Aggregation Services Router (ASR) 1000 Family

Aidan Marks

CCIE, CISSP

aidan@cisco.com

Well, we are ready now. Why is this a big deal?

- Innovation is necessary - you can't continue to play this game with generic off-the-shelf processors.
- It took us 5 years and a decent chunk of R&D.
- Cisco's first generation programmable silicon chip is here.
- ASR1000 is the first platform to leverage this silicon.
- Price/performance just shifted gear
- Very large scalability possible without necessarily high cost, heat, power and space

ASR1000 Topics Covered

- Overview and Positioning
- Places in the Network
- Hardware & System Architecture
- Performance
- Broadband on the ASR
- Software Architecture, HA and Packaging
- Summary and Roadmap

ASR1000 Overview and Positioning

ASR1000 Overview

- Next-generation of Midrange router family
  - 2RU / 4RU / 6RU chassis
  - 5 / 10 / 20 / 40 + Gbps forwarding
  - “Next-Gen” 7200 supporting same feature set at different price performance points
- ASR1000 Differentiators
  - QFP – QuantumFlow Processor
  - Carrier-class design - H/W and S/W HA
  - On-chip services (SBC, FPM, Security..) and QoS
  - Feature velocity
  - Built on Linux! :-)

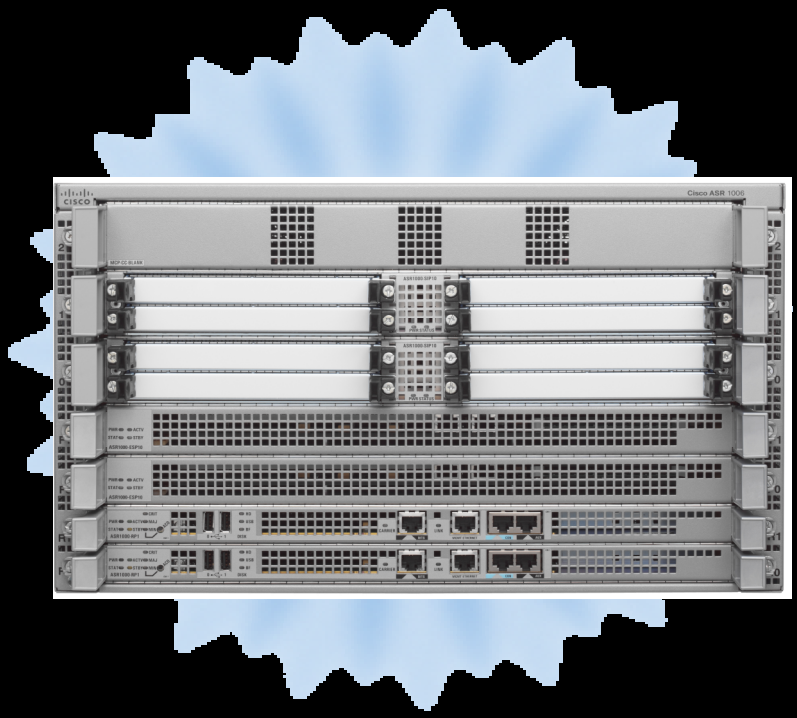

ASR1000 – Product Positioning

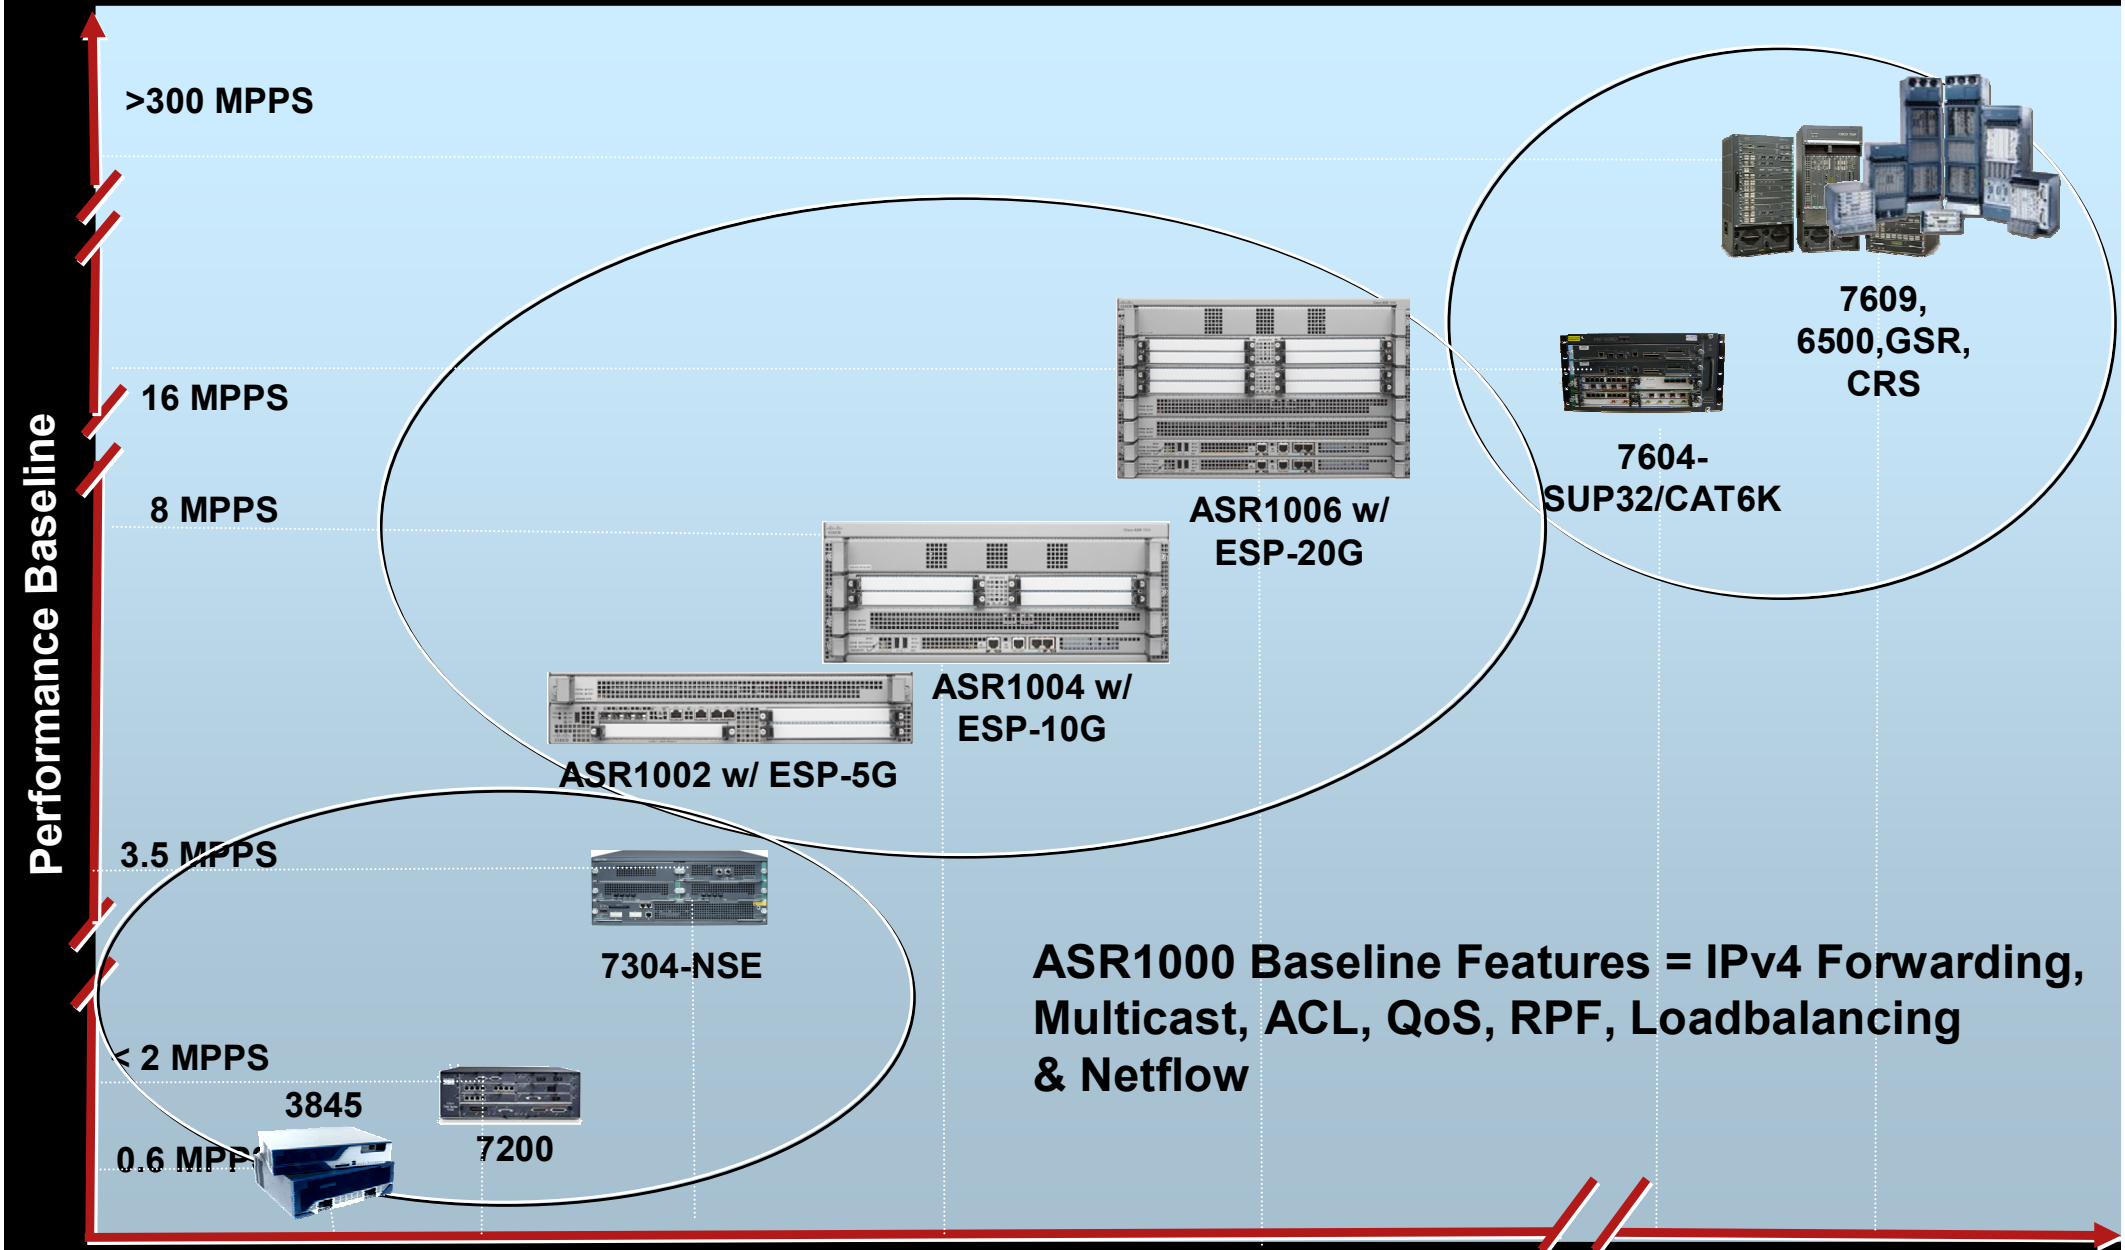

Product Positioning - Broadband

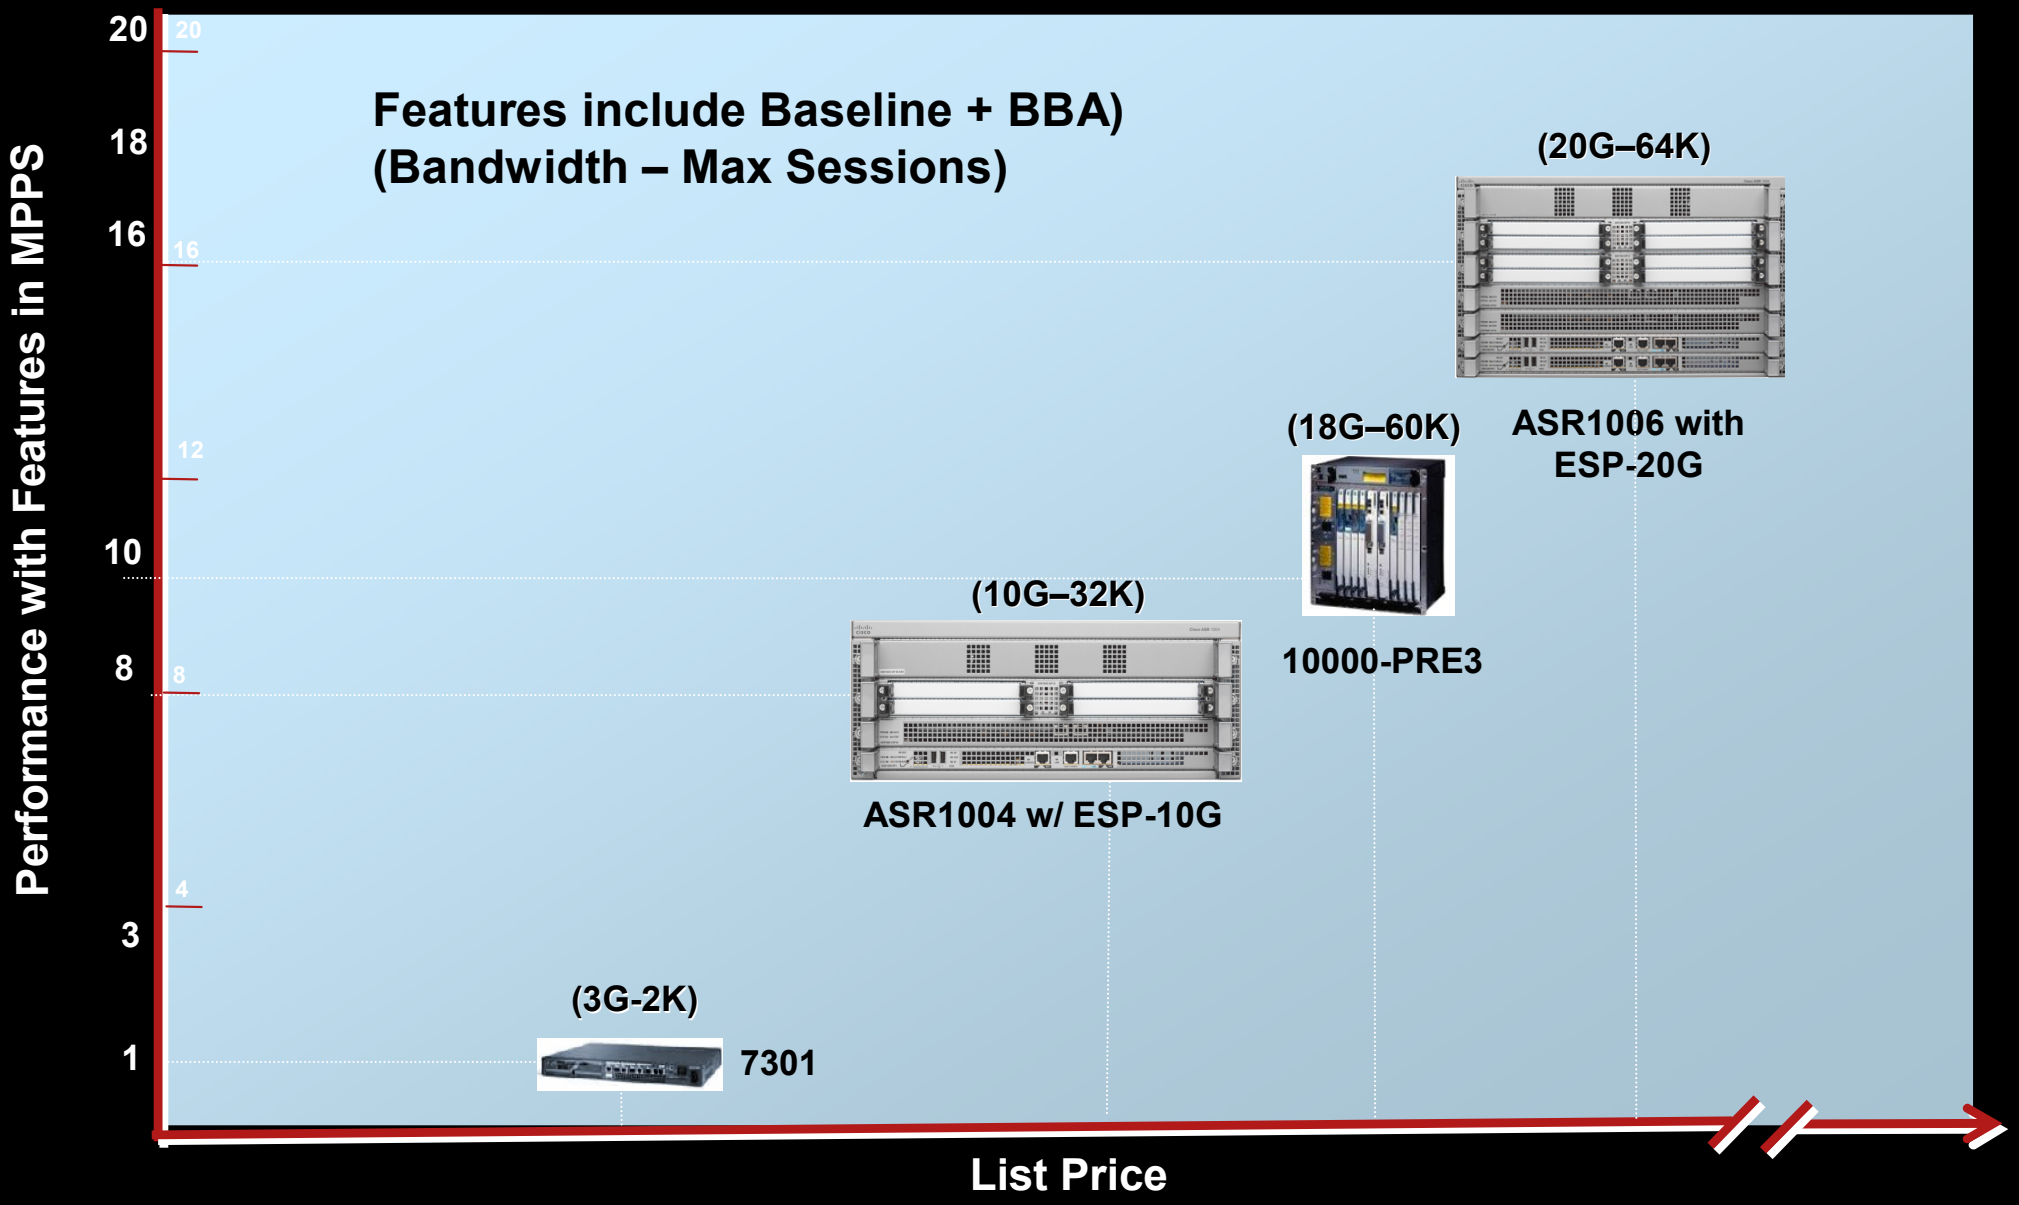

ASR1000 Product Family

SPA Slots	2 RU ASR1002	4 RU ASR1004	6 RU ASR1006
	3-slot	8-slot	12-slot
# of ESP Slots	1	1	2
# of RP Slots	Integrated (RP1)	1	2
# of SIP Slots	Integrated (SIP-10)	2	3
IOS Redundancy	Software Option	Software Option	Hardware
Built in GigE	4 x GE SFP	N/A	N/A
Height	3.5” (2RU)	7” (4RU)	10.5” (6RU)
Bandwidth	5-10 Gbps	10-40+ Gbps	10-40+ Gbps
Performance	4(5G)-8(10G) Mpps	8(10G) - 16+ Mpps	8(10G) - 16+ Mpps
Air Flow	Front to Back	Front to Back	Front to Back
Power Supply (Watts)	470	765	1275

Aggregated Services & Scale

ASR Places In The Network

ASR 1000 in Service Provider IP Next Generation Network

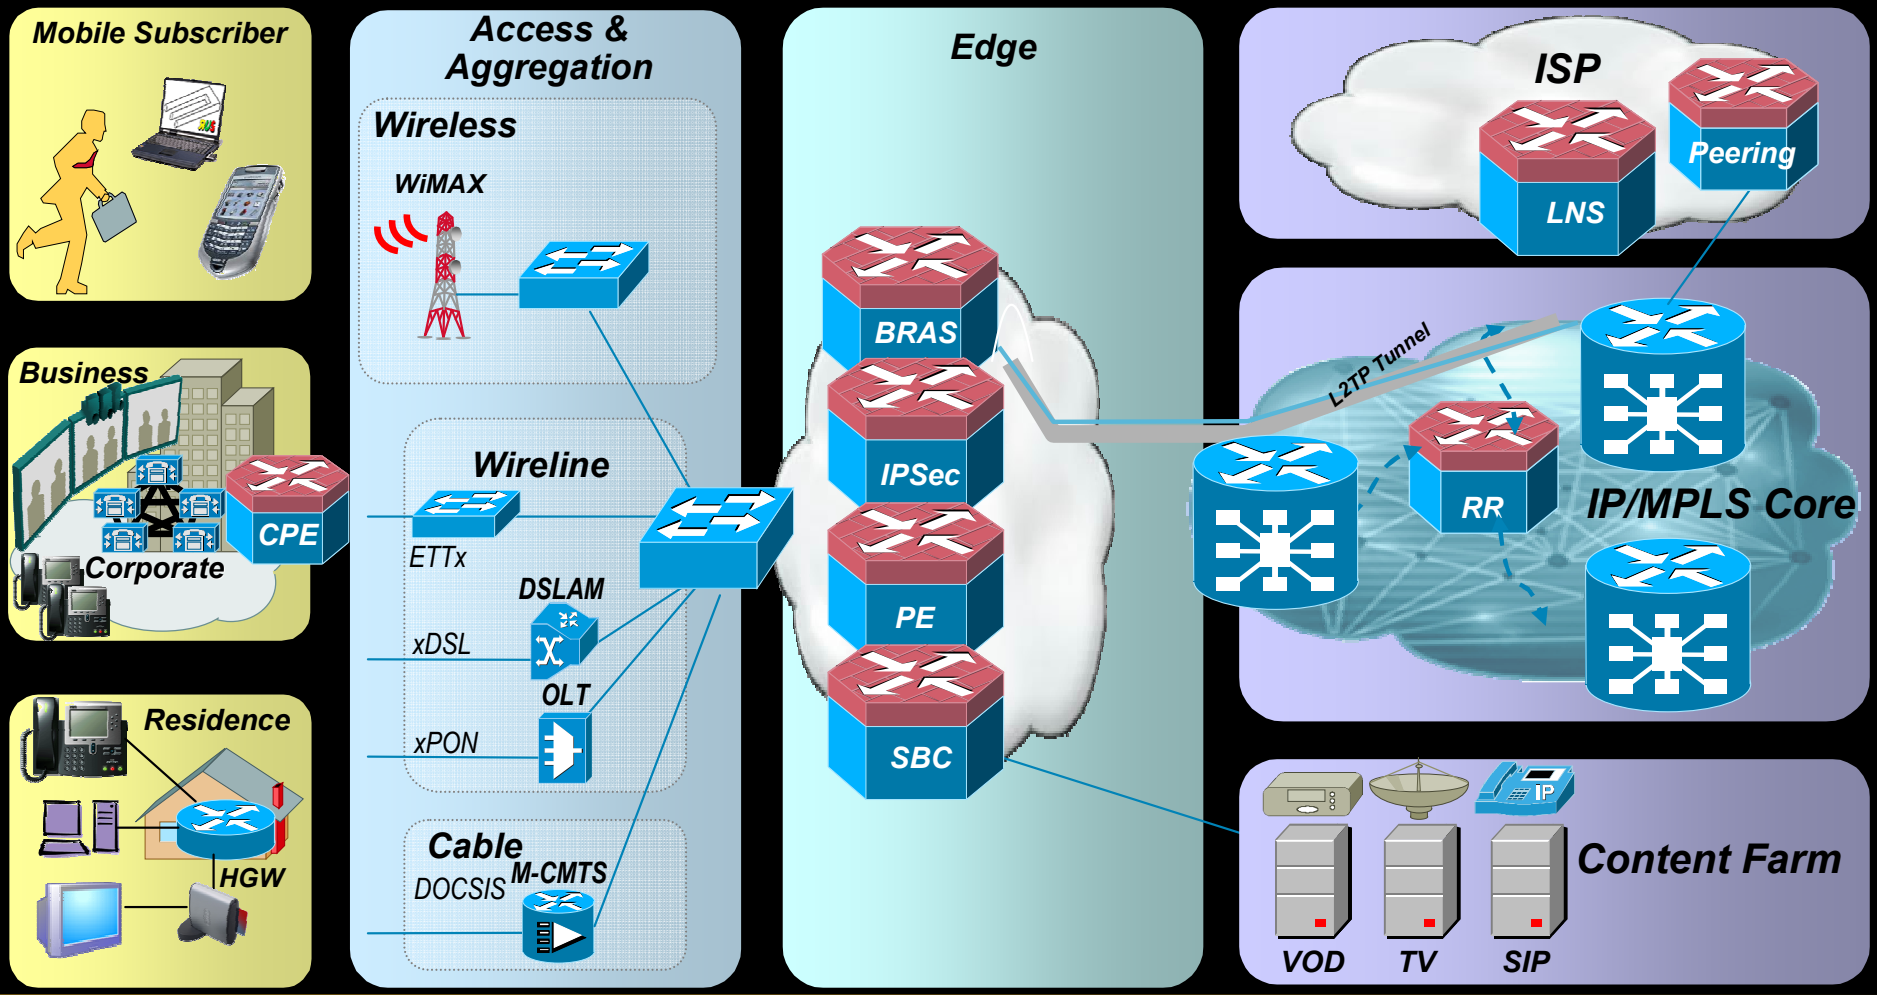

- High Speed CPE

- BRAS-PPPoE
- LAC, PTA, ISG
- IPSec Aggregator
- VoIP SBC
- PE (L3VPN PE)

