Multicast Network Management

Multicast Architecture

Andy Kessler

Cisco Networkers 2007

Agenda

- IP Multicast MIBs
- IP Multicast Syslogs
- IP Multicast NetFlow
- Network Management Systems for IP Mcast
- IP SLAs for IP Multicast
- Monitoring and Trouble shooting Examples

What Is Network Management ?

Practically, many people say:

- **Fault Detection and Isolation**
  - **Are any WAN links flapping ?**
  - **How long has that been happening ?**
- **Monitoring**
  - **Do you know what your network is doing right now ?**
  - **Do you know where your packets are ?**
- **Configuration Management**
  - **Which routers are included in that routing domain ?**
  - **Are there any obvious misconfigurations ?**

IP Multicast Net Mgmt Challenges

- **Multicast forwarding state is dynamic**
- **Best Effort Delivery**
- **No Congestion control**
  - **Requires External Monitoring**

What Is Network Management for IPmc?

- **Some people may say....**
  - **How any active mroutes do we have now ?**
  - **What data rates are they running at ?**
  - **Where are the receivers for that group ?**
  - **Is the traffic behaving as expected ?**
  - **Which RP supports that group?**
  - **How does the multicast traffic flow affect other traffic ?**
- **What do you think it includes ?**

Multicast MIBs

Multicast MIBS

IGMP	**IGMP-MIB.my** **IGMP-STD-MIB.my**
IGMP Snooping	**CISCO-IGMP-SNOOPING-MIB.my (CatOS only)**
Mroute	**IPMROUTE-MIB.my** **IPMROUTE-STD-MIB.my** **CISCO-IPMROUTE-MIB.my**
PIM	**PIM-MIB.my** **CISCO-PIM-MIB.my**
MSDP	**MSDP-MIB.my**
mVPN	**CISCO-MVPN-MIB.my**

Multicast MIBS IOS Support

	12.1E	12.2SX	12.3	12.4	12.0S
IGMP-MIB	Yes	Yes	No	No	Yes
IGMP-STD-MIB	No	No	Yes	Yes	No
MROUTE-MIB	Yes	Yes	No	No	No
MROUTE-STD-MIB	No	No	Yes	Yes	Yes
CISCO-IPMROUTE-MIB	Yes	Yes	Yes	Yes	Yes
PIM-MIB	Yes	Yes	Yes	Yes	Yes
CISCO-PIM-MIB	Yes	Yes	Yes	Yes	Yes
MSDP-MIB	No	Yes	Yes	Yes	Yes
CISCO-MVPN-MIB	No	No	No	Yes	Yes

IP-MROUTE-STD-MIB

- Based on RFC 2932
- Contains information about the status of multicast routing
- Traffic statistics
  - Packet counters per mroute
  - Packet counters per mroute, per outbound interface
    - NextHopPkts
  - Octet counters per mroute
  - Octet counters per interface – in/out

Traffic Reporting on 6500/7600

- Cat6500 traffic statistics are collected by hardware counters and updated periodically to MSFC
- Native IOS updates 25% of mroutes every 25 seconds
  - worse case stats can be 100 seconds old
  - in 12.2(18)SX this was changed to 10% with a default of 9 seconds – worse case 90 secs
- The stat update time can be adjusted with mls ip multicast flow-stat-timer <secs>
- May cause increase in CPU utilization depending on number of mroutes. Use with care.

Multicast Notifications (Traps)

Mroute	**ciscoIpMRouteMissingHeartBeats**
PIM	**pimNeighborLoss** **ciscoPimRPMappingChange** **ciscoPimInvalidRegister** **ciscoPimInvalidJoinPrune** **ciscoPimInterfaceUp** **ciscoPimInterfaceDown**
MSDP	**msdpEstablished[1]** **msdpBackwardTransition**
mVPN	**ciscoMvpnMvrfChange**

[1]Supported in latest images – CSCek00661

More Info

- **For more information about IP Multicast MIBs:**

  Search on CCO for "Multicast Network Management"

Or

- **http://www.cisco.com/go/ipmulticast**

  White Papers

  IP Multicast Network Management

Multicast Syslog Messages

Multicast Syslogs

- **There are dozens of multicast Syslog messages in these categories:**
  - **Mroute Messages**
  - **MDS Messages**
  - **PIM Messages**
  - **AUTORP Messages**
  - **MDT Messages**
  - **MSDP Messages**
  - **DVMRP Messages**
  - **MCAST Messages - Layer 2 Multicast**
- **Many customers use a correlation engine to collect and process Syslog messages – such as CNS Notification Engine**

Useful Multicast Syslogs

Invalid RP Register Syslog:

%PIM-1-INVALID_RP_REG: Received Register from 210.0.1.202 for 239.3.3.3 not willing to be RP

This message indicates that an edge router is configured with the wrong RP address. DR addr is 210.0.1.202

Some users confuse the DR addr with the source addr. New format will make the message more readable. Adding address of RP from Reg msg:

%PIM-1-INVALID_RP_REG: Received Register from router 210.0.1.202 for group 239.3.3.3, 210.1.1.3 not willing to be RP

Multicast NetFlow

NetFlow Origination

- **Developed by Darren Kerr and Barry Bruins at Cisco Systems in 1996**
  - **US Patent 6,243,667**
- **The value of information in the cache was a secondary discovery**
  - **Initially designed as a switching path**
- **NetFlow is now the primary network accounting technology in the industry**
- **Answers questions regarding IP traffic: who, what, where, when, and how**

Principle NetFlow Benefits

Service Provider

- Peering arrangements
- Network planning
- Traffic engineering
- Accounting and billing
- Security monitoring

Enterprise

- Internet access monitoring (protocol distribution, where traffic is going/coming)
- User monitoring
- Application monitoring
- Charge back billing for departments
- Security monitoring

What is a Traditional IP Flow ?

