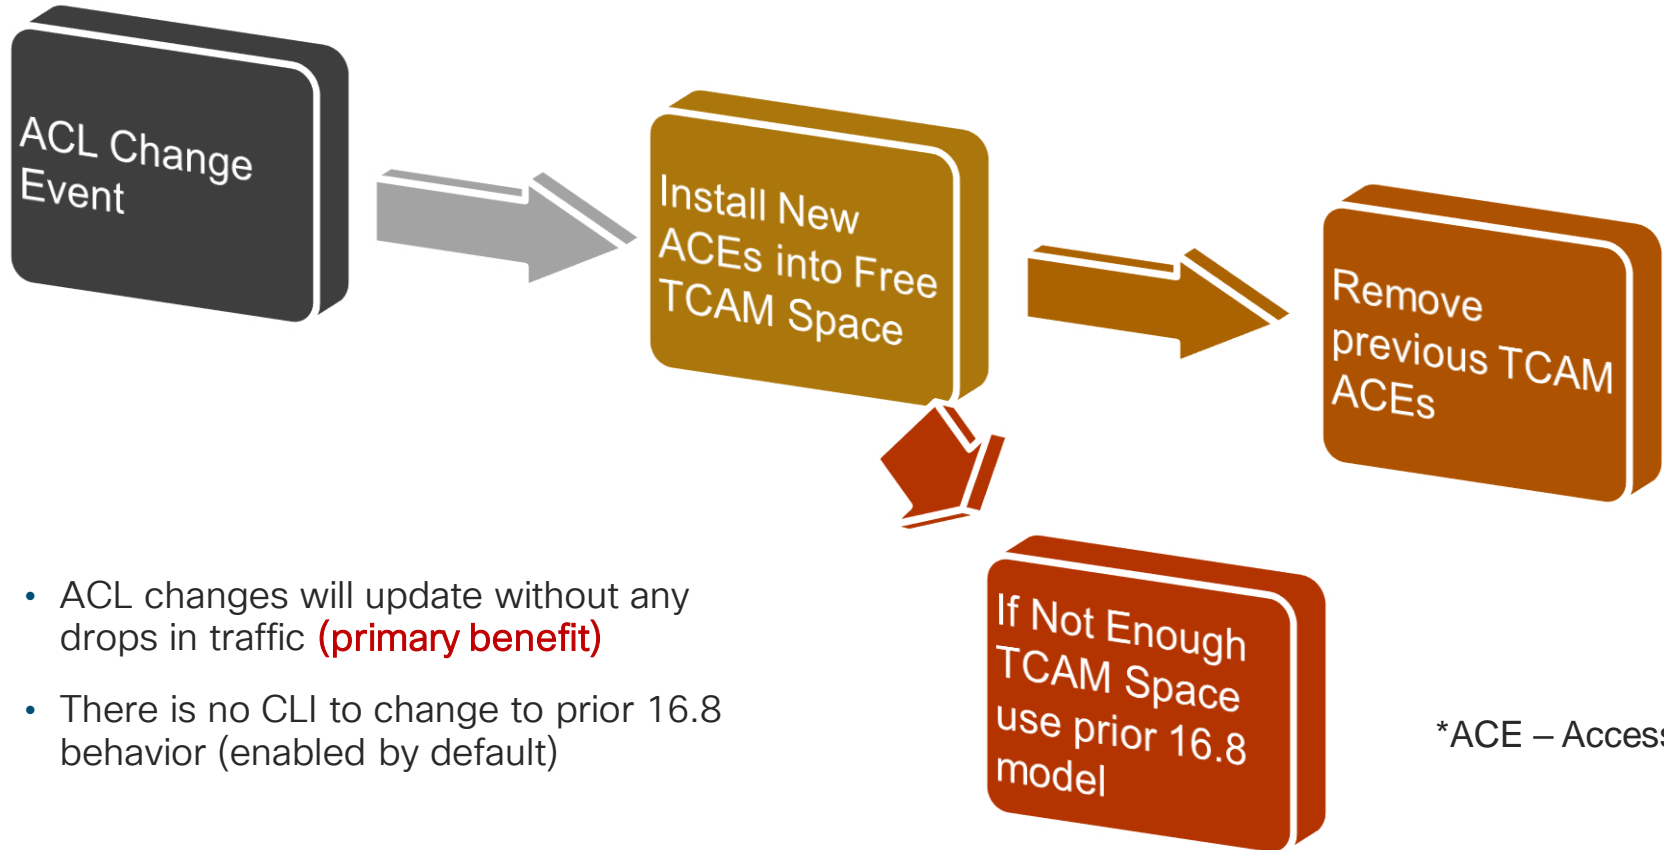
ACL Update Sequence



- In case of ACL change event, all traffic is dropped until the changes are applied
- In case of the new changes cannot fit the TCAM the traffic is also dropped until rollback is completed

*ACE – Access-list Entry

Hitless ACL Update – After 16.8 Sequence



- ACL changes will update without any drops in traffic (**primary benefit**)
