Cisco Catalyst 6500
IP Multicast Architecture and Troubleshooting

RST-3262

Session Goal

To provide you with a thorough understanding of the Catalyst® 6500 IP multicast architecture and packet flow, as well as provide key approaches and tools for troubleshooting IP multicast in the Catalyst 6500 switches

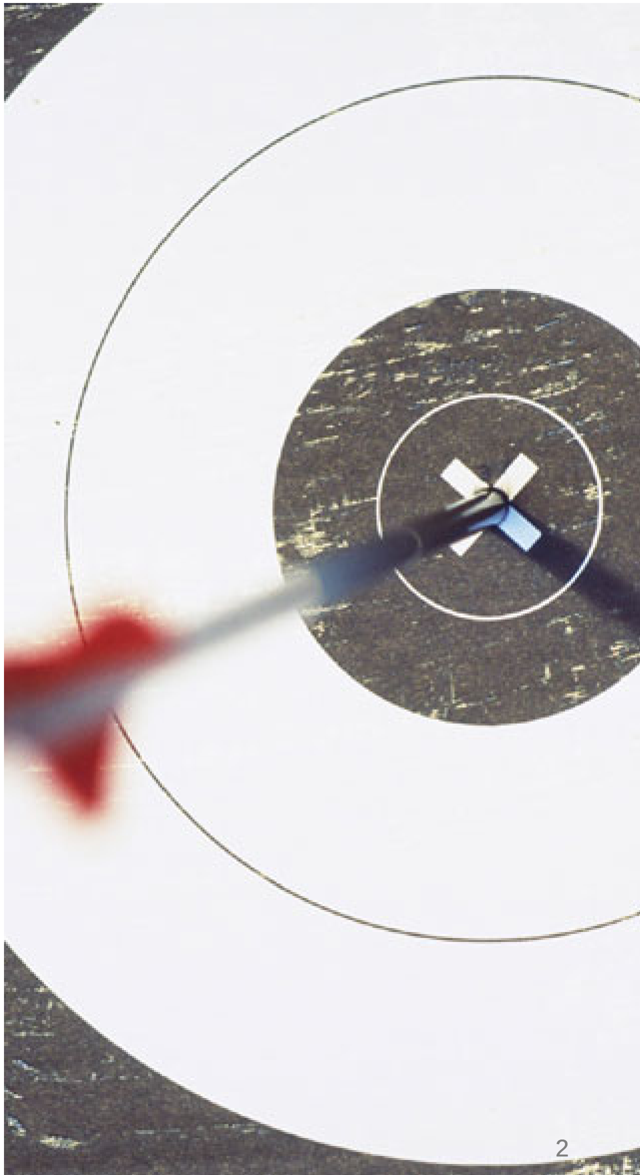

Session Assumptions

- **Working understanding of IPv4 multicast**
- **Working understanding of Catalyst 6500 platform architecture and operation**

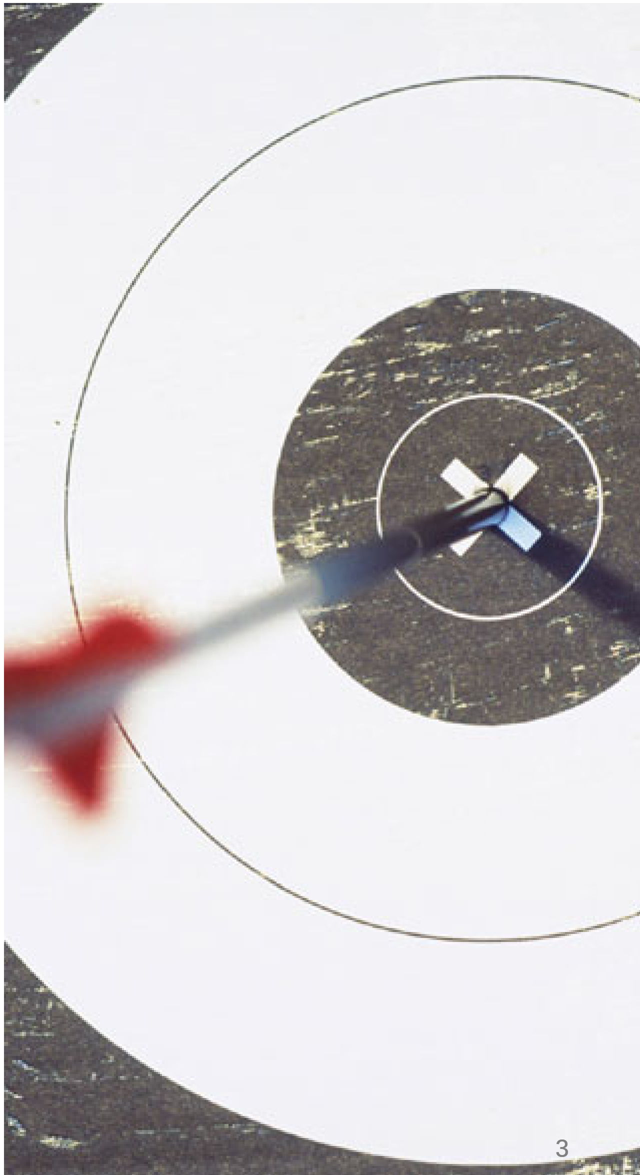

General Catalyst 6500 architecture covered in:
RST-3465: Cisco Catalyst 6500 Switch Architecture

Agenda

- IP Multicast Overview
- IP Multicast Hardware Architecture
- IP Multicast Hardware Forwarding
- IP Multicast Replication
- IP Multicast Packet Flows
- IGMP and IGMP Snooping
- Multicast Troubleshooting

IP Multicast Overview

Multicast Protocols and Components

Interdomain Multicast

- PIM-SM
- PIM-SSM
- IPv6 multicast
- PIM snooping
- mVPN
- MBGP
- MSDP

Enterprise Multicast

- IGMP
- IGMP snooping
- PIM-SM
- Bidir-PIM
- MLD
- MLD snooping
- IPv6 multicast

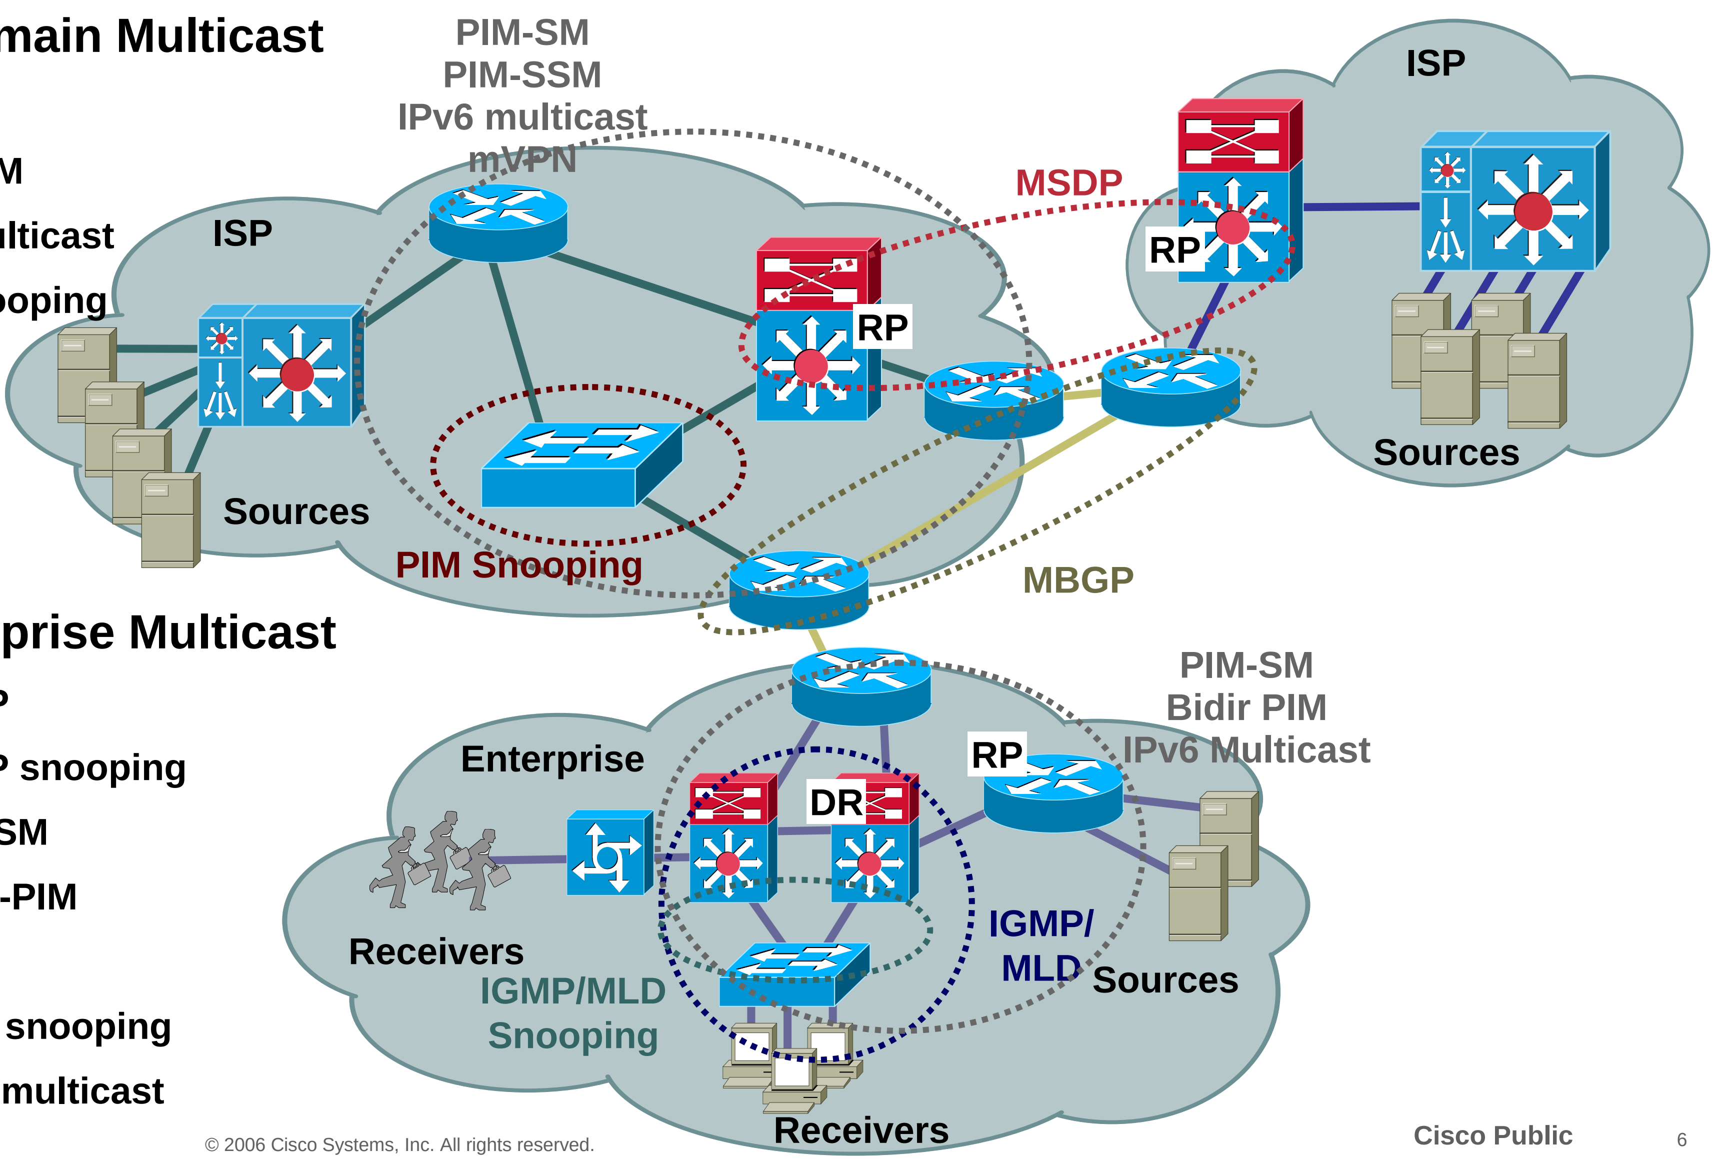

Catalyst 6500 Hardware IPv4 Multicast

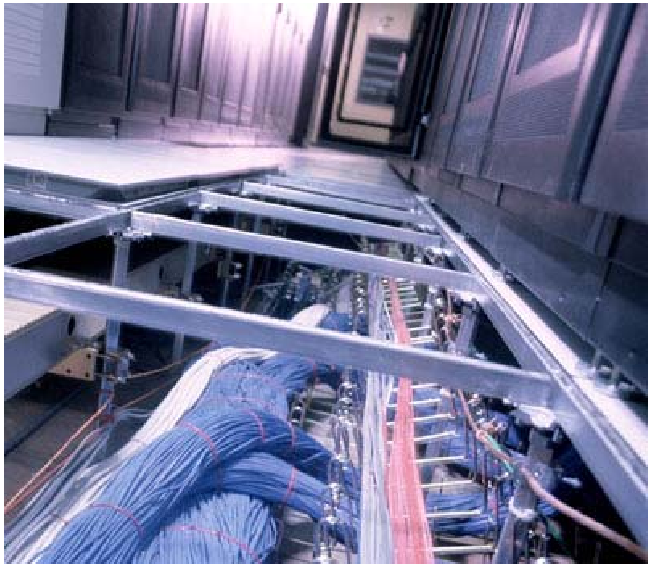

Protocols:

- **Most common:**
  - PIM-SM
  - IGMPv2
- **Emerging:**
  - PIM-SSM
  - IGMPv3
  - Bidir-PIM

Hardware:

- **Most common:**
  - Supervisor Engine 2
  - Supervisor Engine 720
  - CEF256 and CEF720 modules
- **Emerging:**
  - Supervisor Engine 32

Software:

- **Core/Distribution/DC:**
  - Cisco IOS®
- **Access/distribution:**
  - Catalyst OS/Hybrid

Catalyst 6500 IPv4 Multicast Overview

- **Implements centralized and distributed IPv4 multicast hardware switching**
  - **Off-loads majority of forwarding tasks from RP CPU**
- **Supports PIM-SM (*,G) mroute forwarding in hardware**
- **Supports PIM-SM and PIM-SSM (S,G) mroute forwarding in hardware**
- **Supervisor 720 and Supervisor 32 support Bidir (*,G) forwarding in hardware**
- **Supports IGMPv1/v2/v3 and v1/v2/v3 snooping**
- **Supervisor 2 and Supervisor 720 support distributed multicast packet replication**

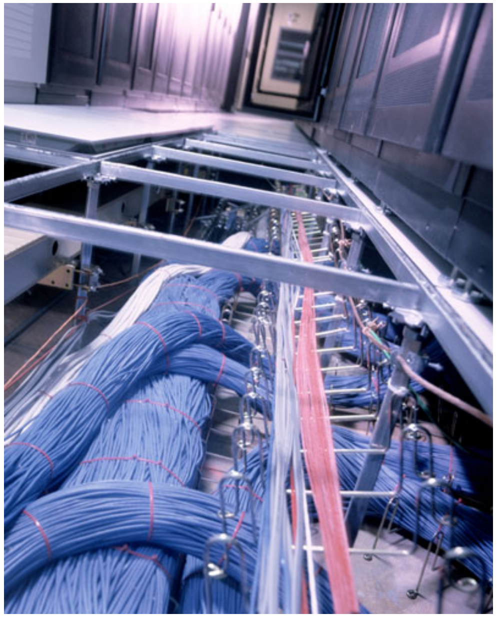

PIM Sparse Mode

PIM-SM

- General purpose multicast routing protocol
- Automatic source discovery
- Efficient packet delivery (on-demand, not flood and prune)
- Uses both shared and source-based trees
  - Distribution trees are unidirectional
- Can support arbitrary source and receiver distribution
- Group membership tracked via IGMPv1, v2, or v3
- Supported in hardware in Supervisor 2, Supervisor 32, and Supervisor 720

PIM Source-Specific Multicast

PIM-SSM

- **Simplifies one-to-many multicast delivery—uses source trees only**
- **Assumes one-to-many model**
  - Most Internet multicast fits this model
  - Video distribution also fits this model
- **Hosts responsible for source discovery—**
  - Typically via some out-of-band mechanism (web page, content server, etc.)
  - Eliminates need for RP and shared trees
  - Eliminates need for MSDP
- **Group membership tracked via IGMPv3 or combination of IGMPv2 and SSM mapping**
- **Supported in hardware in Supervisor 2, Supervisor 32, and Supervisor 720**
  - Hardware implementation of PIM-SSM virtually identical to PIM-SM

Bidirectional PIM

Bidir-PIM

- **Massively scalable—ideal for many-to-many applications**
- **Data independent—no registers, asserts, non-RPF issues**
- **Drastically reduces network mroute state**
  - Eliminates ALL (S,G) state in the network for Bidir groups
  - Shortest path trees from sources to RP eliminated
  - Source traffic flows both up and down shared RP tree
  - Permits virtually unlimited sources
- **Group membership tracked via IGMPv1, v2, or v3**
- **Hardware support on Supervisor 32 and Supervisor 720**
  - Somewhat different hardware implementation for Bidir versus PIM-SM or SSM
  - Support for up to four Bidir RPs per VRF in hardware, using PIM RP-to-DF interface mapping table

IGMP and IGMP Snooping

- **IGMP support through Cisco IOS software**
  - IGMP v1/v2/v3 protocol support for PIM-SM and Bidir PIM
  - IGMP v3 protocol support for PIM-SSM
  - Option for SSM mapping to translate IGMPv2 joins to PIM-SSM joins
- **IGMP snooping support leveraging both hardware and software**
  - Snooping support for all IGMP versions
  - PFC performs hardware redirection of IGMP packets to SP CPU for analysis

Agenda

- IP Multicast Overview
- IP Multicast Hardware Architecture
- IP Multicast Hardware Forwarding
- IP Multicast Replication
- IP Multicast Packet Flows
- IGMP and IGMP Snooping
- Multicast Troubleshooting

IP Multicast Hardware Architecture

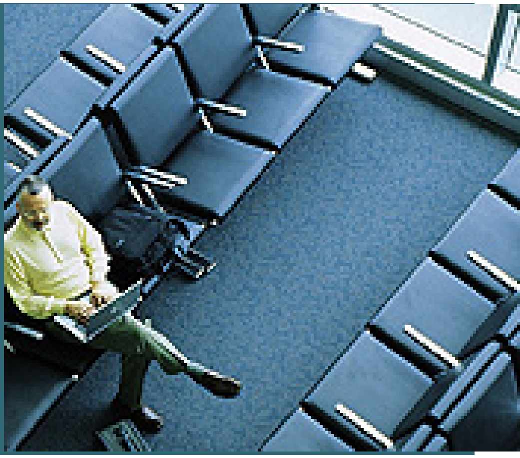

Supervisor Engine Multicast Components

Key supervisor engine components that relate to IP multicast forwarding:

- **Route Processor (RP) and Switch Processor (SP) CPUs**
- **PFC Daughter Card**
- **Multicast Replication Engine**