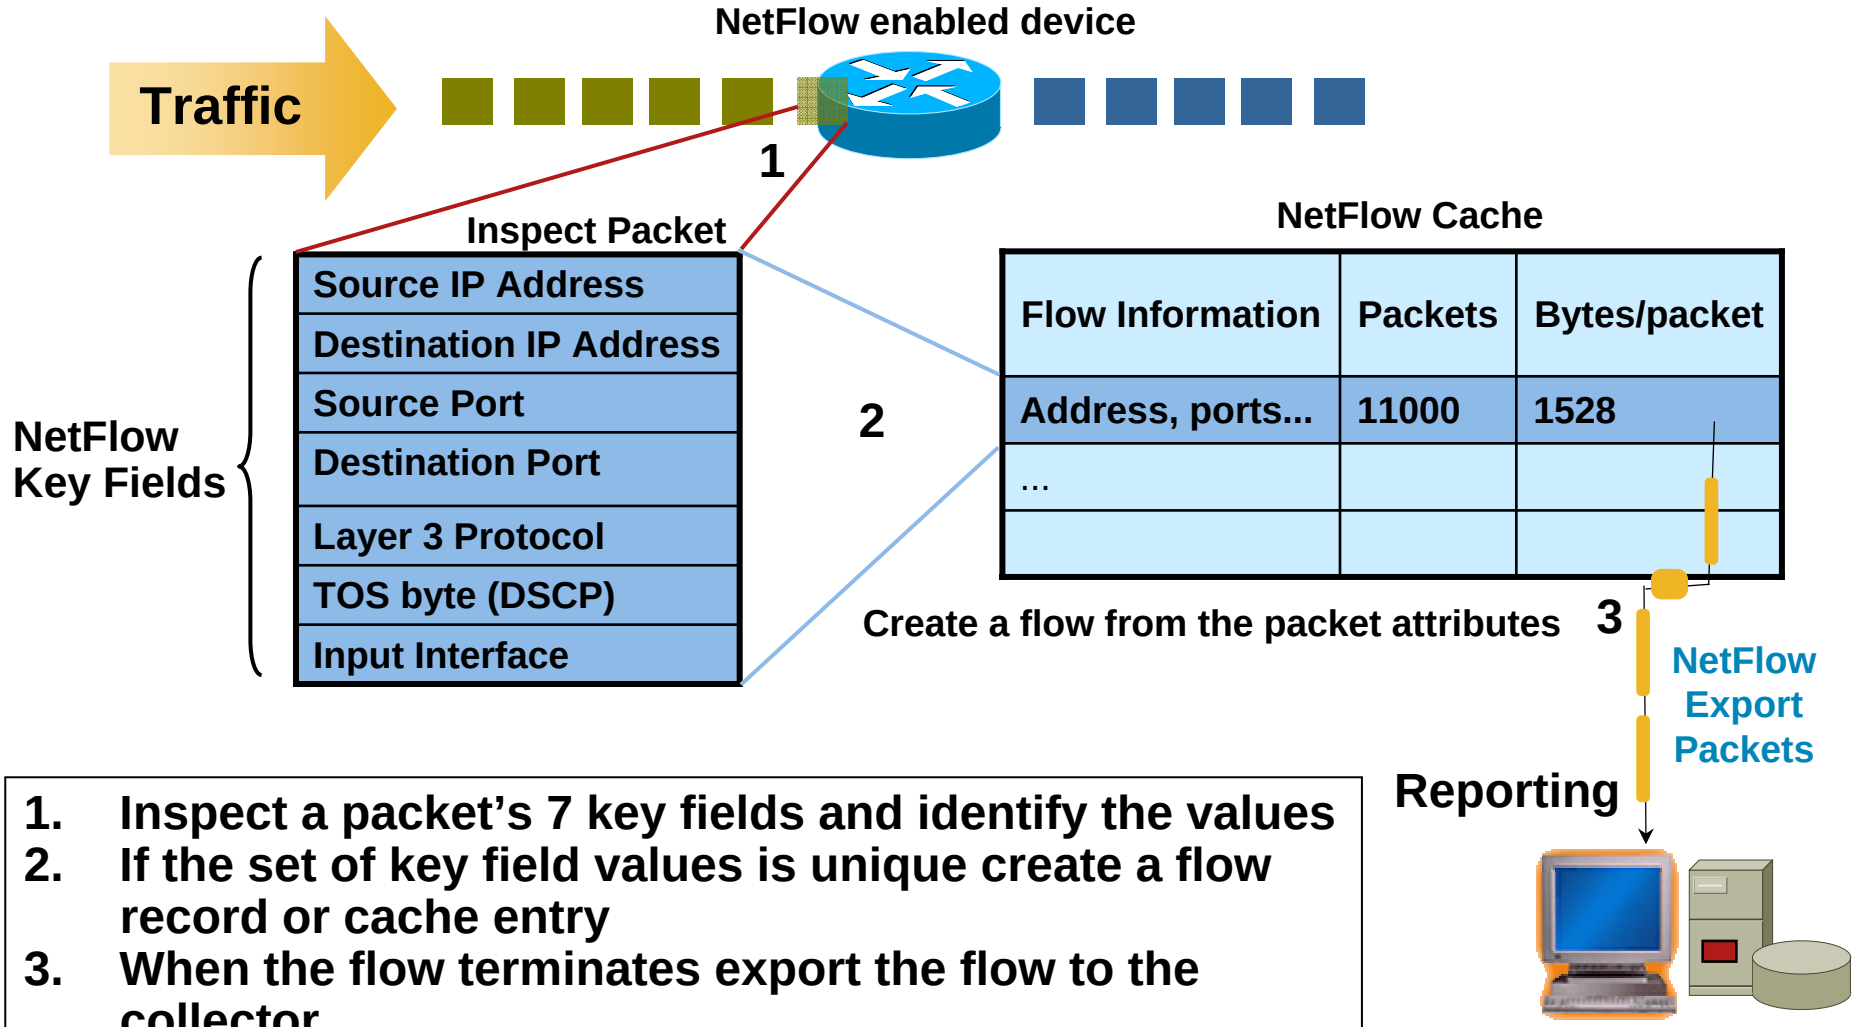

1. Inspect a packet's 7 key fields and identify the values
2. If the set of key field values is unique create a flow record or cache entry
3. When the flow terminates export the flow to the collector

