Protection and Fault Recovery at Internet Peering Points using 802.1ag CFM

Rahul Vir
Product Line Manager
Foundry Networks
Oct 15, 2007

Agenda

- Peering point diagnostic challenges
- Current OAM options
- Overview of IEEE 802.1ag Connectivity Fault Management (CFM)
- Protection and Fault Recovery at Peering points using CFM
- Troubleshooting Example
- Advantages of CFM over current OAM options
- Q&A

Peering Point Diagnostic Challenges
Potential Issues

- Fiber failure, laser or electronics failure
- Card failure / Node failure
- OAM trace and loopback path does not match data path
- Difficulty in separating Exchange Point issues from peer issues
- Insufficient tools for diagnostics and fault isolation
- Want to know more?
  - Check out travails of people using co-location facilities at http://peeringforum.com/presentations/gpf-colo-preso.ppt

Peering Point Diagnostic Challenges
Desirable Features of OAM tools

- Proactive monitoring of critical links
- Provide visibility in Layer 2 network
- Ability to debug networks when component networks belong to different operators
- Diagnostic capabilities during network design and testing phase
- Troubleshooting capabilities on fault detection

Current OAM Options (1)

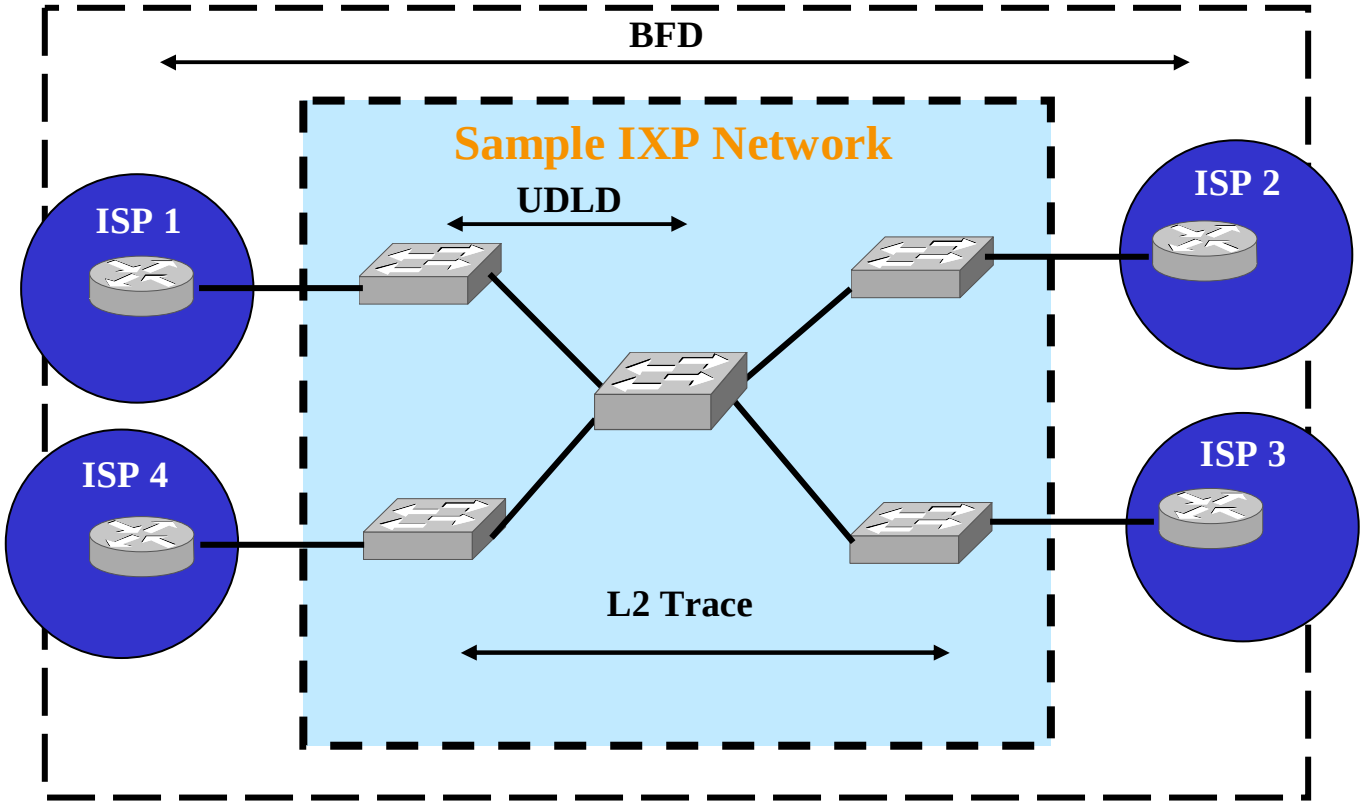

OAM Tools

- Ping/Traceroute
- Proprietary Uni-Directional Link Detection (UDLD)
- Bi-directional Forwarding Detection (BFD)
- Proprietary L2 Trace

Current OAM options (2)

- Layer 3 OAM options
  - Ping
  - Traceroute

- Uni-directional Link Detection (UDLD)
  - Provides fast detection of link failures by exchanging periodic health exchange packets

```
NetIron(config)# show link-keepalive ethernet 8/1
Current State : up                          Remote MAC Addr : 00e0.52d2.5100