- LNS
- Route Reflector
- Internet Peering

ASR 1000 as Distributed Integrated Service Edge in IP NGN

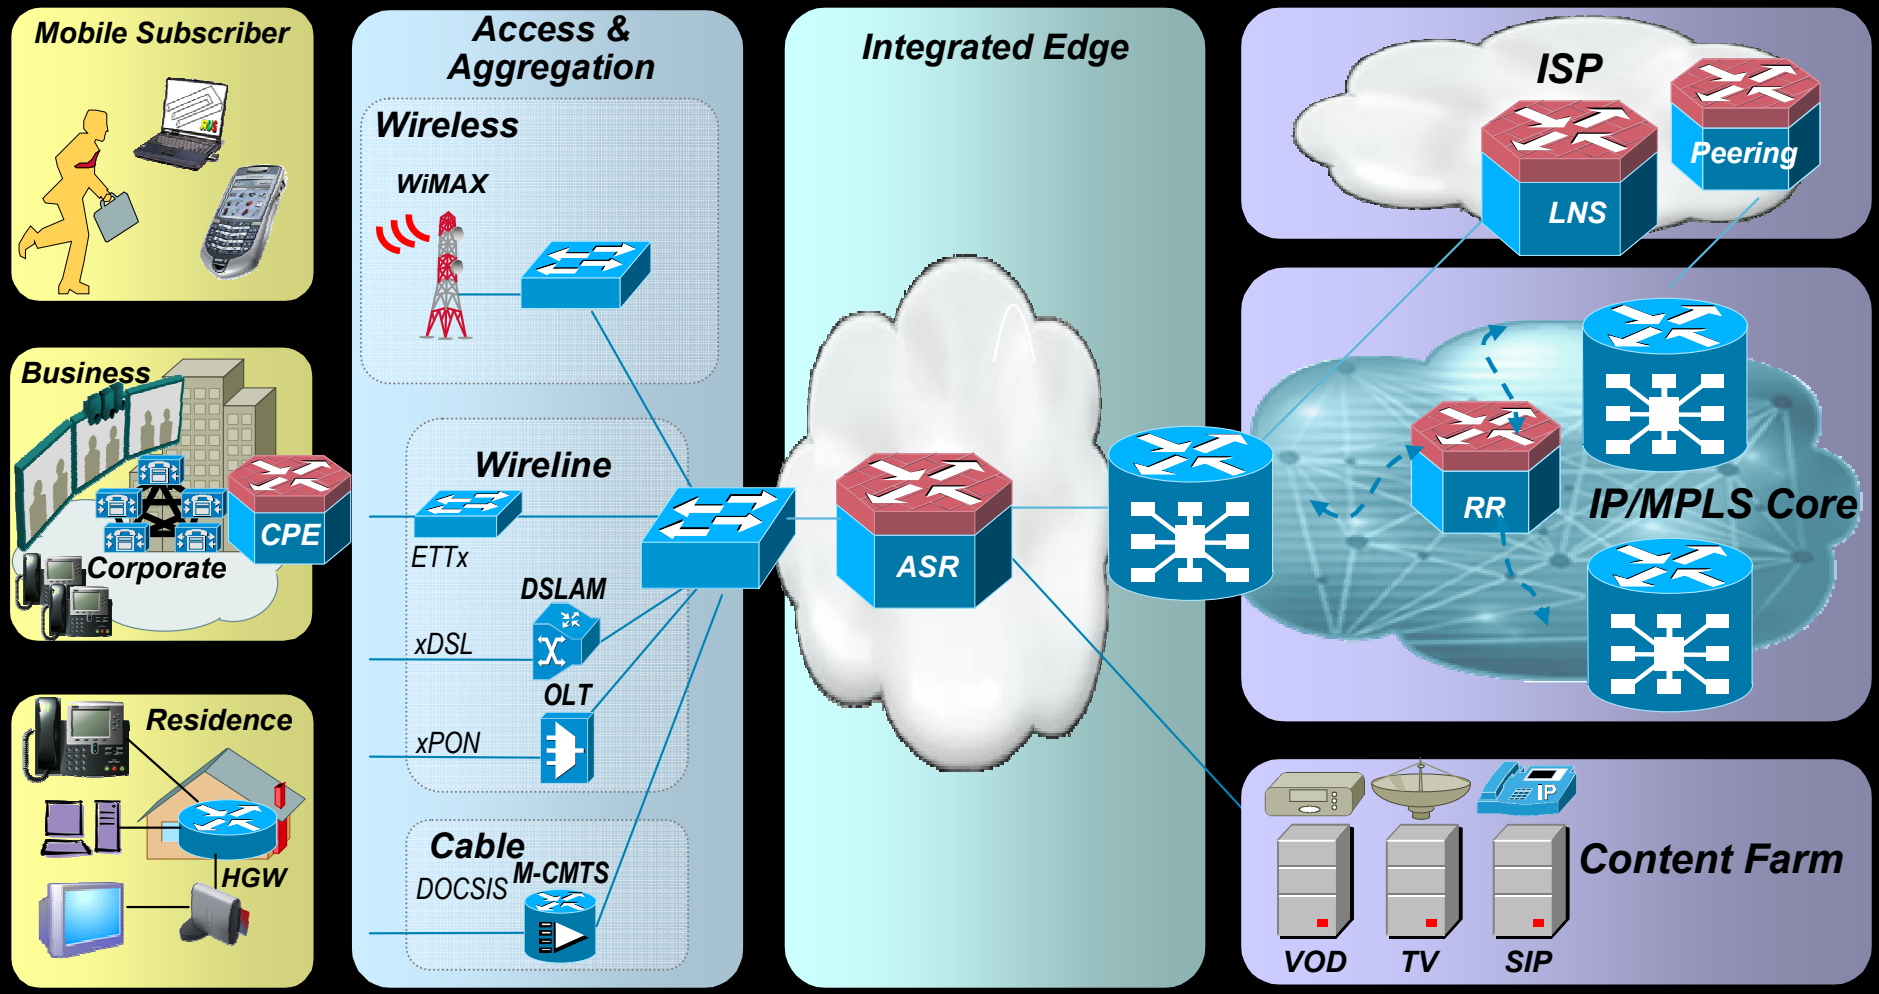

ASR1000 Hardware & System Architecture

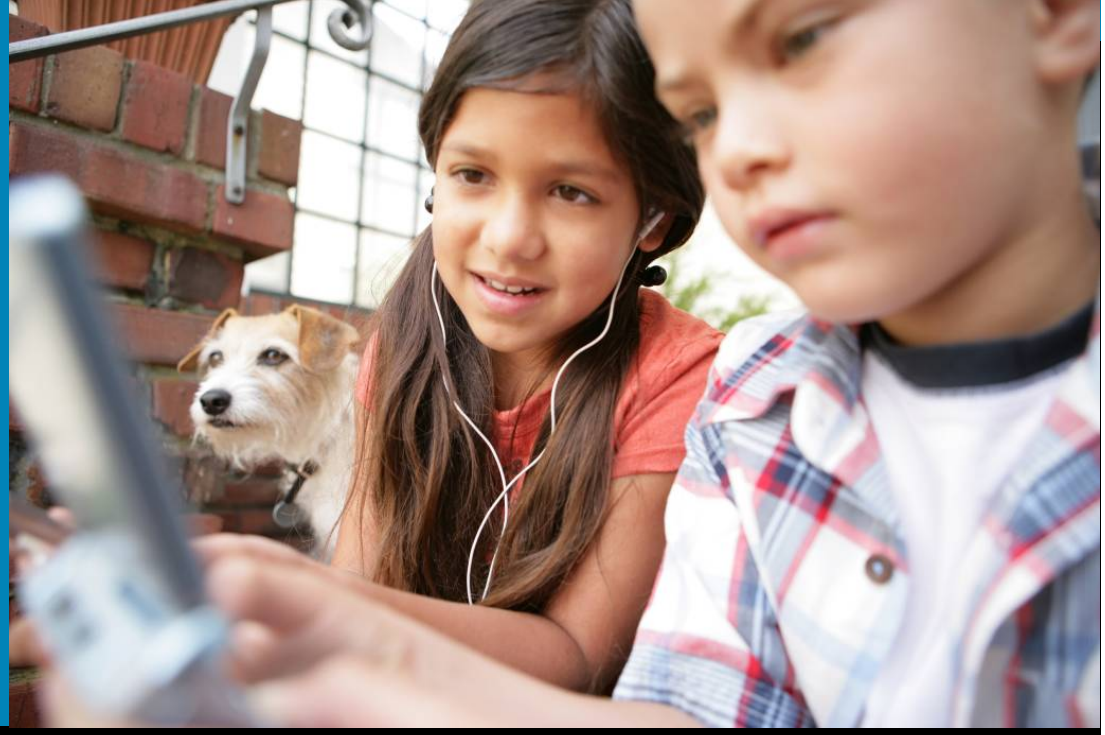

ASR1000 Building Blocks

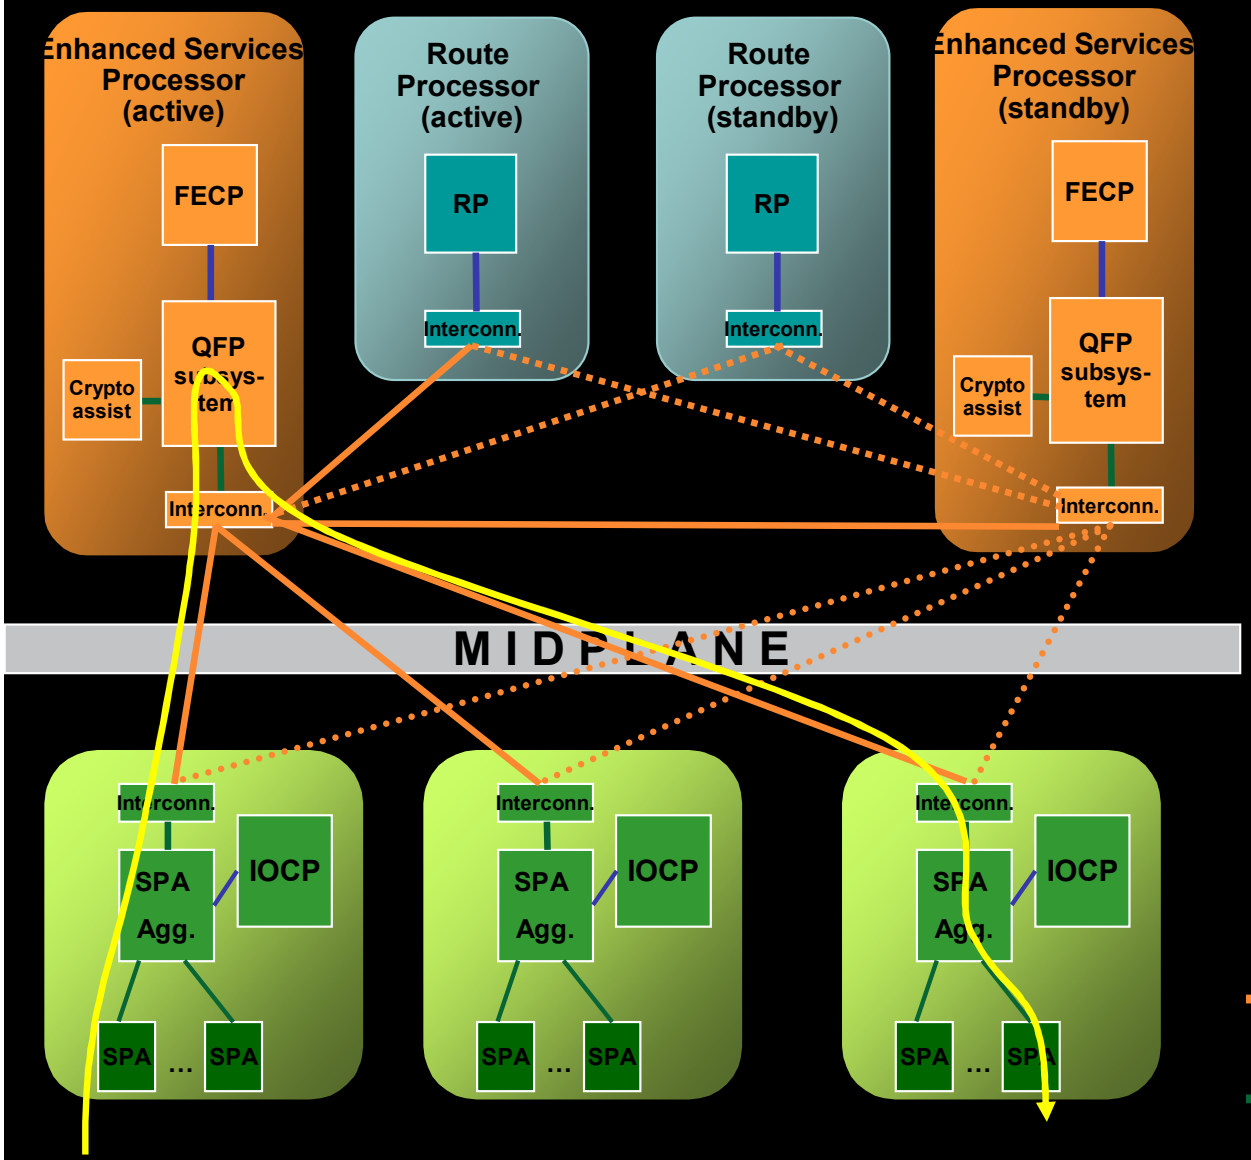

- **RP (Route Processor)**
  - Handles control plane traffic
  - Manages system
- **ESP**
  - Handles forwarding plane traffic
- **SIP**
  - Houses the SPAs
- **SPAs**
  - Provide interface connectivity
- **Centralised Forwarding Architecture**
  - All traffic flows through the ESP

ESI, (Enhanced Serdes Interface) 11.5Gbps

SPA-SPI, 11.2Gbps

Hypertransport, 10Gbps

Quantum Flow Processor Architecture
1st generation

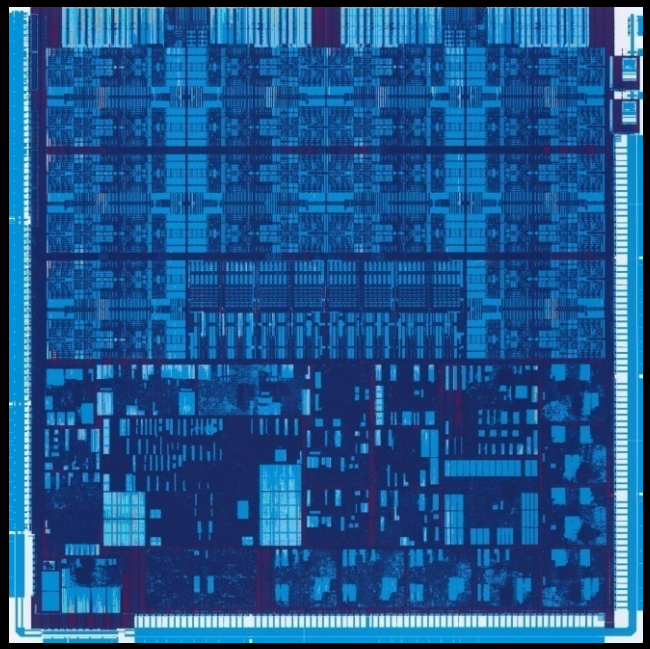

Multi-Core (40) Packet Processor

+

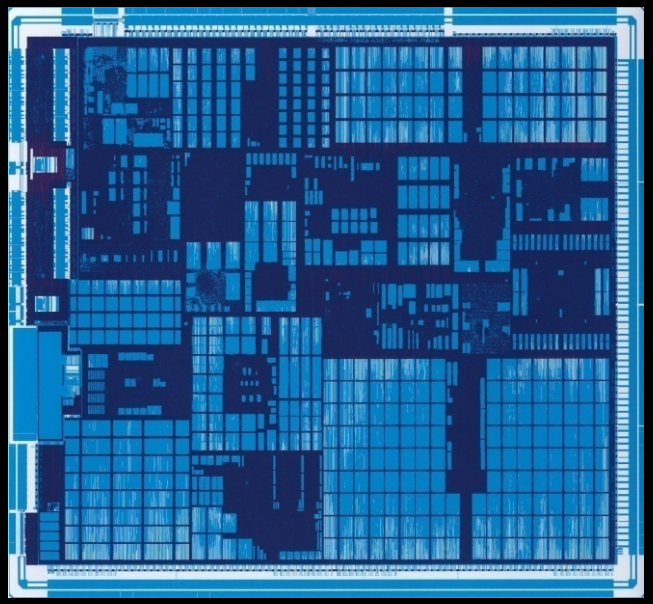

Traffic Manager (BQS)

+ Quantum Flow Processor Software

- Scale - 100s of resources & massive feature scale
- Performance - Designed to deliver 5-40+ Gbps
- Feature Velocity - Software designed to deliver a common forwarding plane for multiple systems.
- Multi-Generational - This is only the 1st Generation!

QFP Highlights

- **Packet Processing Engine (QFP-PPE)**
  - 40 Packet Processors – 4 Contexts (threads) each
  - Up to 1.2GHz (Tensilica ISA) processors + DRAM packet memory
  - Single TCAM4 I/F (can cascade 1-4 devices)
  - C-language for feature development (extensive development support tools)
  - HW assist for flow-locks, look-ups, stats, WRED, policers, range lookup, crypto, CRC
- **Buffer/queue subsystem (QFP-Traffic Manager)**
  - HW hierarchical 3-parameter (min, max & excess) scheduler
  - Fully configurable # of layers based on HQF
  - Priority propagation through the multiple layers

System Architecture – Data Plane

- All data forwarding is through ESP
- Exception: Punt path for Legacy protocols – handled by the RP
- Interconnect ASIC in each of the functional elements provides the backplane connection through ESI links
- ESI (Enhanced Serdes Interconnect) links are used for Data forwarding
- SPA-SPI links connect to the backplane through the SPA-Agg ASIC

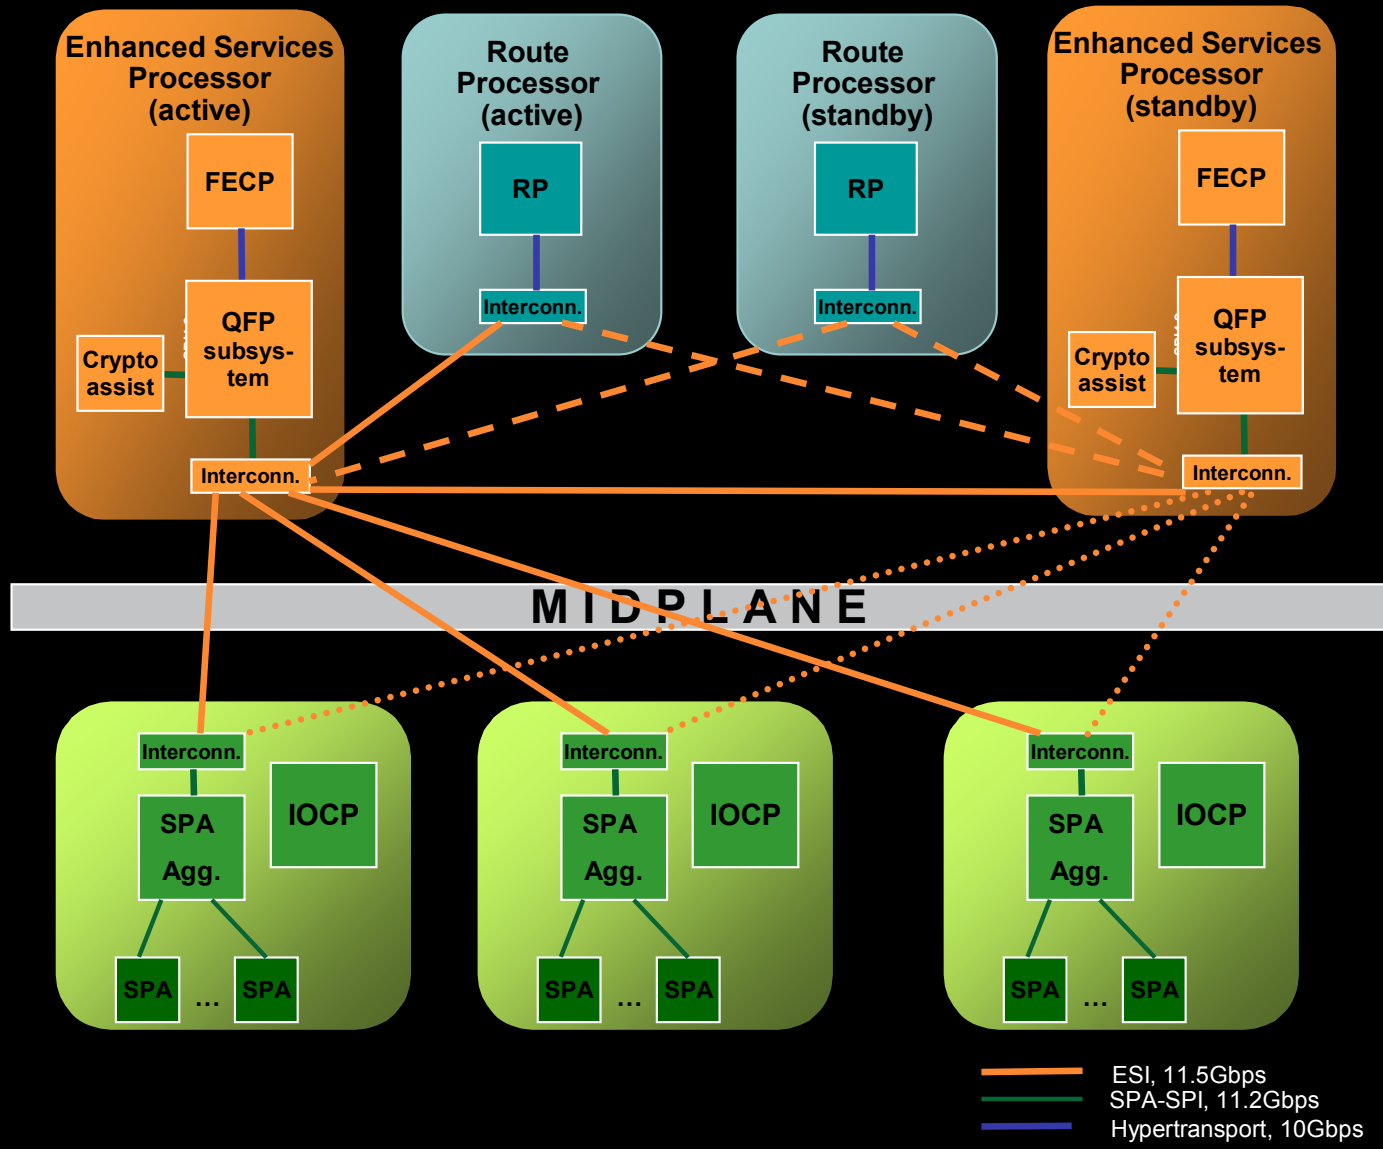

System Architecture - Control Plane

- Two different control plane links separate from the dataplane links
  - Ethernet out-of-band Channel (EOBC).
  - I²C - Monitor health of hardware components
- SPA control links