Traditional Layer 3 NetFlow Cache

1. Create and update flows in NetFlow cache

Key Fields in Yellow
Non-Key Fields white

SrcIf	SrcIPadd	DstIf	DstIPadd	Protocol	TOS	Flgs	Pkts	Src Port	Src Msk	Src AS	Dst Port	Dst Msk	Dst AS	NextHop	Bytes/ Pkt	Active	Idle
Fa1/0	173.100.21.2	Fa0/0	10.0.227.12	11	80	10	11000	A2	/24	5	A2	/24	15	10.0.23.2	1528	1745	4
Fa1/0	173.100.3.2	Fa0/0	10.0.227.12	6	40	0	2491	15	/26	196	15	/24	15	10.0.23.2	740	41.5	1
Fa1/0	173.100.20.2	Fa0/0	10.0.227.12	11	80	10	10000	A1	/24	180	A1	/24	15	10.0.23.2	1428	1145.5	3
Fa1/0	173.100.6.2	Fa0/0	10.0.227.12	6	40	0	2210	19	/30	180	19	/24	15	10.0.23.2	1040	24.5	14

2. Expiration

- Inactive timer expired (15 sec is default)
- Active timer expired (30 min (1800 sec) is default)

SrcIf	SrcIPadd	DstIf	DstIPadd	Protocol	TOS	Flgs	Pkts	Src Port	Src Msk	Src AS	Dst Port	Dst Msk	Dst AS	NextHop	Bytes/ Pkt	Active	Idle
Fa1/0	173.100.21.2	Fa0/0	10.0.227.12	11	80	10	11000	A2	/24	5	A2	/24	15	10.0.23.2	1528	1800	4

4. Export version

Non-Aggregated Flows—Export Version 5 or 9

5. Transport protocol

30 Flows per 1500 byte export packet

Export Packet

Multicast NetFlow — Timers

- IP Multicast uses UDP
- UDP flows do not terminate like TCP flows with a RST or a FIN
- UDP flows depend on the aging timers to be exported
- On SW platforms this is controlled by the active timer

```
ip flow-cache timeout active 1
```

Minimum setting is 1 minute

- On 6500/7600 this is controlled by long aging timer

```
mls aging long 64
```

Minimum setting is 64 seconds

NetFlow Export Versions

NetFlow Version	Comments
1	Original
5	Most Common
7	Specific to Cisco C6500 and 7600 Series Switches Similar to Version 5, but Does Not Include AS, Interface, TCP Flag and ToS Information
8	Choice of Eleven Aggregation Schemes Reduces Resource Usage
9	Flexible, Extensible Export Format to Enable Easier Support of Additional Fields and Technologies e.g. MPLS, Multicast, BGP Next Hop, and IPv6. Defined by RFC 3954.

Multicast NetFlow

Three Types of NetFlow Implementations for Multicast Traffic:

1. **Traditional Ingress NetFlow**
2. **Multicast NetFlow Ingress**
3. **Multicast NetFlow Egress**

Multicast: Traditional NetFlow

Traditional NetFlow Configuration

```
Interface Ethernet 0
    ip route-cache flow
    or
    ip flow ingress

ip flow-export version 9

ip flow-export destination 10.255.1.1 9995
```

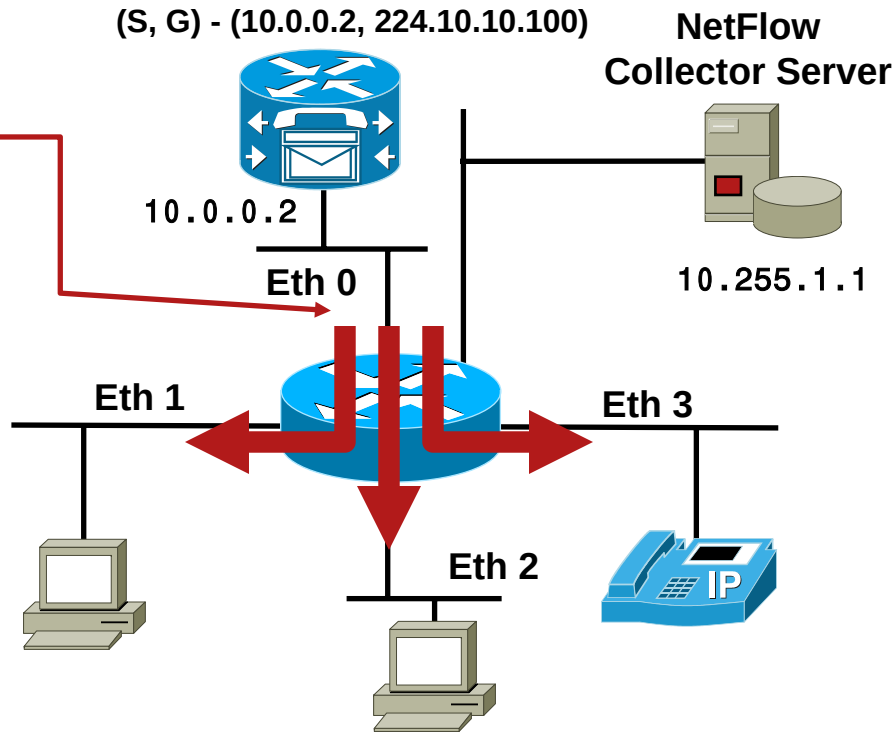

Flow Record Created in NetFlow Cache

SrcIf	SrcIPadd	DstIf	DstIPadd	Protocol	TOS	Flgs	SrcPort	DstPort	Bytes	Packets		
Eth0	10.0.0.2	Null	224.10.10.100	11	80	10	00A2	00A2	23100	21		