Supervisor 2 Multicast Architecture

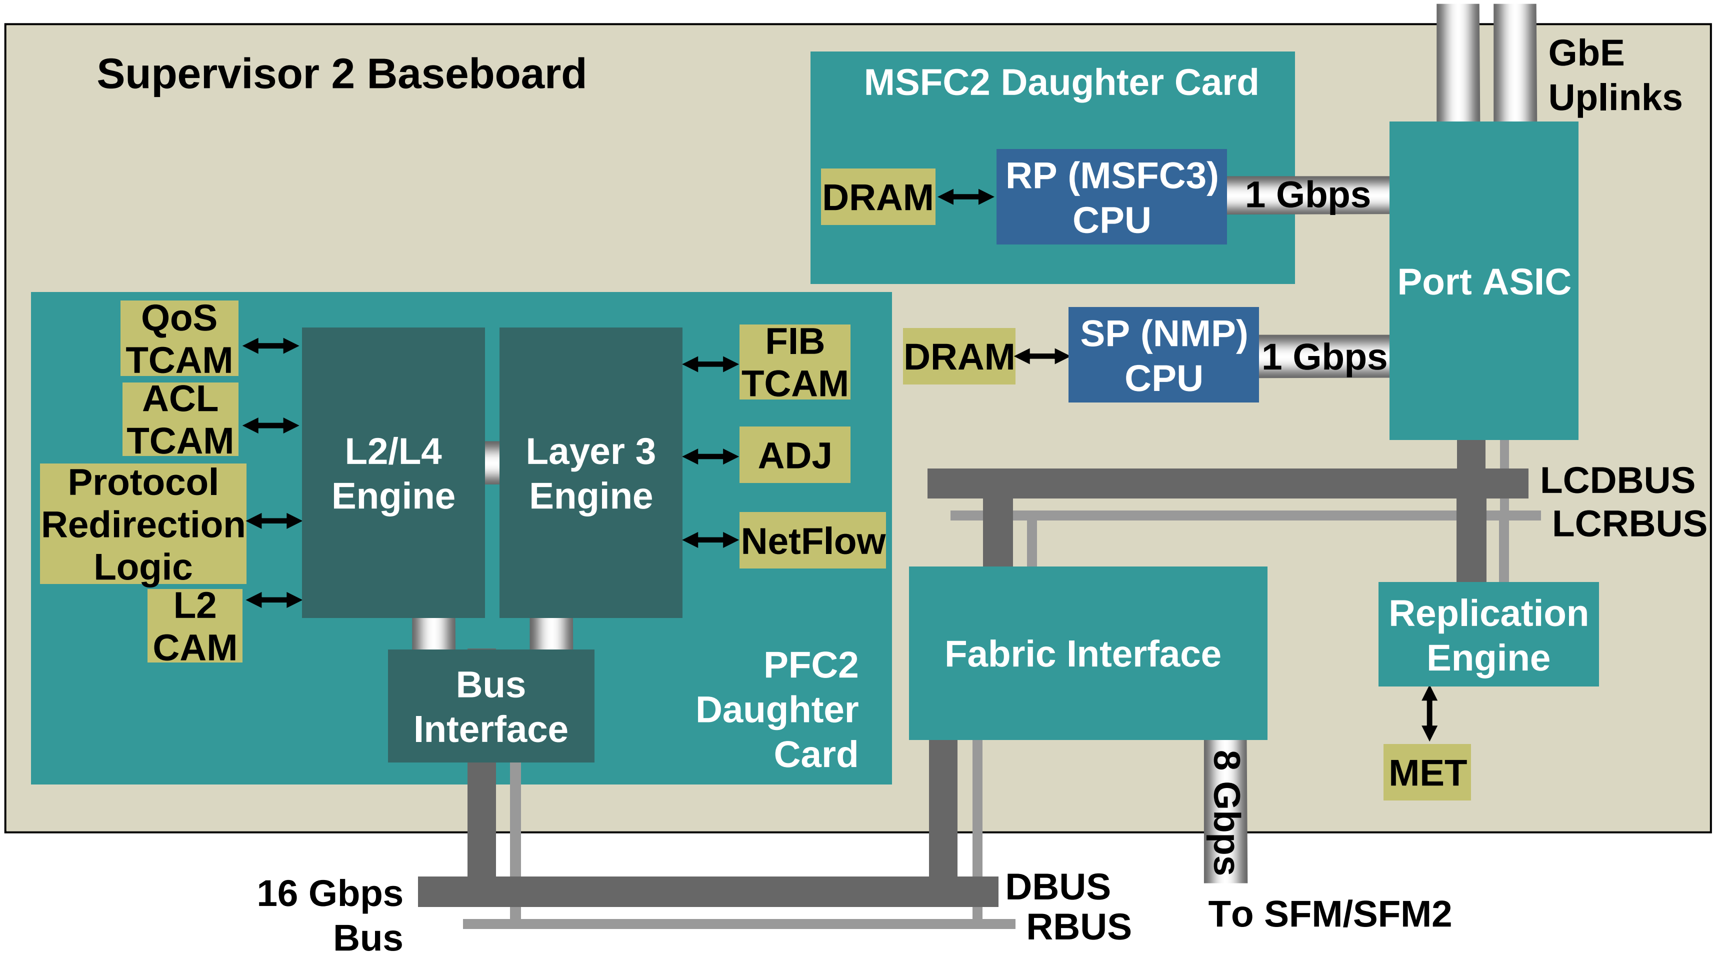

Supervisor 2 Multicast Architecture

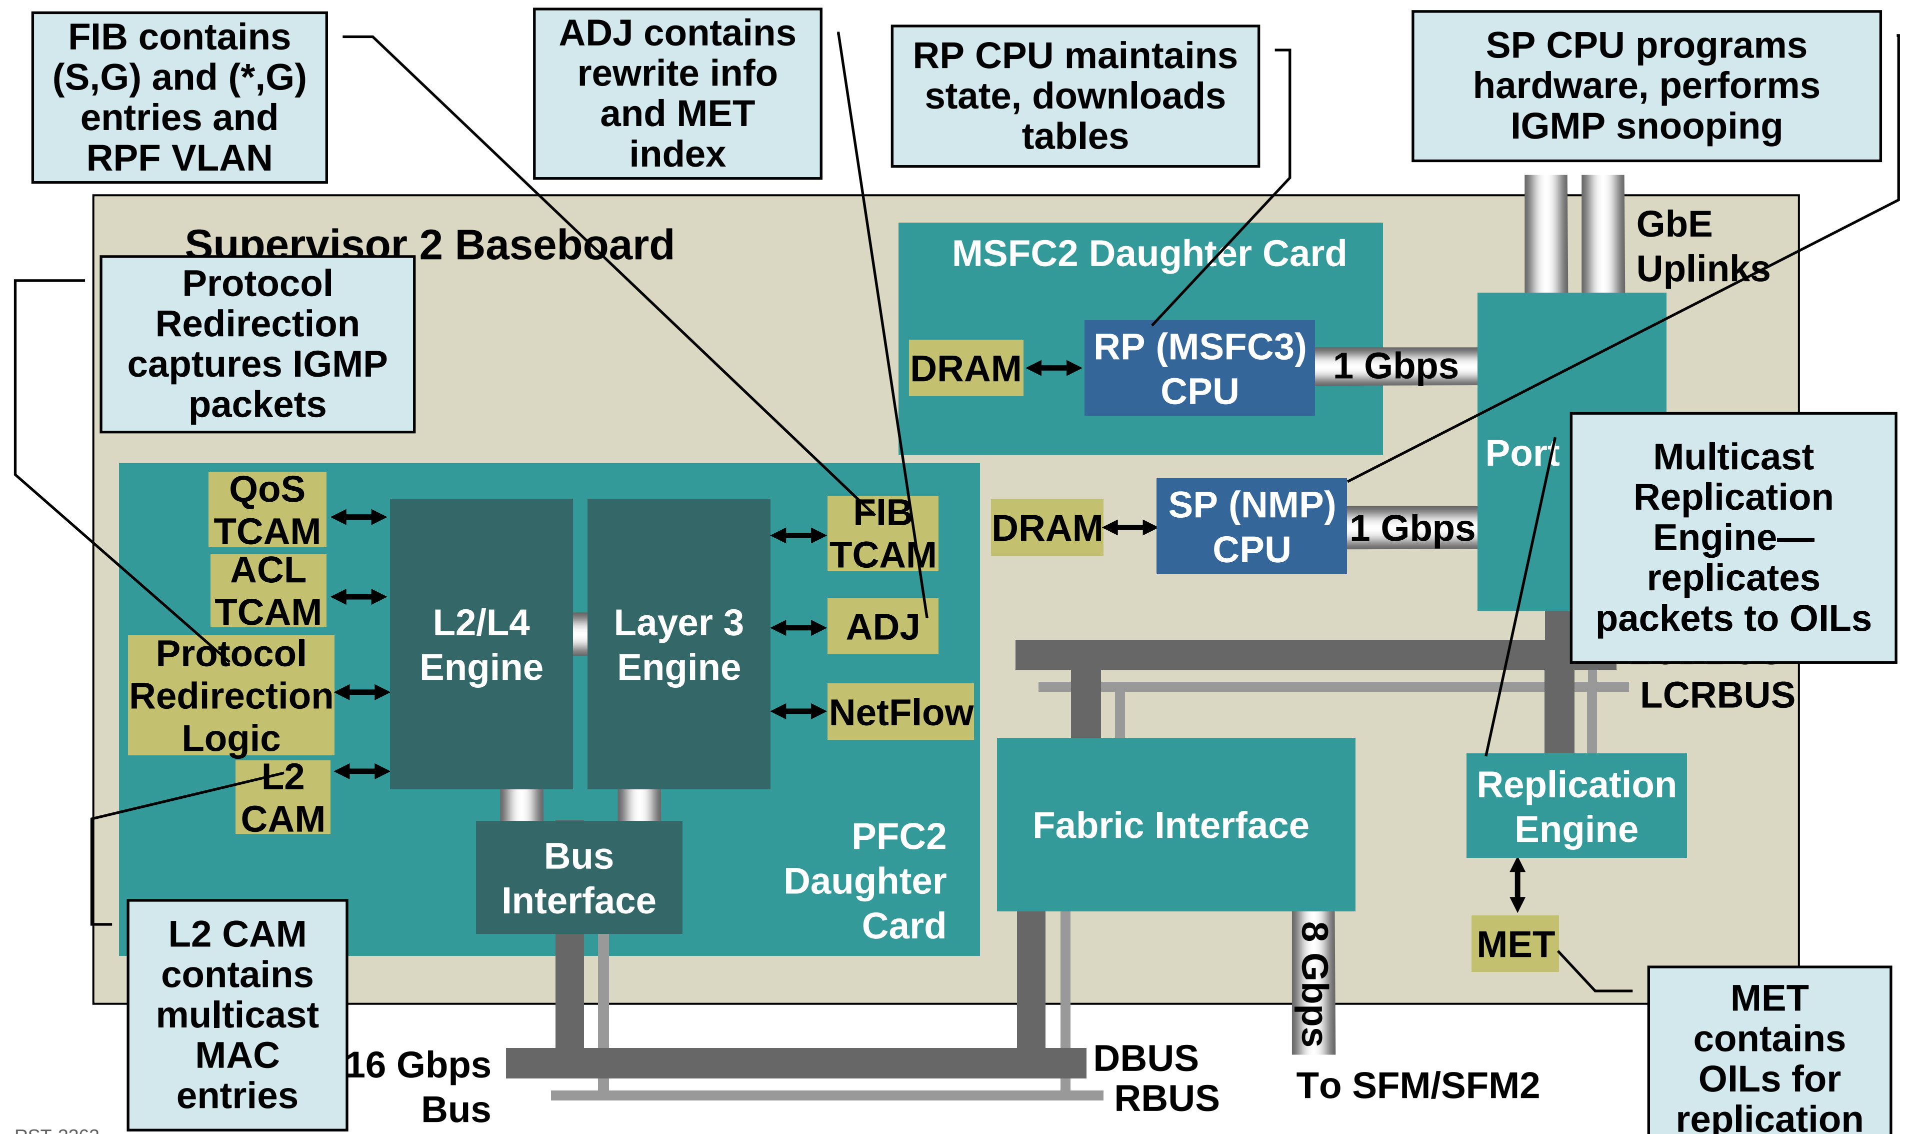

Supervisor 720 Multicast Architecture

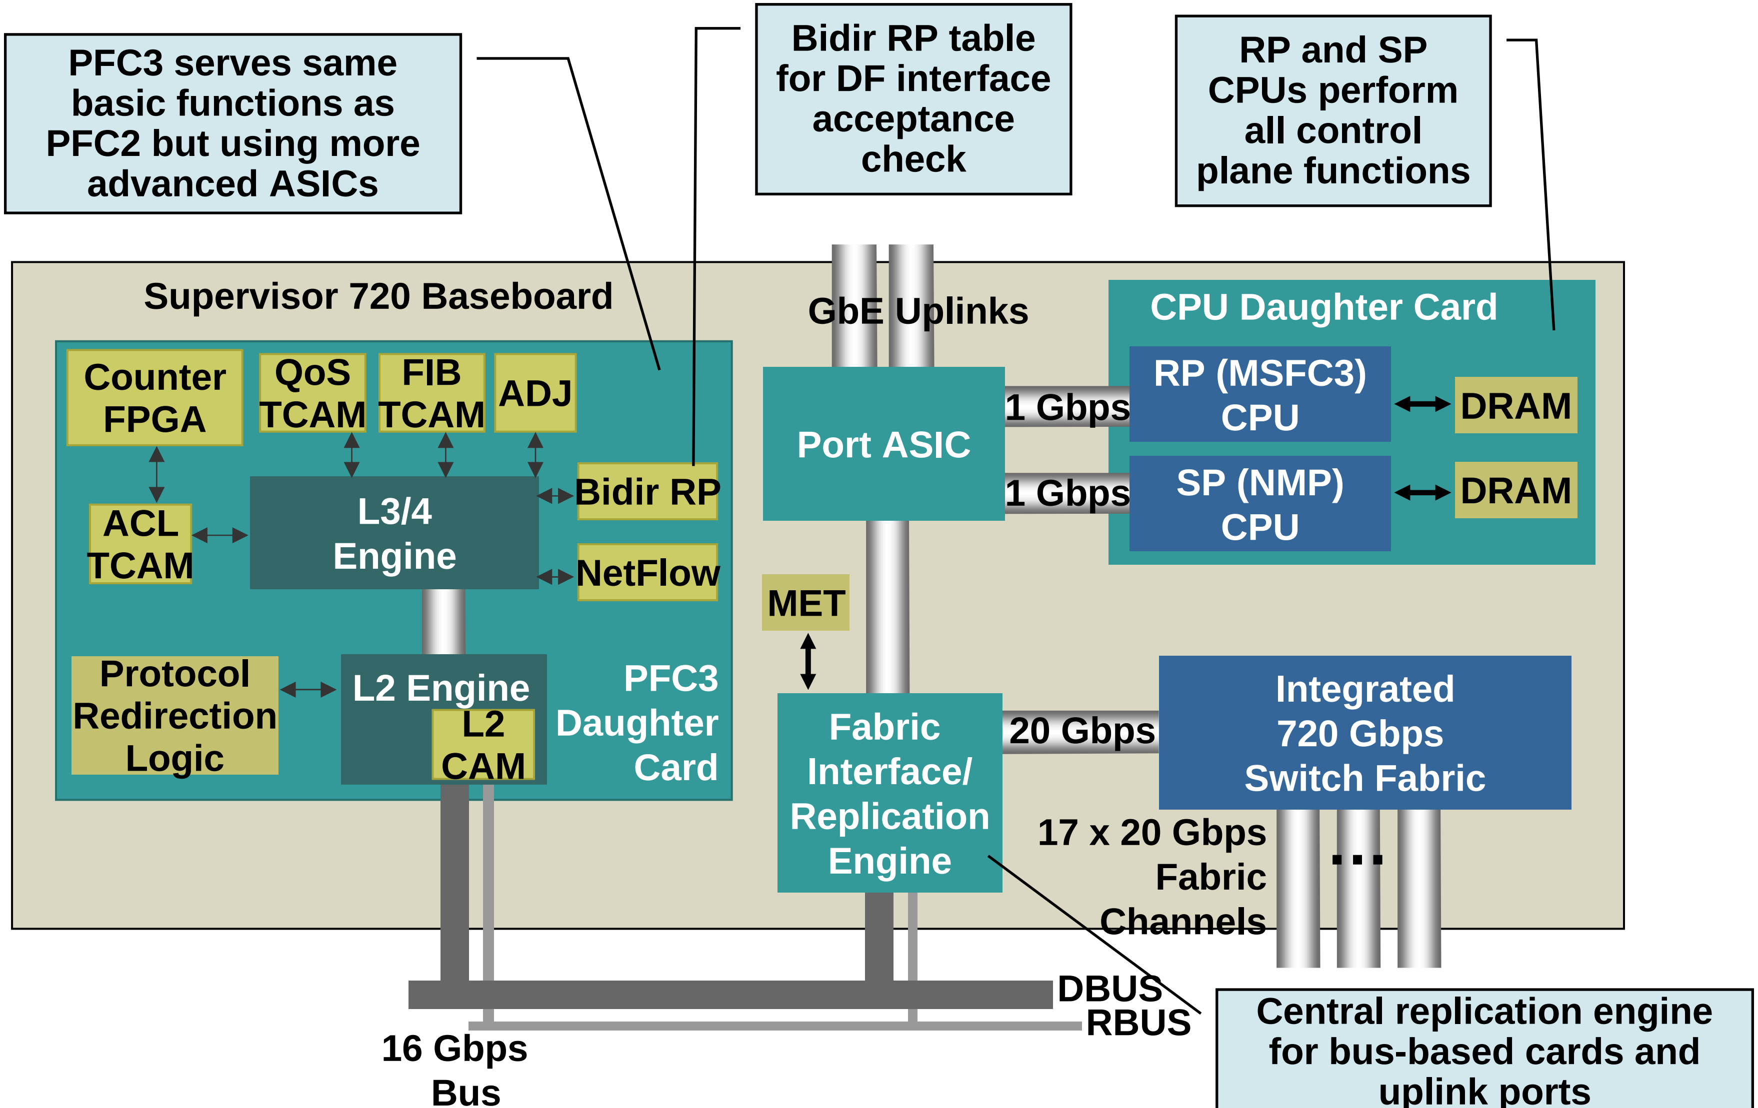

Supervisor 32 Multicast Architecture

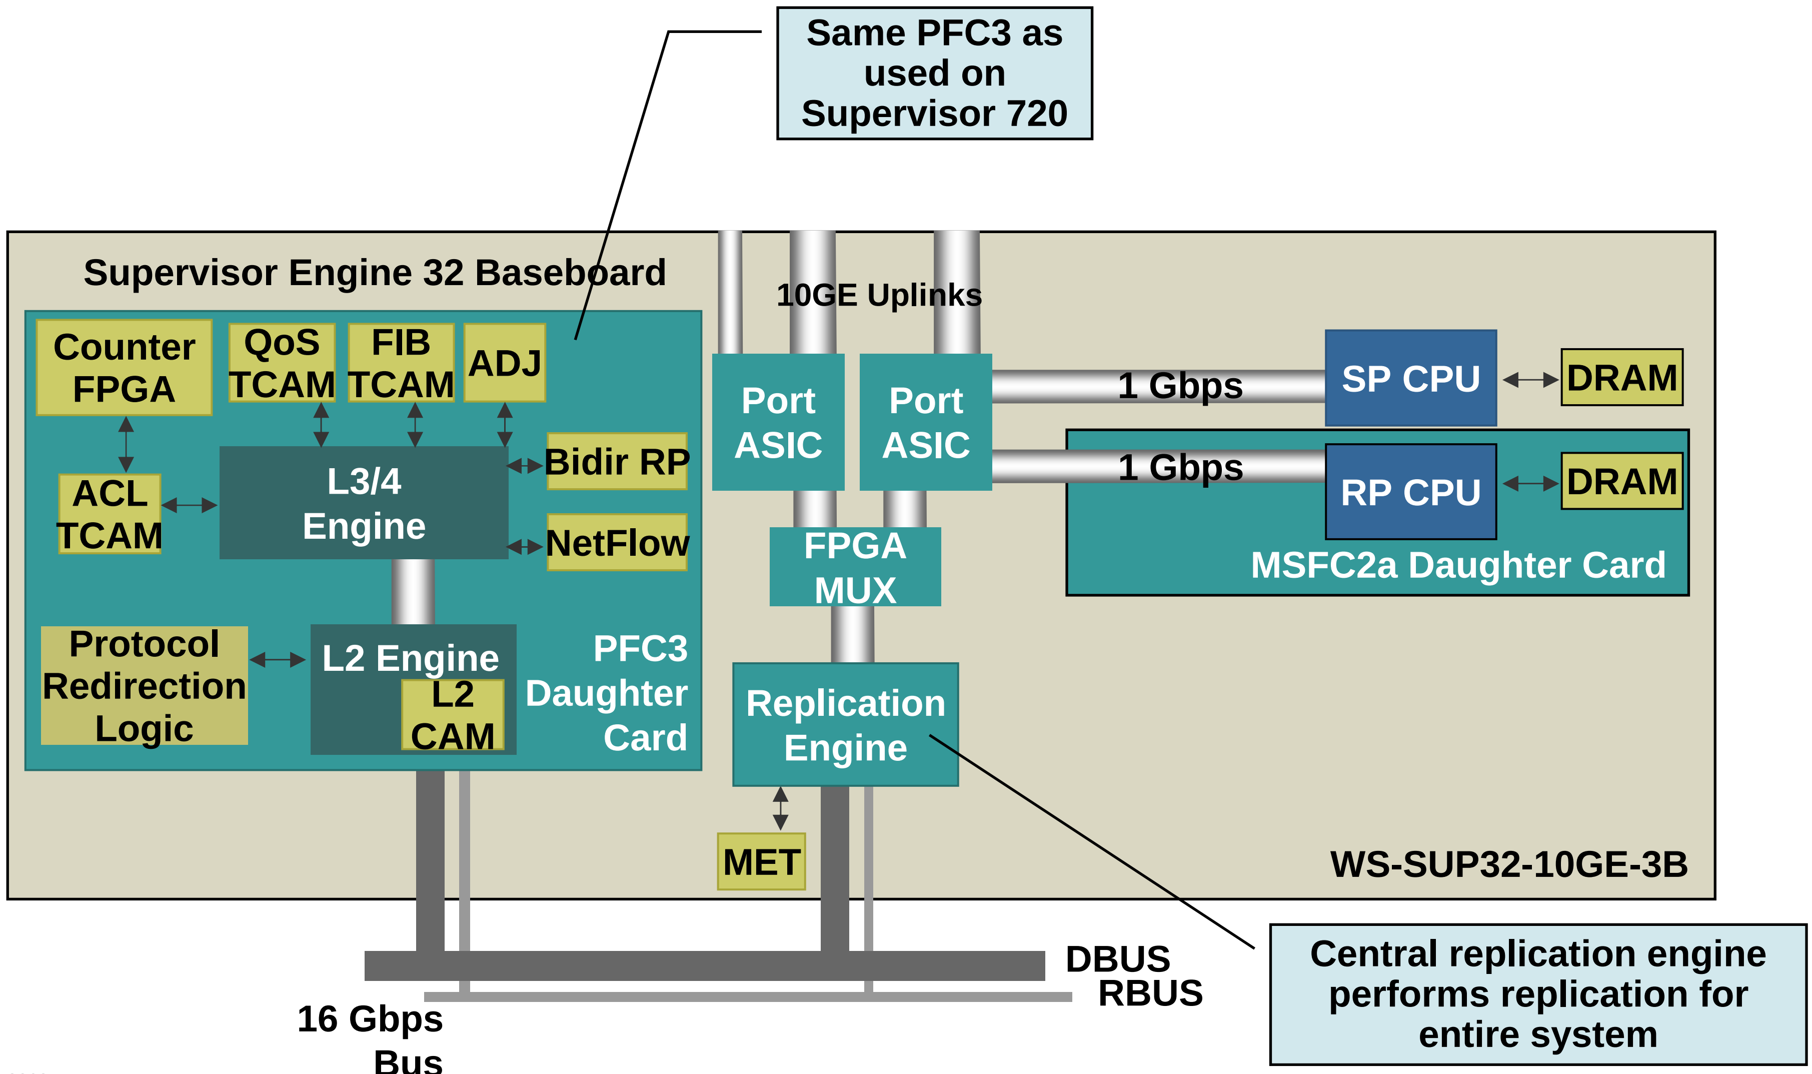

Switching Module Multicast Components

Key components of switching modules that relate to IP multicast forwarding:

- **Multicast Replication Engine**
- **DFC Daughter Card**

Classic Module Architecture

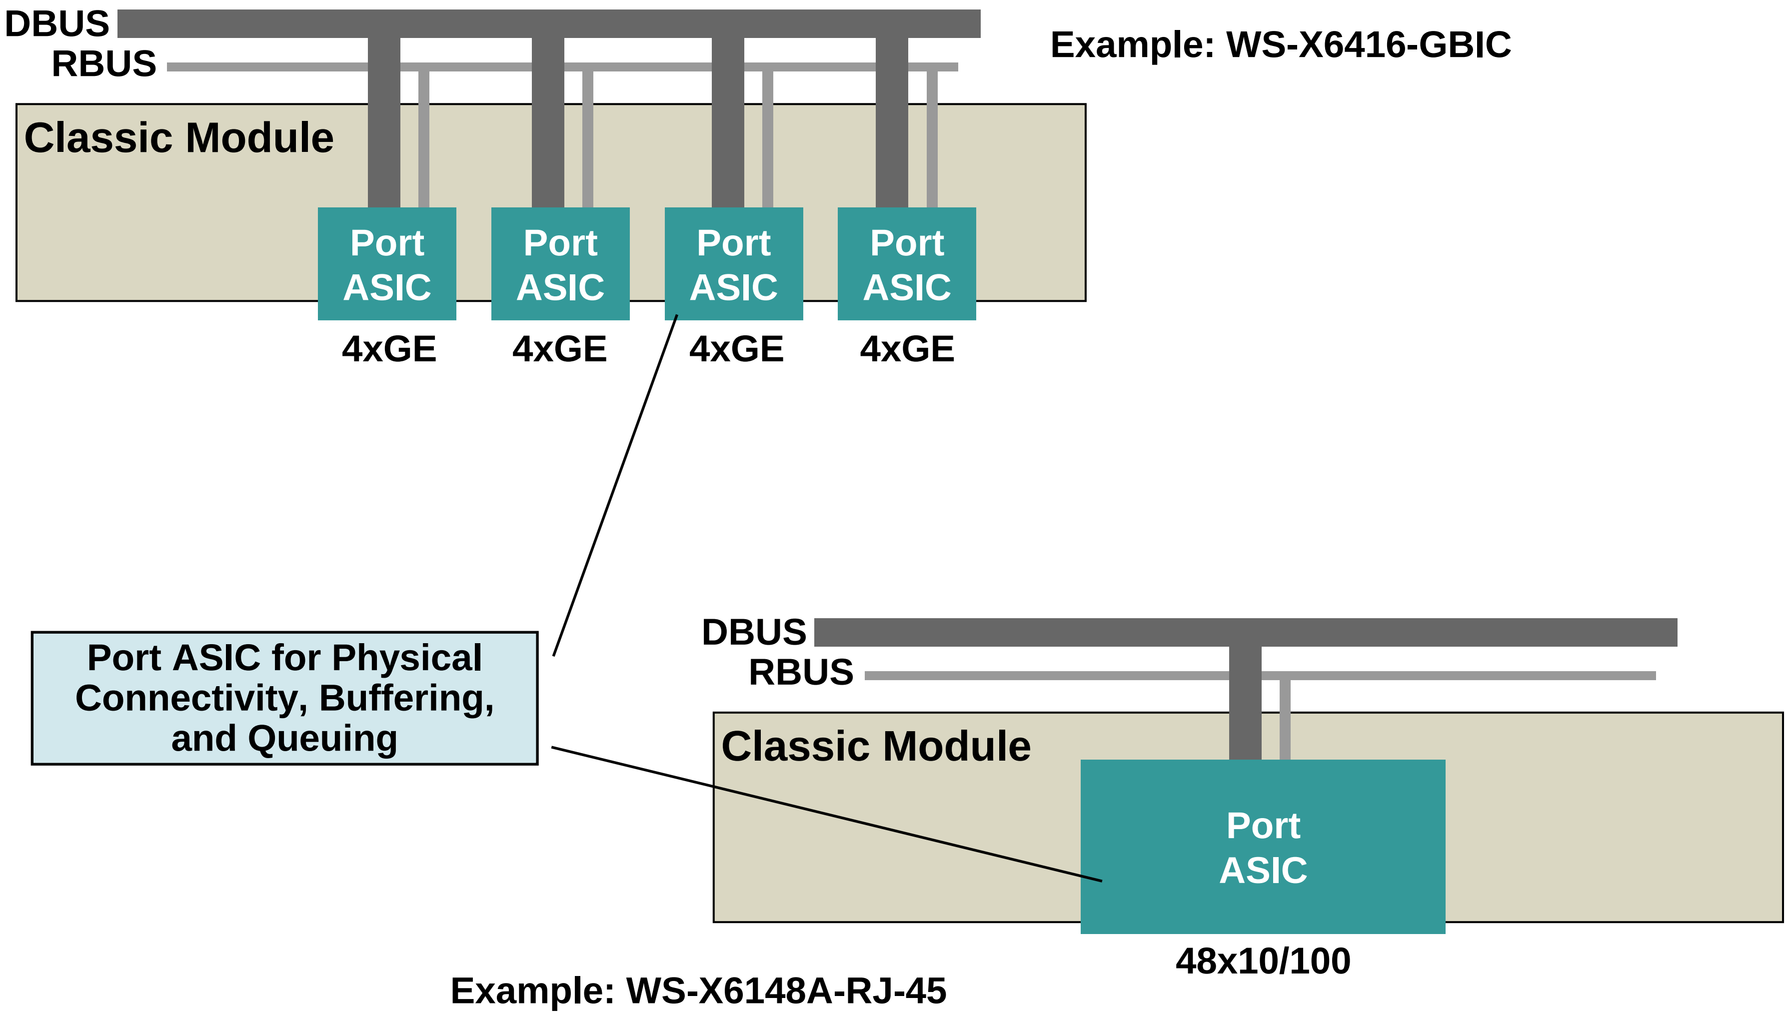

CEF256 Module Architecture

Example: WS-X6516-GBIC

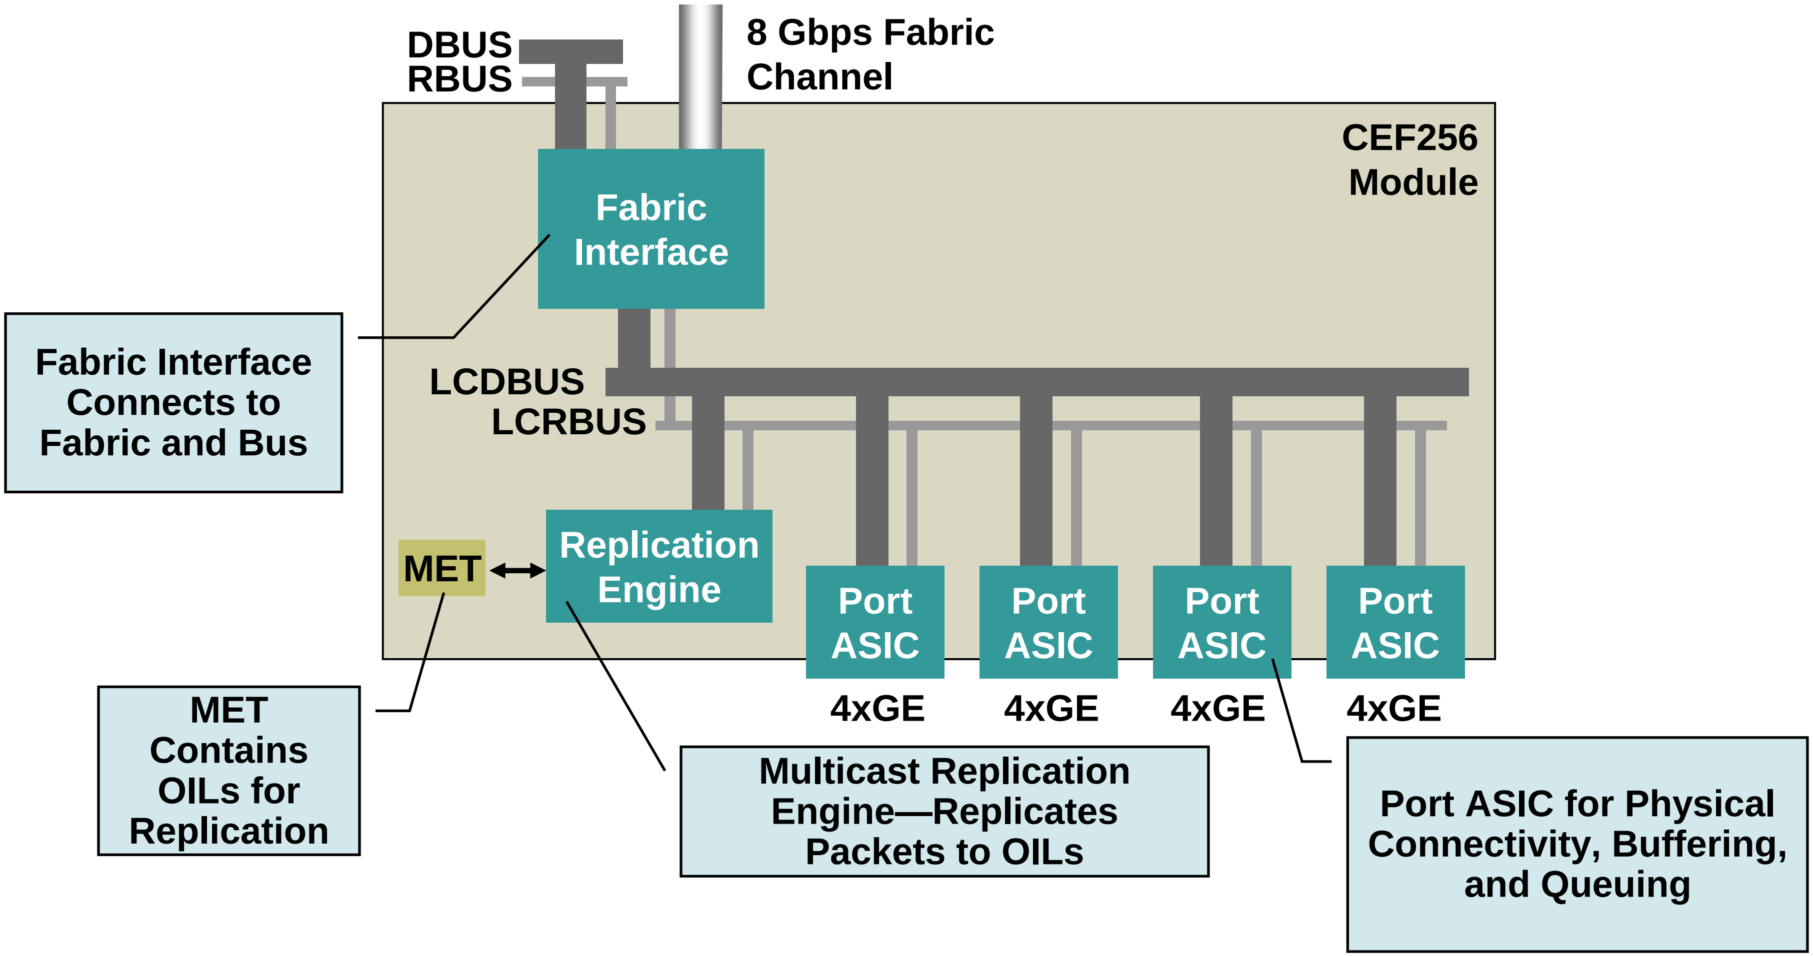

CEF256 Module Architecture—with DFC

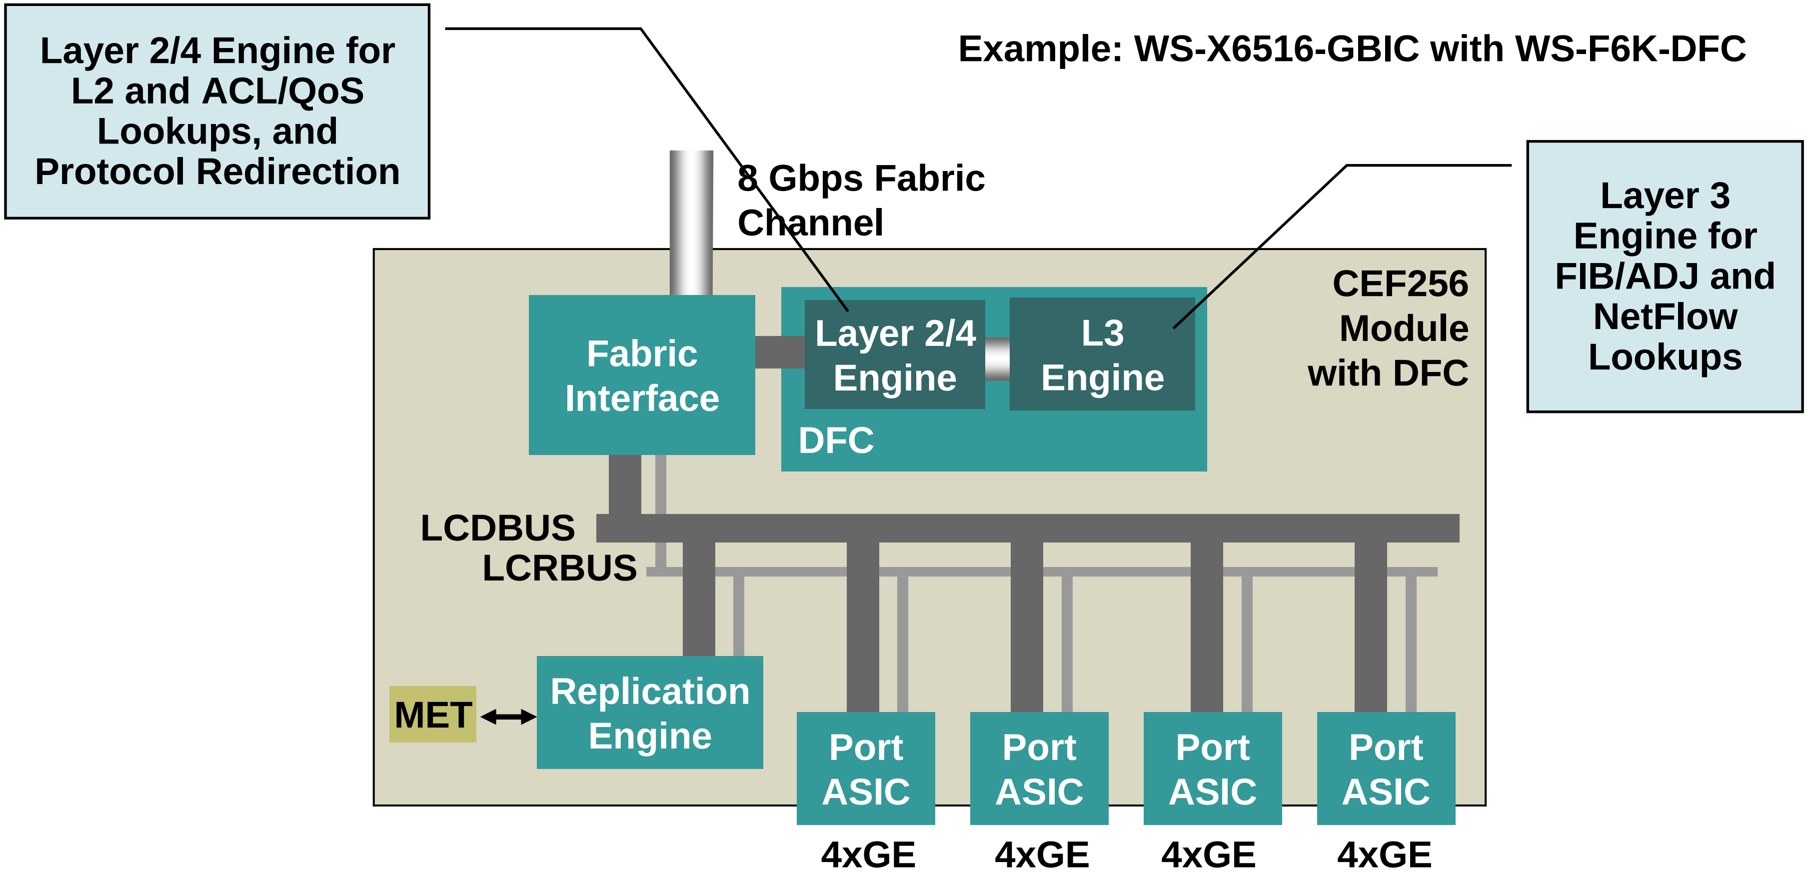

CEF720 Module Architecture

Example: WS-X6748-SFP

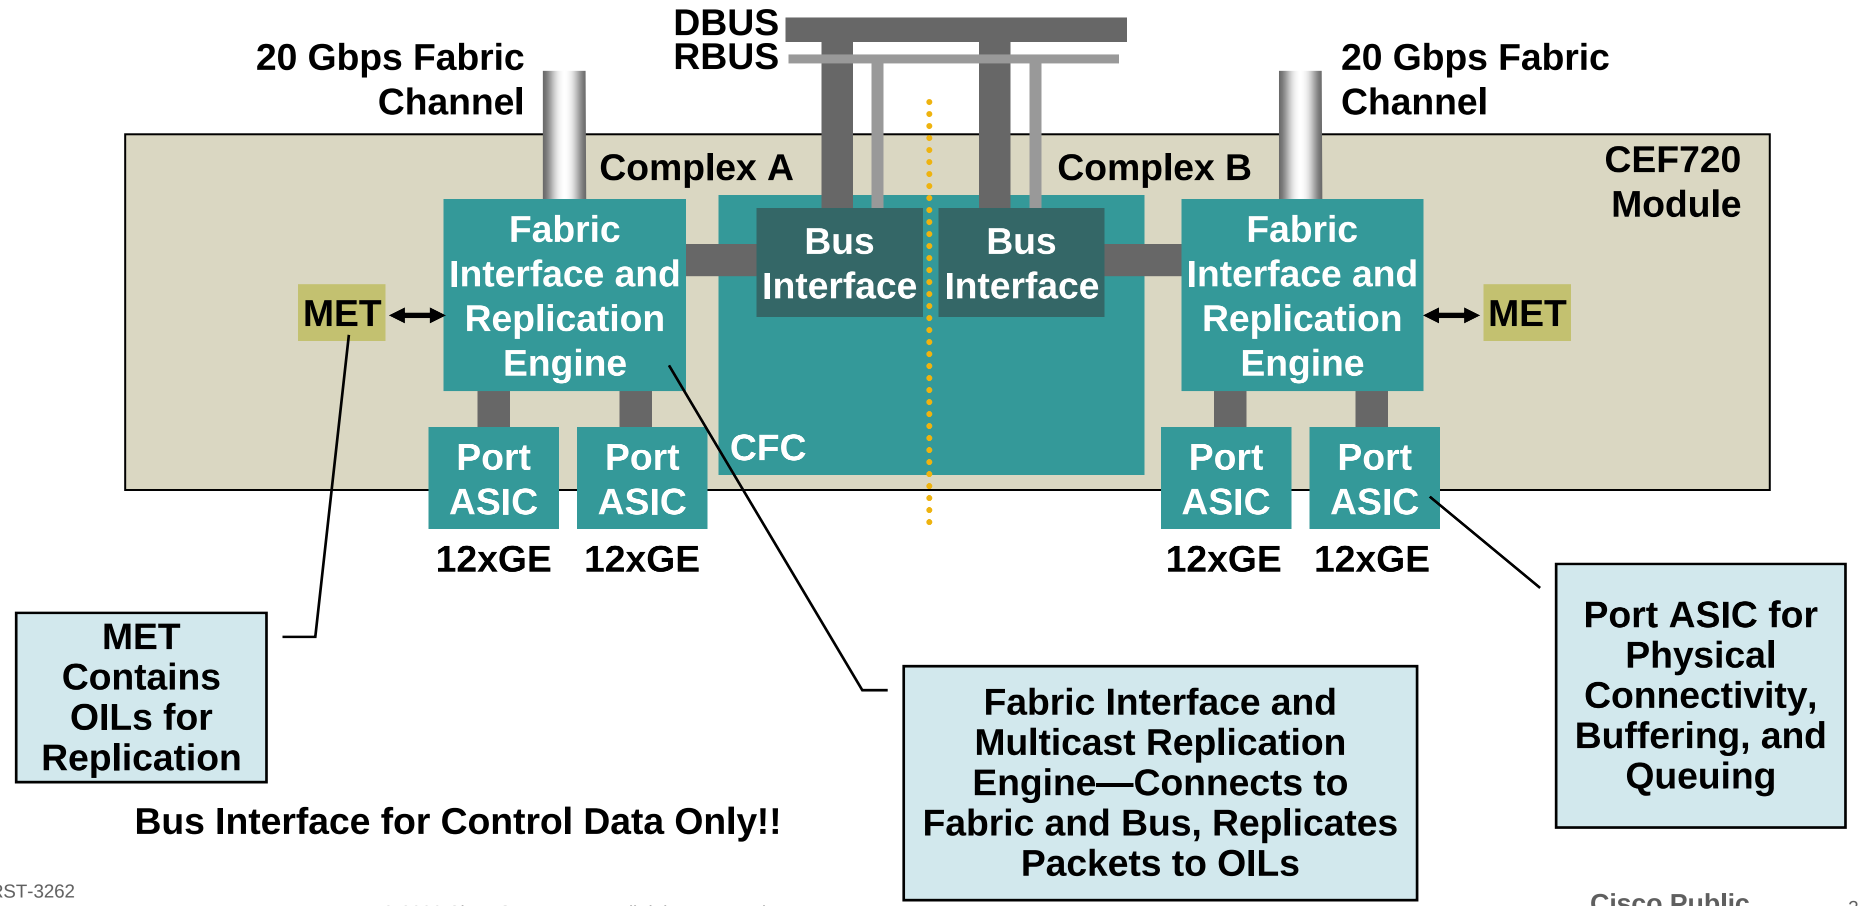

Bus Interface for Control Data Only!!

CEF720 Module Architecture—with DFC3

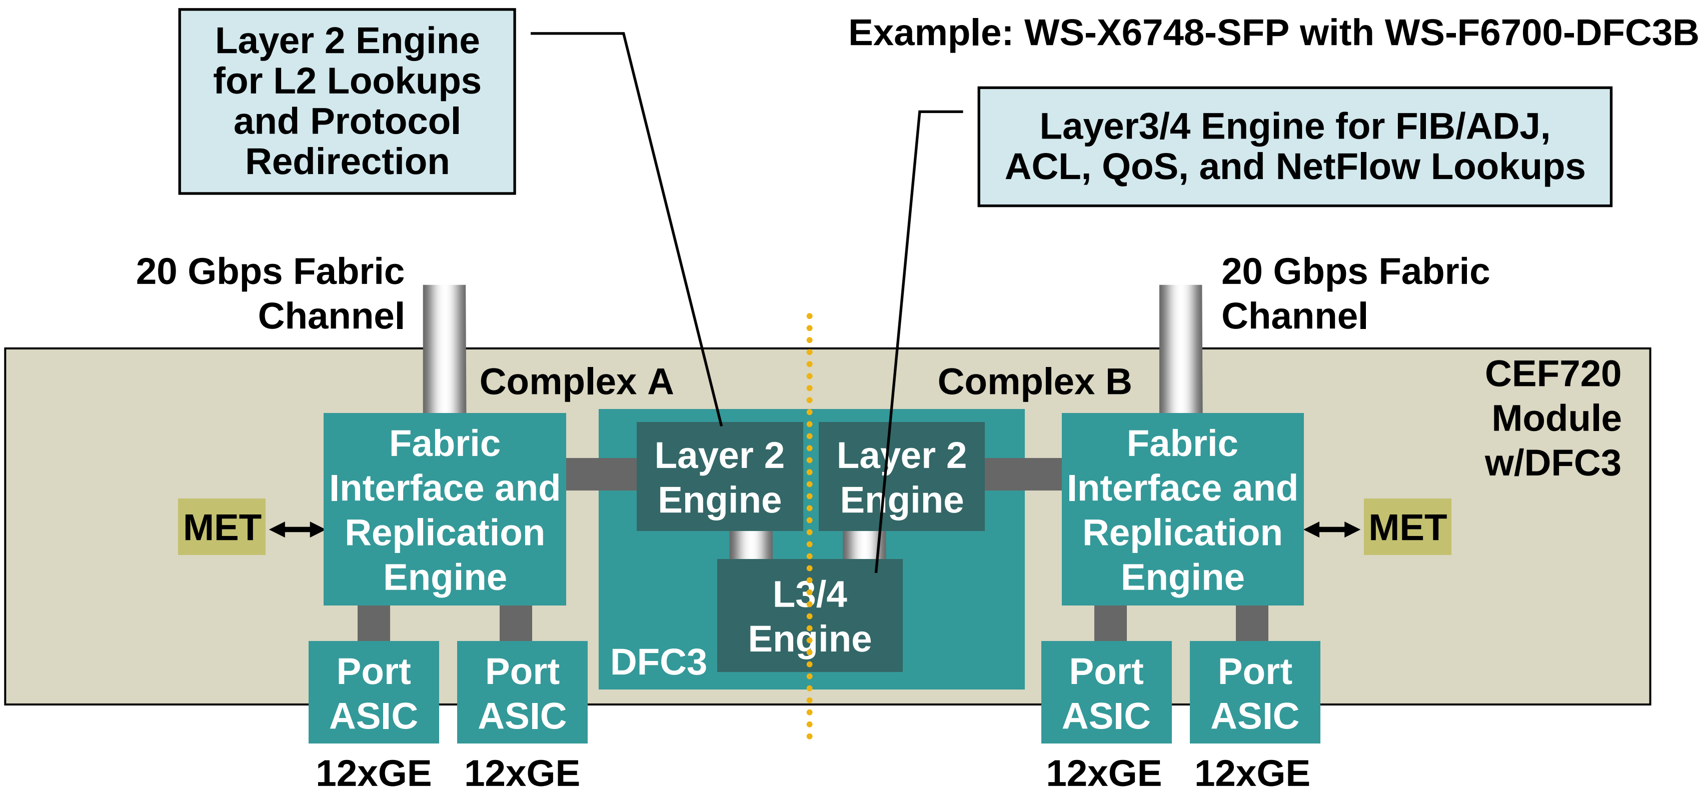

Agenda

- IP Multicast Overview
- IP Multicast Hardware Architecture
- IP Multicast Hardware Forwarding
- IP Multicast Replication
- IP Multicast Packet Flows
- IGMP and IGMP Snooping
- Multicast Troubleshooting