Local Port : 8/1                            Remote Port : 5/1
Local System ID : e0927400                  Remote System ID : e0d25100
Packets sent : 254                          Packets received : 255
Transitions : 1
```

- Bidirectional Forwarding Detection (1-hop) for BGP and IGPs
  - Provides ability to quickly track connectivity between two directly-connected systems

```
NetIron# show bfd neighbor
Total number of Neighbor entries: 2
NeighborAddress     State   Interface   Holddown   Interval   RH
12.14.1.1           UP      eth 1/1     300000     100000     1
12.2.1.1            UP      eth 2/1     300000     100000     1
```

Current OAM options (3)
Proprietary L2 Trace

- Probe Layer 2 Topology

  NetIron # trace-l2 vlan 10
  Vlan 10 L2 topology probed, use "trace-l2 show" to display

- Display results

NetIron # trace-l2 show
Vlan 10 L2 topology was probed 6 sec ago, # of paths: 2
path 1 from e2/7, 1 hops:

hop	input	output	IP and/or MAC address	microsec	comment
1	e1/3		1.1.1.1 0004.8057.0d00	383	

path 2 from e2/5, 2 hops:

hop	input	output	IP and/or MAC address	microsec	comment
1	e2/7	e2/6	1.1.1.3 00e0.8052.ea00	657	
2	e2/8		1.1.1.4 00e0.803f.c400	296	

Overview of IEEE 802.1ag CFM

- IEEE 802.1ag Connectivity Fault Management (CFM)
- Standard for detecting, isolating and reporting connectivity faults in a network
- Facilities for multiple nested maintenance domains over a Bridged network
- Ability to cross networks maintained by different administrative organizations
- Intended for detecting and isolating faults across link layer
- Designed to be transparent to customer traffic that is transported by the network
- CFM functions that are facilitated by 802.1ag:
  - Path discovery
  - Fault detection
  - Fault verification and isolation
  - Fault notification
  - Fault recovery