  Runs between IOCP and SPAs

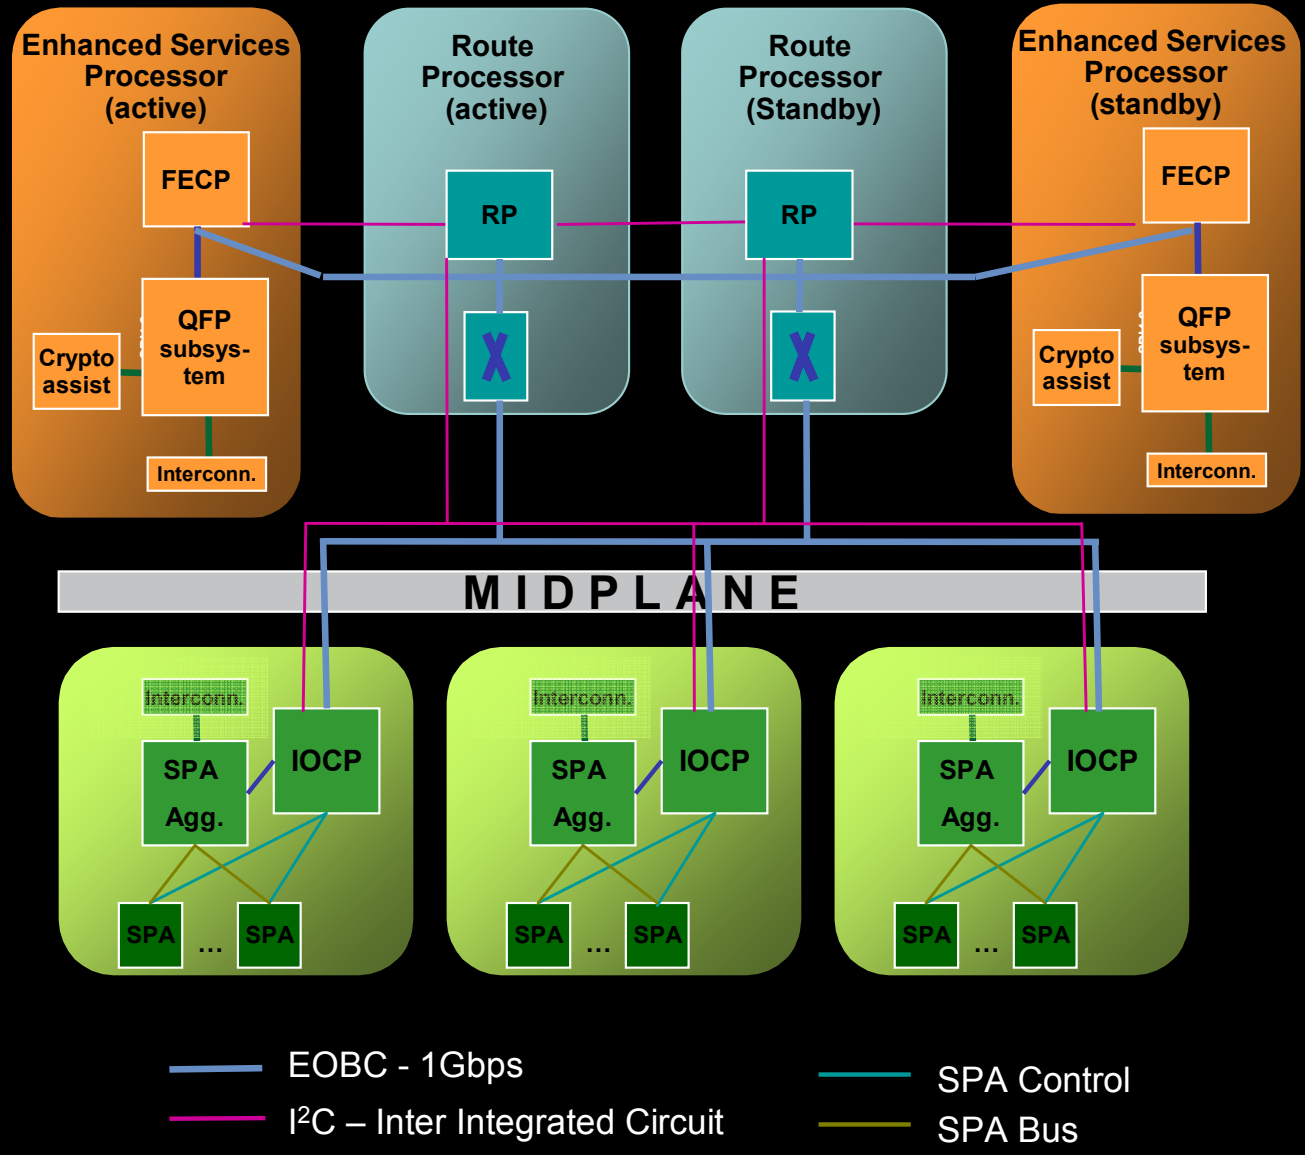

2RU

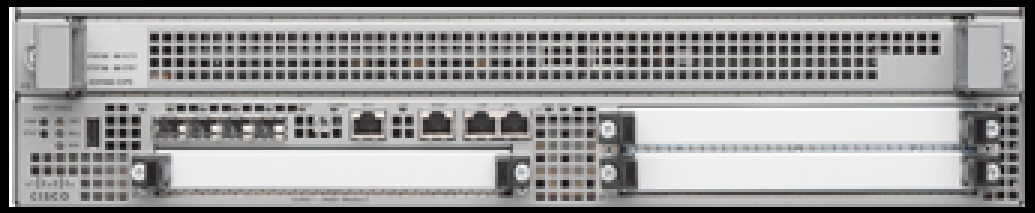

- Integrated RP1. Exact same hardware features as RP1 with the following exception
  - DRAM – 4G default
  - No Hard Drive for Mass storage
  - eUSB of 8GB
    - 2 x 32MB used as NVRAM
    - Remaining used for mass storage for code storage, boot, config, logs, etc
- Modular ESP; supports ESP-5G and ESP-10G
- Built-in 4 X GE SFP ports – feature/performance parity to GE SPA v2
- Copper SFP will support 10/100/1000
- Software redundancy feature supported with IOS processes