IP Multicast Hardware Forwarding

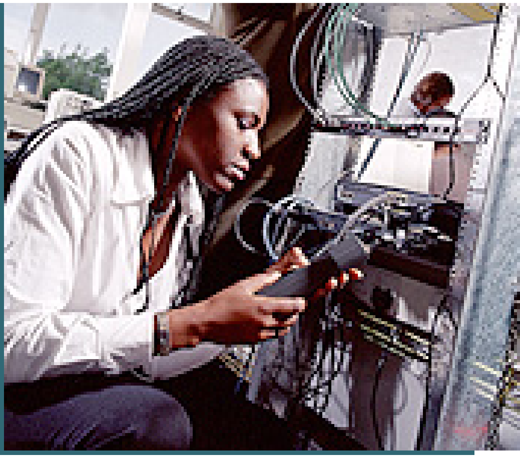

Multicast Control Plane

- **RP CPU multicast control plane functions:**
  - Control plane protocols—PIM, IGMP, AutoRP, BSR, MSDP, routing protocols
  - Calculating RPF interfaces
  - Managing software IP mroute table
  - Downloading IP mroute table entries to SP for installation in the PFC hardware
- **SP CPU multicast control plane functions:**
  - Managing PFC hardware tables
  - IGMP snooping packet processing
  - PIM snooping/RGMP packet processing
  - IGMP Querier function

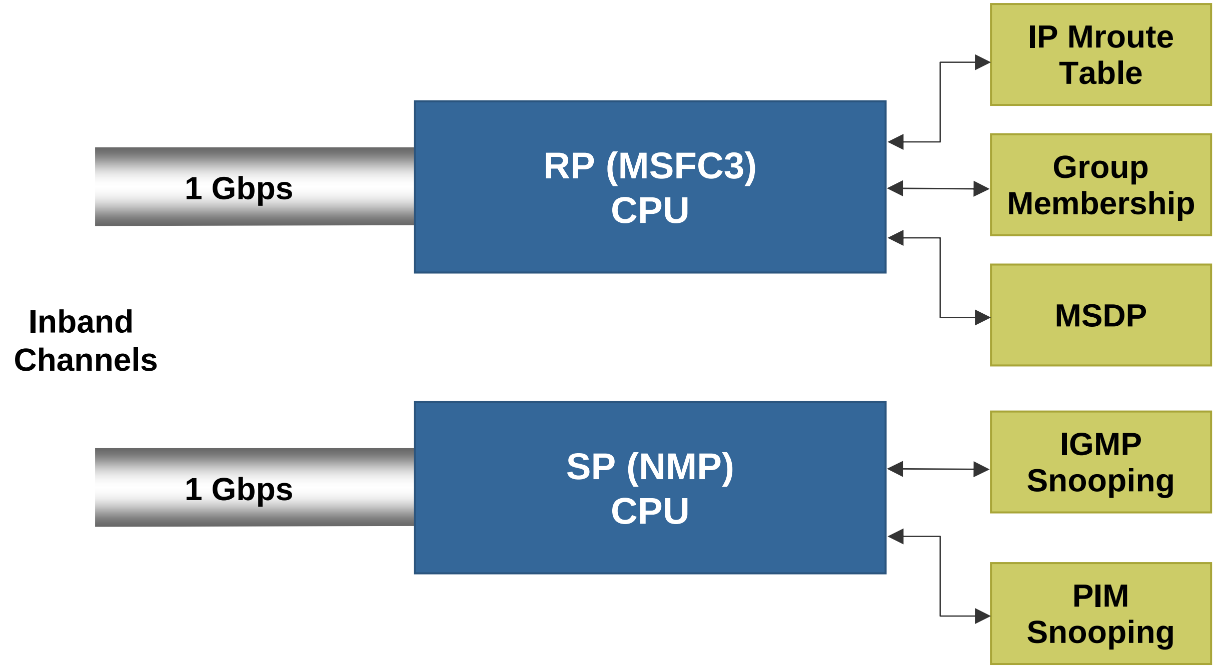

Hardware Multicast Switching

- **RP CPU derives 3 key data structures from multicast routing table**
  - **Multicast FIB—Consists of (S,G) and (*,G) entries, and RPF VLAN or Bidir-PIM RP index**
  - **Adjacency table—Contains rewrite MAC and MET index**
  - **Multicast Expansion Table (MET)—Contains output interface lists (OILs), i.e., lists of interfaces requiring replication**
- **RP CPU downloads tables to SP CPU**
- **SP CPU installs tables in the appropriate hardware**
  - **Multicast FIB and adjacency tables installed in PFC/DFC hardware**
  - **MET installed in replication engines**
- **SP CPU also maintains L2 table for IGMP/PIM snooping**

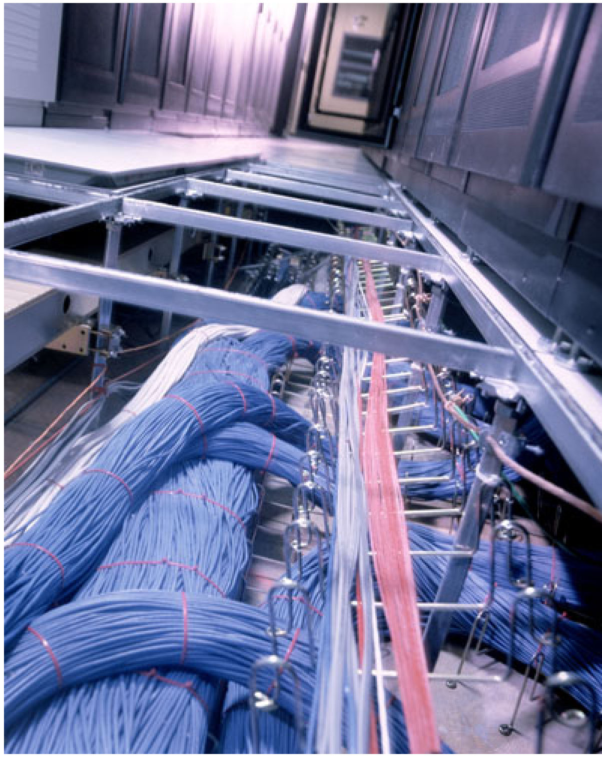

Multicast FIB TCAM Lookup

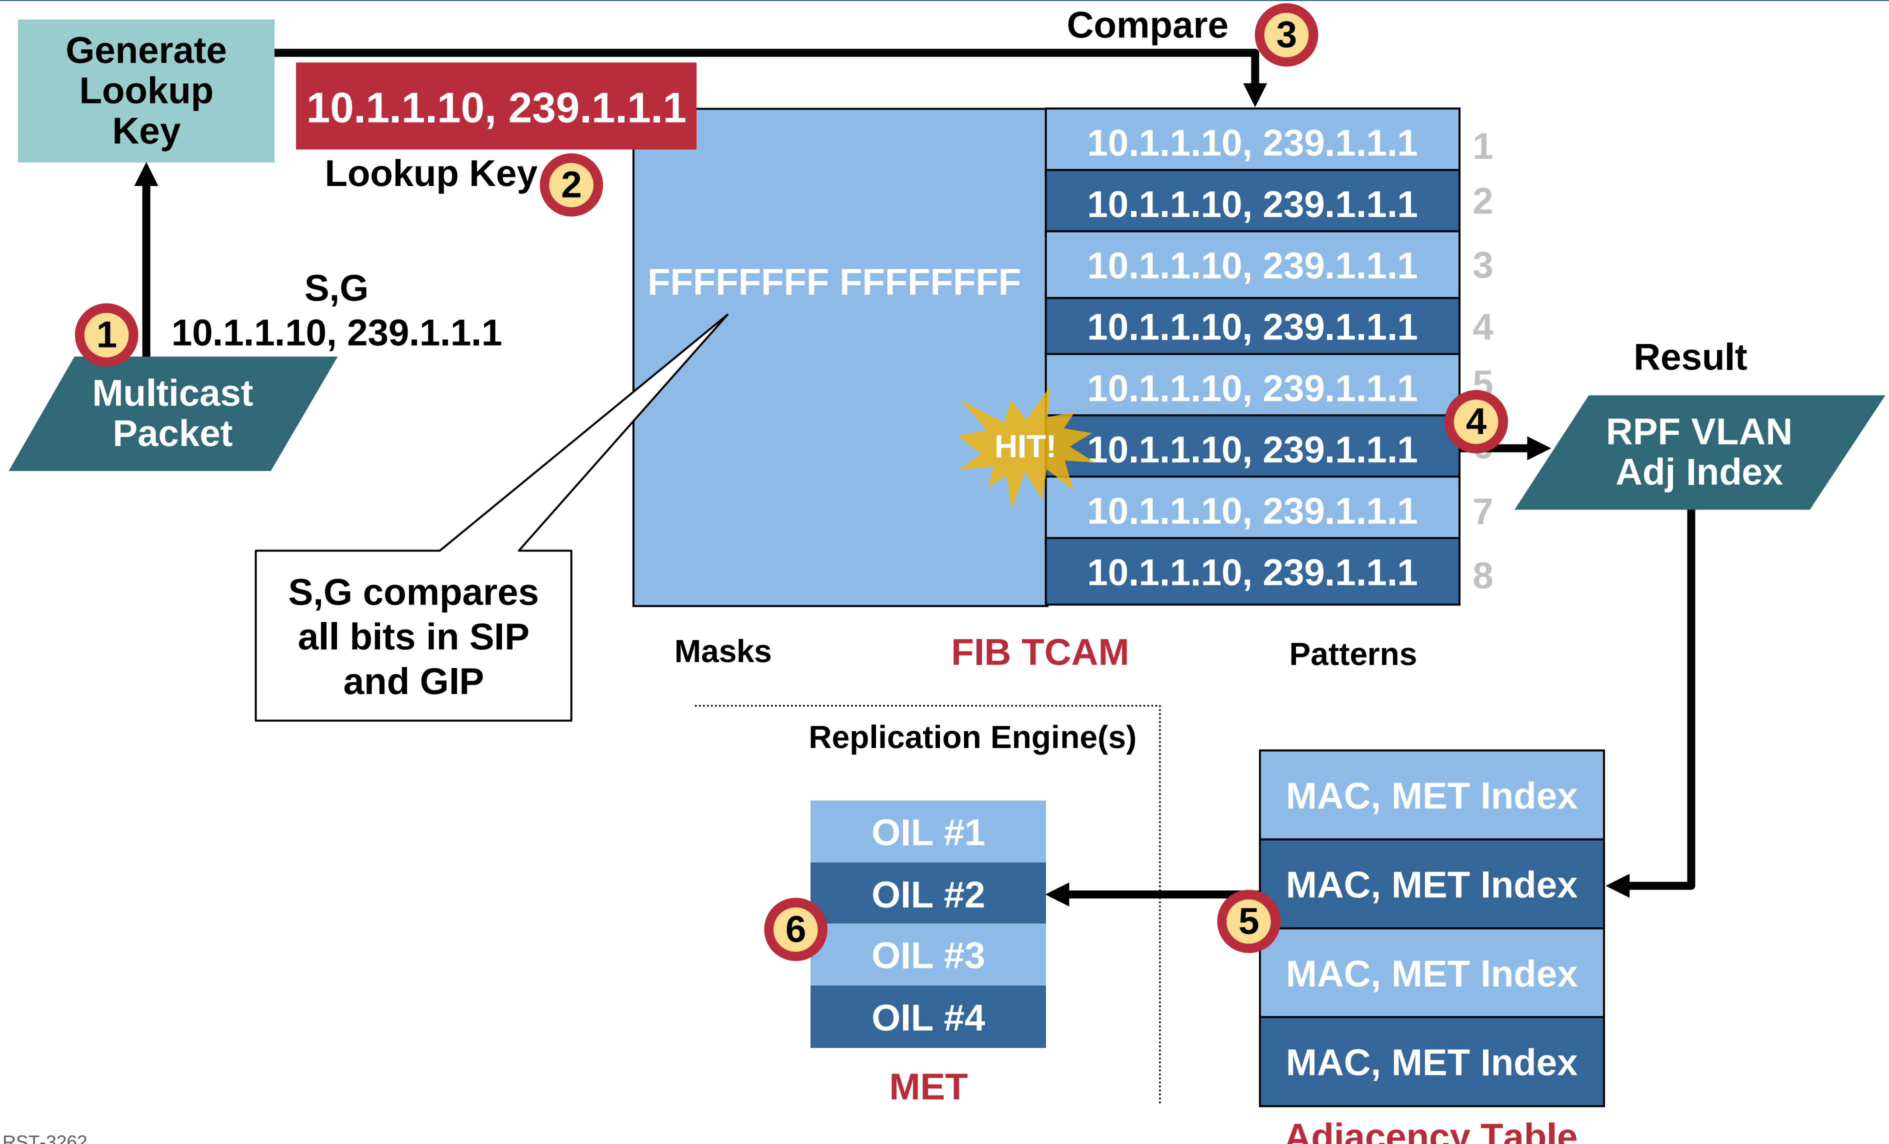

Multicast Expansion Table (MET)

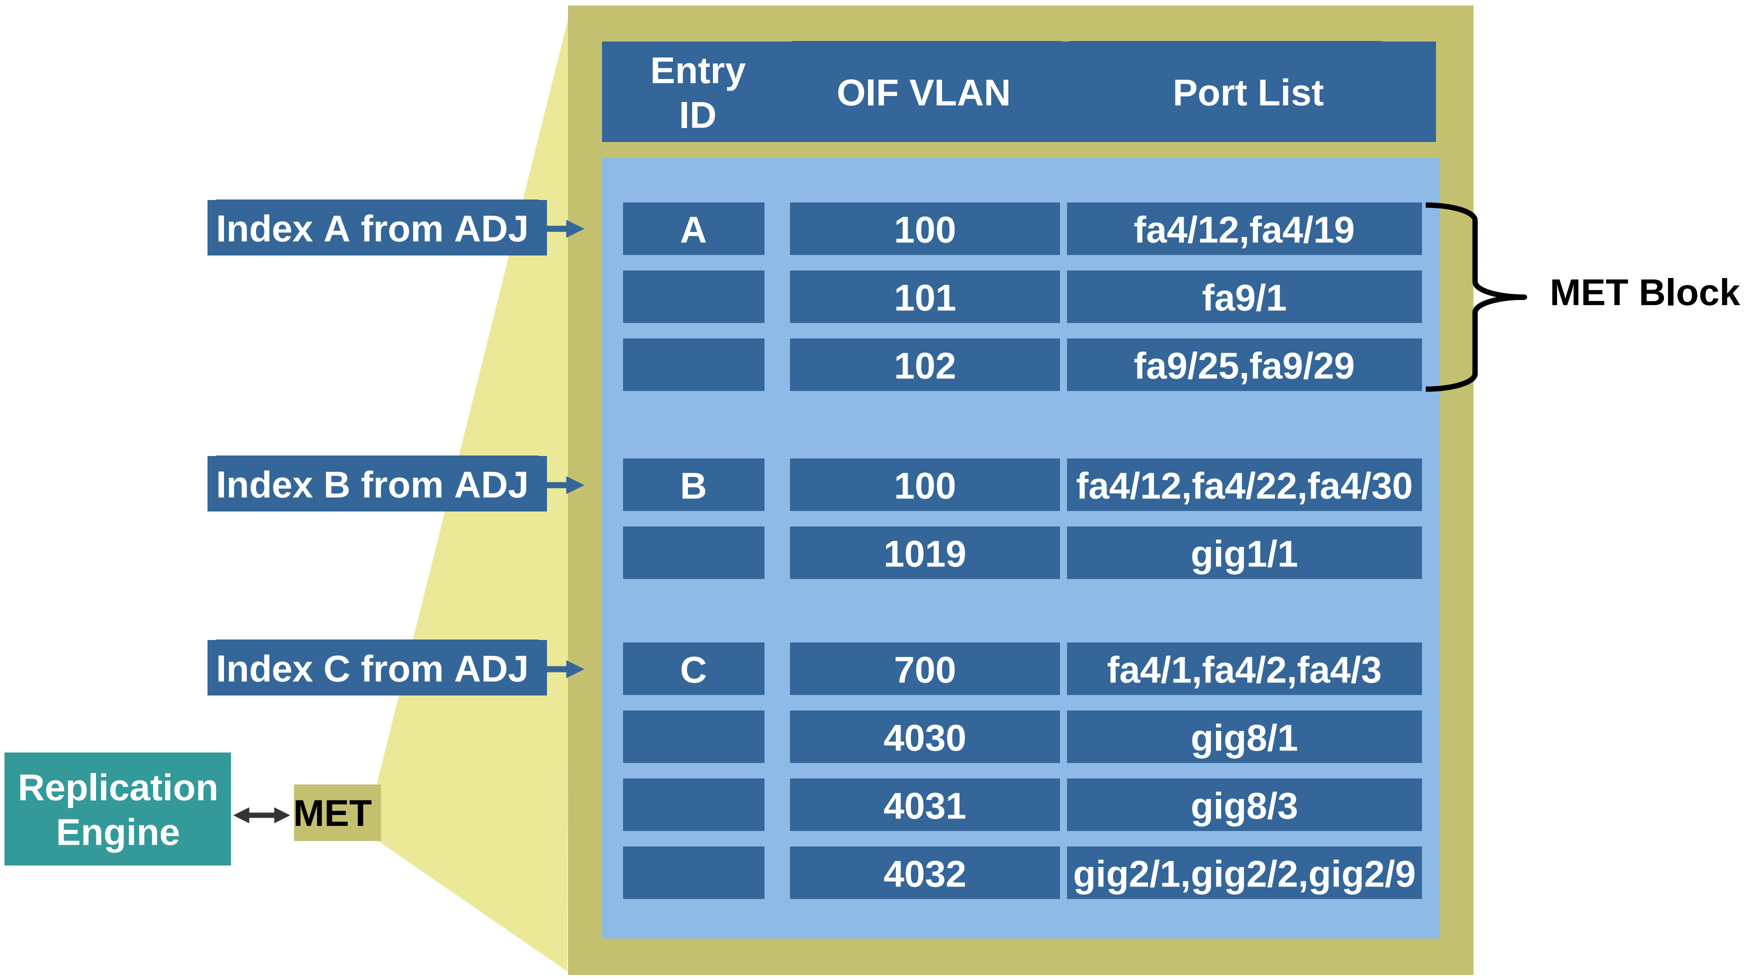

Entry ID	OIF VLAN	Port List
A	100	fa4/12,fa4/19
	101	fa9/1
	102	fa9/25,fa9/29
B	100	fa4/12,fa4/22,fa4/30
	1019	gig1/1
C	700	fa4/1,fa4/2,fa4/3
	4030	gig8/1
	4031	gig8/3
	4032	gig2/1,gig2/2,gig2/9