Concepts and Definitions

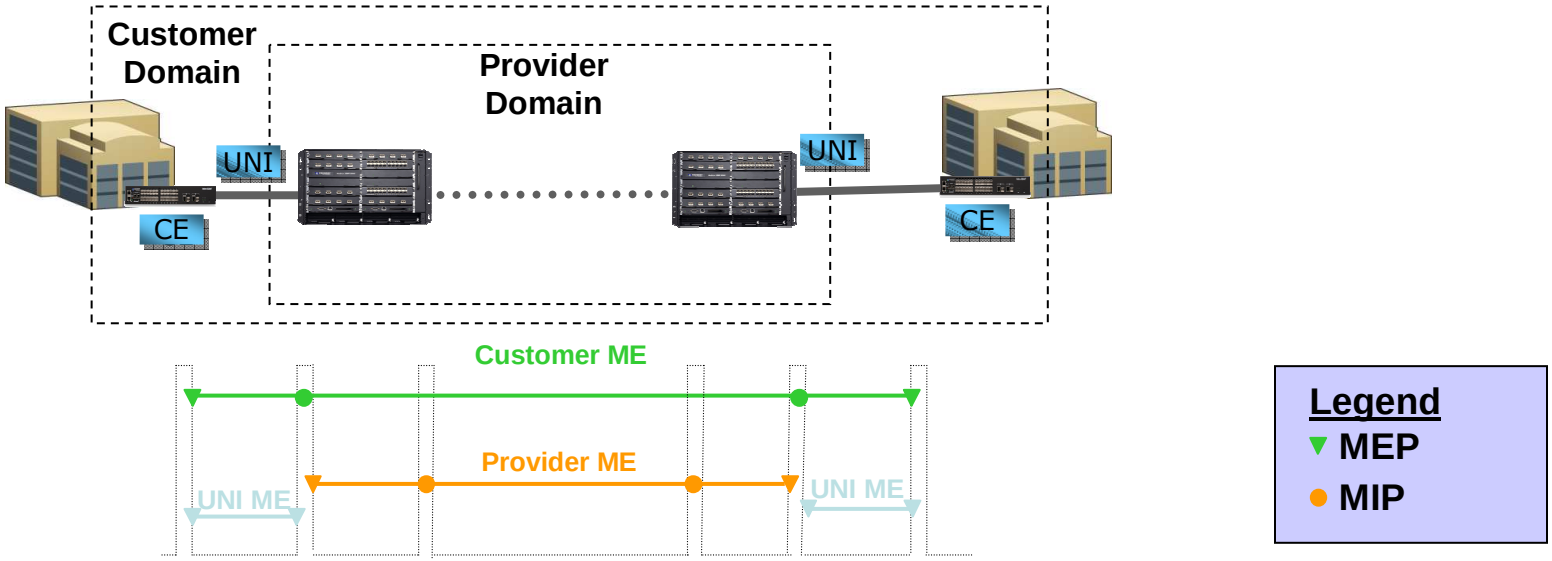

- Concepts:
  - Maintenance Entity (ME) – An OAM entity that needs management
  - Maintenance Association (MA) – MEs that belong to the same service in an OAM domain
  - MA End Point (MEP) – A provisioned reference point that can initiate/terminate proactive OAM frames
  - Maintenance Domain (MD) – A network controlled by an operator that supports connectivity between MEPs
  - MD Intermediate Point (MIP) – A provisioned reference point that can respond to diagnostic OAM frames initiated by a MEP
  - MD Level – It determines the MPs that are interested in the contents of the CFM frame and through which the CFM frame is allowed to pass.

Types of CFM messages

- Ethernet CFM messages have a special EtherType (8902). E.g.:

Destination MAC Address	Source MAC Address	8100	C-VLAN	8902	802.1ag frame data

- There are different types of CFM messages:
  a) Continuity Check Message (CCM)
  b) Loopback Message (LBM)
  c) Loopback Response (LBR)
  d) LinkTrace Message (LTM)
  e) LinkTrace Response (LTR)
- Each message type is identified by a unique Opcode:

8	5	0
MD Level	Version	
Version		
Opcode		
Flags		
First TLV Offset		
…		
End TLV(0)		

← Contents of a CFM frame

Continuity Checking

- CCM sent periodically by a Maintenance End-Point (MEP) with a multicast destination address
- Transmitted to the network at configurable intervals (3.33 msec to 10 min)
- Receiver can use it to discover the remote end-point or know the health of the transmitting station
- Loss of 3 consecutive CCM messages or receipt of a CCM with incorrect information indicates a fault
- Facility to send Remote Defect Indication (RDI) in CCM to indicate a fault
- Useful for detecting failures, cross-connect misconfigurations etc.