Chassis Options: ASR-1002

Chassis Options: ASR-1004

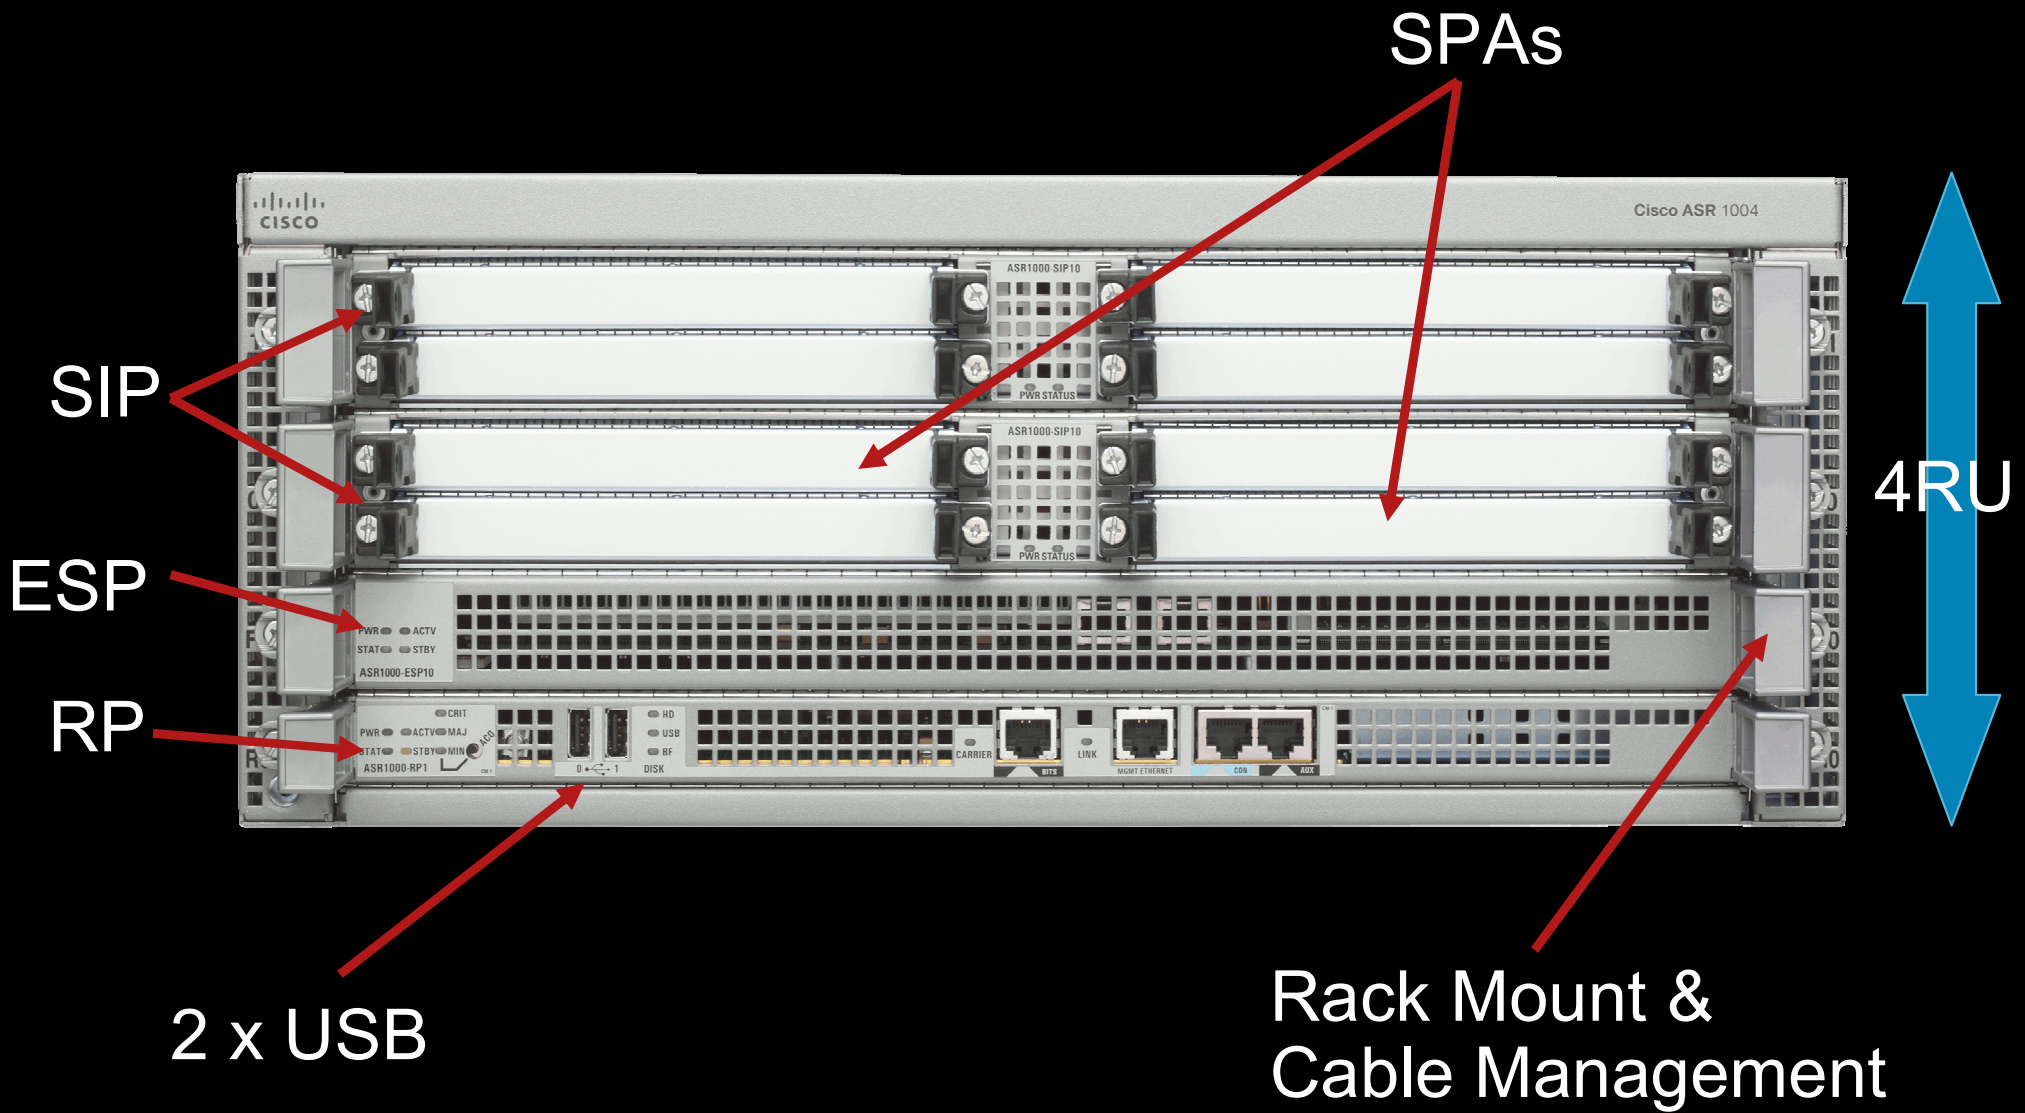

Chassis Options: ASR-1006

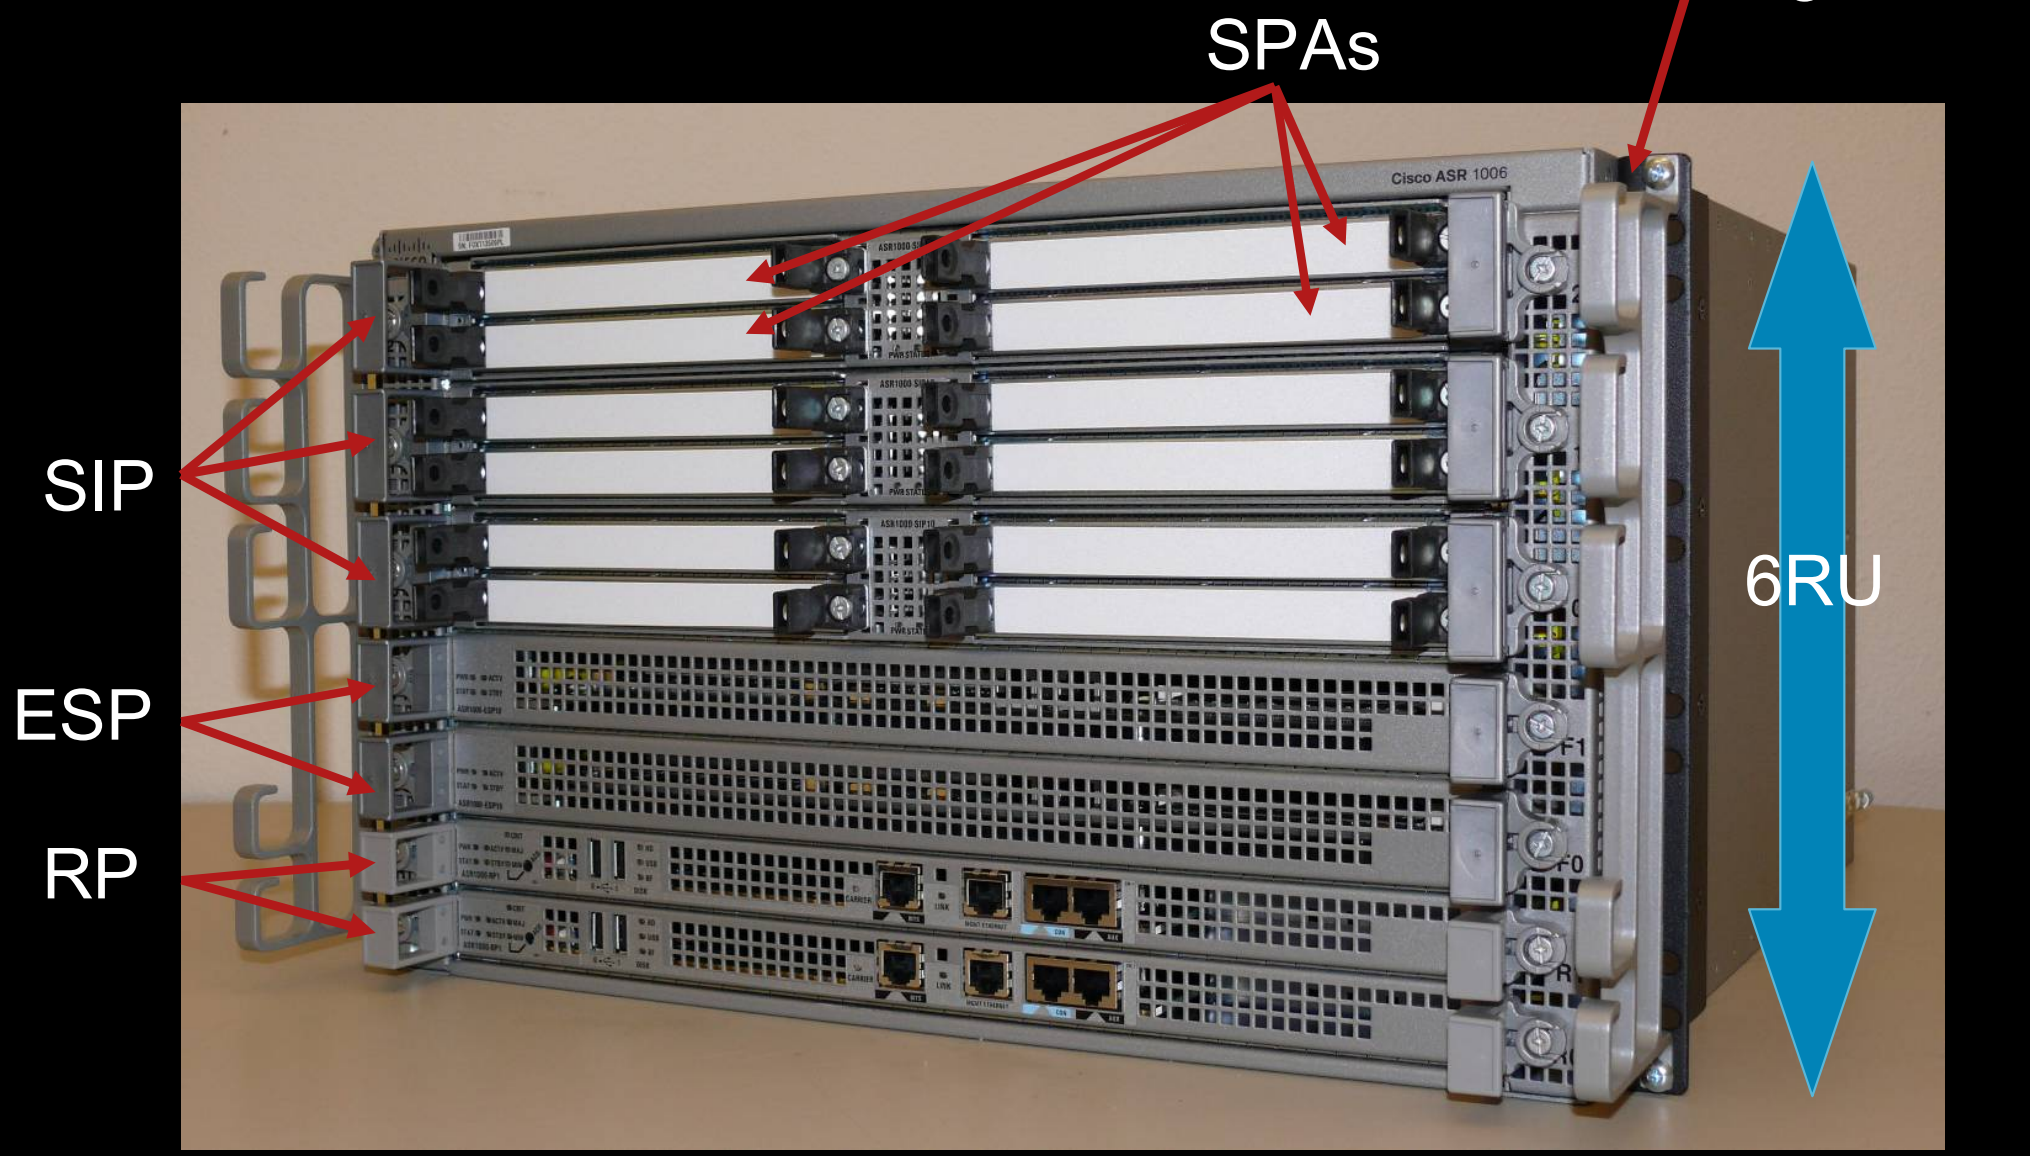

Route Processor – RP1

- **Features**
  - –1.4Ghz PowerPC Freescale 8548 CPU
  - – **2GB** default / **4GB** max optional DRAM, 32MB NVRAM
  - –1GB eUSB ‘On-Board’ Storage for code, boot, crashinfo etc
  - – 40GB rotary HDD, 32GB SSD optional (post FCS) – core dumps, failure capture
  - – EOBC Switch For Inter-Card Command & Control
  - – Stratum 3 network clock circuitry

- **Interfaces**
- Recessed Alarm Cut Off (ACO) button
- 2x USB ports
- RJ-45 BITS port
- RJ-45 100BASE-T Mgmt port
- RJ-45 Console port
- RJ-45 Auxiliary port
- NO Compact Flash!

A word on modern CPUs for ASR…

- Modern routers have a tough time handling the complex control plane
- There is just a lot to compute and scale is going up and up
- IOS XE allows us to use any CPU available as long as there is a Linux port
- This is important for future RPs and ASR platforms
- Cisco can piggy back on the mainstream server market's requirement for faster, mostly x86 multicore CPUs that also mainly run Linux

Enhanced Services Processor – ESP-5G, ESP-10G, ESP-20G, ESP-40G

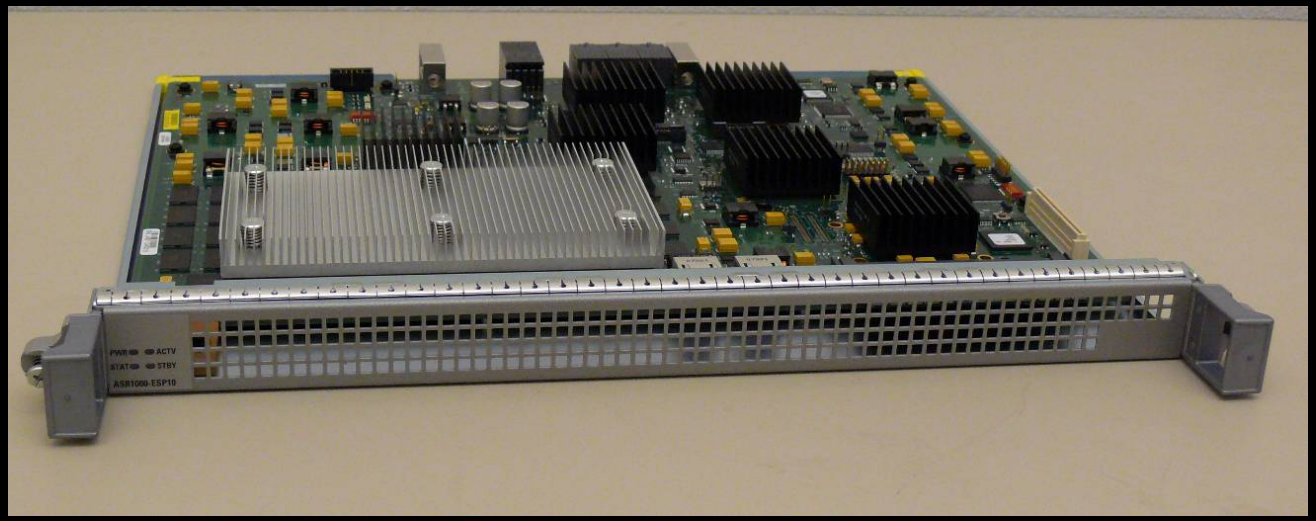

- Centralised, programmable forwarding engine (i.e. QFP subsystem (PPE) and crypto engine) providing full-packet processing
- Packet buffering and queuing/scheduling (BQS)
  - For output traffic to carrier cards/SPA's
  - For special features such as input shaping, reassembly, replication, punt to RP, etc.
- Interconnect providing data path links (ESI) to/from other cards over midplane
  - Transports traffic into and out of QFP
  - Input scheduler for allocating QFP BW among ESI's
- FECP CPU managing QFP, crypto device, midplane links, etc

ESP Generations

Roadmap

	ESP-5G	ESP-10G	ESP-20G	ESP-40G
Bandwidth	5Gbps	10Gbps	20Gbps	40Gbps
Based on	QFP	QFP	QFP	QFP