Bidir FIB TCAM Lookup

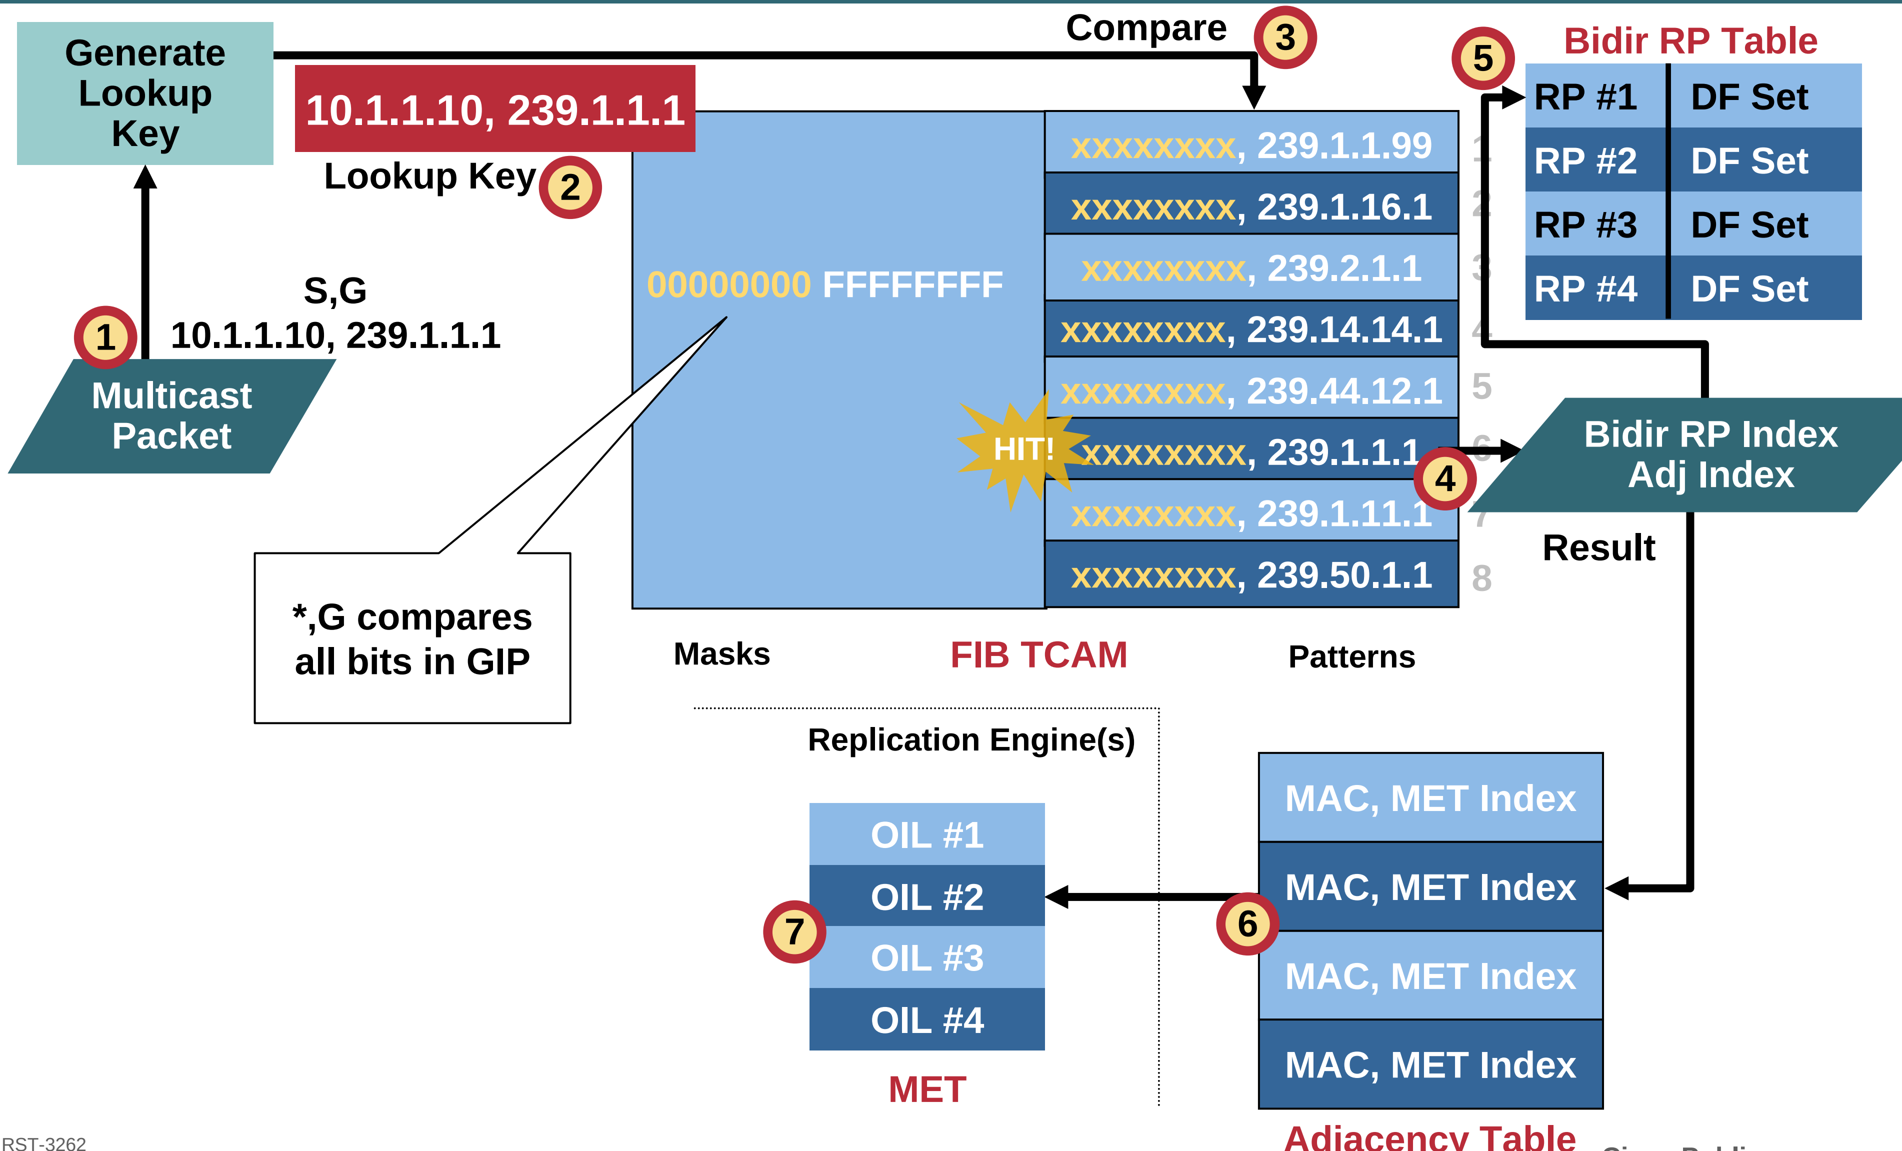

Bidir RP-to-DF Mapping Table

RP Index	DF Interface Set								
0	1	2	3	4	5	6	...	4093	4094
1	1	2	3	4	5	6	...	4093	4094
2	1	2	3	4	5	6	...	4093	4094
3	1	2	3	4	5	6	...	4093	4094

4K VLAN IDs

RP #1	DF Set
RP #2	DF Set
RP #3	DF Set
RP #4	DF Set

Bidir RP Table
(PFC/DFC)

RP Index	RP IP Address	DF Interfaces
0	10.1.1.1	VLAN 1 2 3
1	10.2.2.2	VLAN 3 4 5 6
2	10.3.3.3	VLAN 3 4093
3	10.4.4.4	VLAN 200 201

Mapping Table
(Software Data Structure)

Agenda

- IP Multicast Overview
- IP Multicast Hardware Architecture
- IP Multicast Hardware Forwarding
- IP Multicast Replication
- IP Multicast Packet Flows
- IGMP and IGMP Snooping
- Multicast Troubleshooting

IP Multicast Replication

Multicast Replication

- **Process of creating copies of multicast packets on each Layer 3 OIF**
  - **Example: (S,G) with three OIFs—multicast replication creates three copies of every packet received from source (S) destined to group (G)**
- **Replication on Catalyst 6500 occurs in one or more “replication engine” ASICs**
- **Supervisor engine always has a replication engine**
- **Fabric-enabled switching modules have local replication engines**

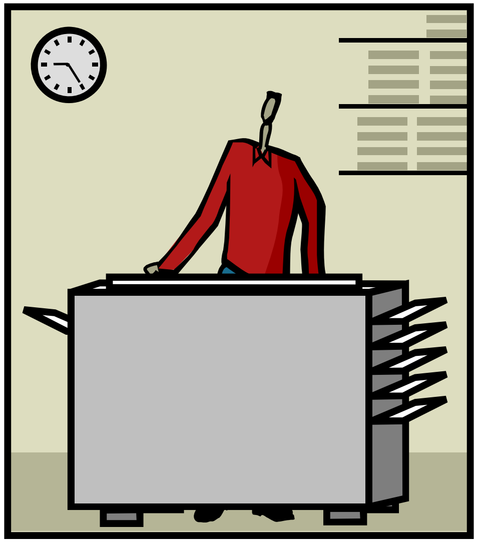

Multicast Replication Modes

- Replication mode refers to where in the system multicast replication occurs
- In classic system, replication always occurs centrally on the supervisor engine
- In fabric-enabled system, two possible replication modes:
  - Ingress replication
  - Egress replication

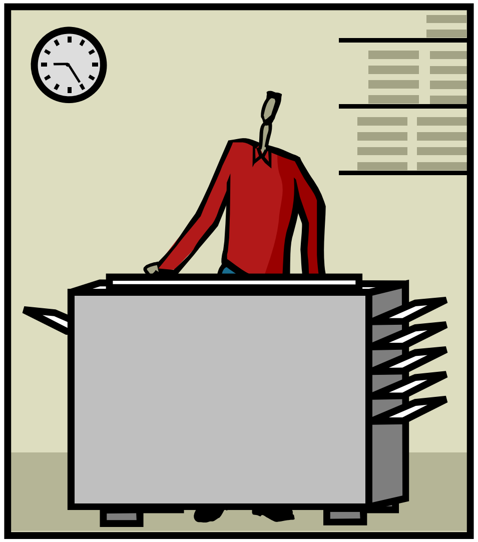

Replication in Classic System

- Supervisor engine performs multicast replication for all modules
- All input packets pass on switching bus
- All replicated copies pass on switching bus
- PFC performs lookups for input and all replicated packets
- Only one MET in system—the MET on supervisor replication engine

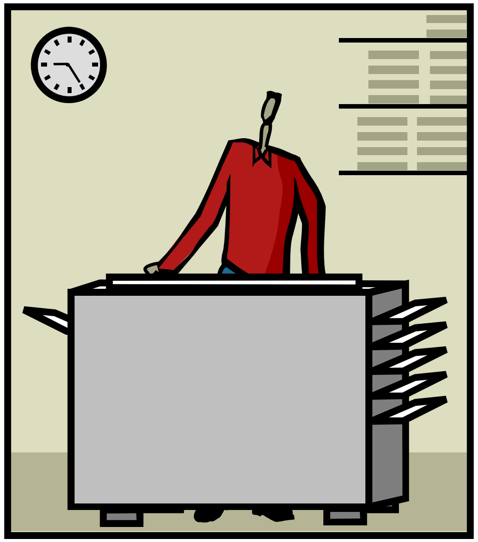

Ingress Replication

- Supported on Supervisor 2 (with fabric) and Supervisor 720
- Requires fabric-enabled modules
- Replication load distributed—Supervisor and switching modules perform replication
- Replication engine on ingress module performs replication for all OIFs
- Input and replicated packets get lookup on PFC or ingress DFC
- Replicated copies pass over fabric to egress modules
- Multiple MET tables, but MET on all replication engines synchronized

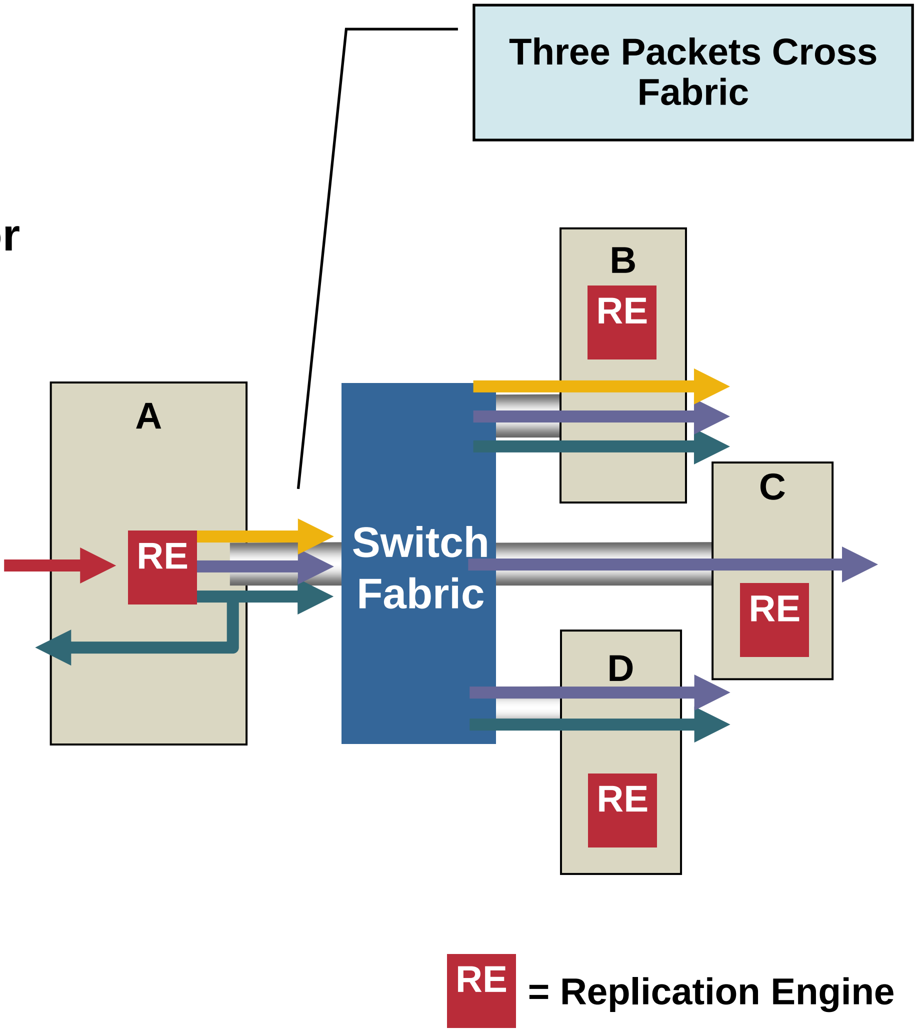

Egress Replication

- Supported on Supervisor 720 with certain switching modules only (CEF720, 6516A, 6548-GETX, SIPs)
- Replication load distributed—Supervisor and switching modules perform replication
- All modules in chassis must be egress-capable
- Egress mode not optimized unless DFCs present on modules
- Input packets get lookup on ingress DFC, replicated packets get lookup on egress DFC
- For OIFs on ingress module, local engine performs the replication
- For OIFs on other modules, ingress engine replicates a single copy of packet over fabric to all egress modules
- Engine on egress module performs replication for local OIFs
- MET tables on different modules can be asymmetric

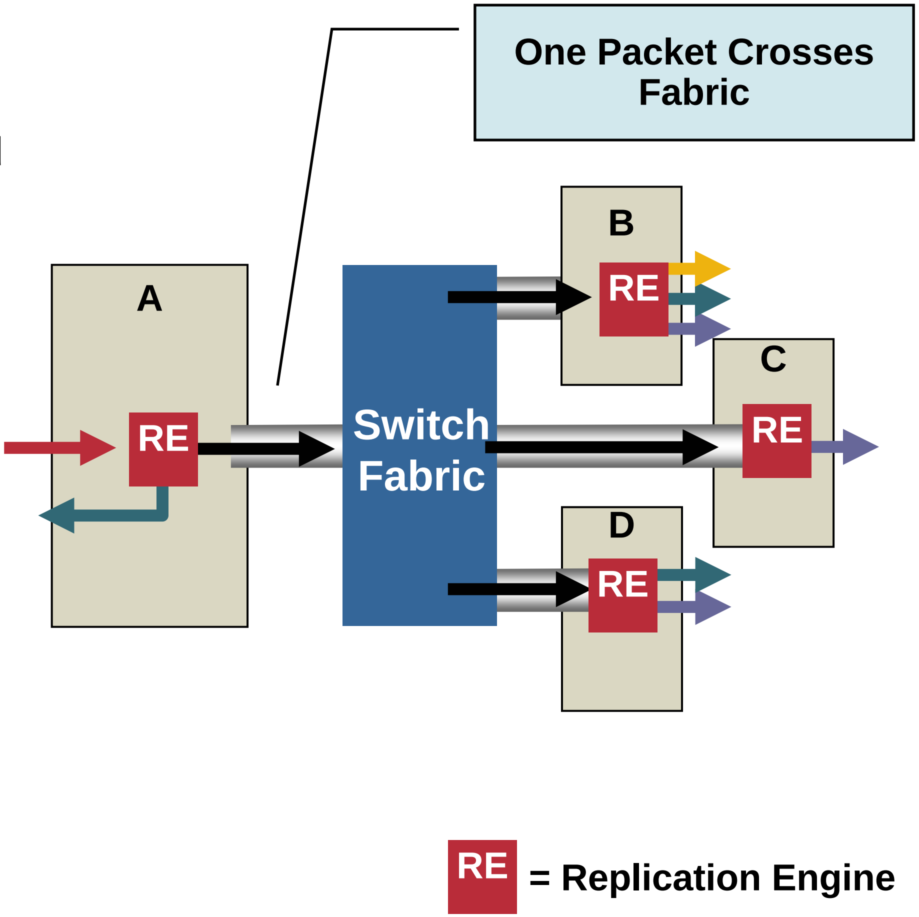

Agenda

- IP Multicast Overview
- IP Multicast Hardware Architecture
- IP Multicast Hardware Forwarding
- IP Multicast Replication
- IP Multicast Packet Flows
- IGMP and IGMP Snooping
- Multicast Troubleshooting

IP Multicast Packet Flows

Centralized Replication (1)

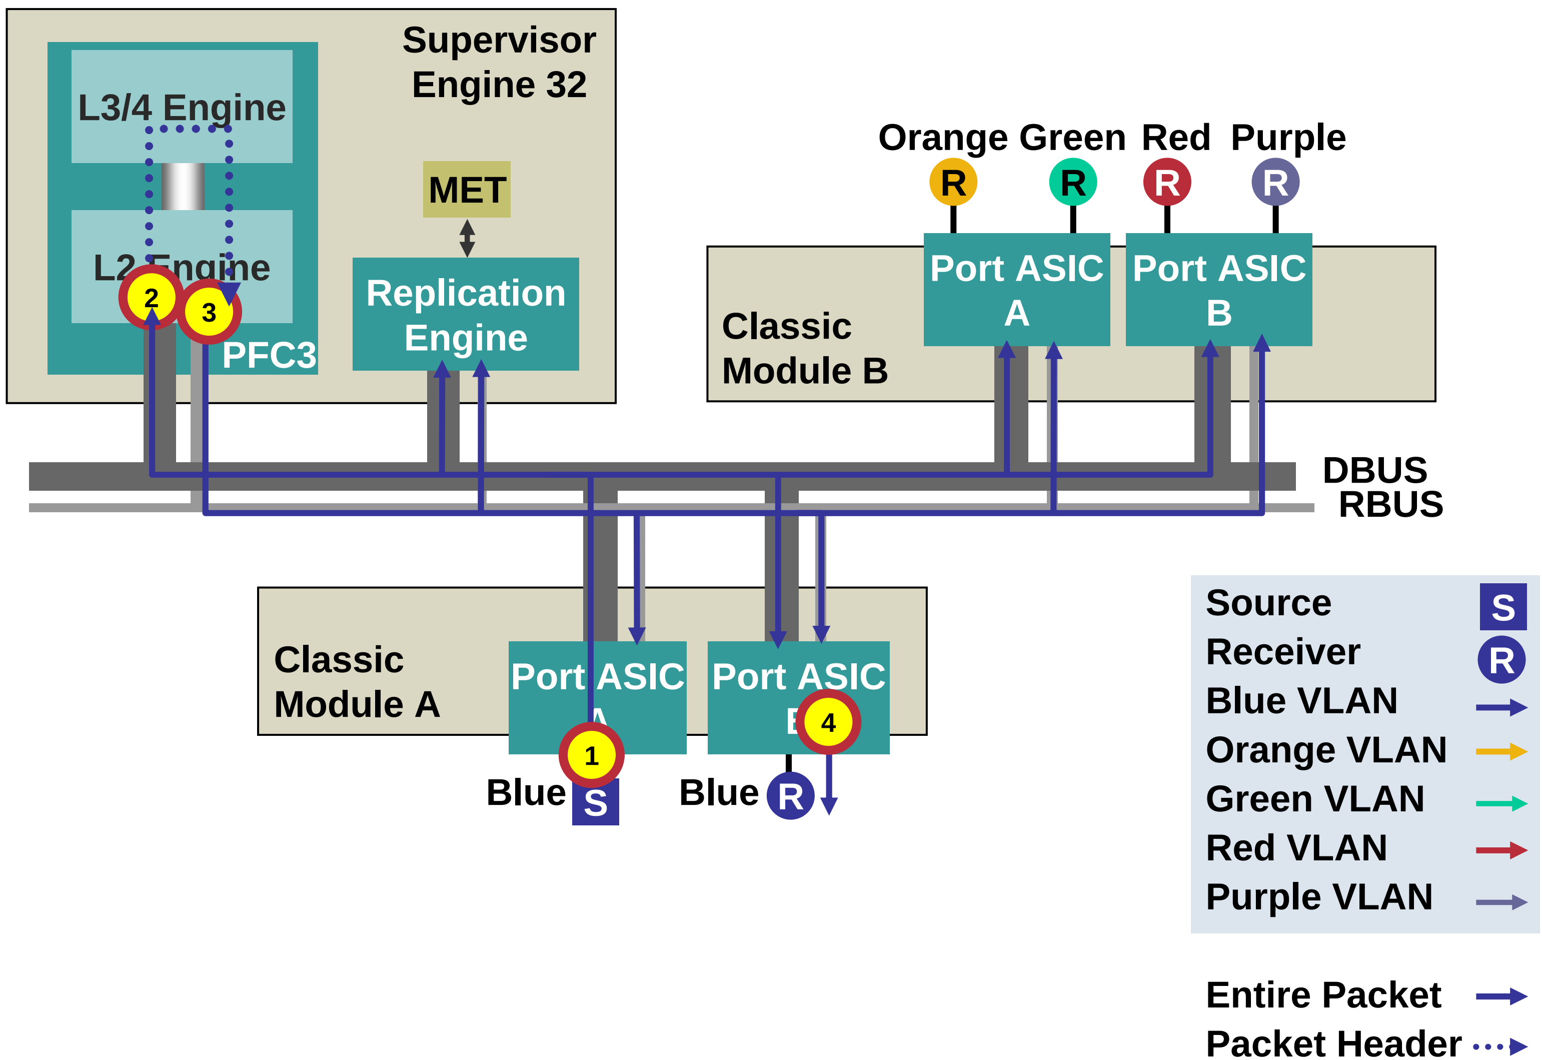

Centralized Replication (2)

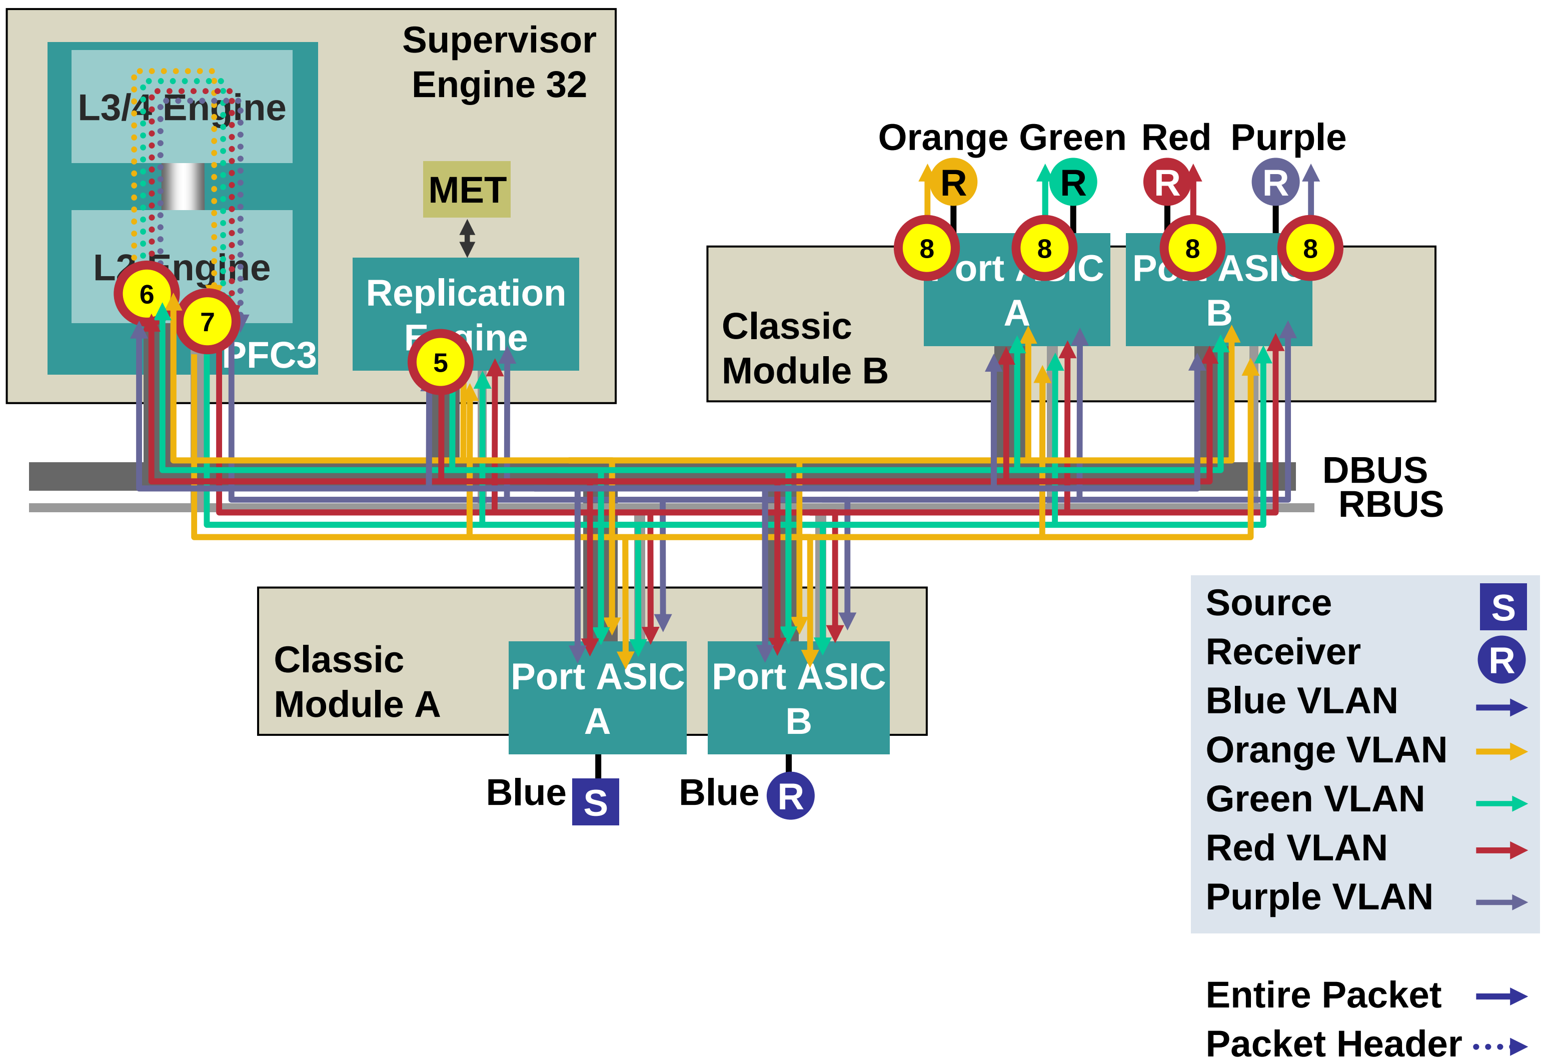

Ingress Replication (1)