Loopback Operation

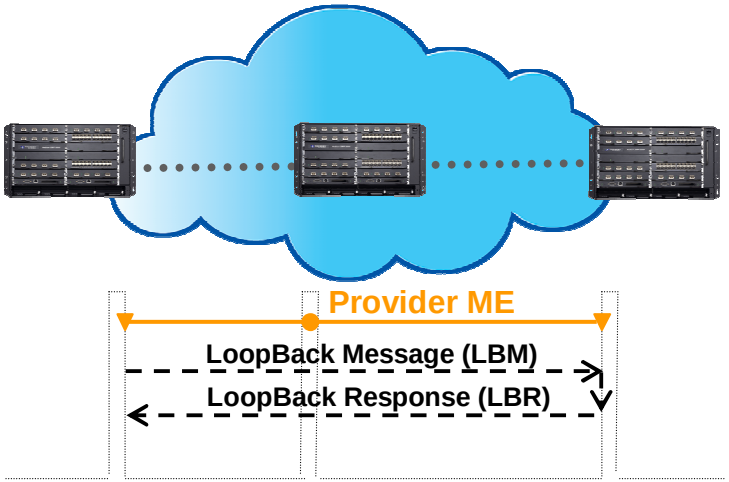

Legend
- MEP
- MIP

- A Loopback Message (LBM) is sent to a unicast destination MAC address.
- MEP at the Destination MAC address responds to the LBM message with an LBR
- Either a MEP or a MIP can respond to LBM if Destination MAC address in LBM matches that of the MAC address corresponding to the MEP/MIP
- Similar to ICMP Echo/Response (but happens at Layer 2)
- Useful for verifying connectivity with a specific Layer 2 destination

Tracing a Layer 2 Topology

- LinkTrace Message (LTM) and LinkTrace Response (LTR)