# of QFP Processors	20	40	40	40
Clock Rate	900 Mhz	900 Mhz	1.2 Ghz	1.2 Ghz
Crypto Engine BW	~1Gbps	3Gbps	8Gbps	10Gbps
QFP Memory	256MB	512MB	1GB	1GB
Packet Buffer	64MB	128MB	256MB	256MB
TCAM	10Mb	10Mb	40Mb	40Mb

SPA Interface Processor (SIP) – SIP-10, SIP-40

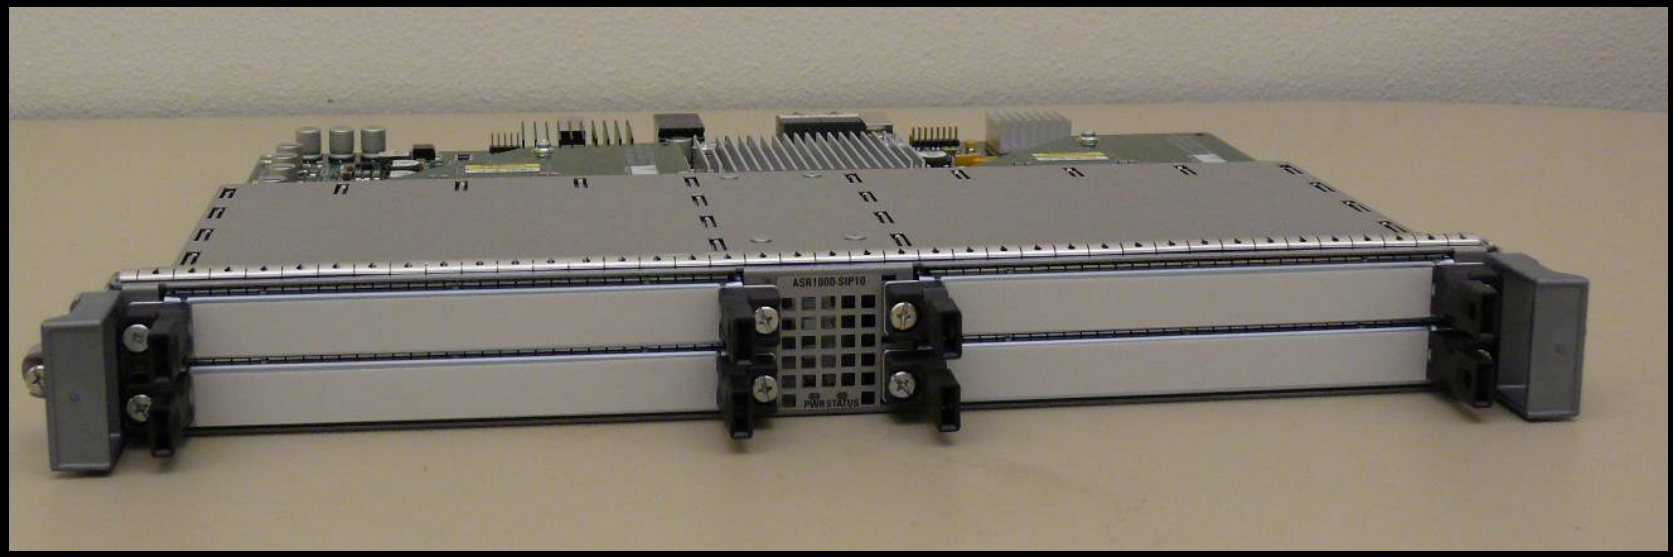

- Physical termination of SPA
- Supports up to 4 SPA's
  - 4 half-height, 2 full-height, 2 HH+1FH
  - full OIR support
- Does not participate in forwarding
- Capture stats on dropped packets
- Network clock distribution to SPA's, reference selection from SPA's
- Limited QoS
  - Ingress packet classification – high/low
  - Ingress over-subscription buffering (low priority) until ESP can service them. Up to 128MB of ingress oversubscription buffering
- IOCP manages Midplane links, SPA OIR, SPA drivers

What does the ESP Bandwidth mean?

- ESP bandwidth denotes the total 'OUTPUT' bandwidth of the system, regardless of the direction
- High priority traffic (as long as it is not over-subscribed - Example: <=10G for ESP-10G) will not be affected by this bandwidth limit
- ESP-10G Examples:

5G Unicast in each direction
Total Output bandwidth 5+5=10

1G Multicast with 8X replication in one direction
2G unicast in the other direction
Total Output bandwidth 8+2=10G

5G Unicast in one direction and
6G Unicast in the other direction
Total output bandwidth (5+6=11) exceeds 10G; Only 10G will go through

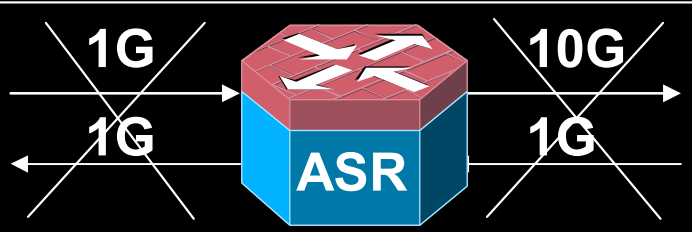

1G Multicast with 10X replication in one direction
1G Unicast in the other direction
Total bandwidth (10+1=11) exceeds 10G; only 10G will go through

Data Packet Flow: Through ESP10

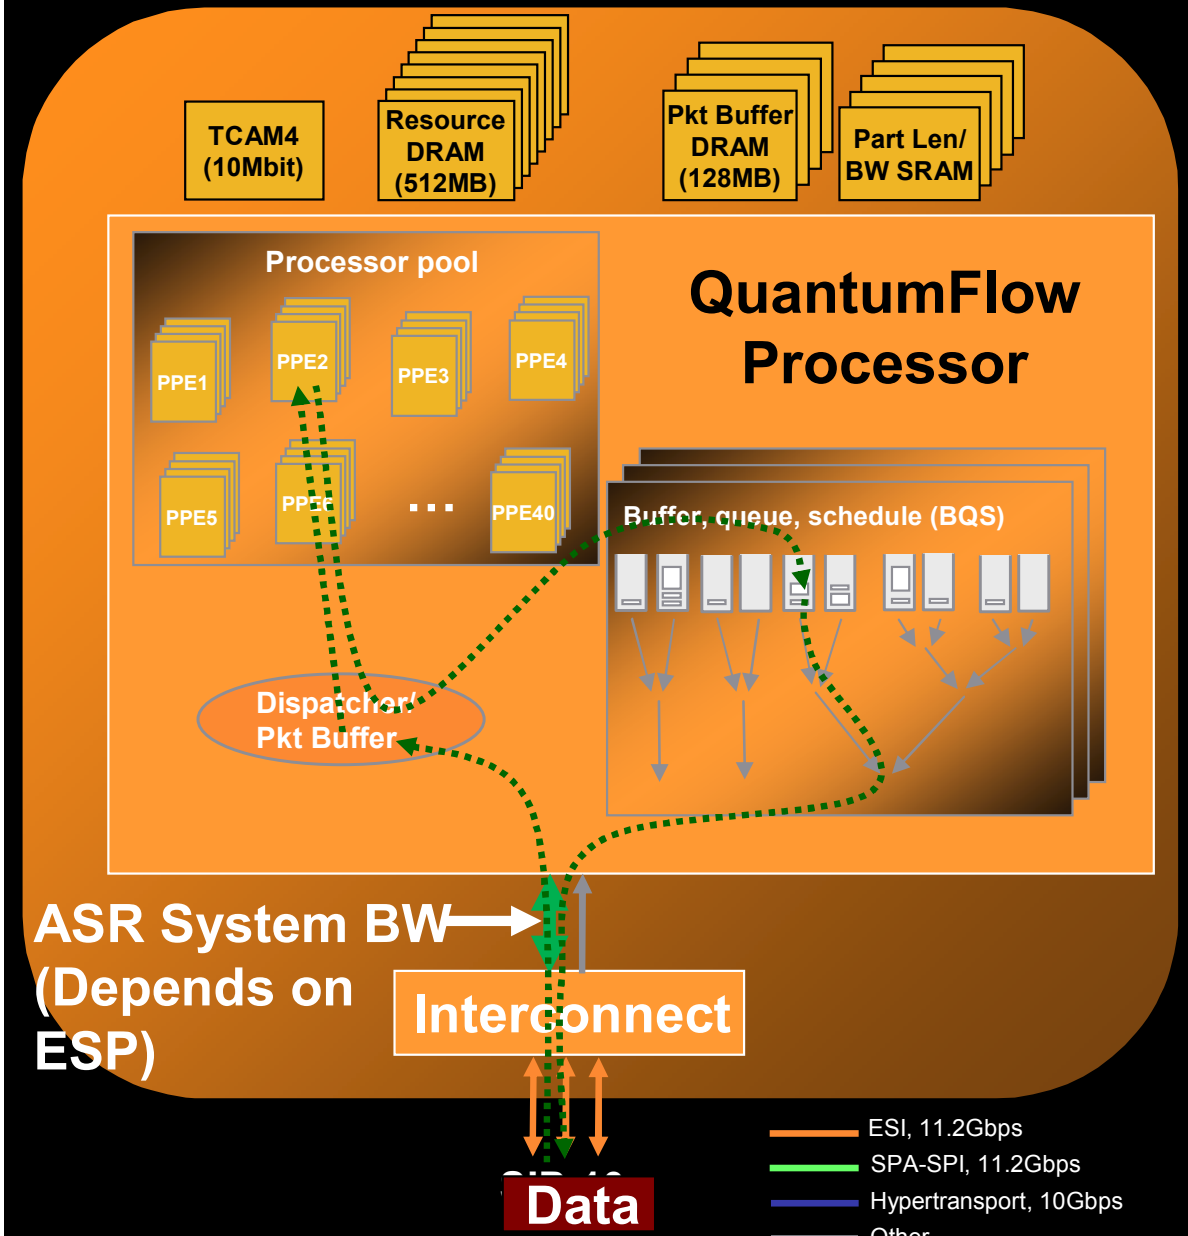

1. Packet arrives on QFP
2. Packet assigned to a PPE thread.
3. The PPE thread processes the packet in a feature chain similar to 12.2S IOS (very basic view of a v4 use case):

Input Features applied

Netflow, MQC/NBAR Classify, FW, RPF, Mark/Police, NAT, WCCP etc.