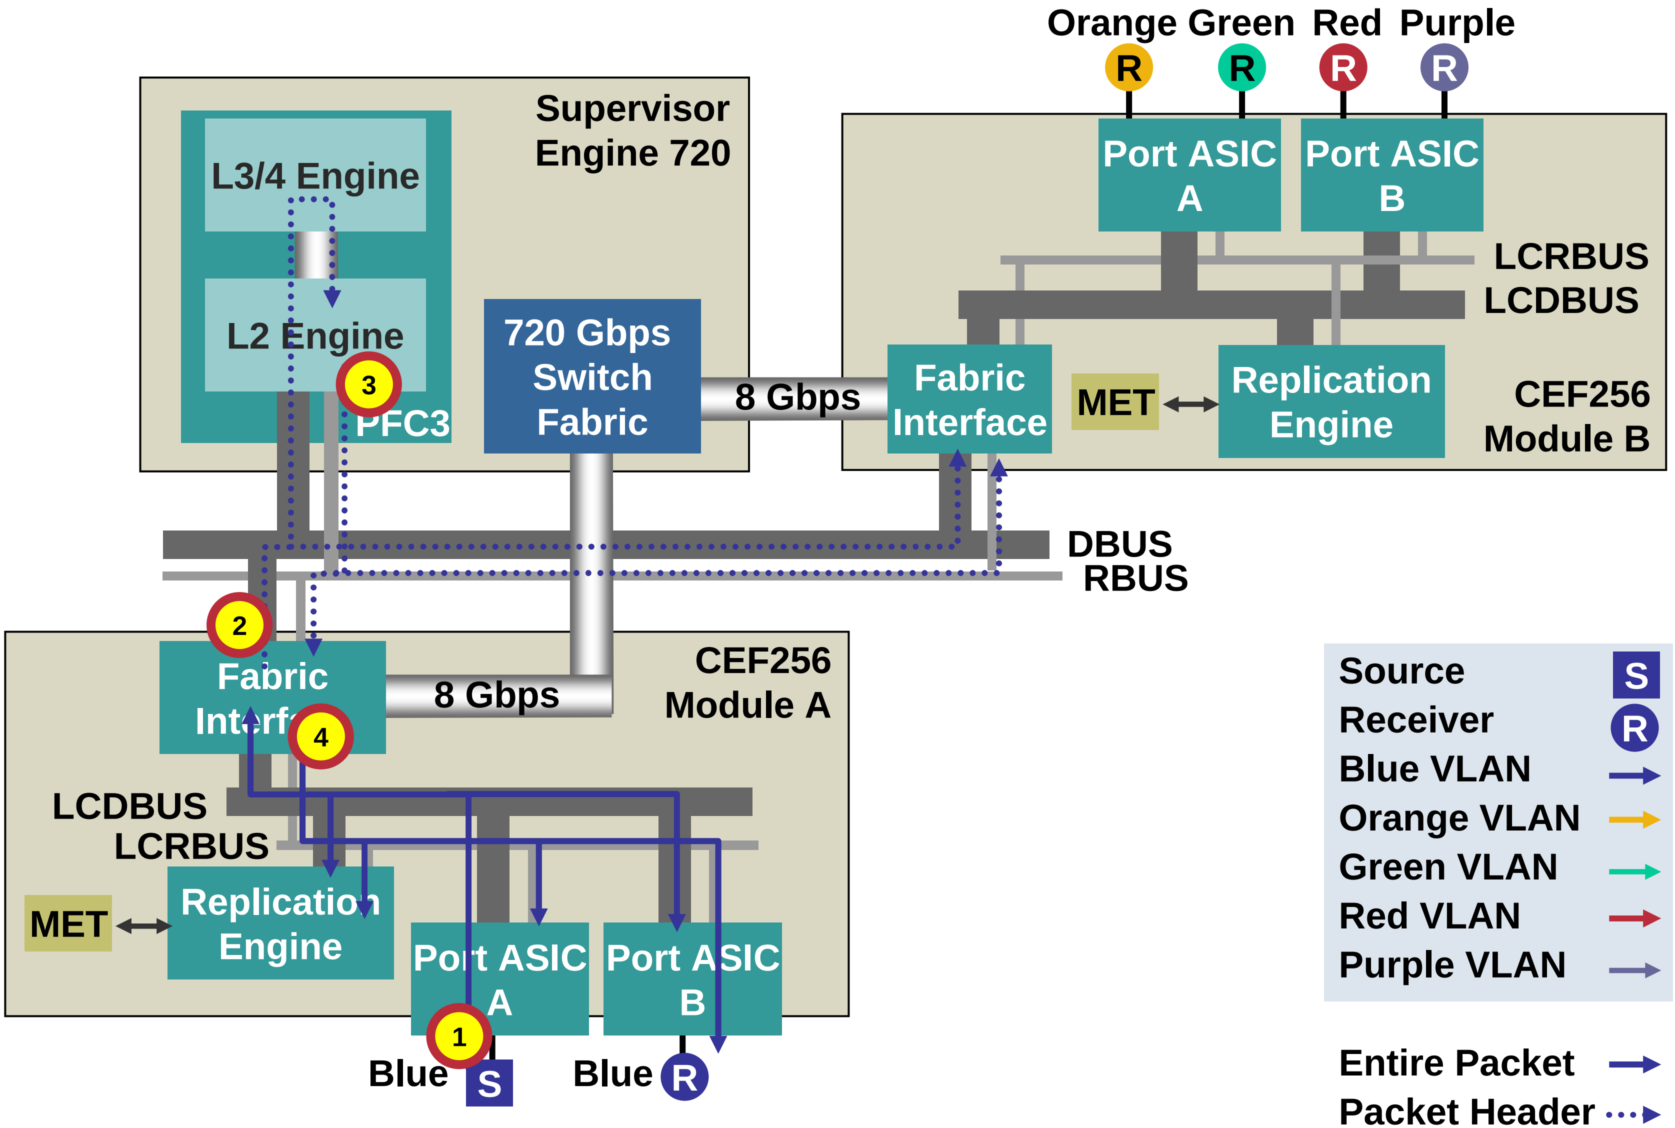

Ingress Replication (2)

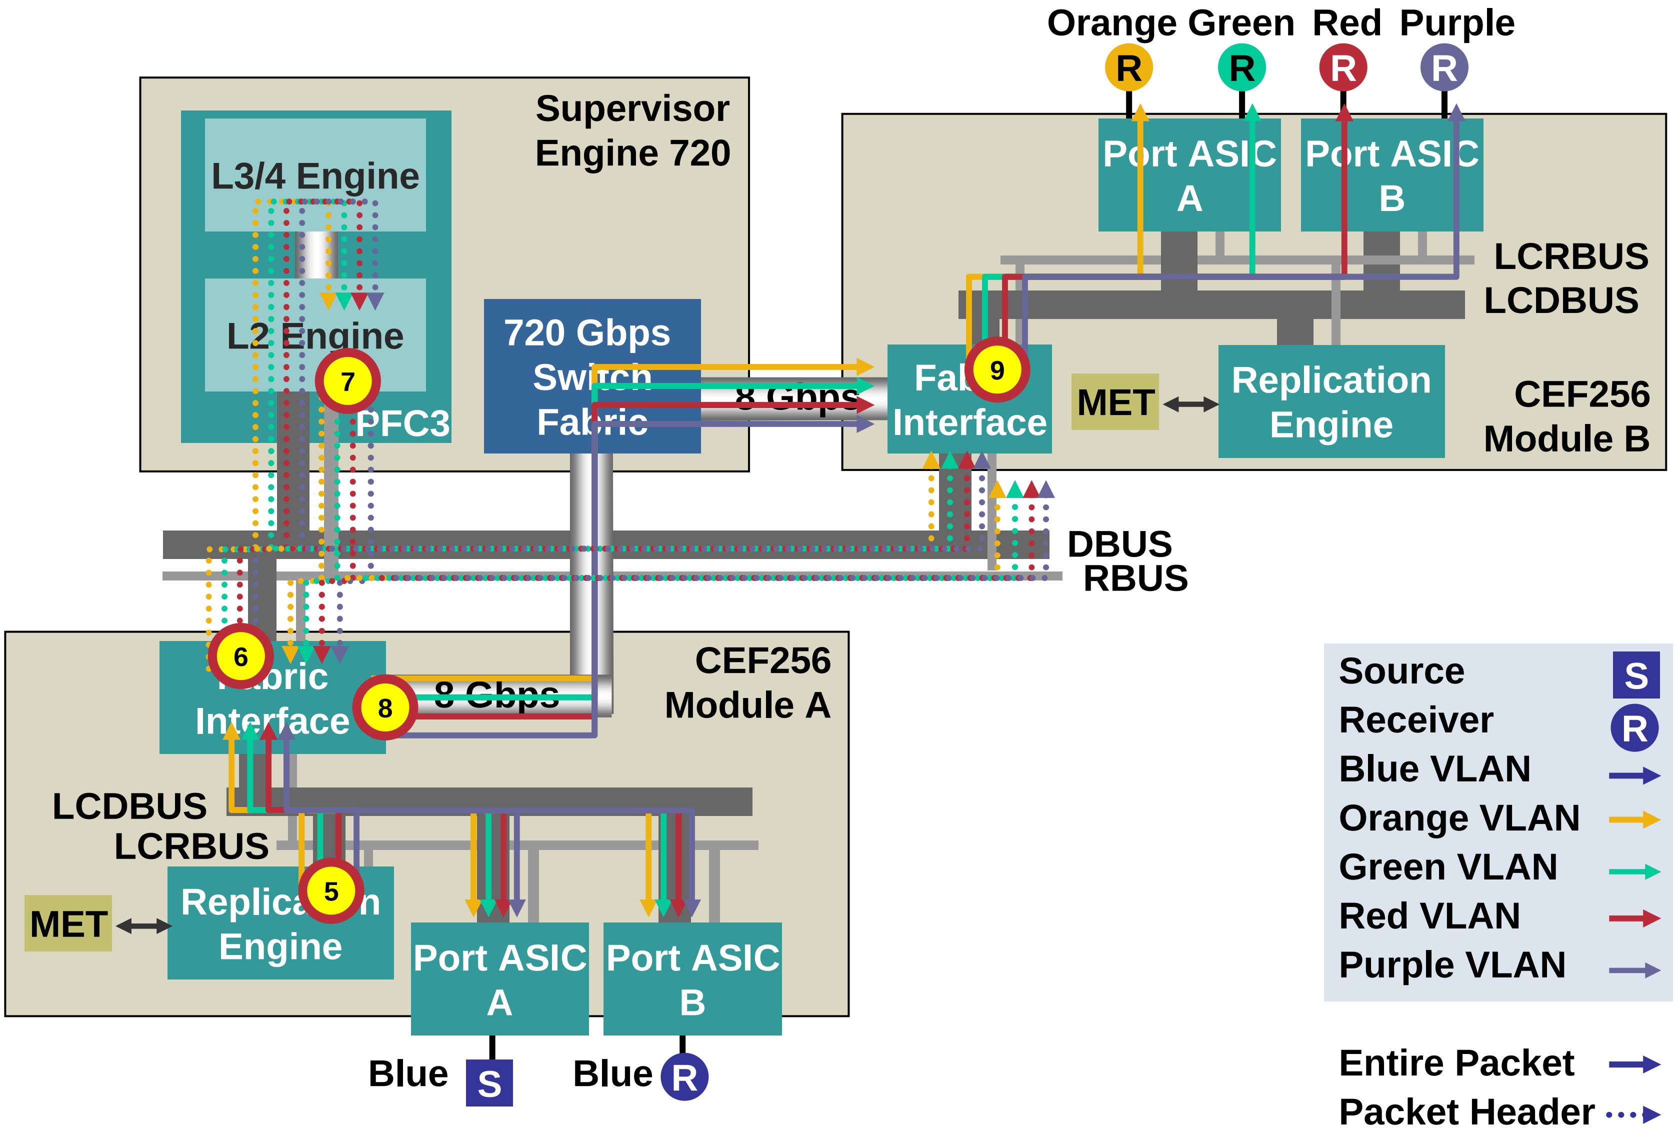

Egress Replication (1)

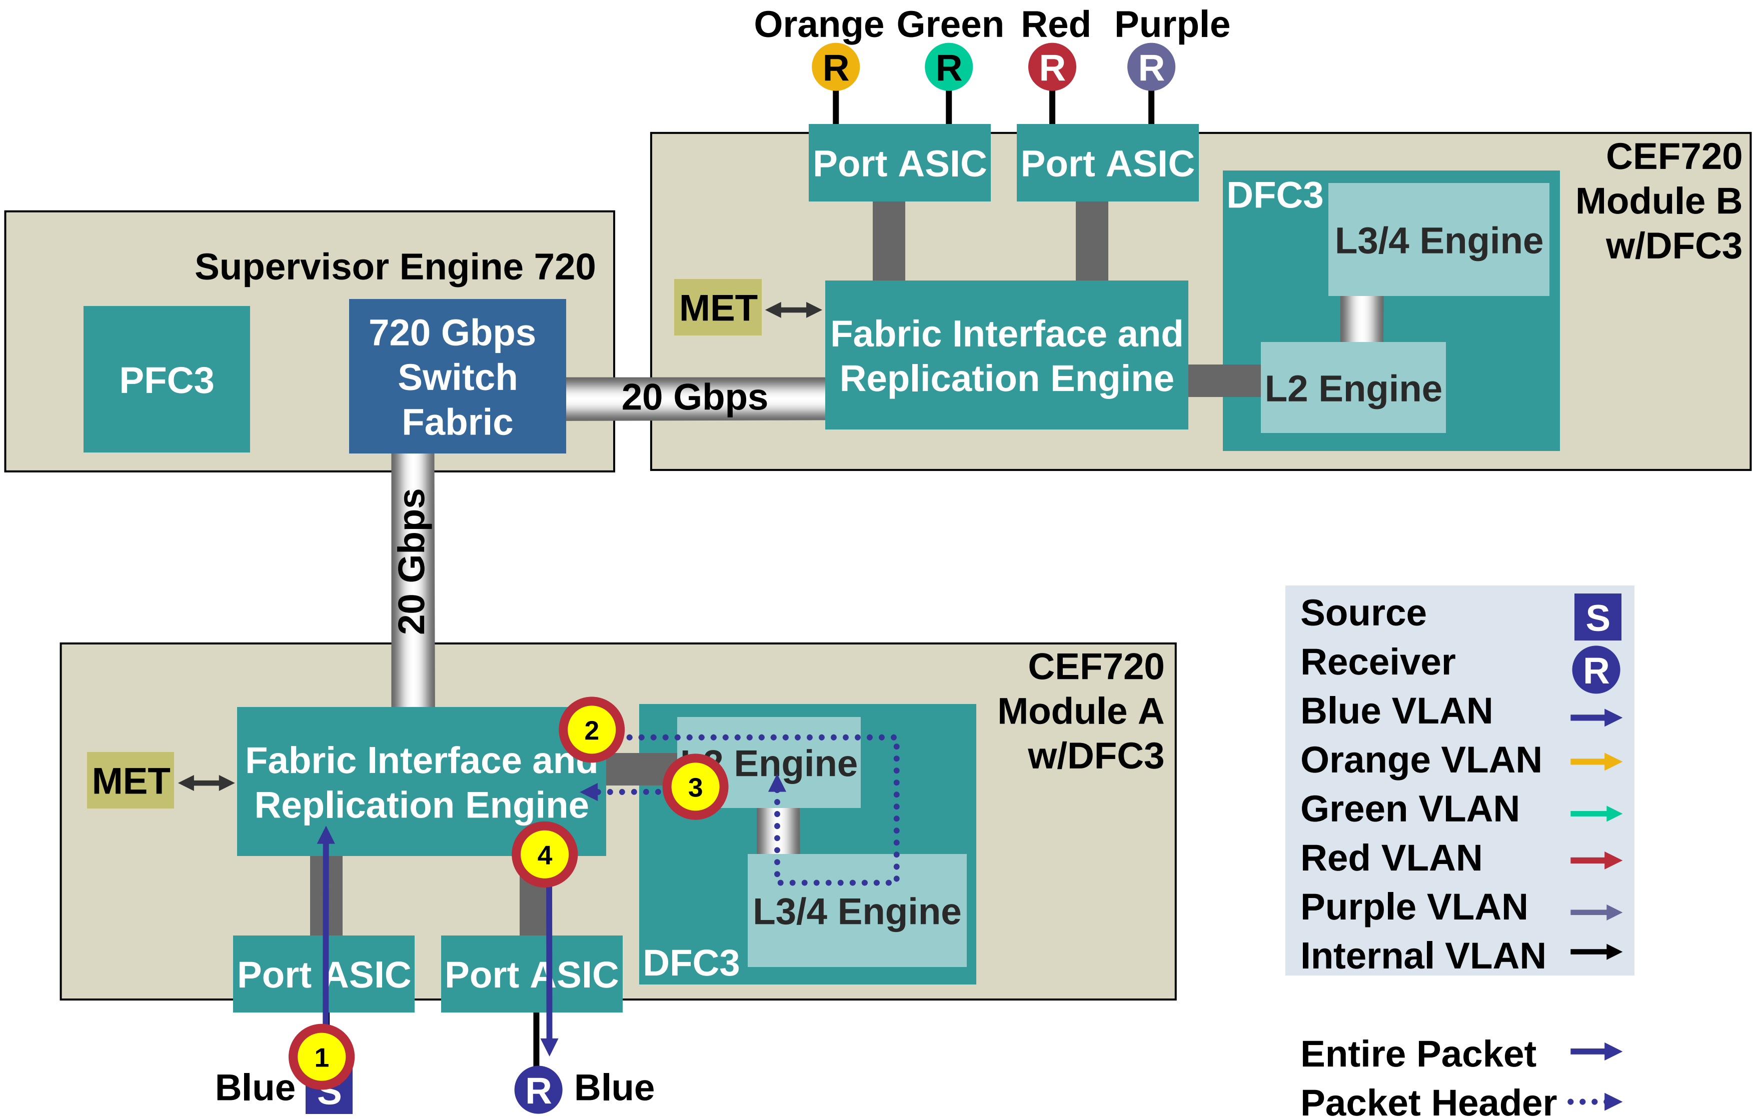

Egress Replication (2)

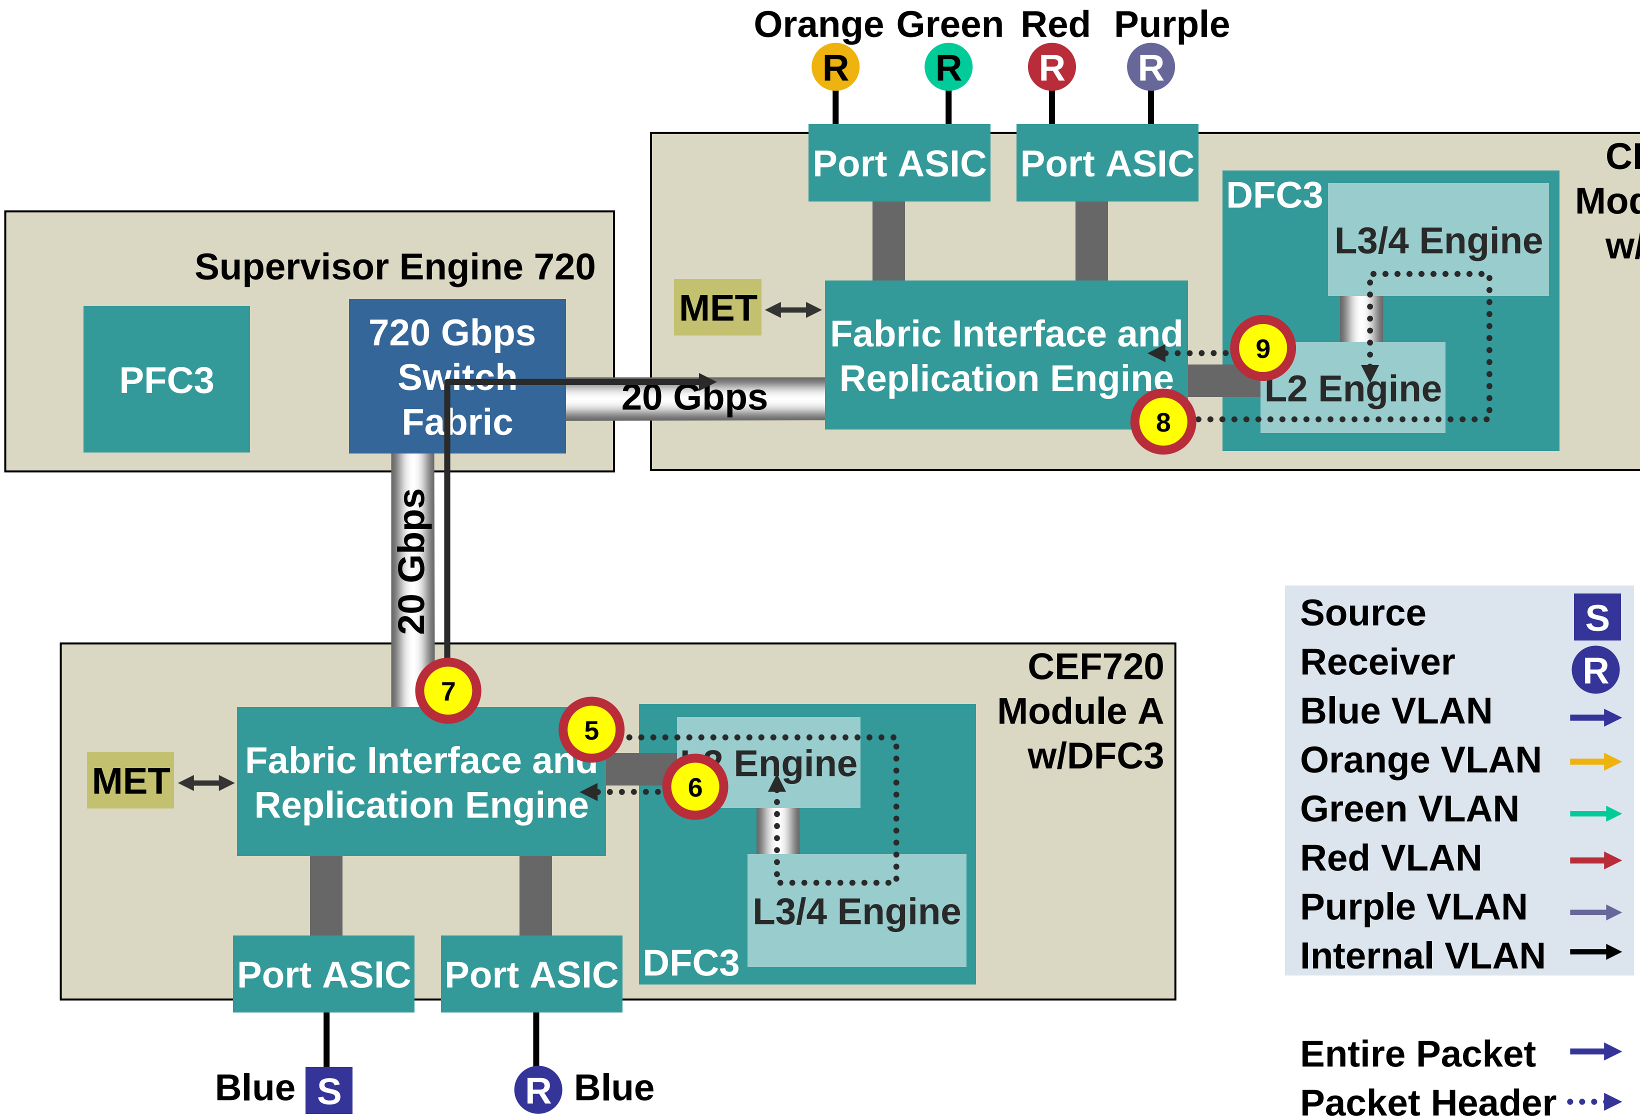

Egress Replication (3)