- There is no CLI to change to prior 16.8 behavior (enabled by default)

*ACE – Access-list Entry

Application Visibility and Control

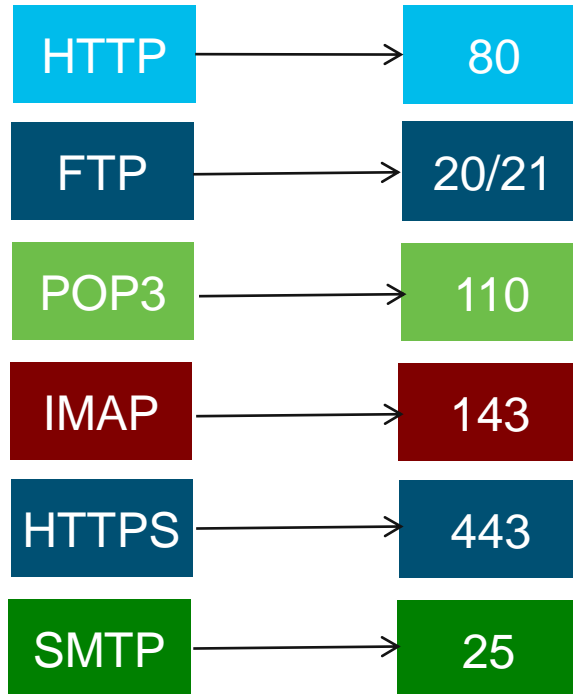


You make customer experience **possible**

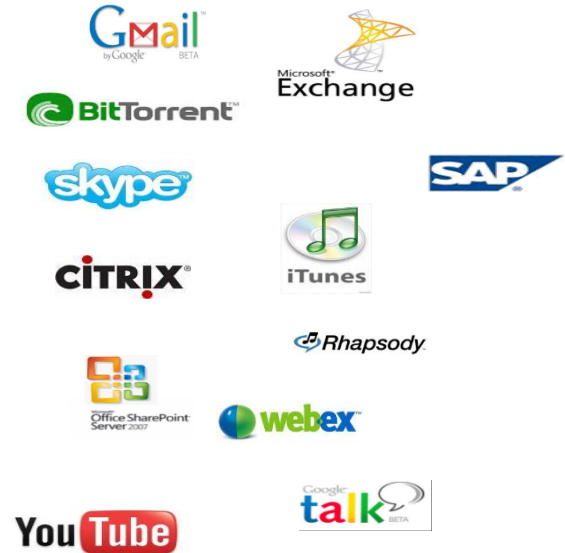
Challenges of Today's Network



Yesterday's Applications



Today's Applications



Enabling and Monitoring AVC – CLI

CLI

```
switch# show run int g1/0/23
```

Building configuration...