Forwarding Decision is made

Ipv4 FIB, Load Balance, MPLS, MPLSoGRE, Multicast etc.

Output Features applied

Netflow, FW, NAT, Crypto, MQC/NBAR Classify, Police/Mark etc.

Finished

4. Packet released from on-chip memory to Traffic Manager (Queued)
5. The Traffic Manager schedules which traffic to send to which SIP interface (or RP or Crypto Chip) based on priority and what is configured in MQC
6. SIP can independently backpressure ESP via ESI control message to pace the packet transfer if overloaded.

Possible Ingress Oversubscription Ratios

ESP Version	Max SPAs per SIP	Max BW / SPA (Gbps)	Max ESI BW / SIP (Gbps)	SIP Oversubscription Ratio (worst case)	6RU System Oversubscription Ratio (worst case)
ESP10 (with SIP10)	4	10	~10	4:1	12:1
ESP20 (with SIP40)	4	10	~20	2:1	6:1
ESP40 (with SIP40)	4	10	~40	1:1	3:1

ESP and SIP Ingress QOS functions are integrated into the ASR design to deal with oversubscription.

I/O – Shared Port Adapters

Ethernet:	• FE • GE • 10GE
POS: POS/DPT/RPR:	• OC3 • OC12 • OC48 • OC192
RPR	• GE • 10GE

- SPAs currently supported in other Cisco Platforms will also be supported on ASR1000

CEOP:	• OC3 • T3/E3 • T1/E1

Clear Chan.:	• T3/E3
Channelised:	• T1/E1 • T3 • STM1 • OC12 • OC48
ATM:	• T3/E3 • OC3 • OC12

Single-Height SPA

Double- Height SPA

ASR1000 Maximum Physical Interface Termination Capacity

	2RU	4RU	6RU	Comment
# SPAs (single-height)	3	8	12	
10GE	3	8	12	1-port 10GE
GE	28	64	96	8-port GE SPA;2RU has 4 built-in GE ports
FE	24	64	96	8-port FE
STM-4	3	8	12	1-port STM4 POS
STM-1	12	32	48	4-port STM1 POS
T3/E3	12	32	48	4-port T3/E3
ChT3 @T1	336	896	1344*	4-port Channelised T3
ChT3 @DS0	3072*	4096*	4096*	4-port Channelised T3
ChT1 / ChE1 @DS0	576/768	1526/2048*	2304/3072*	8-port Channelised T1/E1
V.35/X.21/EIA-232…	12	32	48	4-port Serial (12in1)
ChSTM1 @ T3 / E3	9/9	24/24	36/36	1-port Channelised STM1
ChSTM1 @ T1 / E1	252/189	672/504	1008* / 756	1-port Channelised STM1
ChSTM1 @ DS0	3066*	4096*	4096*	1-port Channelised STM1
STM-64	1	4	6	1-port OC192 (double-height)
STM-16	12	32	48	4-port OC48

- Physical interface termination capacities only
  - Throughput performance depends on ESP used (ESP-5G / ESP-10G / ESP-20G / ESP-40G)
  - Assumes all SPA slots are filled with the respective SPA

* Requires FP20G

ASR1000 Performance

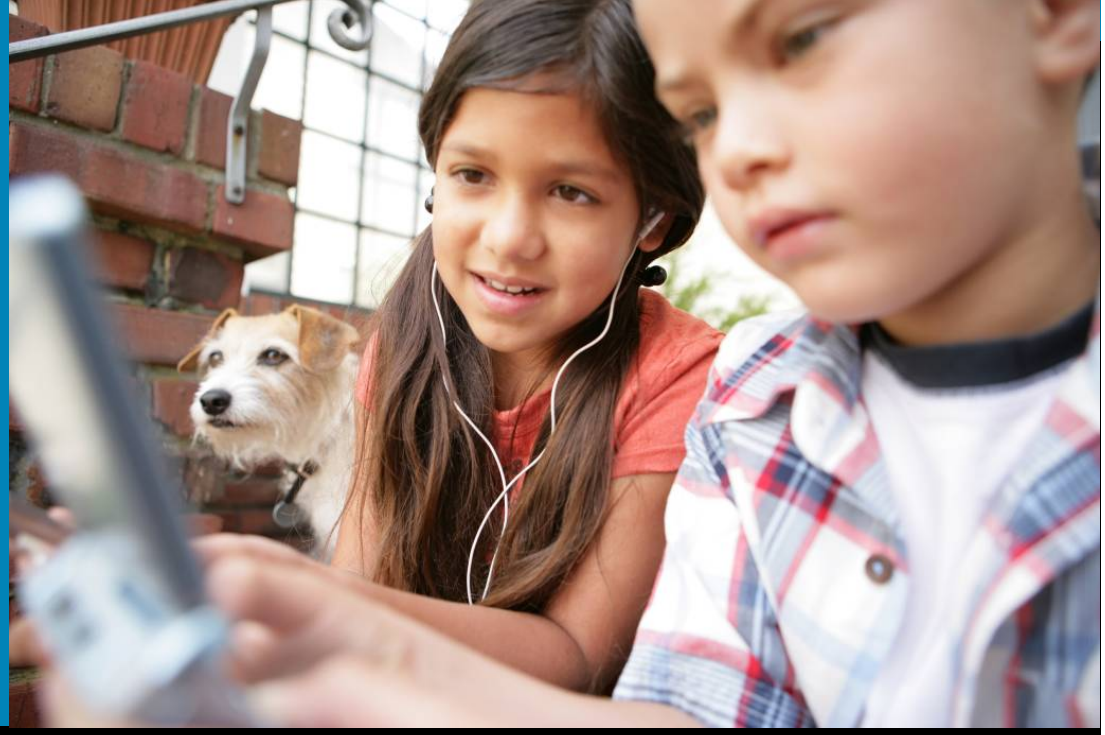

What affects ASR/QFP Performance?

- More features = more instructions = less MPPS
  - Feature Invocation Array (FIA)
  - Fancy name for per interface, per protocol feature list
- Use of multi-pass processing (recycling)
  - Multicast and IPsec processing are examples
  - Packet fragmentation

ASR Performance and Scaling

For performance information, please contact your local Cisco representative.

Broadband on the ASR

Broadband Feature Summary

- Based on 12.2S feature set
- LAC and LNS functionality
- Per session QoS (Post-FCS)
- Broadband HA: SSO/ISSU
- Subscriber-aware: ISG (Post-FCS)

ASR Broadband HA

- Preserves PPPoL2TP sessions for both LAC and LNS
- Preserves L2TP tunnel between LAC and LNS
- Basis for ISSU
- Standby RP:
  - Synchronises state information between Active RP & Standby RP
- Standby ESP:
  - Active RP will program the standby ESP while programming the active ESP. This ensures the state consistency of active and standby ESP
- ESP CPU maintains the per-session state in QFP – this makes possible the different types of redundancy (RP vs. ESP)

Broadband SSO/ISSU Supported Features

1. PPPoE and PPP SSO Session Features
   - BBA per-Session Features for both terminated [VAI] and forwarded [SSF] sessions
   - IP Address Management/Local Pools
   - Virtual Template Manager
2. L2TP Silent Mode Switchover: LAC/LNS
3. DHCPv4 & DHCPv6 SSO
4. AAA SSO
5. MIBs/SNMP
   - Session info
   - Statistics will be reset

Note: Statistics are NOT SSO'ed.

L2TP Silent Switchover for RP and ESP

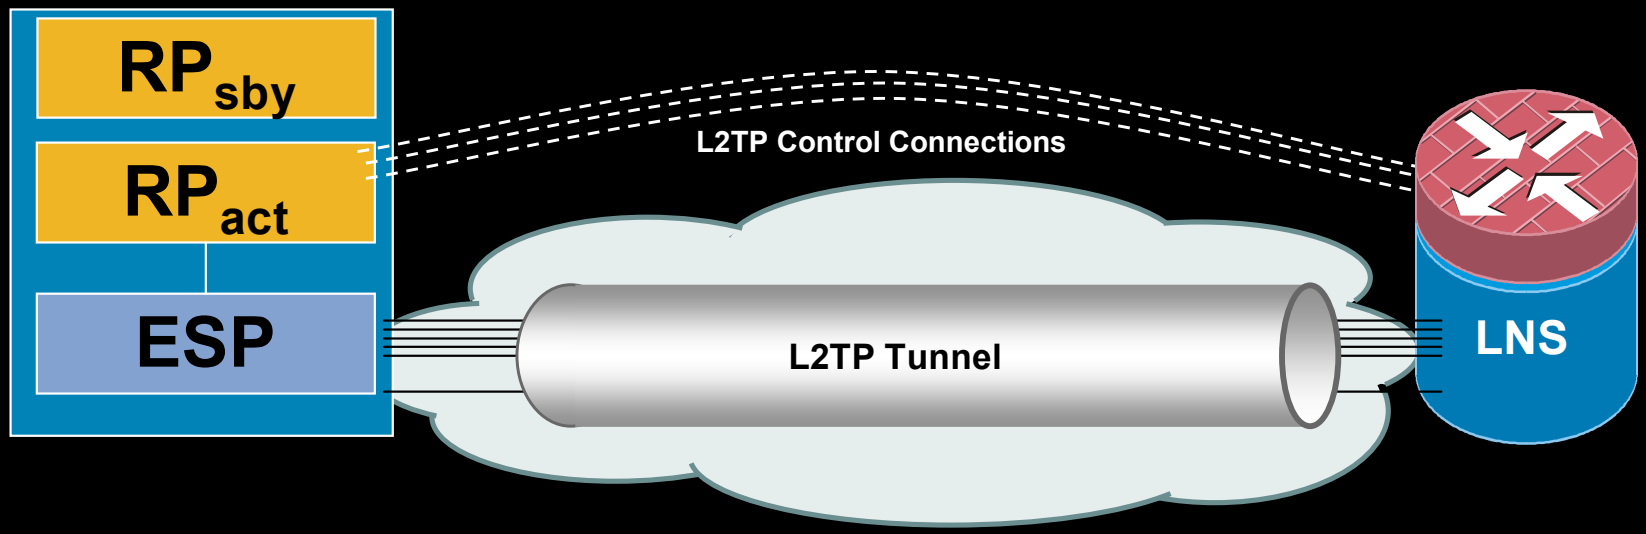

- RP_{act} synchronises state with RP_{sby}

  State includes configuration, PPP session IDs, L2TP CC sequence numbers etc.

  Sequence numbers (Ns, Nr) for L2TP Control Connections (CC) are only synched once for a packet window of X (i.e. once every X L2TP control packets)

- The standby ESP will take over forwarding as it is sync'd (except transient packets on the previously active ESP may be dropped)

Broadband QoS

- Leverages ASR/QFP QoS infrastructure
  - Sessions in QFP use the same data structures as interfaces and sub interfaces
- Per session QoS:
  - Session level classification (match IPv4 precedence / IPv4 DSCP, match Access list, match QoS group)
  - Session level marking (IPv4 precedence / IPv4 DSCP)
  - Session level policing
  - Session level queuing/shaping (Post FCS)
  - QoS config in vpdn-group and virtual-template
  - QoS config from Radius (Post FCS)
- Other Broadband related QoS features:
  - Automatic Mapping of L3 QoS to L2 CoS on Egress (Post FCS)
  - Marking L2 CoS for PPP/PPPoE control packets

ASR Broadband Scalability

	ESP-5G	ESP-10G	ESP-20G TARGET
Max # of Subscribers per System	16,000 (ASR limit)	32,000 (ASR limit)	64,000 (ASR limit)
Max # of L2TP Tunnels per System	4,000 (ASR limit)	8,000 (ASR limit)	16,000 (ASR limit)
Max # of VLANs per SPA	8,000 (SPA limit)		
Max # of Subscribers per SPA	8,000 (1 subscriber / VLAN)		
Max # of Virtual Template Interfaces	4,000 (IOS limit)		
Max # of bba-group (LAC)	4,000		
Max # of vpdn-group (LNS)	4,000		

Broadband Forwarding Performance

	ESP-10G	ESP-20G TARGET
BBA Performance: Raw Forwarding rate	**~ 8Mpps**	**~ 12Mpps**
Throughput bandwidth	**10 Gbps**	**20 Gbps**

BB QFP packets do not normally require packet recycling avoiding a performance tax

Broadband Roadmap

ASR FCS Broadband Features

Subscriber Access

- PPPoE over VLAN
- PPPoEoVLAN tunneled into L2TP
- PPPoE Security
- PPPoE Session Limiting
- Logical Line ID (LLID) and LLID Blocking
- VPDN Logging
- VPDN Session Rate Compute

IP Address Assignment

- PPP IPCP Subnet
- Local IP Pool
- IP Pools downloaded from RADIUS
- IP Address assignment via RADIUS (Attr 8 Framed IP address)
- Overlapping IP Address Pools
- DHCP server
- DHCP option 82
- DHCP relay unnumbered

L2TP

- L2TPv2
- L2TPv2 over IPv4
- TCP MSS Adjust
- L2TP ToS Reflection
- L2TP IP UDP checksum ignore
- L2TP LNS Random Session Balancing
- Managed LNS (VRF termination)
- L2TP Extended Failover
- Session/Idle Timeout
- L2TP Calling Station ID Suppression

High Availability

- 802.1Q SSO/ISSU
- PPPoE SSO/ISSU
- AAA SSO/ISSU
- DHCPv4 SSO/ISSU

ASR FCS Broadband Features

AAA

- RADIUS
- PAP/CHAP
- RADIUS Attribute Screening
- VPDN group-session limiting
- VPDN Group Templates
- RADIUS attribute 8 in access request
- RFC-2869 attribute 52 and 53 Gigaword support
- Radius round-robin failover
- Method List/Server Group Scalability
- Accounting of VPDN disconnect cause
- NAS-Port ID (for VLAN)

- IP VRF Attribute VSA Support
- Radius Attribute 77 for DSL
- L2TP Disconnect Cause Information
- Virtual Template Limit Expansion
- Per-User Access-List
- TACACS+
- AAA Method Lists Enhancement
- Start/Stop Messages
- Interim Messages
- Packet of Disconnect
- RADIUS Accounting

ASR FCS Broadband Features

MIBS

- IF-MIB (RFC-2863) support for PPPoEoVLAN interfaces
- CISCO-AAA-SESSION-MIB
- CISCO-VPDN-MGMT-MIB
- IF-MIB (RFC-2863) support for L2TP tunnel interfaces
- TUNNEL-MIB (RFC-4087)
- L2TP-MIB (RFC-3371)
- CISCO-L2TP-MIB
- CISCO-CONTEXT-MAPPING-MIB
- CISCO-CDP-MIB (Context aware)
- CISCO-PING-MIB (Context aware)
- CISCO-AAA-SERVER-MIB
- CISCO-AAA-SERVER-EXT-MIB
- CISCO-IP-LOCAL-POOL-MIB
- CISCO-DHCP-MIB

ASR Post FCS Features

- ISG – RLS2
- Per Session Queuing – RLS2
- PPPoA – RLS5
- L2TP SSO/ISSU – RLS2 Silent Mode
- PPPoEoQinQ – RLS3