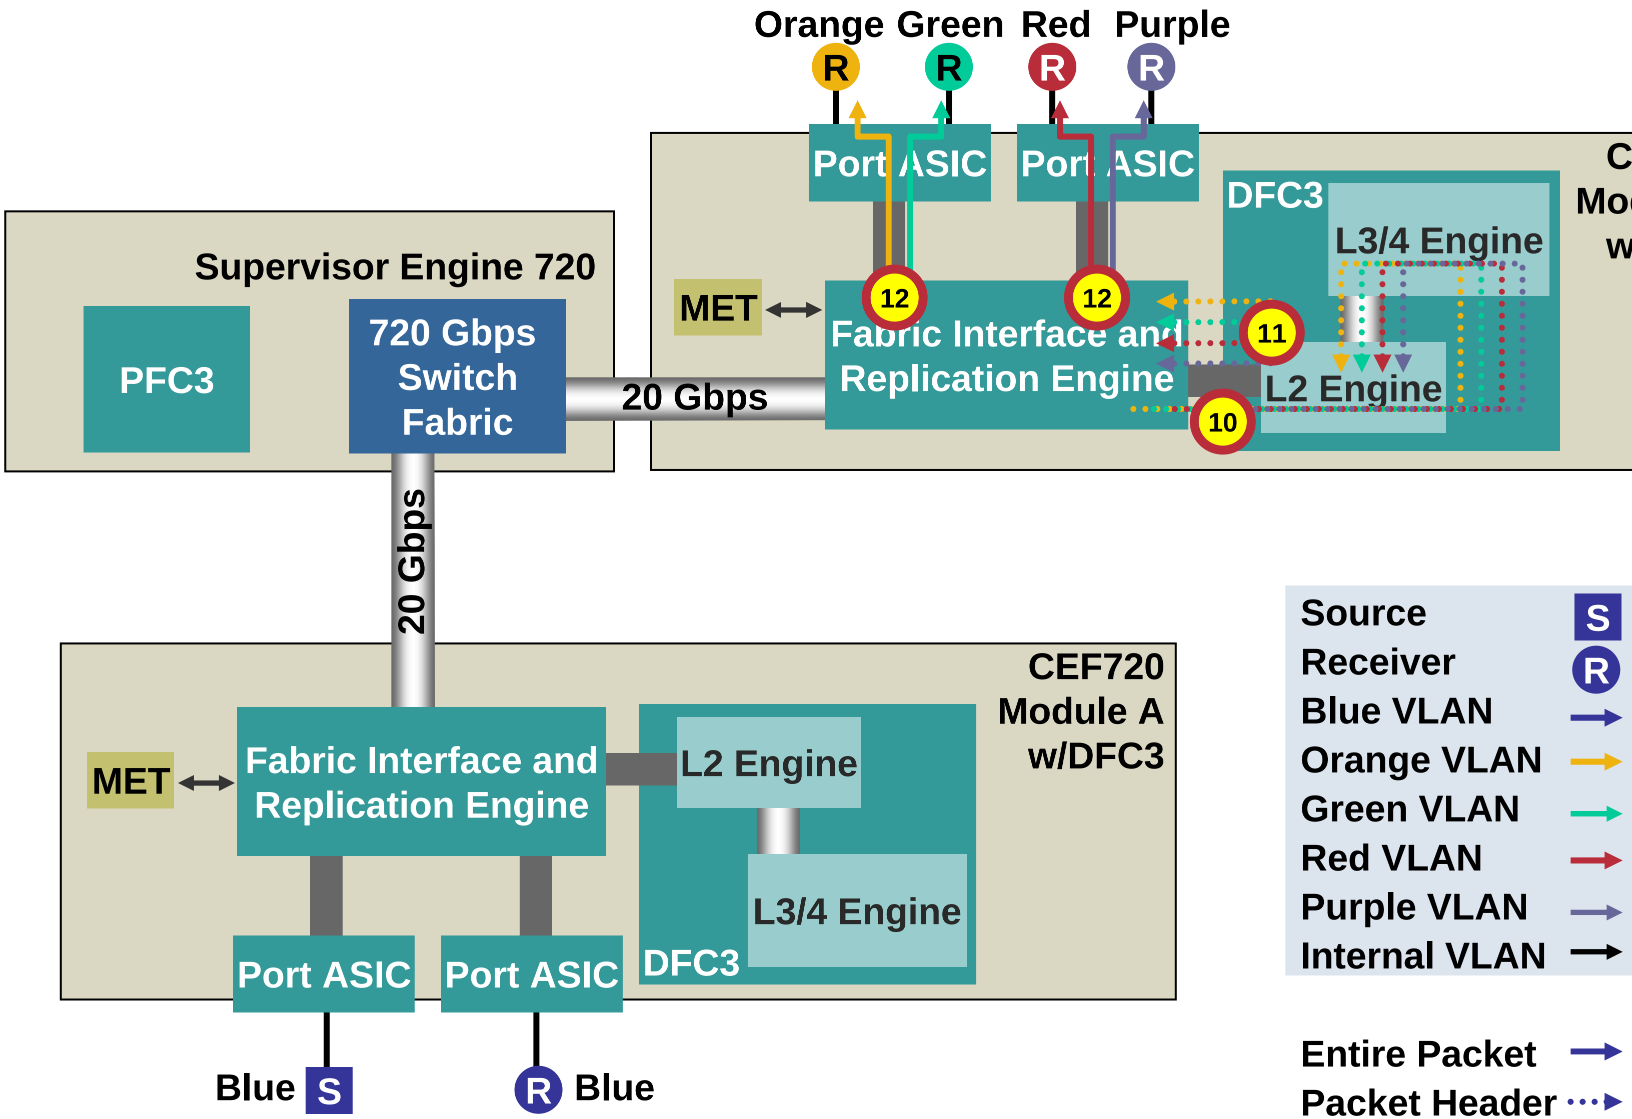

Agenda

- IP Multicast Overview
- IP Multicast Hardware Architecture
- IP Multicast Hardware Forwarding
- IP Multicast Replication
- IP Multicast Packet Flows
- IGMP and IGMP Snooping
- Multicast Troubleshooting

IGMP and IGMP Snooping

IGMP

- Purpose—Signal and refresh group membership on receiver segments
- IGMP support through Cisco IOS software
- IGMP v1/v2/v3 protocol support for PIM-SM and Bidir-PIM
- IGMP v3 protocol support for PIM-SSM

  Option for SSM mapping to translate IGMPv2 joins to PIM-SSM joins

IGMP Snooping

- **Purpose—Constrain multicast flooding on Layer 2 ports**
- **Implementation leverages both hardware and software**

  PFC ASICs recognize IGMP packets and redirect them to SP CPU ("protocol redirection logic")

  Switch "snoops" contents of IGMP packets

  Switch installs static Layer 2 forwarding entries for each multicast group MAC

  Multicast data traffic forwarded to appropriate interfaces according to MAC address table entries (per VLAN)

- **Does not affect performance for multicast data traffic**

  Protocol redirection ONLY redirects IGMP packets, not UDP (data) packets

IGMP Snooping Packet Flow

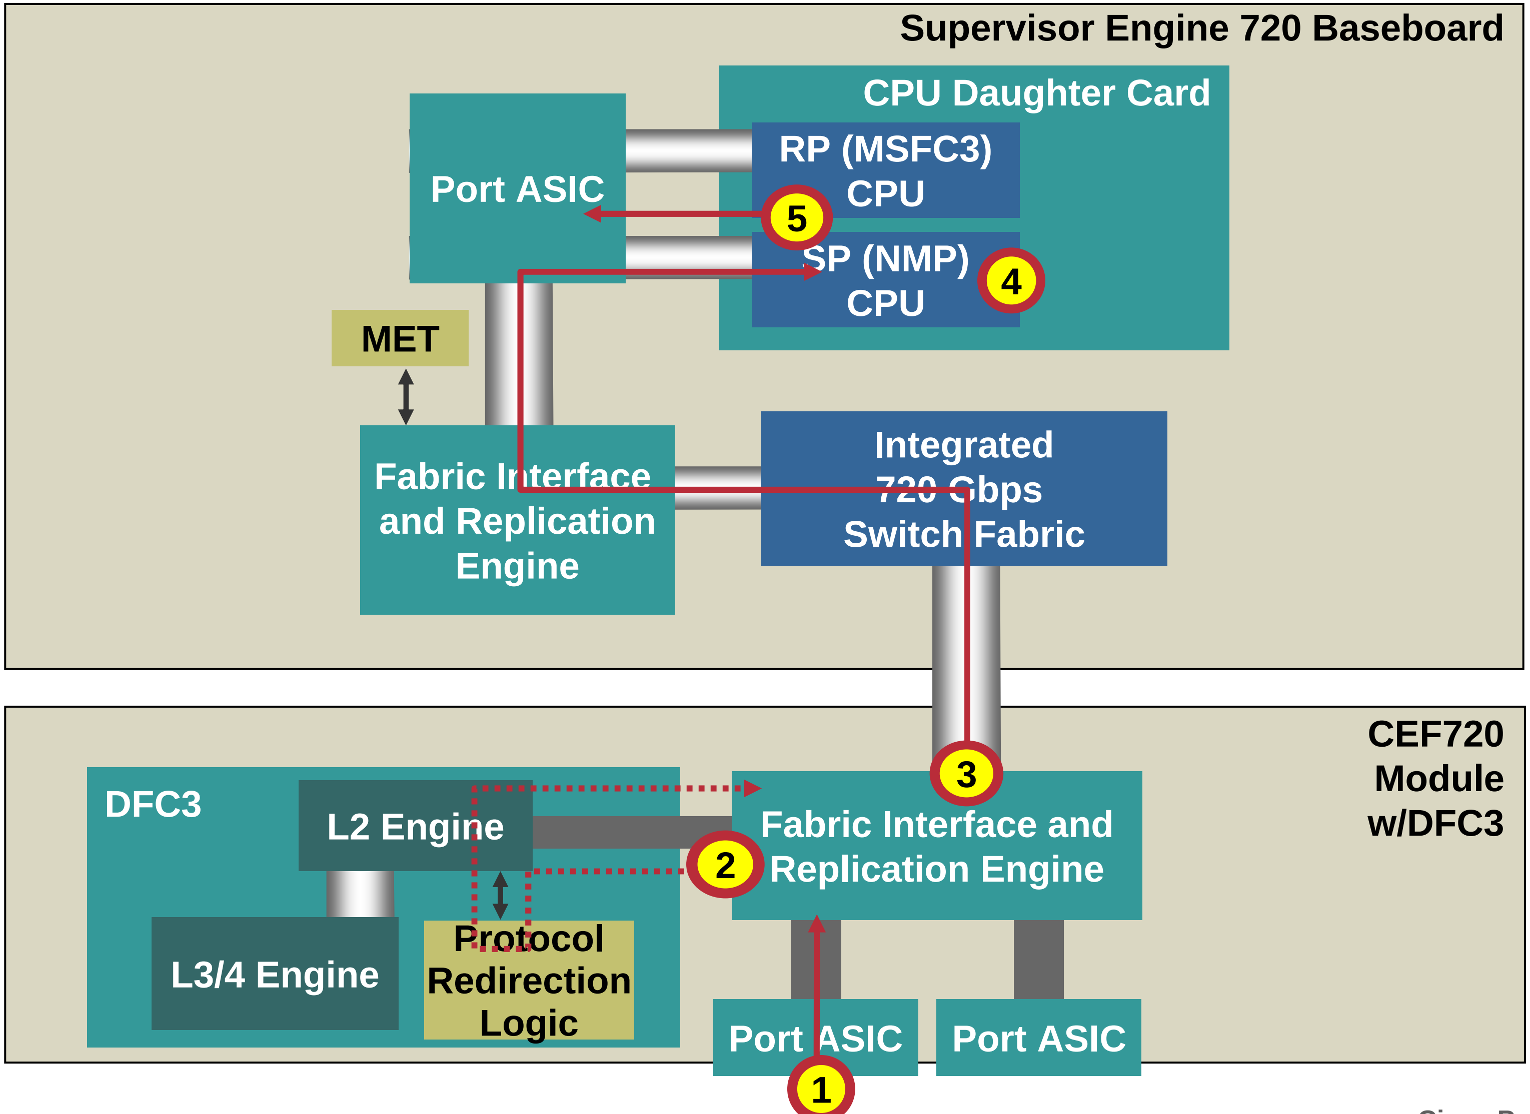

Agenda

- IP Multicast Overview
- IP Multicast Hardware Architecture
- IP Multicast Hardware Forwarding
- IP Multicast Replication
- IP Multicast Packet Flows
- IGMP and IGMP Snooping
- Multicast Troubleshooting

IP Multicast Troubleshooting

Catalyst 6500 Multicast Troubleshooting Overview

Key Problem Areas:

- Configuration
- Software and hardware multicast forwarding state
- Software and hardware scalability limits
- Other issues

  Port ASICs, fabric, switching bus, network events

Verify the Multicast Configuration

- `ip multicast-routing` **enabled!!**
- **PIM on all the interfaces (remember the loopbacks)**
- **RP configuration (AutoRP, BSR, AnycastRP, Phantom RP, static)**
- **Hardware MMLS and IGMP snooping enabled (on by default)**
- **Watch out for TTL thresholds, multicast boundary, security ACLs, VACLs, policers, etc.**
- **Watch for proper SSM or Bidir address range configuration—SSM and Bidir require coordination between Network and Application groups**
- **Unicast routing—attend another session for this one ☺**

Troubleshooting Forwarding State

- **Verify RP and DR/DF state**
- **Verify software IP mroute state**
- **Verify hardware multicast forwarding tables**

Verifying RP and DR/DF State

- Check RP IP addresses and group-to-RP mappings
- Verify RP RPF/upstream information
- Verify DR or DF interface state

Verifying Group-to-RP Mappings

```
tstevens-6506#show ip pim rp mapping
PIM Group-to-RP Mappings
Group(s): 224.0.0.0/4, Static
    RP: 10.255.255.3 (tstevens-6509.cisco.com)
tstevens-6506#show ip pim rp
Group: 239.1.1.10, RP: 10.255.255.3, v2, uptime 00:01:10, expires never
Group: 239.1.1.11, RP: 10.255.255.3, v2, uptime 00:01:10, expires never
Group: 239.1.1.12, RP: 10.255.255.3, v2, uptime 00:01:10, expires never
Group: 239.1.1.13, RP: 10.255.255.3, v2, uptime 00:01:10, expires never
Group: 224.0.1.40, RP: 10.255.255.3, v2, uptime 00:01:10, expires never
tstevens-6506#
```

Verifying IP RPF Information

```
tstevens-6506#show ip rpf 10.255.255.3
RPF information for tstevens-6509.cisco.com (10.255.255.3)
  RPF interface: GigabitEthernet1/5
  RPF neighbor: tstevens-6513 (10.20.1.2)
  RPF route/mask: 10.255.255.3/32
  RPF type: unicast (ospf 10)
  RPF recursion count: 0
  Doing distance-preferred lookups across tables
tstevens-6506#
```

Identifying the DR for a Segment

`show ip pim interface` **identifies DR for each interface**

```
tstevens-6513#show ip pim interface

Address          Interface              Ver/   Nbr    Query  DR     DR
                                        Mode   Count  Intvl  Prior
10.255.255.2     Loopback0              v2/S   0      30     1      10.255.255.2
10.255.254.1     Loopback1              v2/S   0      30     1      10.255.254.1
10.10.1.2        GigabitEthernet4/3     v2/S   1      30     1      10.10.1.2
10.30.1.1        GigabitEthernet4/16    v2/S   1      30     1      10.30.1.2
10.100.1.1       Vlan100                v2/S   1      30     1      10.100.1.2
10.101.1.1       Vlan101                v2/S   1      30     1      10.101.1.2
10.200.1.1       Vlan200                v2/S   1      30     1      10.200.1.2
10.201.1.1       Vlan201                v2/S   1      30     1      10.201.1.2
tstevens-6513#
```

This router's interface addresses

Compare with DR address

Identifying Bidir DF Interfaces

```
tstevens-6506#show ip pim interface df
Interface               RP               DF Winner        Metric     Uptime
Loopback0               10.255.255.3     10.255.255.1     2          02:21:07
GigabitEthernet1/5      10.255.255.3     10.20.1.2        0          02:21:08
GigabitEthernet1/13     10.255.255.3     10.13.1.1        2          02:21:07
GigabitEthernet2/13     10.255.255.3     10.13.2.1        2          02:21:07
GigabitEthernet2/14     10.255.255.3     10.14.2.1        2          02:21:07
GigabitEthernet2/24     10.255.255.3     10.2.24.1        2          00:35:15
Vlan100                 10.255.255.3     10.100.1.2       2          00:03:39
Vlan101                 10.255.255.3     10.101.1.2       2          00:01:13
Vlan200                 10.255.255.3     10.200.1.2       2          00:01:10
Vlan201                 10.255.255.3     10.201.1.2       2          00:01:07
tstevens-6506#
```

PIM enabled interfaces

Bidir RP address

IP address of current DF winner

Winner's routing metric to RP

Viewing the Bidir RP-to-DF Interface Mapping Table

```
tstevens-6506#show mls ip multicast rp-mapping df-cache
State: H - Hardware Switched, I - Install Pending, D - Delete
  Pending,
       Z - Zombie

RP Address       State       DF              State
10.255.255.3       H         Gi2/13              H
10.255.255.3       H         Gi1/13              H
10.255.255.3       H         Gi2/14              H
10.255.255.3       H         Gi2/24              H
10.255.255.3       H         Vl100               H
10.255.255.3       H         Vl101               H
10.255.255.3       H         Vl200               H
10.255.255.3       H         Vl201               H
tstevens-6506#
```

All DF interfaces should be in H state

DF interfaces for specified Bidir RP

Verifying Software IP Mroute State

- **Ensure IP mroute exists in software—** `show ip mroute`

  Does (*,G) and/or (S,G) exist in software mroute table?

  Does hardware state information contained in Cisco IOS `show ip mroute` output appear correct?