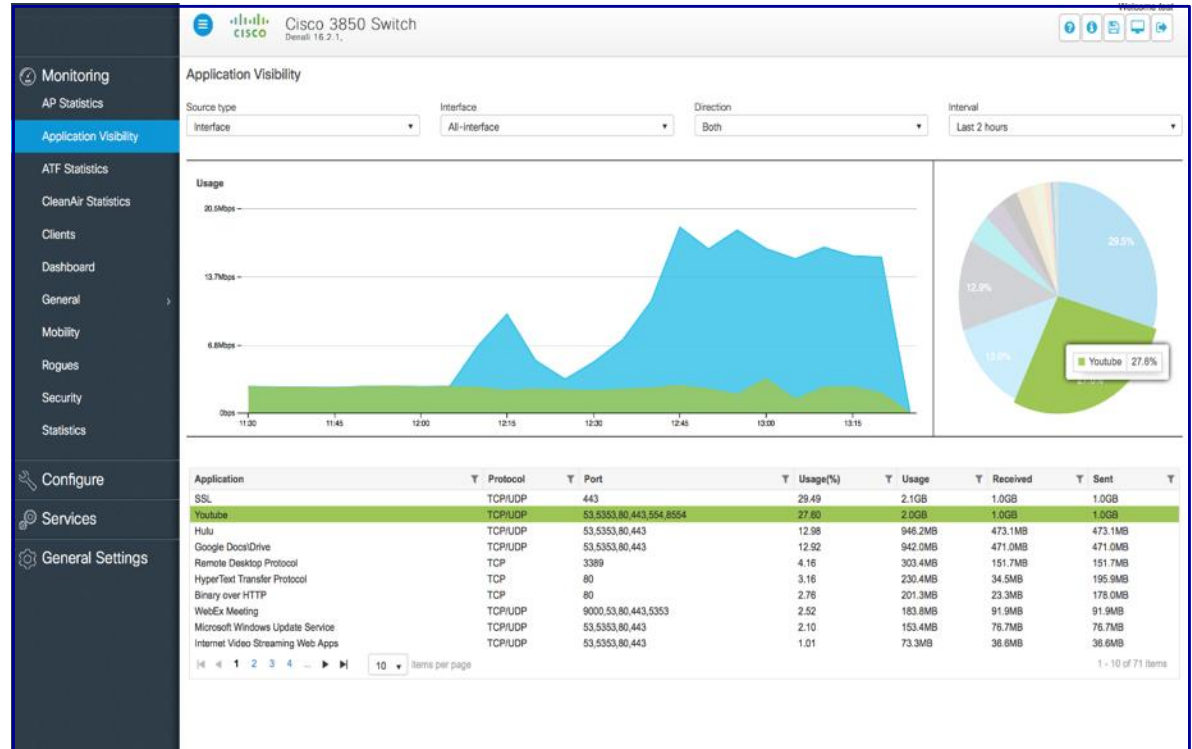
```
interface GigabitEthernet1/0/23
switchport access vlan 193
ip nbar protocol-discovery
end
```

```
switch# show ip nbar protocol-discovery top-n
```

```
GigabitEthernet1/0/23
      Input          Output
      ----          -
Protocol  Packet Count      Packet Count
          Byte Count      Byte Count
          5min Bit Rate (bps)  5min Bit Rate (bps)
          5min Max Bit Rate (bps)  5min Max Bit Rate (bps)
-----
youtube   356             187
          264713          25603
          0              0
          6000            3000
bing      2741            2384
          493258          423925
          0              0
          3000            3000
```

WebUI - Monitoring AVC

- Filter Monitoring Over Ingress/Egress interfaces and direction
- Identify Top Talkers
- Monitor Data over 2, 24 or 48 hours
- Monitor percentage Bandwidth usage



AVC with WEBUI Demo

Cisco WS-C3650-24PD 16.9.2

Welcome cisco

Dashboard Switch View

Search Menu Items

- Dashboard
- Monitoring
- Configuration
- Administration
- Troubleshooting

CPU & Memory Pressure Graph

Last Updated: 5/30/2019, 7:22:38 PM

Slot: 1-RP0

CPU Utilization

CPU: 0

Process	CPU (%)
User	6.90
System	1.10
Idle	91.90

[Advanced CPU View](#)