Software Architecture

Software Architecture – IOS XE

- IOS XE = IOS + IOS XE Middleware + Platform Software
- Operational Consistency - same look and feel as IOS Router
- IOS runs as its own Linux process for control plane (Routing, SNMP, CLI etc). 32bit and 64bit capable.
- Linux kernel with multiple processes running in protected memory for
  - Fault containment
  - Re-startability
  - ISSU of individual SW packages
- ASR HA Innovations
  - Zero-packet-loss RP Failover
  - <50ms ESP Failover
  - "Software Redundancy"

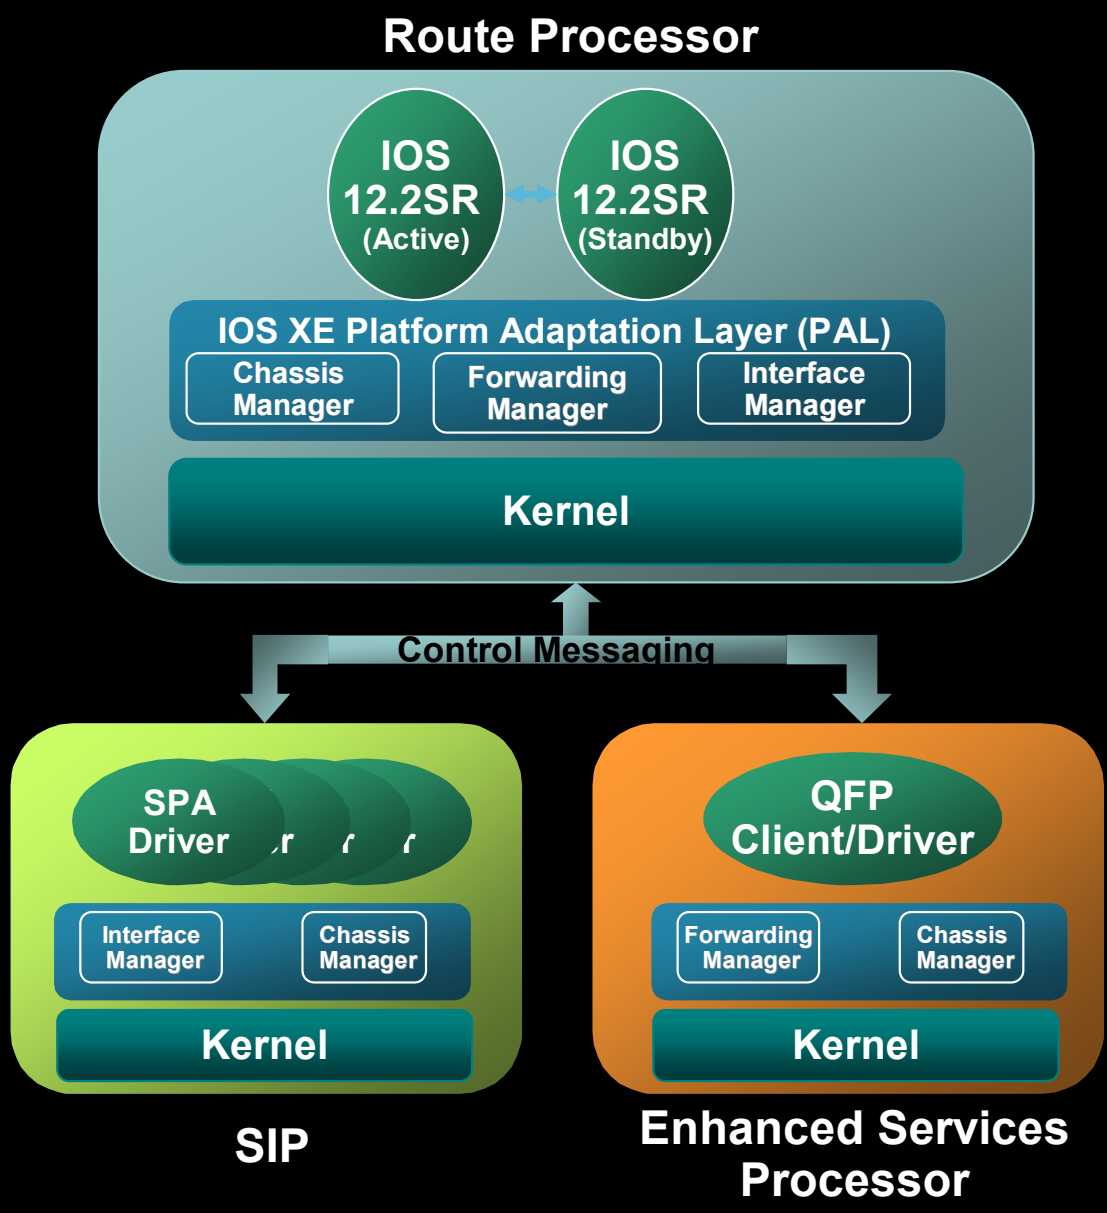

High Availability

ASR 1000 HA Highlights

- ASR 1000 leverages Cisco IOS HA infrastructure – NSF/SSO, ISSU
- 1+1 redundancy option for RP and ESP
  - Active and standby
  - No load balancing
- RP's are separate from ESP's
  - Switchover of ESP does not result in switchover of RP
  - Switchover of RP/IOS does not result in switchover of ESP
- Single RP may be configured with dual IOS for SW redundancy (single RP only)
- No redundancy for SIP or other I/O cards
  - SPA plugs into a single SIP
- Protection against SPA or SIP failure is via APS or Y-cable redundancy feature (Future: requires SPA support)

System Architecture – RP failover (control plane)

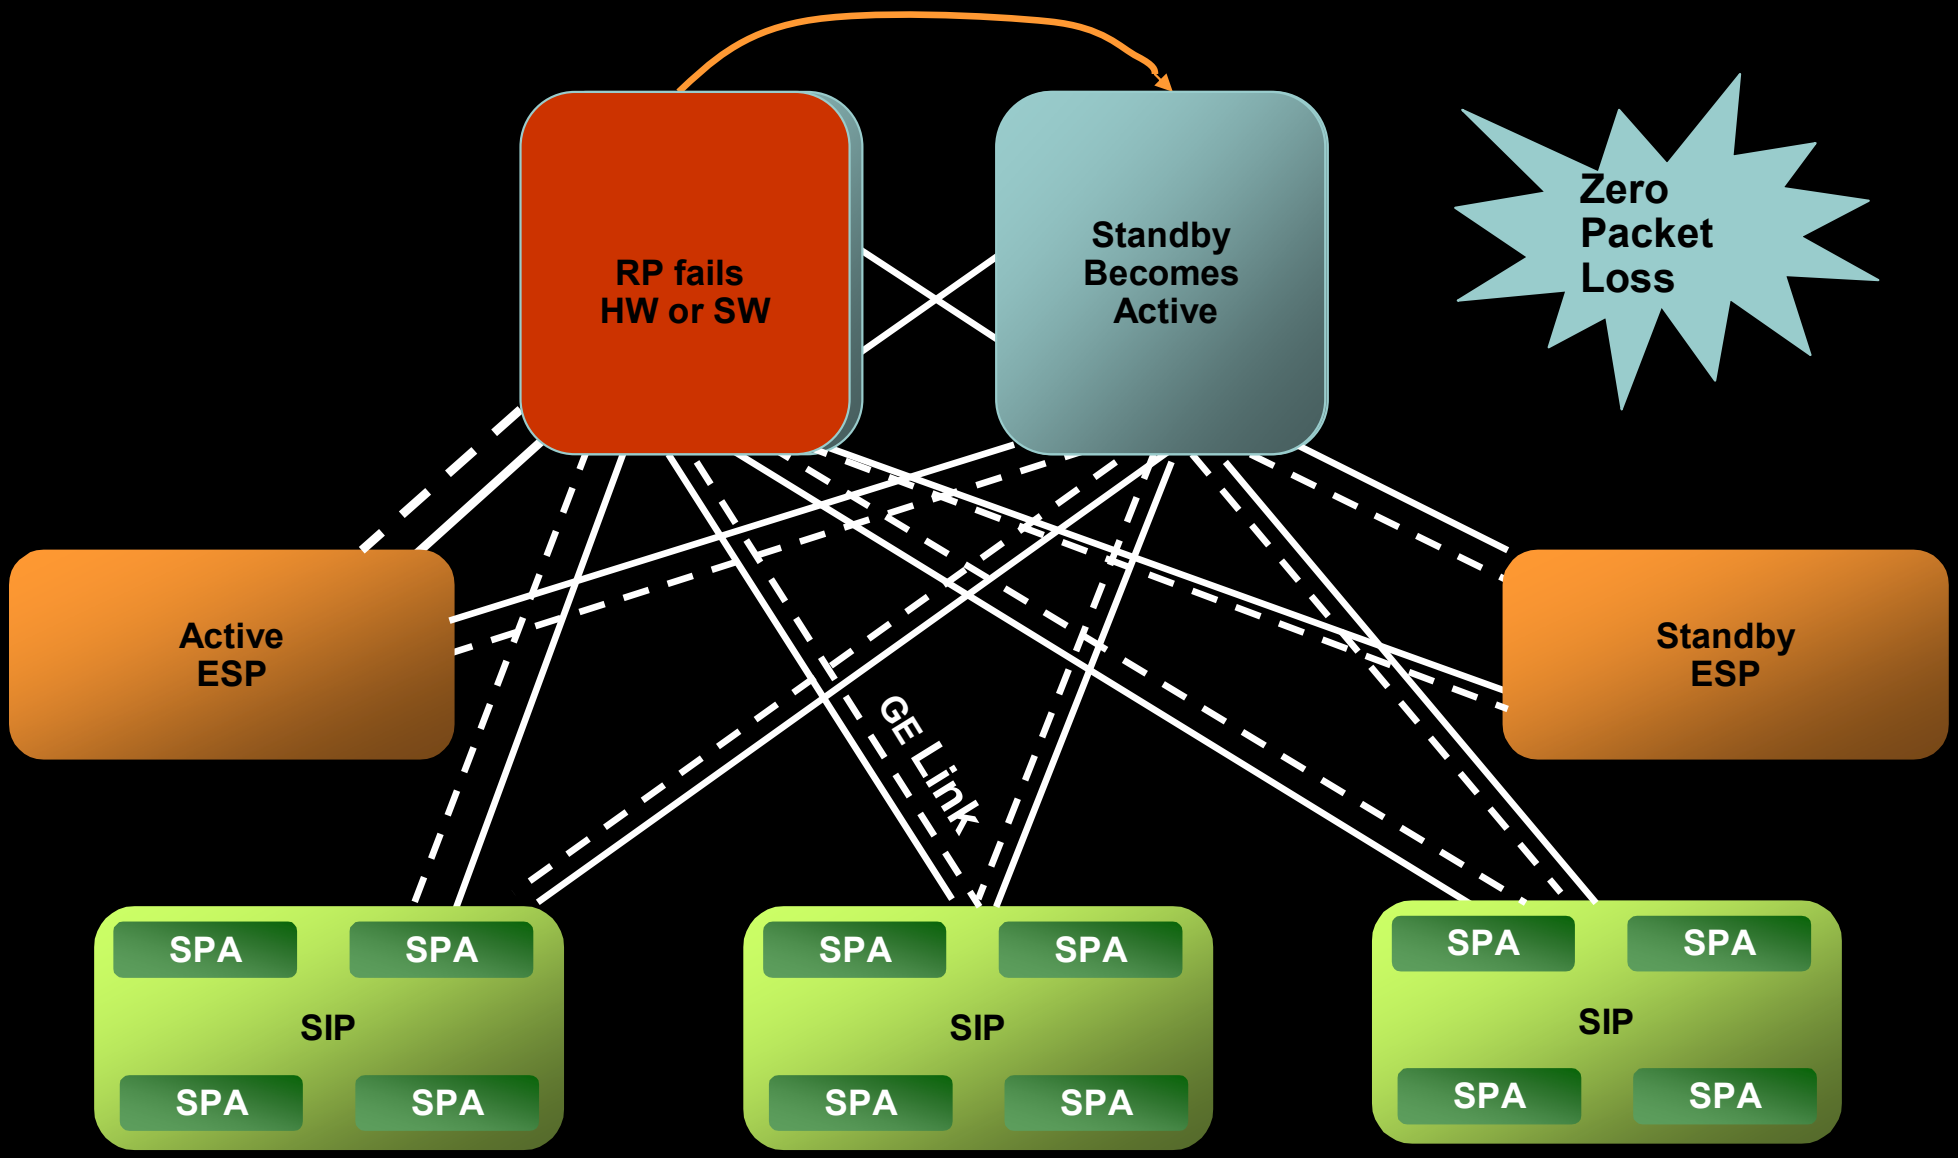

Separate and independent internal communication link for control plane (GE)

System Architecture – ESP Failover (data plane)