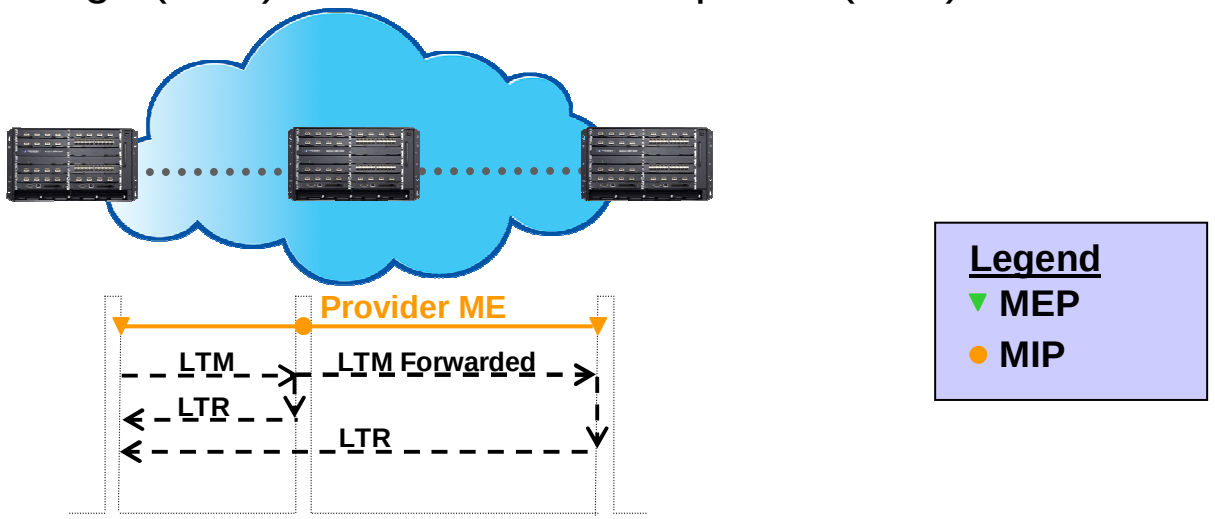

- A LinkTrace Message (LTM) is sent to a multicast MAC address; payload contains a target unicast MAC address
- Each MIP at the same MD level responds with a LinkTrace Response (LTR). Message is then forwarded to the next hop until it reaches the destination MAC
- Originating MEP collects all the LTR messages to determine path through the network
- Similar to a Layer 3 Traceroute (but happens at Layer 2)
- Useful for tracing the Layer 2 path to a specific Layer 2 destination

Protection and Fault Recovery at Peering points

- CFM tools provide
  - Path discovery using linktrace protocol
  - Fault detection using continuity check protocol
  - Fault verification and isolation using loopback and linktrace protocol
  - Fault notification provided by MEP due to loss or errors in continuity check messages
- Helps determine service or network connectivity in a Layer 2 domain
- Facilitates rapid troubleshooting and isolation of faults in an Ethernet network
- Provides visibility into Layer 2 network
- Promotes proactive detection of faults without waiting on customers to report a defect
  - Net result: Improves SLA offered to end-customer