- There is only one flow per NetFlow configured input interface
- Destination interface is marked as “Null”
- Bytes and Packets are the incoming values

Multicast NetFlow Ingress (v9)

Multicast NetFlow Ingress Configuration

```
Interface Ethernet 0
    ip flow ingress
    ip multicast netflow ingress

ip flow-export version 9

ip flow-export destination 10.255.1.1 9995
```

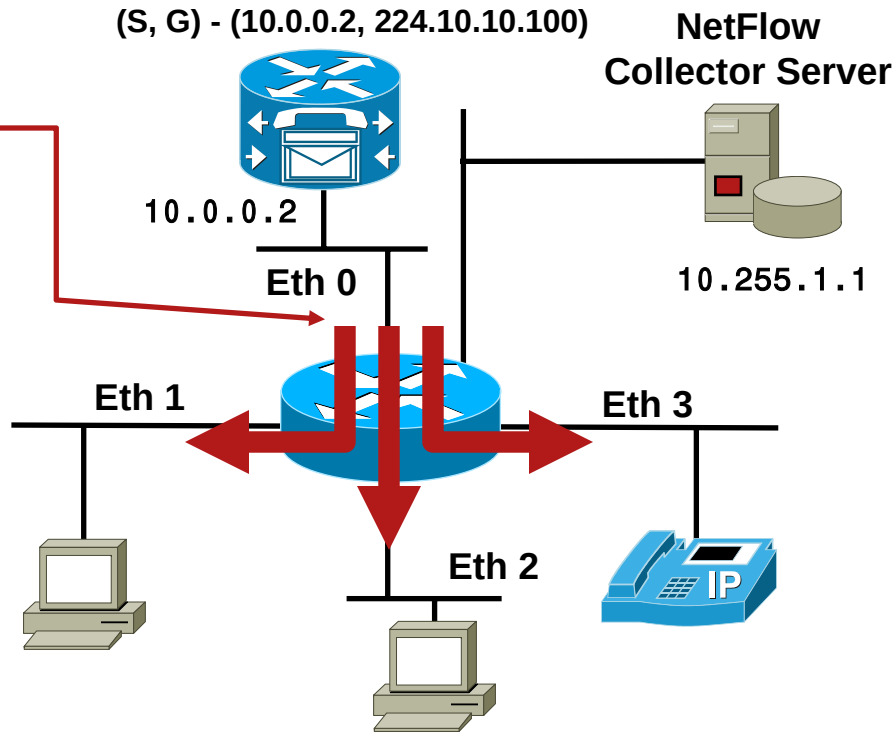

Flow Record Created in NetFlow Cache

SrcIf	SrcIPadd	DstIf	DstIPadd	Protocol	TOS	Flgs	SrcPort	DstPort	Bytes	Packets	Obytes	Opackets
Eth0	10.0.0.2	Null	224.10.10.100	11	80	10	00A2	00A2	23100	21	69300	63

- There is only one flow per NetFlow configured input interface
- Destination interface is marked as “Null”
- Bytes and Packets are the incoming values
- Obytes and Opackets are outgoing values across all interfaces – sw based routers only

Multicast NetFlow Egress (v9)

Multicast NetFlow Egress Configuration

```
Interface Ethernet 1
    ip multicast netflow egress
Interface Ethernet 2
    ip multicast netflow egress
Interface Ethernet 3
    ip multicast netflow egress


ip flow-export version 9
ip flow-export destination 10.255.1.1 9995
```

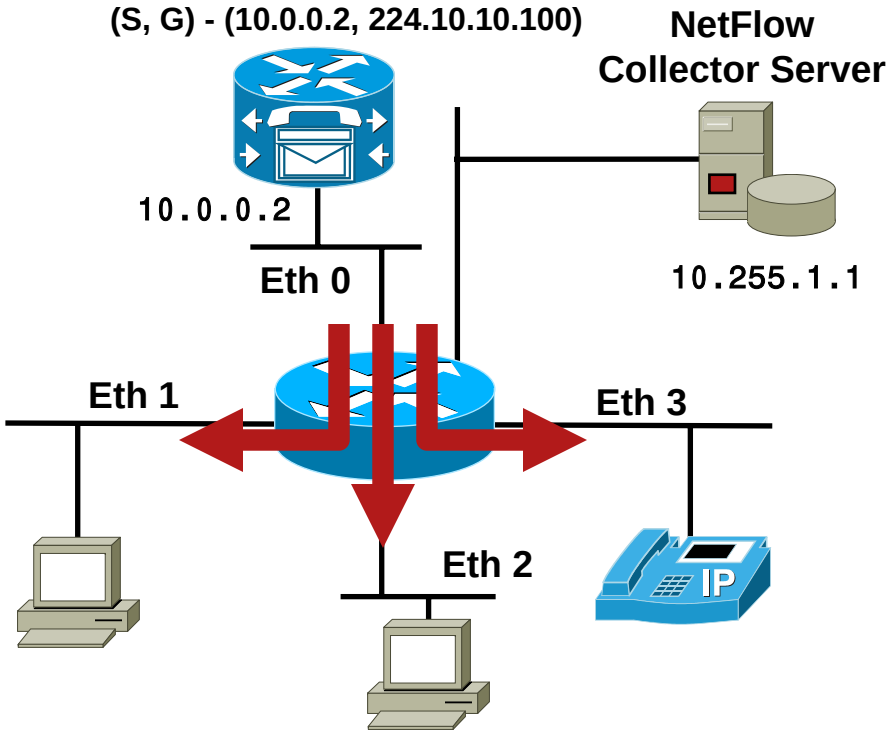

Flow Records Created in NetFlow Cache

SrcIf	SrcIPadd	DstIf	DstIPadd	Protocol	TOS	Flgs	SrcPort	DstPort	Bytes	Packets		
Eth0	10.0.0.2	Eth1*	224.10.10.100	11	80	10	00A2	00A2	23100	21		
Eth0	10.0.0.2	Eth2*	224.10.10.100	11	80	10	00A2	00A2	23100	21		
Eth0	10.0.0.2	Eth3*	224.10.10.100	11	80	10	00A2	00A2	23100	21		