CPU (%) vs Device Time

Time	User (%)	System (%)	Idle (%)
02:27	~15	~5	~80
02:27	~10	~5	~85
02:28	~10	~5	~85

Memory Utilization

Memory Details	Size (KB)
Total	4008352
Used	2504636
Free	1503716
Committed	3456592

[Advanced Memory View](#)

Memory Used (%) vs Device Time

Time	Memory Used (%)
02:27	~62
02:27	~62
02:28	~62

Temperature

Last Updated: 5/30/2019, 7:21:40 PM

System Temperature : 27°C

System Information

Last Updated: 5/30/2019, 7:21:48 PM

- Hostname:** C3650
- Device Uptime:** 1 hour, 24 minutes
- System Time:** 02:27:34.912 UTC Fri May 31 2019
- Device Type:** WS-C3650-24PD
- Boot Image:** flash:cat3k_caa-universalk9.16.09.02.SPA.bin
- FIPS Mode:** Disabled
- Last Reload Reason:** Reload Command

PoE Power Consumption

Last Updated: 5/30/2019, 7:21:44 PM

Total Power Supported : 390 W


- Unused
- PoE
- UPoE

MPLS

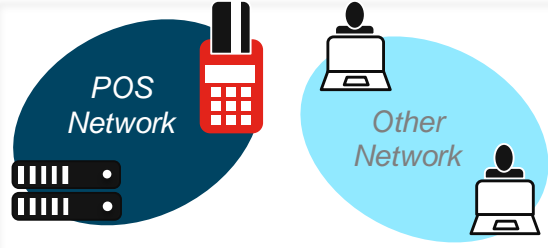


You make the power of data **possible**


MPLS Enables Network Segmentation in Campus




Line of business – BU segmentation



Payment Card Industry




Hospital Network



Bring-Your-Own-Device (BYOD)



Mergers and Acquisitions



Multi-Tenancy

Conclusion



You make networking **possible**

Do you feel you know the Car in and out Now?





Catalyst 3850/3650 is built on Robust Architecture



UADP



IOS-XE 16.x



The Combination of UADP and IOS-XE 16.x Makes your Network Ready and Future proofed