Troubleshooting Example
Simplified Peering Network

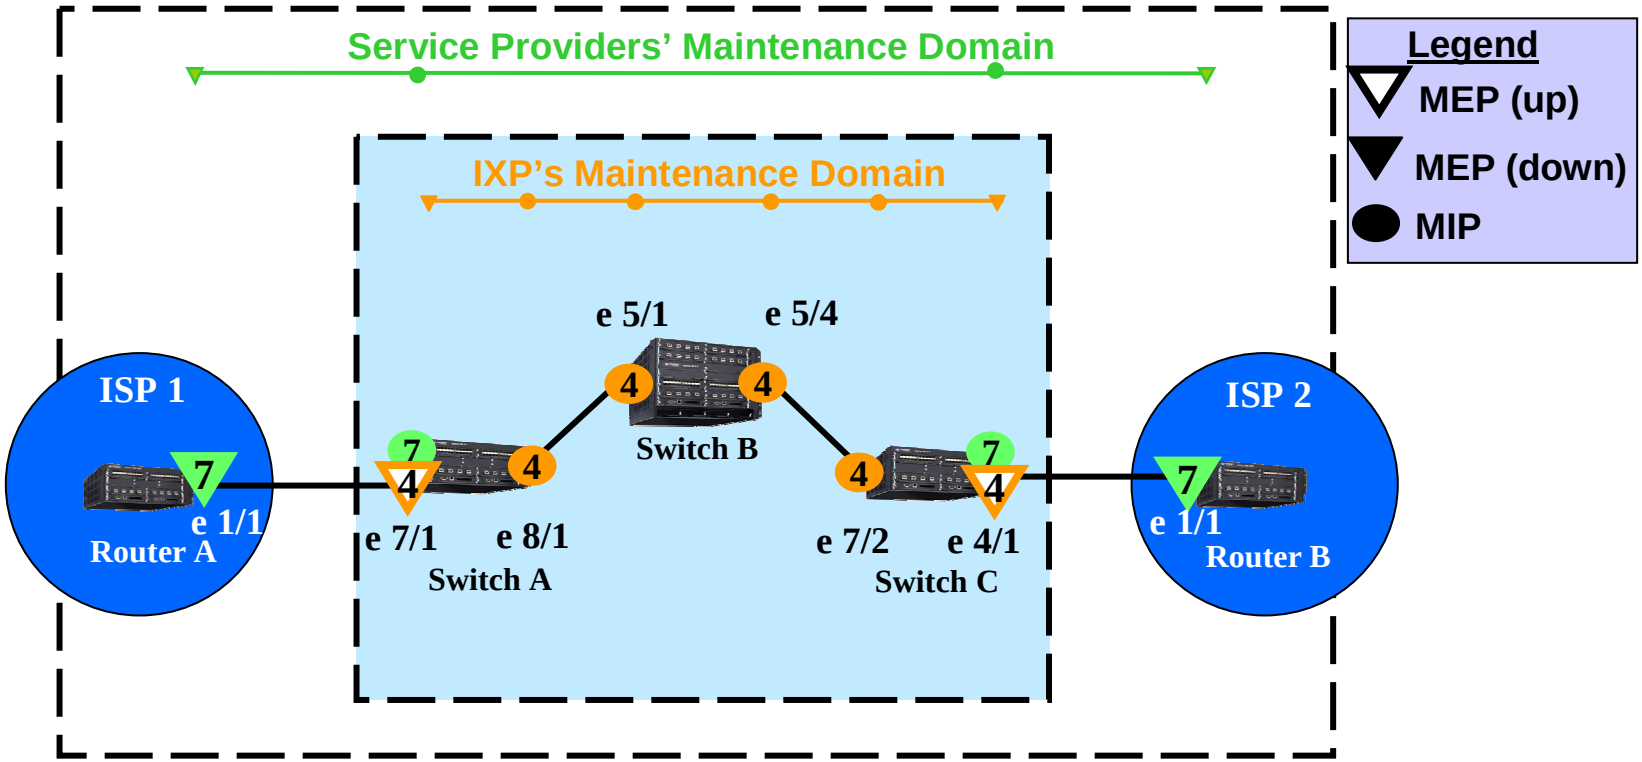

- Set MD level 4 for IXP operator, and MD level 7 for ISPs
- Configure ISPs peering interfaces and IXP customer interfaces as MEPs
- Configure MIPs in the IXP network
- MEP generates alerts on connectivity fault detection
- Both ISPs and IXPs can quickly detect faults
- Linktrace is used for fault isolation and Loopback for connectivity verification

Example of Fault Management
ISP View - Router A

- Shows local MEPs, MIPs

Router A# show cfm
Domain: md1 Level: 7
Maintenance association: ma1
CCM interval: 10
VLAN ID: 20
Priority: 1

MEP	Direction	MAC	PORT
====	=========	=========	====
22	DOWN	000c.dbf3.a700	ethe 1/1

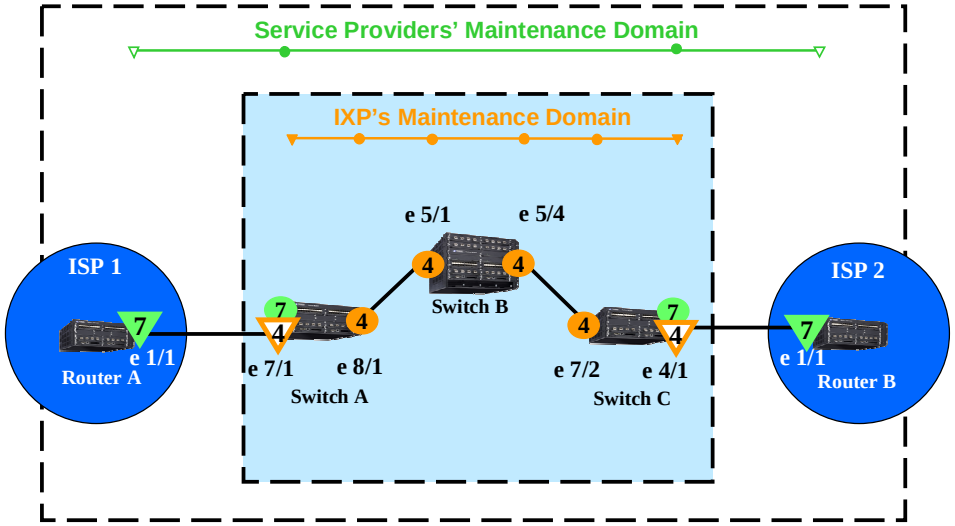

- Shows remote MEPs, MIPs

Router A# show cfm conectivity
Domain: md1 Level: 7
Maintenance association: ma1
CCM interval: 10
VLAN ID: 20
Priority: 1