- **Ensure RPF interface is known and correct**

  Make sure `show ip mroute` and `show ip rpf` show correct RPF interface for (*,G) or (S,G)

  PIM reliance on unicast routing means multicast issues often caused by unicast routing issues

Verifying Software IP Mroute State (2)

OIFs Are Known and Correct—OIF Inclusion Driven by PIM or IGMP

- Ensure PIM neighbors active and stable (`show ip pim neighbor`)
- IGMP dictates connected receiver membership—ensure joins/leaves sent by receivers and received by RP CPU
- Might need to verify PIM and IGMP packet exchange using SPAN/sniffer and/or debugs

IP Mroute Table with Complete Shortcut

- **show ip mroute**

```
tstevens-6506#show ip mroute 10.1.1.100 239.1.1.1
<…>
(10.1.1.100, 239.1.1.1), 00:00:38/00:02:53, flags: TA
  Incoming interface: GigabitEthernet1/5, RPF nbr 10.20.1.2, RPF-MFD
  Outgoing interface list:
    Vlan100, Forward/Sparse, 00:00:38/00:02:21, H
    Vlan200, Forward/Sparse, 00:00:38/00:02:21, H
    Vlan101, Forward/Sparse, 00:00:38/00:02:21, H
    GigabitEthernet1/13, Forward/Sparse, 00:00:38/00:02:58, H
    GigabitEthernet2/13, Forward/Sparse, 00:00:38/00:02:58, H
    GigabitEthernet2/14, Forward/Sparse, 00:00:38/00:02:58, H

tstevens-6506#
```

Characteristics of Complete Shortcuts

- **RPF-MFD flag is set for the entry**
  - "RPF-Multicast Fast Drop"—fancy way of saying "the router does not need to receive this traffic, so don't punt a copy to the RP CPU"
- **RPF-MFD flag is only set when every OIF associated with the forwarding entry is hardware switched**
  - Hardware-switched OIFs have the "H" flag set
- **Gotcha: if at least one OIF has MTU smaller than the RPF interface, packets exceeding the MTU of that OIF are punted to the RP CPU for software replication for all OIFs**
  - Mroute entry may NOT be flagged as a partial shortcut in this case!!
  - Make sure OIFs have MTU >= the RPF interface

IP Mroute Table with Partial Shortcut

- **show ip mroute**

```
tstevens-6506#show ip mroute 10.1.1.100 239.1.1.1
<…>
(10.1.1.100, 239.1.1.1), 00:02:08/00:02:53, flags: TA
  Incoming interface: GigabitEthernet1/5, RPF nbr 10.20.1.2, Partial-SC
  Outgoing interface list:
    Vlan100, Forward/Sparse, 00:02:08/00:02:44, H
    Vlan200, Forward/Sparse, 00:02:08/00:02:51 (ttl-threshold 8)
    Vlan101, Forward/Sparse, 00:02:08/00:02:11, H
    GigabitEthernet1/13, Forward/Sparse, 00:02:08/00:02:58, H
    GigabitEthernet2/13, Forward/Sparse, 00:02:08/00:02:58, H
    GigabitEthernet2/14, Forward/Sparse, 00:02:08/00:02:58, H

tstevens-6506#
```

Characteristics of Partial Shortcuts

- **Partial-SC flag set for the entry**
- **Partial-SC flag set when at least one OIF associated with the forwarding entry requires software switching**
  - **OIFs requiring software switching clearly labeled with the reason**
- **Partial shortcut typically due to hardware-unsupported configuration (TTL threshold, ip igmp join-group, etc.)**
- **RP CPU must receive all traffic for that mroute to apply the feature and replicate the traffic to the unsupported interface(s)**
- **Other OIFs still get hardware switching**
- **Partial-SC traffic can be rate-limited in PFC3 using the `mls rate-limit multicast partial <rate>` command**
  - **Caps aggregate rate for partial-shortcut traffic, but can result in suboptimal delivery for software forwarded OIFs**

Reasons for Partial Shortcuts

- Packets require PIM register encapsulation
- `ip igmp join-group` command present on RPF interface or on an OIF
- `ip multicast ttl-threshold` command present on an OIF
- `ip multicast helper-map` command present on RPF interface or an OIF
- An OIF is a tunnel interface (PFC2, PFC3 pre-18SXE)
- Hardware switching disabled on an OIF
- NAT configured on an OIF and source address translation required
- (*,G) entry on last-hop leaf router, if shortest path threshold not set to infinity

Displaying Hardware Forwarding Entries

- Cisco IOS: `show mls ip multicast`
- Catalyst OS: `show mls multicast entry`

```
tstevens-6506#show mls ip multicast group 239.1.1.1 source 10.1.1.100
Multicast hardware switched flows:
(10.1.1.100, 239.1.1.1) Incoming interface: Gi1/5, Packets switched: 293919
Hardware switched outgoing interfaces:
        Vlan100 Vlan101 Gi1/13 Gi2/13 Gi2/14 Vlan200
RPF-MFD installed


Total hardware switched flows : 1
tstevens-6506#
```

- **Applies to PIM-SM, PIM-SSM, and Bidir (*,G) hardware entries**

Bidir-PIM (*,G/m) Entries

- **Source-only traffic must reach RP (could be receivers on other branches)**
- **For efficiency, system installs (*,G/m) hardware forwarding entry/entries to transport this traffic**
  - **Entries based on Bidir-PIM RP ACL configuration**
- **At RP, assuming no receivers on other branches, packets dropped in hardware**
- **NOTE: (*,G/m) entries not shown in software mroute table today**

Viewing Bidir (*,G/m) Forwarding Entries

```
tstevens-6506#show mls ip multicast rp-mapping gm-cache
State: H - Hardware Switched, I - Install Pending, D - Delete Pending,
       Z - Zombie

RP Address     State     Group          Mask          State       Packet/Byte-count
10.255.255.3     H       224.0.0.0      240.0.0.0       H         1183799/1754389822
tstevens-6506#
```

Statistics

Group IP and Mask (224/4)

Bidir RP IP Address

Troubleshooting Group Membership

- **Verify IGMP and IGMP snooping configuration status**
- **Make sure IGMP snooping requirements are met**
- **Make sure the Layer 3 and Layer 2 entries exist and interface/port membership is correct**

Verifying IGMP and IGMP Snooping Configuration Status—Cisco IOS

```
tstevens-6506#show ip igmp interface vlan 100
Vlan100 is up, line protocol is up
  Internet address is 10.100.1.2/24
  IGMP is enabled on interface
  Current IGMP host version is 2
  Current IGMP router version is 2
  IGMP query interval is 60 seconds
  IGMP querier timeout is 120 seconds
  IGMP max query response time is 10 seconds
  Last member query count is 2
  Last member query response interval is 1000 ms
  Inbound IGMP access group is not set
  IGMP activity: 3 joins, 2 leaves
```

IGMP Configuration State

IGMP Packet Statistics

Verifying IGMP and IGMP Snooping Configuration Status—Cisco IOS (Cont.)

```
  Multicast routing is enabled on interface
  Multicast TTL threshold is 0
  Multicast designated router (DR) is 10.100.1.3
  IGMP querying router is 10.100.1.2 (this system)
  No multicast groups joined by this system
  IGMP snooping is globally enabled
  IGMP snooping is enabled on this interface
  IGMP snooping fast-leave is disabled and querier is disabled
  IGMP snooping explicit-tracking is enabled
  IGMP snooping last member query response interval is 1000 ms
  IGMP snooping report-suppression is disabled
tstevens-6506#
```

DR and Querier Information

IGMP Snooping Configuration State

Verifying IGMP and IGMP Snooping Configuration Status—CatOS

- Use `show ip igmp interface` on MSFC to verify IGMP configuration status
- Use `show multicast protocols status` to verify IGMP snooping configuration status

```
tstevens-6503> (enable) show multicast protocols status
IGMP enabled
IGMP fastleave enabled
IGMP V3 processing disabled
IGMP V3 fastblock feature disabled
RGMP disabled
GMRP disabled
tstevens-6503> (enable)
```

IGMP Snooping Configuration State

IGMPv3 Snooping Configuration State

Additional Protocols Configuration State

Verifying IGMP Snooping Requirements

- **IGMP querier must be present in the VLAN**
  - **Could be multicast router or switch configured as querier**
- **Make sure all multicast router ports known**
  - **Switch tracks location of all multicast routers on per-VLAN basis**
  - **Detection based on IGMP queries and PIM hellos**
  - **Snooping switch uses list of mrouter ports to flood certain traffic—e.g., proxied joins/leaves**
  - **Loss of multicast router port will impact traffic flow**

Viewing Multicast Routers (Cisco IOS)

```
tstevens-6509-neb#show ip igmp snooping mrouter

vlan            ports
-----+----------------------------------------
 100  Gi3/7,Gi3/16
 101  Gi3/7,Gi3/16
 200  Gi3/7,Gi3/16
 201  Gi3/7,Gi3/16
tstevens-6509-neb#
```

List of Multicast Router Ports for VLAN

VLAN for Which Specified Ports Are Multicast Router Ports

Viewing Multicast Routers (CatOS)

```
tstevens-6503> (enable) show multicast router
'*' - Configured
'+' - RGMP-capable
'#' - Channeled Port
'$' - IGMP-V3 Router
'@' - IGMP-Querier Router
Port               Vlan
-------------- -----------------
 1/3            @ 50,60

Total Number of Entries = 1
tstevens-6503> (enable)
```

VLANs for Which Port Is Multicast Router Port

Flags (Here, @ = IGMP Querier)

Multicast Router Port

Verifying IGMP Membership

- **IGMP enabled when PIM configured on an interface**
- **For receiver segments, IGMP drives OIF presence in mroute table**
  - **Use `show ip igmp groups` to see IGMP join status on mrouter**
  - **Use `debug ip igmp <group>` to monitor IGMP packet reception at router**
  - **Be aware of effect of IGMP snooping on Layer 3 IGMP behavior (e.g. not all joins/leaves seen by router)**
- **Verify IGMP querier consistency and group membership for routers on shared segments**
- **Watch for IGMP access groups, multicast boundary**
- **Verify IGMP snooping entries**
  - **Make sure entry exists and port membership correct**

Verifying IGMP Entries

```
tstevens-6506#show ip igmp groups

IGMP Connected Group Membership
Group Address     Interface                  Uptime    Expires   Last Reporter
239.1.2.0         Vlan200                    00:00:12  00:00:00  10.200.1.101
239.1.2.1         Vlan200                    00:00:12  00:00:00  10.200.1.101
239.1.2.2         Vlan200                    00:00:12  00:00:00  10.200.1.101
239.1.2.3         Vlan200                    00:00:12  00:00:00  10.200.1.101
239.1.1.10        Vlan201                    00:00:13  00:00:00  10.201.1.101
224.0.1.40        Loopback0                  1w4d      00:02:18  10.255.255.1
tstevens-6506#
```

Interface with Connected Receiver

Multicast Group Joined

IP Address of Last Receiver to Report for the Group on the Interface

Verifying IGMP Snooping Entries—Cisco IOS

```
tstevens-6509-neb#show mac-address-table multicast igmp

 vlan   mac address     type    learn qos            ports
-----+---------------+--------+-----+---+--------------------------------
 200  0100.5e01.0101    static  Yes          -   Gi3/1,Gi3/7,Gi3/16
 100  0100.5e01.0101    static  Yes          -   Gi3/5,Gi3/7,Gi3/16
 101  0100.5e01.0101    static  Yes          -   Gi3/6,Gi3/7,Gi3/16

tstevens-6509-neb#
```

List of Receiver Ports (Also Includes Any Multicast Router Ports in VLAN)

Group Destination MAC

VLAN Where Group MAC Is Learned

Verifying IGMP Snooping Entries—CatOS

```
tstevens-6503> (enable) show multicast group
* = Static Entry. + = Permanent Entry.

                                                Destination Ports or
VLAN  Dest MAC/Route Des    [CoS]     Age       VCs / [Protocol Type]
----  ------------------    -----  ----------   ---------------------
50    01-00-5e-01-01-01                          1/1,1/3,3/1

Total Number of Entries = 1
tstevens-6503> (enable)
```

Group Destination MAC

List of Receiver Ports (Also Includes Any Multicast Router Ports in VLAN)

VLAN Where Group MAC Is Learned

Expiring TTL Traffic

- TTL expiring in multicast traffic will drive up RP CPU
- Use `show ip traffic` to see if “Bad Hop Count” incrementing

```
tstevens-6509#show ip traffic | include bad hop
         0 format errors, 0 checksum errors, 441855 bad hop count
tstevens-6509#show ip traffic | include bad hop
         0 format errors, 0 checksum errors, 448633 bad hop count
tstevens-6509#
```

- Enable CPU rate limiter for expiring TTL traffic (PFC3 only)

  `mls rate-limit all ttl-failure <rate> <burst>`

  Affects both unicast and multicast traffic

Establishing New Multicast State

- **Many new multicast sources/groups will drive RP CPU utilization**
- **Traffic punted to RP CPU on first-hop router to establish new multicast state and perform RP register encapsulation**
- **Only a problem with PIM-SM/SSM, not with Bidir**
- **Use `mls rate-limit multicast ipv4 connected <rate> <burst>` to control rate of traffic on first-hop routers (PFC3 only)**
- **Can impact convergence time (longer time to establish new state)**

Source-Only Flooding

- **New multicast source handling in VLAN environment**
- **Switch installs "source-only" MAC entries to prevent flooding of streams with no receivers**
- **Lag in time to install source-only entries (software driven)**
  - **Initial flooding unavoidable without using static group MACs**
  - **Dependent on number of new groups (~5ms for a single new sender)**
- **In Cisco IOS, two-stage aging prevents further flooding as long as source is active**
- **In CatOS, source-only entries removed periodically (every ~5m), causing continuous periodic flooding**
  - **Periodic flooding can be avoided by configuration: `set igmp flooding disable`**

Platform-Dependent Debugs/Traces

- **Generally recommend debugs/traces as last resort**
  - Beware excessive screen output, locking up console, and pegging CPU
  - Consider sending debugs/traces to logging buffer, not console/terminal
  - Enable `service timestamps debug datetime msec`!!!
  - Use caution! Recommend using only under TAC direction
- **Multicast MLS debugs in Cisco IOS (native and hybrid)**
  - `debug mls ip multicast…`
- **IGMP snooping debugs in Cisco IOS (all from SP console):**
  - `debug mls_mcast {igmp-event | igmp-pak}`
- **IGMP snooping traces in CatOS:**
  - `set trace mcast <level>`
  - Level "2" produces least output, will see high-level packet events
  - Level "5" more verbose, will see internal process triggers etc.
  - Level "10" very verbose, use caution, will see full packet decodes etc.
  - `set trace monitor enable` enables traces to terminal session (otherwise only seen on console)
  - `set trace all 0` disables all traces

Scalability Limits

- Ensure Layer 3 and Layer 2 entry capacity not exceeded
- Monitor MET utilization
- Monitor RP and SP CPU utilization
- Monitor fabric and bus utilization

Checking Layer 3 Entry Capacity

- **Multicast entries share FIB TCAM entries and hardware adjacency entries with other protocols (unicast, IPv6, MPLS)**
- **Syslogs printed when hardware capacity exceeded (total FIB exhaustion, or exceeded max-routes configuration)**
- **Check current entry status:**

  `show mls ip multicast summary`

  `show mls cef maximum-routes` (PFC3)
- **Maximum capacity by forwarding engine documented earlier in your handouts**

show mls ip multicast summary

- **Cisco IOS:**

  `show mls ip multicast summary`

- **Catalyst OS:**

  `show mls multicast`

```
6506#show mls ip multicast summary
21210 MMLS entries using 3394656 bytes of memory
Number of partial hardware-switched flows: 0
Number of complete hardware-switched flows: 21210

Directly connected subnet entry install is enabled
Hardware shortcuts for mvpn mroutes supported
Current mode of replication is Ingress
Auto-detection of replication mode is enabled
Consistency checker is enabled
Bidir gm-scan-interval: 10
6506#
```

Monitoring MET Utilization

- **Recall that MET in replication engines contains OIFs for mroutes**
- **MET is limited hardware resource (64K entries total)**
  - Allocated in fixed size blocks
  - CatOS always allocates in four-entry blocks
  - Cisco IOS allocates in variable block sizes based on need (one, two, four, eight, or 16 entries)
- **Monitor current MET utilization**
  - Inexact science—some internal “overhead” (pointers, etc.) and “wasted” space
- **Syslog generated if MET capacity reached**

```
%MMLS-SP-6-MET_LIMIT_EXCEEDED: Failed to allocate MET entry,
exceeded system limit of (65536) entries. Number of times MET
limit is exceeded in the last 1 min : 7
```

Monitoring MET Utilization (Cisco IOS)

```
tstevens-6513-sp#show mmls met
met free entries: 61462
met total entries: 65526
<…>
tstevens-6513-sp#
```

Total Free Entries

64K Total MET Entries

Monitoring MET Utilization (CatOS)

```
tstevens-6503> (enable) show mls multicast met utilization
Total blocks: 16383
Available blocks: 15915
Used blocks: 468

Table utilization: 3%
tstevens-6503> (enable)
```

Total Blocks = 16K
16k * 4 = 64k

Number of Four-Entry Blocks Still Available

Number of Four-Entry Blocks in Use

Percent MET Utilization

Layer 2 Entry Capacity

- **Limit is 15,488 Layer 2 multicast groups**

  Lower in earlier CatOS releases (3,072)

- **In Cisco IOS, limit is configurable using `ip igmp snooping l2-entry-limit` command**

- **Syslog posted when limit exceeded or hash collision occurred**

```
%MCAST-SP-6-GC_LIMIT_EXCEEDED: IGMP snooping was trying to allocate more
Layer 2 entries than what allowed (15488)
```

```
%MCAST-SP-6-L2_HASH_BUCKET_COLLISION: Failure installing (G,C)->index:
(0100.5e01.1465,1017)->0x82C
```

- **When limit exceeded, flooding occurs for additional group MACs**

Monitoring Layer 2 Entry Capacity

- `show mac-address-table multicast igmp count` (Cisco IOS)

```
tstevens-6509-neb#show mac-address-table multicast igmp count

Multicast MAC Entries for all vlans:    455
tstevens-6509-neb#
```

- `show multicast group count` (CatOS)

```
tstevens-6503> (enable) show multicast group count

Total Number of Entries = 257
tstevens-6503> (enable)
```

Port ASIC Issues

- **Port ASICs**
  - **Handle physical interface, queuing/scheduling, trunking, channeling, etc.**
- **Port ASIC problems frequently independent of multicast**
  - **Input/output errors**
  - **Port oversubscription**
  - **Use the usual commands for port problems**
    - **show interface counters errors/show counters interface (Cisco IOS)**
    - **show port counters/show mac/show counters (CatOS)**
- **Heavy OutDiscards/Xmit-Err, especially in conjunction with poor multicast performance (pixelization, etc.)**
  - **Typical of multicast flooding or other oversubscription**
  - **Speed mismatch may cause output discards (GE can transmit 100M much quicker than 100M link can process it)**

Fabric and Switching Bus Considerations

- **As with port ASICs, fabric and bus issues frequently independent of multicast**
- **Check for drops and errors on bus or fabric**

  Bus sync and CRC errors reported in syslog, e.g.:

  ```
  %SYS-3-SYS_LCPERR3:Module 13: Pinnacle #1, Frames with Bad Packet CRC
  Error (PI_CI_S_PKTCRC_ERR - 0xC7) = 29
  ```

  Check for fabric errors using `show fabric errors all`

  Monitor fabric utilization with `show fabric utilization`

  Make sure switching bus is not over-utilized using `show catalyst6000 traffic-meter` (Cisco IOS) or `show traffic` (CatOS)

Monitoring Fabric Utilization

- Monitor fabric utilization to ensure adequate bandwidth
- Ingress replication will increase fabric load
- **show fabric utilization**

```
6506#show fabric utilization
 slot    channel      speed    Ingress %     Egress %
    1          0         8G           22           23
    2          0         8G            4            9
    3          0        20G            0            1
    3          1        20G           11           12
    4          0        20G            0            1
    4          1        20G           10           13
    6          0        20G            0            1
6506#
```

Network Events

- **Network events can impact multicast forwarding**
- **Spanning tree TCNs**
  - **For IGMP snooping, no impact**
  - **But CGMP devices will purge L2 group MAC entries**
- **Route flaps and other routing issues affect multicast**
  - **Constant RPF interface changes, etc. affect stability and consistency of packet delivery**
- **Network congestion**
  - **Dropped protocol and control packets (PIM packets, IGMP, PIM registers/register stops) will negatively impact stability and reliability**

Conclusion

You should now have a thorough understanding of the Catalyst 6500 IP multicast switching architecture and packet flow, as well as key tools for troubleshooting the platform…

ANY QUESTIONS?

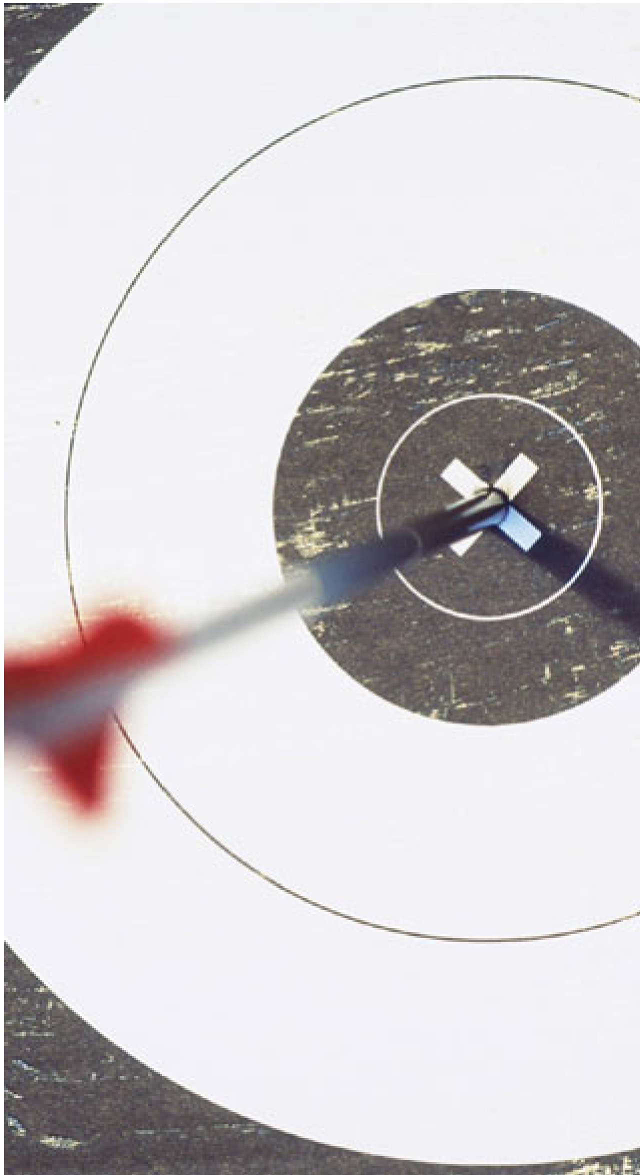

Related Networkers Sessions

- **RST-3465: Cisco Catalyst 6500 Switch Architecture**
- **RST-1261: Introduction to IP Multicast**
- **RST-2261: Deploying IP Multicast**
- **RST-2262: Multicast Security**
- **RST-2263: Multicast Network Management**
- **RST-3261: Advanced IP Multicast**
- **RST-3143: Troubleshooting Catalyst 6500 Series Switches**
- **TECRST-1008: Enterprise IP Multicast**
- **TECRST-3101: Troubleshooting Cisco Catalyst Switches**

Recommended Reading

- **Continue your Cisco Networkers learning experience with further reading from Cisco Press**
- **Check the Recommended Reading flyer for suggested books**

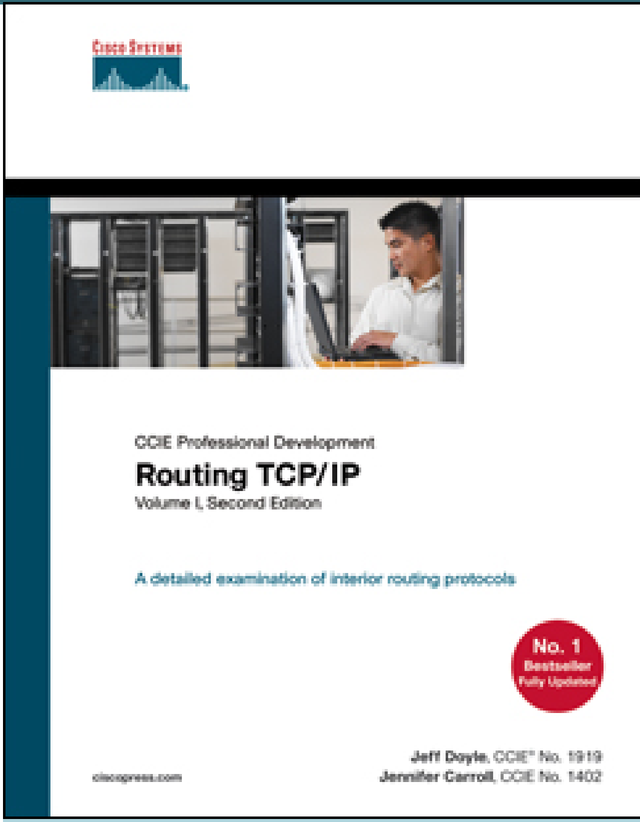

Available Onsite at the Cisco Company Store

Complete Your Online Session Evaluation

- **Win fabulous prizes; Give us your feedback**
- **Receive ten Passport Points for each session evaluation you complete**
- **Go to the Internet stations located throughout the Convention Center to complete your session evaluation**
- **Drawings will be held in the World of Solutions**

  Tuesday, June 20 at 12:15 p.m.

  Wednesday, June 21 at 12:15 p.m.

  Thursday, June 22 at 12:15 p.m. and 2:00 p.m.