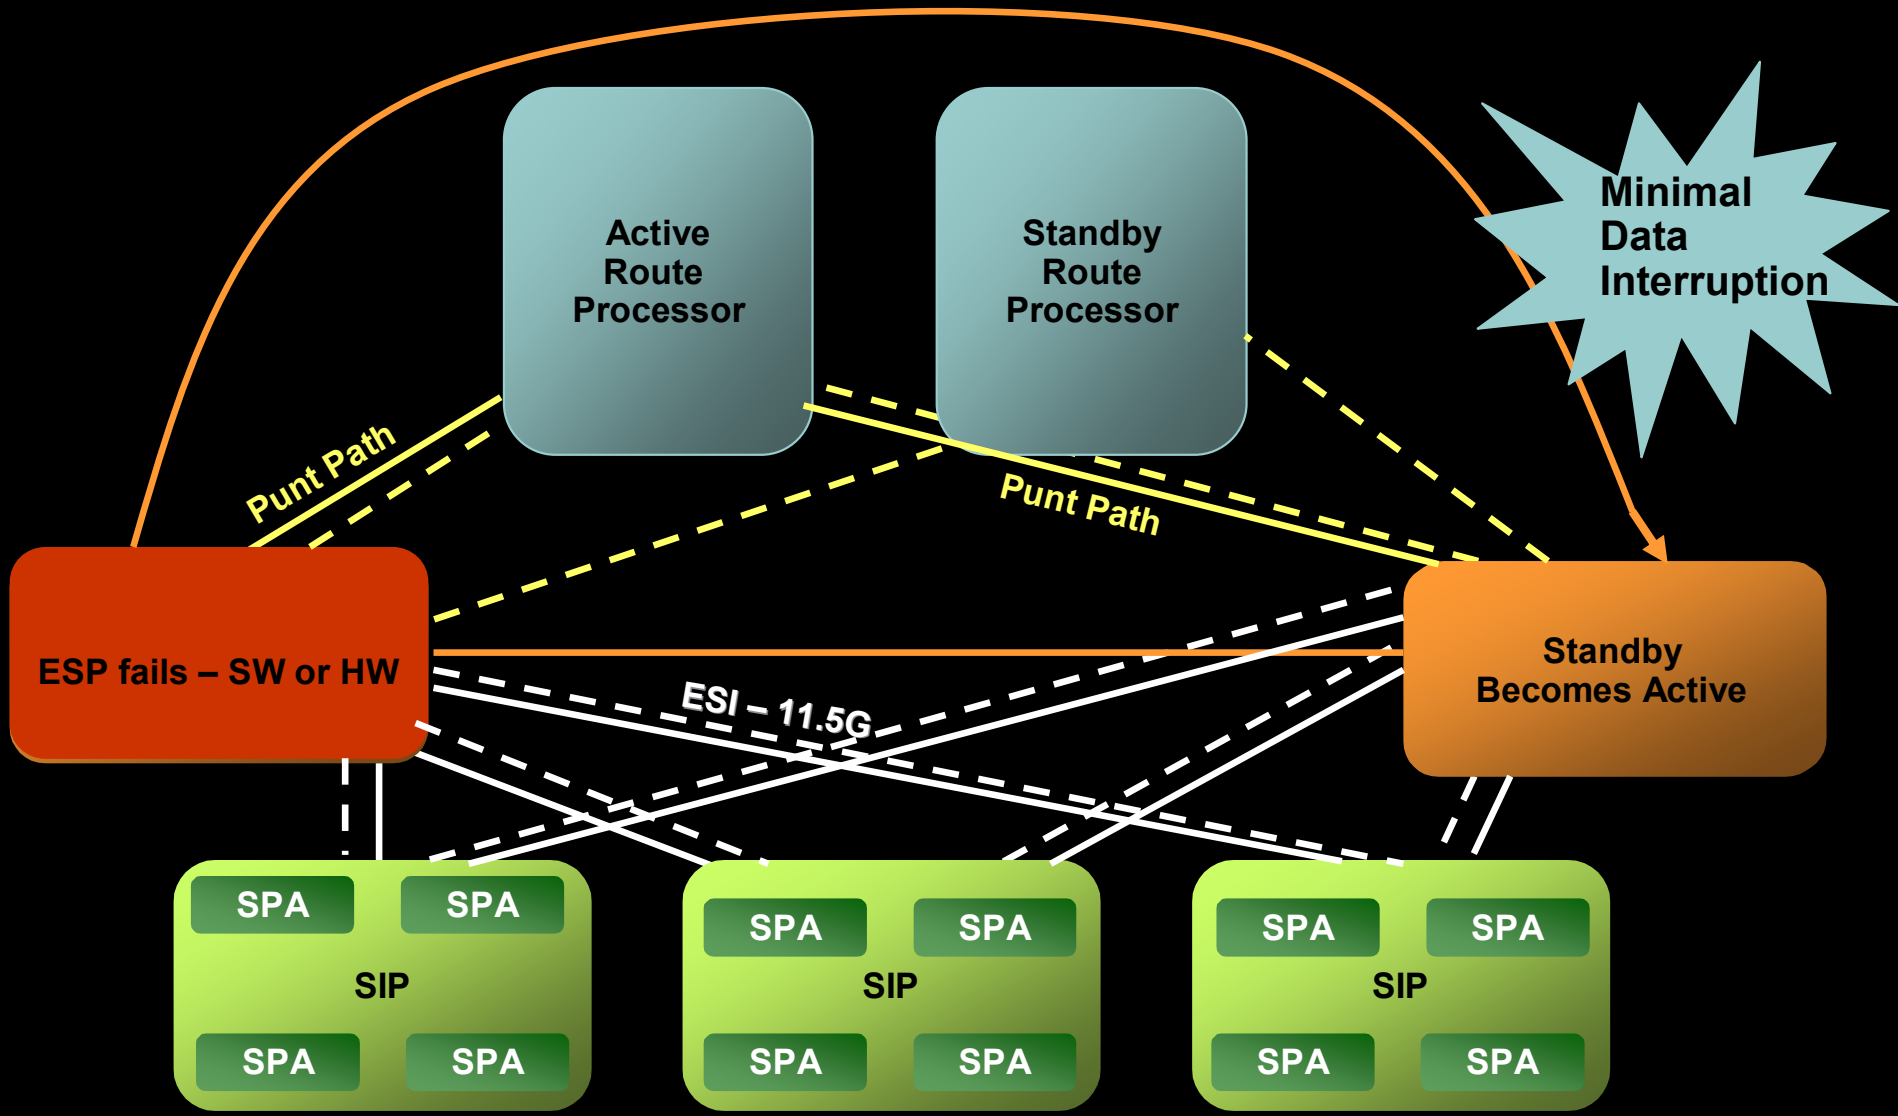

- All packets processed by QFP for forwarding
- Separate and Independent links for Data Plane communication (ESI 11.5G)

Software Redundancy on 2RU/4RU Single RP/ESP

- Stand-by IOS process in RP in the single-engine 4RU/2RU system
- Two IOS process in a single RP function similar to different processes on separate RP
- Support all NSF/SSO features supported by dual-RP systems
- Requires additional RP memory – 4G

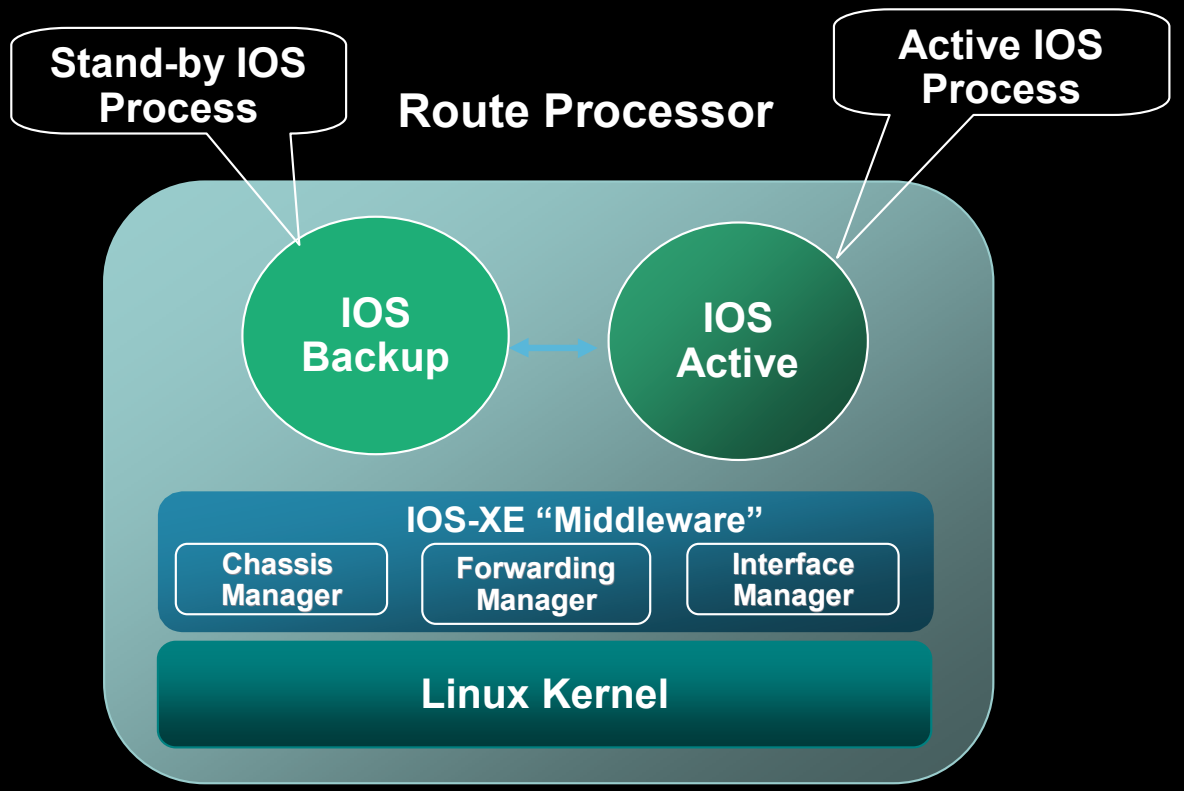

ASR Software Packaging

- Each functional element of ASR will support different modular software packages
- In total, 7 different software packages will be available
- The packages are designed to maximise the 'In Service Software Upgrade' capability
- At every release of ASR software, all 7 components will be integrated and available as one software package for download from CCO

- RP
  1. **RPBase:** RP OS
  2. **RPControl:** Control Plane processes that interface between IOS and the rest of the platform
  3. **RPIOS**
  4. **RPAccess:** Software required for Router access; 2 versions will be available. One that contains open SSH & SSL and one without (RPAccess and RPAccess-K9)
- SIP
  5. **SIPBase:** SIP OS + Control processes
  6. **SIPSPA:** SPA drivers and FPD (SPA FPGA image)
- ESP
  7. **ESPBase:** ESP OS + Control processes + QFP client/driver/ucode

+ **ROM Monitor:** One ROM Monitor package containing ROMMON for RP, ESP, SIP (released when needed)

ASR IOS Feature Sets

RP1 ADVANCED ENTERPRISE SERVICES	"Kitchen Sink" image includes all features (includes all legacy protocols, but excludes SNA Switching)
RP1 ADVANCED IP SERVICES	Excludes all Legacy Protocols
RP1 IP BASE	Includes only basic IP features with SSH support
RP1 IP BASE W/O CRYPTO	Includes only basic IP features

Note: No BGP in IP BASE

AISK9 and AESK9 images include both LI & SBC

ASR1000 Roadmap

Mid Range Product Strategy

ASR 1000 Architectural Insertion

- Ent: Wan Aggregation w/ FW
- Ent: IPSec Headend
- Ent: Internet Gateway
- SP: Broadband Agg, Service Mesh
- SP: SBC (DBE only)
- SP: Managed Services

- Ent: Full Security: GETVPN, SSL, IPS
- Ent: Voice (SBC+GW)
- SP& Ent: L3VPN, L2VPN
- SP: BB Phase 2 (ISG, L2CP)
- SP: SBC Unified
- SP: MSE

ASR 1000 Product Strategy

- Introduction of QFP
- Introduction of IOS-XE
- Initial SP and Ent Arch insertions

- Scale IOS-XE control plane with multi-core technology
- Expand SP and Ent architectures into adjacent markets (full integrated security, voice)

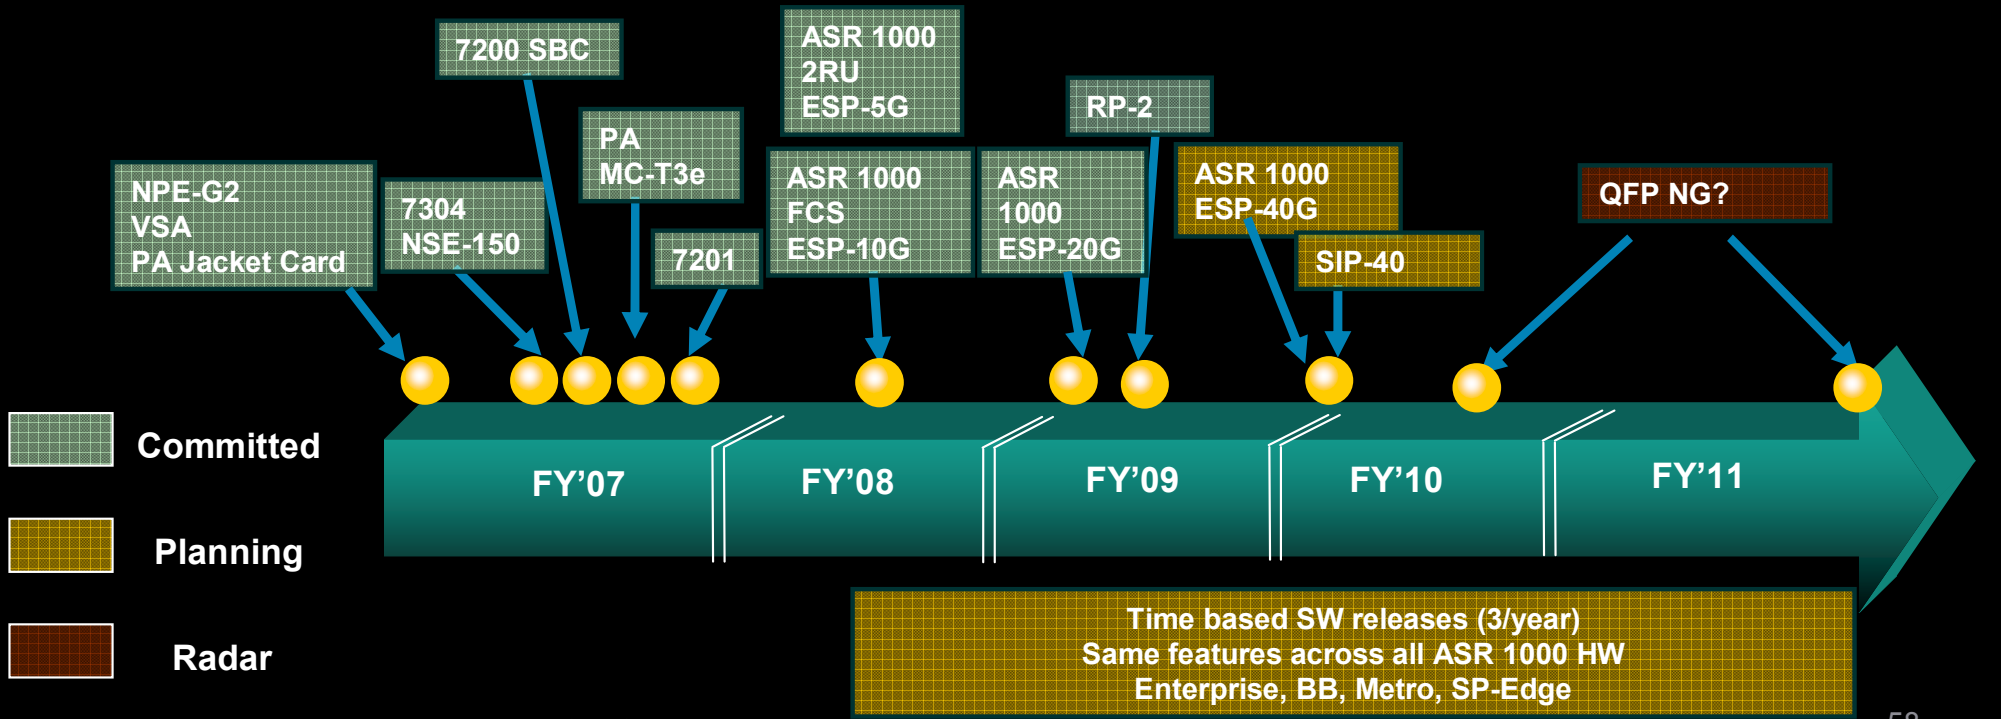

ASR 1000 RLS 1 – April 2008

Software Features – High level

- IPv4 Unicast & Multicast
- IPv6 Unicast and IPv6 Multicast
- High Availability – NSF/SSO & ISSU
- BB Aggregation - Phase I
- QOS
- Security ACLs
- SBC – Session Border Controller – DBE function
- NBAR
- Netflow
- cRTP
- Security Features
  - NAT, IOS FW – Phase I
  - IPSec
- MPLS VPN – Phase I
- GRE

Hardware Features

- 2/4/6 RU Systems
- ESP-10G, SIP-10, RP-1 on 4/6RU
- ESP-5G/ESP-10G 2RU only
- SPAs
  - 8-port GE
  - 1-port 10GE
  - 2-port GE and 5-port GE, 10-port GE
  - 8-port FE
  - 8-port T1/E1
  - 2&4-port T3/E3
  - 2&4-port OC-3/STM1 POS
  - 1-port OC12/STM4 POS
  - 2&4 –port Channelised T3
  - 4-port Serial (12-in-1)

ASR 1000 RLS 1 - Software Feature Details - 1

IPv6

- IPv6 Unicast/Multicast
- IPv6 ACLs
- Data Link Protocols
  - Ethernet, dot1q
- Routing Protocols
  - OSPFv3, Static Routes, BGP extensions for v6
- MIBs & Stats

Broadband Features

- LAC, LNS & PTA
- PPPoE and related features
- L2TP & related features
- AAA support - RADIUS
- DHCP Relay features – v4 & v6
- RA-MPLS

High Availability

- NSF
  - BGP, OSPF, OSPFv3, IS-IS, EIGRP, LDP
- SSO/ISSU:
  - CEF
  - SNMP, ARP, NAT
  - Stateful ISIS
  - MPLS, MPLS VPN, LDP, VRF-lite
  - IPv6 (NDP, uRPF)
  - FR, PPP, MLPPP, HDLC, VLAN
  - Broadband: DHCPv4, DHCPv6 PD
  - BFD
  - IPSec
- Network
  - IP event dampening
  - BGP & SPF optimisations
  - Multicast sub second convergence
  - GLBP, HSRP, VRRP
  - BFD for ISIS, OSPF, OSPFv3 & static v4/v6

ASR 1000 RLS 1 Software Feature Details - 2

QoS

- Ingress/egress classification, marking and action
- Egress traffic shaping
- Dual-rate/Single-rate 3-colour policer
- Bandwidth Remaining Ratio & Equal Bandwidth sharing
- MPQ - Multiple Priority Queue
- BB – Per session policing/Marking

Security

- IPsec VPN – 3DES/AES
- IOS FW/NAT – Phase I
- NBAR
- Network Foundation & Protection – Phase I

MPLS

- Basic MPLS
- MPLS VPN

SBC – Distributed DBE

- Topology Identity hiding
- DoS Protection
- Pinhole/filter control
- SIP Signaling/latching
- NAPT
- Flow-based QoS control
- DBE control interface – H.248, V4 transport, UDP, TCP, etc
- Megaco/H.248

Network Management

- Enterprise LAN Manager (LMS)
- QOS Policy Manager
- Multicast manager
- Netflow collector
- IP Solution Centre (ISC)
- Cisco Information Centre (CIC)
- Bandwidth Quality Manager (BQM)

Summary – ASR 1000 Key Innovations and Advantages

- **QFP as Forwarding Engine**
  - Flexible network services in hardware at high-speed
  - No “service-blades” needed, just a licence
- **Platform Architecture**
  - Separation of control plane and data plane
  - RP scalability to support integration of other Cisco server based s/w products
- **Modular software architecture**
  - “Touch and feel” like IOS
  - Modular software packages for separate upgrades
- **HA capability**
  - ISSU, Redundant HW system , SW Redundancy w/o additional HW