RMEP	MAC	VLAN/VC	PORT
====	=====	=======	====
23	000c.dbf3.9c00	20	1/1

Example of Fault Management

ISP View - Router A

Linktrace

```
Router A# cfm linktrace domain md1 ma ma1 src-mep 22 target-mep 23
Linktrace to 000c.dbf3.9c00 on Domain md1, level 7: timeout 10ms, 8 hops
-------------------------------------------------------------------------------------------------------------------
Hops        MAC                 Ingress          Ingress Action          Relay Action
            Forwarded           Egress           Egress Action           Nexthop
-------------------------------------------------------------------------------------------------------------------
  1         0012.f23b.60f0                                               RLY_FDB
            Forwarded           8/1              EgrOK
  2         000c.dbfb.5378                                               RLY_FDB
            Forwarded           4/1              EgrOK
  3         000c.dbf3.9c00      1/1              IgrOK                   RLY_HIT
            Not Forwarded
Destination 000c.dbf3.9c00 reached
```

Loopback

```
Router A# cfm loopback domain md1 ma ma1 src-mep 22 target-mep 23
DOT1AG: Sending 10  Loopback to 000c.dbf3.9c00, timeout 10000 msec
Reply from 000c.dbf3.9c00: time<1ms
<repeats 10 times … >
A total of 10 loopback replies received.
Success rate is 100 percent (10/10), round-trip min/avg/max=0/0/1 ms.
```