- There is one flow per Multicast NetFlow Egress configured output interface
- One of the Key fields that define a unique flow has changed from source interface to destination interface
- Bytes and Packets are the outgoing values

Multicast NetFlow: Minimum Config - Ingress

Software Based Routers (e.g. 7200)

```
interface Ethernet 0
   ip flow ingress
   ip multicast netflow ingress

ip flow-export version 9
ip flow-export destination 10.255.1.1 9995
```

ip multicast netflow ingress **is not nvgened and not required**

Multicast NetFlow: Minimum Config - Egress

Software based routers (e.g. 7200)

```
interface Ethernet 0
   ip flow ingress
   ip multicast netflow egress

ip flow-export version 9
ip flow-export destination 10.255.1.1 9995
```

Multicast NetFlow: Minimum Config - Ingress

6500/7600 - Ingress

```
mls flow ip interface-full
mls nde sender
!
interface Vlan10
 ip flow ingress
 ip multicast netflow ingress
!
ip flow-export version 9
ip flow-export destination 10.255.1.1 9995
```

ip multicast netflow ingress **is not nvgened and not required**

Multicast NetFlow: Minimum Config - Egress

6500/7600 - Egress

```
mls flow ip interface-full
mls nde sender
!
interface Vlan10
 ip flow ingress   # can be configured on any interface
 ip multicast netflow egress
!
ip flow-export version 9
ip flow-export destination 10.255.1.1 9995
```

Multicast NetFlow: Export Format Summary

Software Based Router (e.g. 7200) – Ingress Accounting

SrcIf	SrcIPadd	DstIf	DstIPadd	Bytes	Packets	Obytes	Opackets
Eth0	10.0.0.2	Null	224.1.1.10	23100	21	69300	63

Software Based Router (e.g. 7200) – Egress Accounting

SrcIf	SrcIPadd	DstIf	DstIPadd	Bytes	Packets
Eth0	10.0.0.2	Eth1	224.1.1.10	23100	21
Eth0	10.0.0.2	Eth2	224.1.1.10	23100	21
Eth0	10.0.0.2	Eth3	224.1.1.10	23100	21

6500/7600 – Ingress Accounting

SrcIf	SrcIPadd	DstIf	DstIPadd	Bytes	Packets
Eth0	10.0.0.2	Null	224.1.1.10	23100	21

6500/7600 – Egress Accounting

SrcIf	SrcIPadd	DstIf	DstIPadd	Bytes	Packets
Null	10.0.0.2	Eth1	224.1.1.10	23100	21
Null	10.0.0.2	Eth2	224.1.1.10	23100	21
Null	10.0.0.2	Eth3	224.1.1.10	23100	21

Multicast NetFlow: Summary

- **Supported via NetFlow version 9 export format**
- **Performance: Ingress vs. Egress**

  Multicast NetFlow Ingress and traditional NetFlow will have similar performance numbers

  Multicast NetFlow Egress will have performance impact that is proportional to the number of interfaces on which it is enabled (include input interfaces)
- **Availability**

  Cisco IOS Software Release 12.3(1)
- **Cisco Cisco Catalyst 6500 Series and Cisco 7600 Series**

  Multicast NetFlow Ingress is supported on the PFC3A, PFC3B or PFC3B-XL in 12.2(18)SXF

  Multicast NetFlow Egress will require a PFC3B or PFC3B-XL

More Info

For more information about netflow:

http://www.cisco.com/go/netflow

Network Mgmt for PIM-SM

- RPs can be discovered through MIBs
- RP Group Ranges can be discovered for Auto-RP and BSR
- RP knows about all active groups
- Mroute MIB can retrieve the entire forwarding table
- MSDP MIB can show which RPs are running MSDP and their peering status
- IGMP MIB can show you which groups have receivers on which interfaces
- Multicast NetFlow can be used for traffic analysis

Network Mgmt for PIM-SSM

- **No RP**
  - **No central place to check for all S,Gs**
- **S,G mroutes can be tracked, measured with IP Mroute MIB**
- **IGMP MIB can give you group membership information**
  - **IGMPv3 is not supported**
  - **No source information**
- **Multicast NetFlow can be used for traffic analysis**

Network Mgmt for PIM-Bidir

- RP knows about all active groups
- No S,G Entries

  Mroute MIB and ‘show ip mroute count’ will not be able to give any info on sources

- *,G still there – MIBs OK

  Traffic info is aggregated on a group

  Source only branches

  Use show mls ip multicast rp-mapping gm-cache

- Need Source info? – Use NetFlow

  Multicast NetFlow will have all S,G info with traffic rates

Network Mgmt for mVPN

- CE routers use same mgmt tools – no change
- On PE routers the CISCO-MVPN-MIB can provide:
  - A list of all active multicast VRFs
  - How many interfaces are configured for each VRF
  - Which default and data MDTs are in use for each VRF
  - Which P Domain S,Gs are being used for each MDT
  - Which P Domain S,Gs are being used for each C Domain mroute
- The P Domain S,G can be looked up in the IPMROUTE-MIB or IPMROUTE-STD-MIB to collect statistics
- P Domain groups can be managed with normal methods
- Data MDT reuse can be tracked with Syslog

Network Management Systems (NMS) for IP Multicast

Some Multicast NMS Products

- **Cisco Multicast Manager**

- **HP OpenView NNM Smart Plug-in for IP Multicast**

- **InCharge™ IP Multicast Manager**

- **SPECTRUM® Multicast Manager**

Cisco NetFlow Applications and Partners

Traffic Analysis

CRANNOGSOFTWARE

Open Source
- **Flow-Tools**
- **FlowMon**
- **Flowd**

Denial of Service

Lancope®

Billing

More info: http://www.cisco.com/warp/public/732/Tech/nmp/netflow/partners/commercial/

Cisco IP SLAs for Multicast

What Is Cisco IP SLAs?