2013



2015



2017

Future proofed for the technologies of tomorrow

Catalyst 3850 / 3650 Related Sessions



For Your
Reference

BRKARC-3146 – Troubleshooting Cisco Catalyst 3650 / 3850 Series Switches

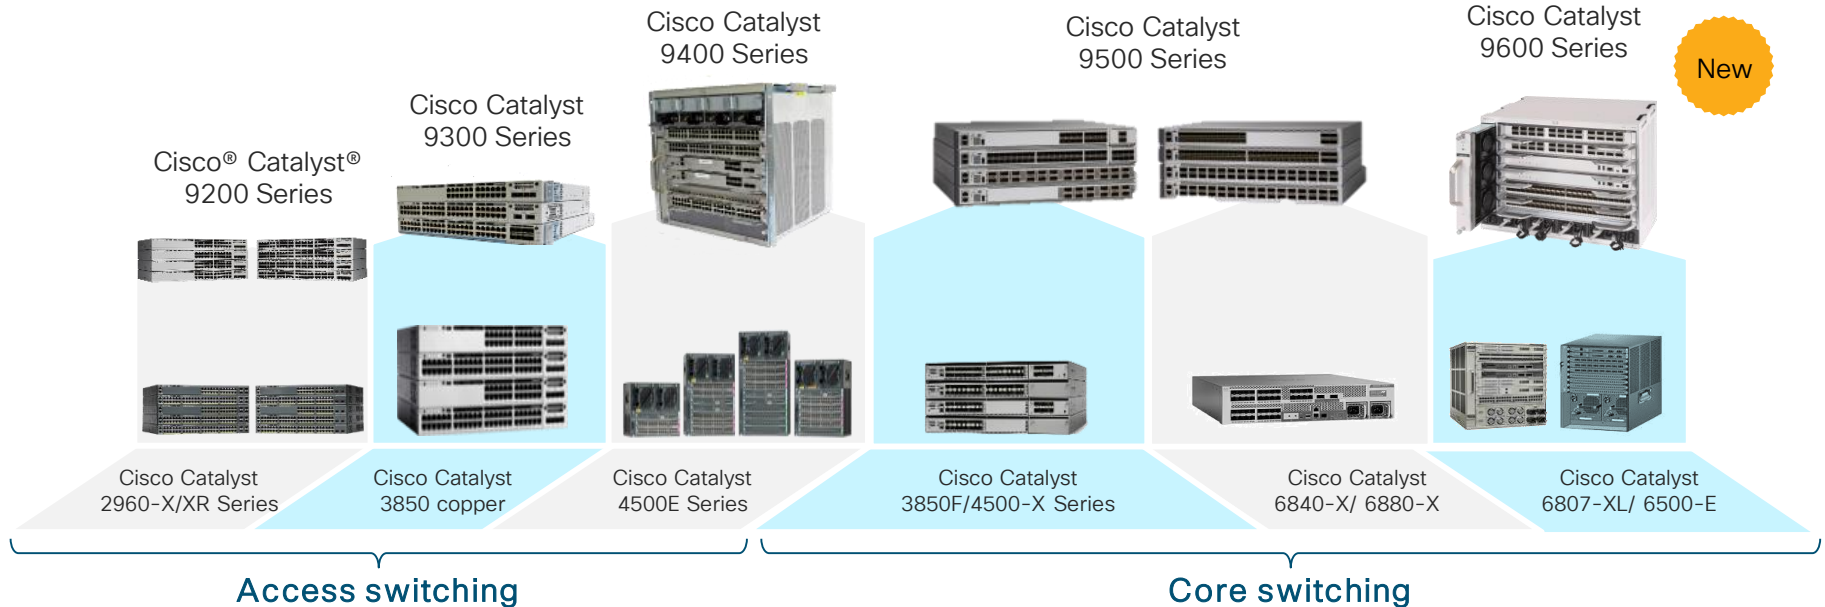
BRKARC-3190 – Troubleshooting Cisco Catalyst 3650, 3850 and 9000 Series Switches

CTHARC-2025 – Cat9300/C3850 Stacking – Tips and tricks

BRKCRS-2901 – Cisco Silicon – The Importance of Hardware in a Software-Defined World

Cisco Catalyst 9000 family switching transitions

Greater flexibility from the branch to business-critical, mission-critical campus core



Catalyst 9000 Switching Related Sessions



For Your
Reference

BRKARC-2035 – The Catalyst 9000 Switch Family - An Architectural View

BRKARC-3863 – Catalyst 9300 Switching Architecture

BRKARC-3873 – Cisco Catalyst 9400 Switch Architecture

BRKARC-2007 – Cisco Catalyst 9500 Switch Architectures

BRKARC-3873 – Cisco Catalyst 9600 Architecture

Continue your education



Demos in the
Cisco campus



Walk-in labs



Meet the engineer
1:1 meetings



Related sessions

NDA Roadmap Sessions at Cisco Live

Customer Connection Member Exclusive

Join Cisco's online user group to ...



Connect online with 29,000 peer and Cisco experts in private community forums



Learn from experts and stay informed about product roadmaps

- Roadmap sessions at Cisco Live
- Monthly NDA briefings



Give feedback to Cisco product teams

- Product enhancement ideas
- Early adopter trials
- User experience insights

Join online: www.cisco.com/go/ccp

CiscoLive!

NETWORKING ROADMAPS	SESSION ID	DAY / TIME
Roadmap: SD-WAN and Routing	CCP-1200	Mon 8:30 - 10:00
Roadmap: Machine Learning and Artificial Intelligence	CCP-1201	Tues 3:30 - 5:00
Roadmap: Wireless and Mobility	CCP-1202	Thurs 10:30 - 12:00

Join at the Customer Connection Booth
(in the Cisco Showcase)

Member Perks at Cisco Live

- Attend NDA Roadmap Sessions
- Customer Connection Jacket
- Member Lounge



Thank you





You make **possible**