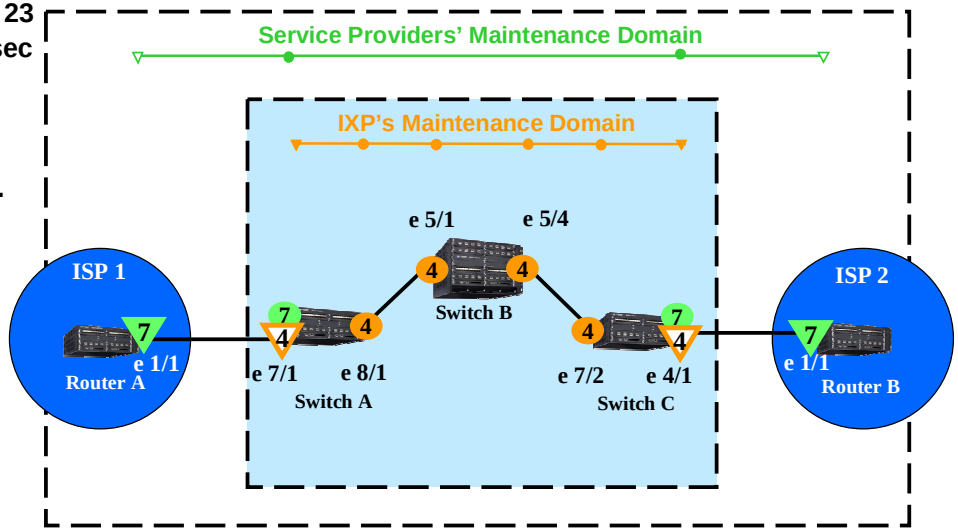

Example of Fault Management
IXP View - Switch A

- Shows local MEPs, MIPs

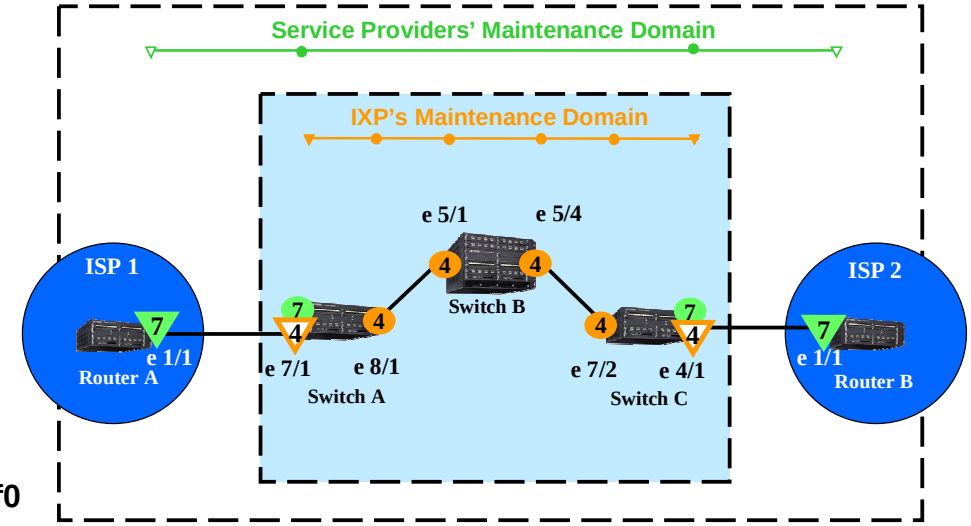

Switch A# show cfm
Domain: md1 Level: 7
Maintenance association: ma1
CCM interval: 10
VLAN ID: 20
Priority: 1

MEP	Direction	MAC	PORT
====	=========	=========	====

MIP	VLAN	Port	MAC
====	====	=====	======
	20	7/1	0012.f23b.60f0

Domain: md2 Level: 4
Maintenance association: ma2
CCM interval: 60
VLAN ID: 20
Priority: 4

MEP	Direction	MAC	PORT
====	=========	=========	====
1	UP	0012.f23b.60f0	ethe 7/1

MIP	VLAN	Port	MAC
====	====	=====	======
	20	8/1	0012.f23b.60f0

Example of Fault Management
IXP View - Switch A

- Linktrace

Switch A# cfm linktrace domain md2 ma ma2 src-mep 1 target-mep 2
Linktrace to 000c.dbfb.5378 on Domain md2, level 4: timeout 10ms, 8 hops

Hops	MAC Forwarded	Ingress Egress	Ingress Action Egress Action	Relay Action Nexthop
1	000c.dbe2.6ea0 Forwarded	5/4	EgrOK	RLY_FDB
2	000c.dbfb.5378 Not Forwarded	7/2	IgrOK	RLY_HIT

Destination 000c.dbfb.5378 reached

- Loopback

Switch A# cfm loopback domain md2 ma ma2 src-mep 1 target-mep 2
DOT1AG: Sending 10 Loopback to 000c.dbfb.5378, timeout 10000 msec
Type Control-c to abort
Reply from 000c.dbfb.5378: time<1ms
<repeats 10 times … >
A total of 10 loopback replies received.
Success rate is 100 percent (10/10), round-trip min/avg/max=0/0/0 ms.

Advantages of CFM over current OAM options

	IEEE 802.1ag	BFD	Proprietary UDLD	Proprietary L2 protocols	IP ping/ traceroute
Fault detection & Isolation	●	◐	◐	?	◐
Standards based	●	●	○	○	●
Visibility in L2 Networks	●	○	●	●	○
Visibility in L3 networks	○	●	○	○	●
Works over 10/100, GE, 10GE (future support for 40GE & 100GE)	●	●	●	●	●
Works with 802.3ad trunk groups	●	●	●	●	●
OAM domain separation to restrict visibility	●	○	○	?	○

○ Good ● Bad

Summary
CFM Advantages

- Single OAM toolset for path discovery, fault detection, fault verification and fault isolation
- Fast detection and recovery leads to improved SLAs
- Provides ability to separate exchange point issues from peer issues
- Nested domains offer ability to restrict visibility in operator's network
- Standards based avoids vendor lock-in


FOUNDRY NETWORKS
XMR 16000
FOUNDRY NETWORKS
XMR 4000
FastIron SuperX

Thank You!!