- Cisco IOS Feature for IP SLA measurement
- SLA Stands for Service Level Agreement
- Includes RTT, One-Way Latency, Inter-Arrival Packet Jitter, Packet Loss, and MOS Measurement Statistics
- Many protocols and applications supported
- TCP, UDP, ICMP, HTTP, RTP, FTP, DNS, DHCP, DLSW, LDP, H.323, SIP

The Idea Behind IP SLA

- **If you have a running Cisco IOS® router, turn it into an active probing device**

  The smart approach

- **Reuse your current equipment and enhance existing network management applications:**

  (CiscoWorks, IP Solution Center (ISC), Cisco InfoCenter)

New Partners 2006

Agilent Technologies

CRANNOGSOFTWARE

Multiprotocol Measurement and Management with Cisco IOS IP SLAs

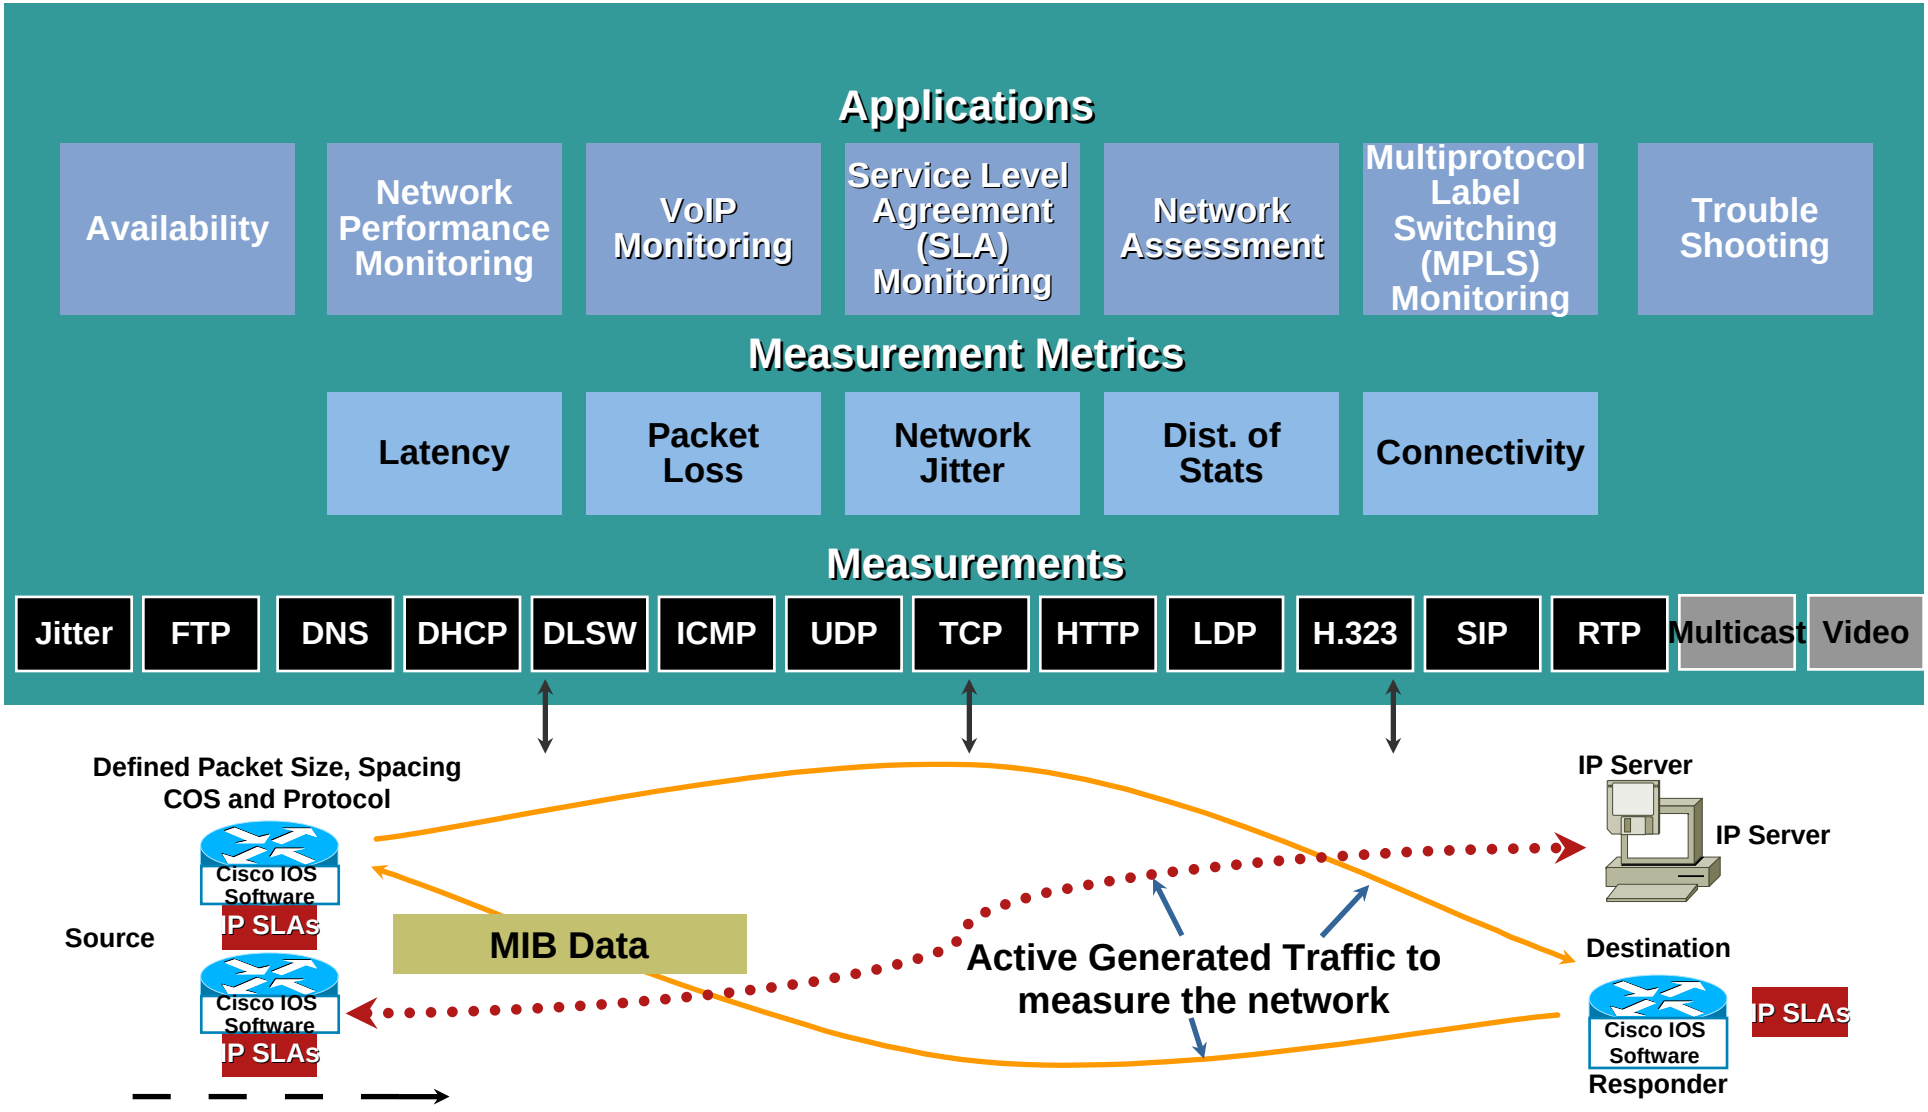

IP SLAs Summary

- Cisco IOS feature
- Provides real-time performance metrics on a wide range of protocols
- Millisecond resolution and trend visibility
- Already in Cisco IOS for unicast and available on most platforms and interfaces
- So what about multicast ?

IP SLAs for Multicast

- Source uses UDP jitter measurement with responder at receiver
- IP SLAs active measurements send stream of packets
- Packets replicated and follow multicast tree
- Receiver's with responders reply via unicast to source
- Source needs to accumulate statistics and group results

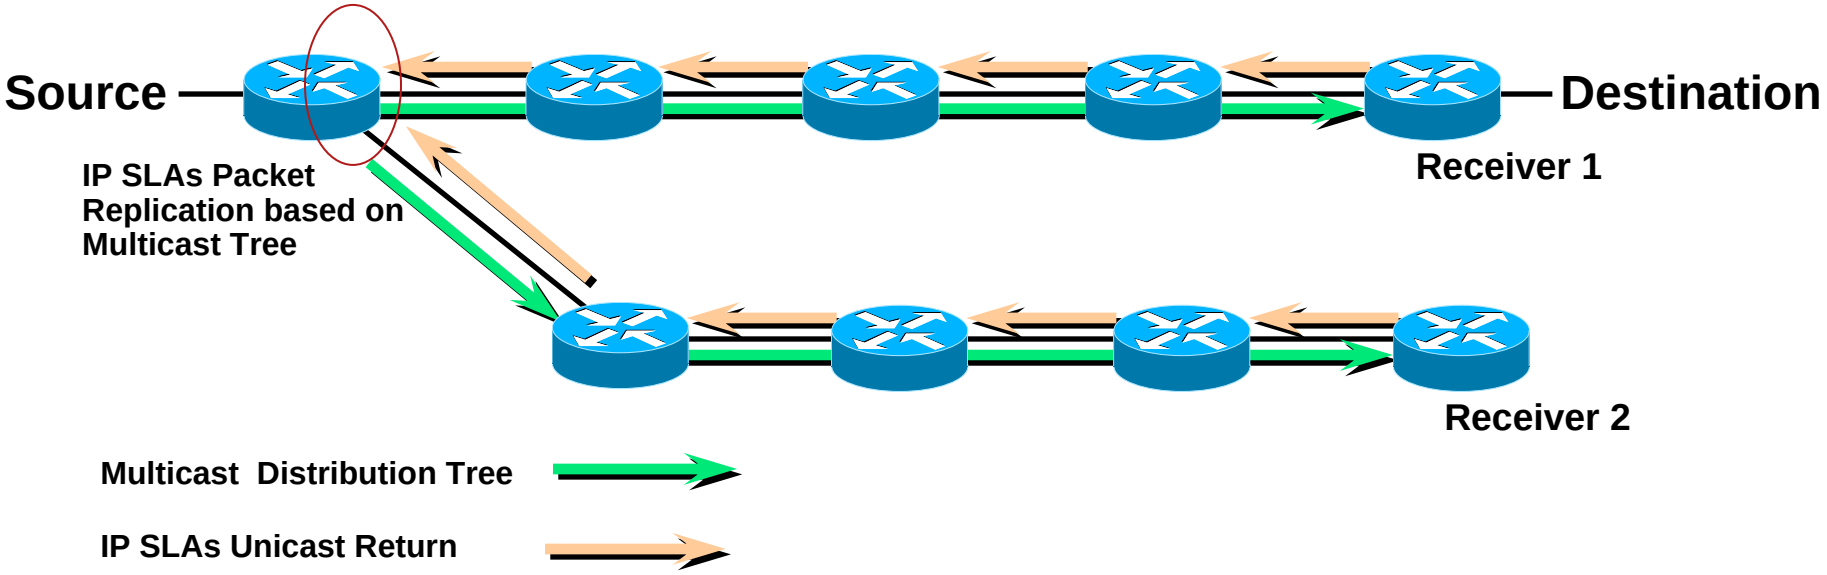

IP SLAs for Multicast: Use Cases

IP SLAs for IP Multicast could be used in these examples:

- **Finance – 30 minutes before Market Open to determine if there any network issues**
- **After a network hw/sw upgrade verify changes did not affect network performance**
- **Throughout the day using using a non-production multicast group**

IP SLAs for Multicast: Summary

IP SLAs will allow you to:

- Create multicast state in the network
- Generate multicast traffic for specific groups
- Verify the end to end multicast distribution network in terms of
  - Connectivity
  - Packet Loss
  - Latency
- Available in a future release

More Info

For more information about IP SLAs:

http://www.cisco.com/go/ipsla

Monitoring and Troubleshooting Examples

Bidir Troubleshooting with NetFlow

- **Problem: A particular Bidir group's traffic levels have jumped dramatically**
- **Might be a misconfigured source**
  - **But Bidir sources can't be seen with MIBs**
- **Solution: Use Multicast NetFlow**
  - **Individual sources can be tracked**
  - **Collectors can point out high traffic source**

Bidir Troubleshooting with NetFlow

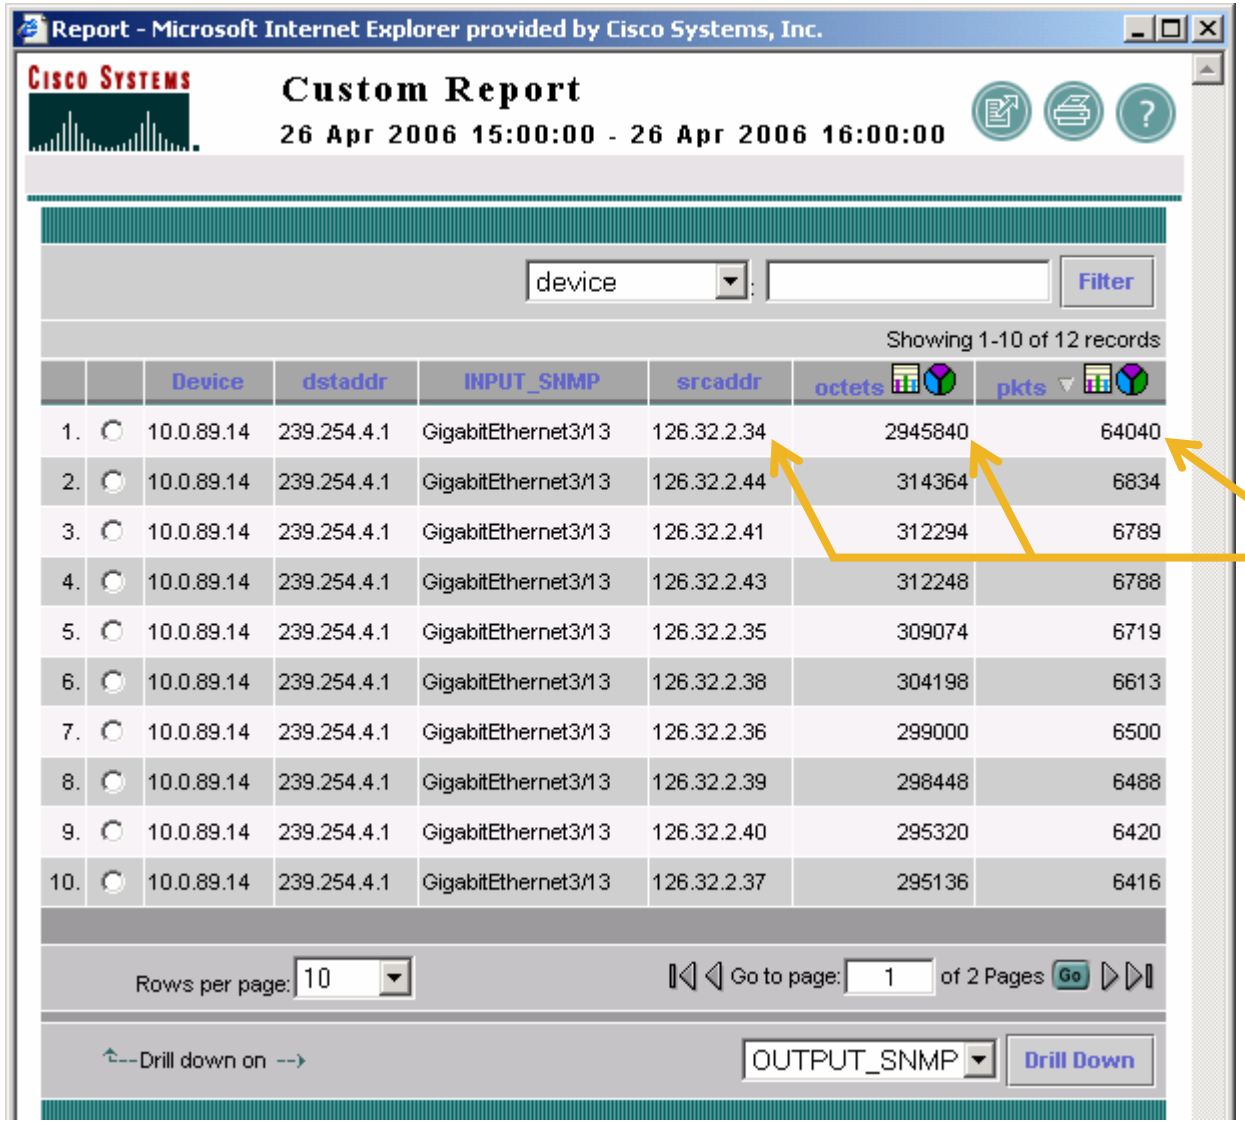

- Cisco NFC can capture all traffic to a mcast dest addr
- All sources for a group can be sorted by data rate

One source is sending significantly more traffic than others

- This host is either misconfigured or its an application problem

Questions?