Understand ICMP Redirect Messages

Contents

Introduction

This document describes the Internet Control Message Protocol (ICMP) packet redirect functionality.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Nexus 7000 platform architecture
- Cisco NX-OS Software configuration
- Internet Control Message Protocol as documented in Request for Comments (RFC) 792

Components Used

The information in this document is based on these software and hardware versions:

- Nexus 7000
- Cisco NX-OS Software

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Background Information

This document discusses packet redirect functionality provided by Internet Control Message Protocol (ICMP). The document explains what presence of ICMP Redirect messages in the network usually indicates, and what can be done to minimize negative side effects associated with network conditions that cause generation of ICMP Redirect messages.

ICMP Redirect Messages

ICMP redirect functionality is explained in RFC 792 Internet Control Message Protocol with this example:

> The gateway sends a redirect message to a host in this situation.
>
> A gateway, G1, receives an Internet datagram from a host on a network to which the gateway is attached. The gateway, G1, checks its routing table and obtains the address of the next gateway, G2, on the route to the datagram Internet destination network, X
>
> If G2 and the host identified by the Internet source address of the datagram are on the same network, a redirect message is sent to the host. The redirect message advises the host to send its traffic for network X directly to gateway G2 as this is a shorter path to the destination.
>
> The gateway forwards the original datagram data to its Internet destination.

This scenario is shown in Image 1. Host and two routers, G1 and G2, are connected to shared Ethernet segment and have IP addresses in the same network 10.0.0.0/24

Image1 - ICMP Redirects in Multi-point Ethernet Networks

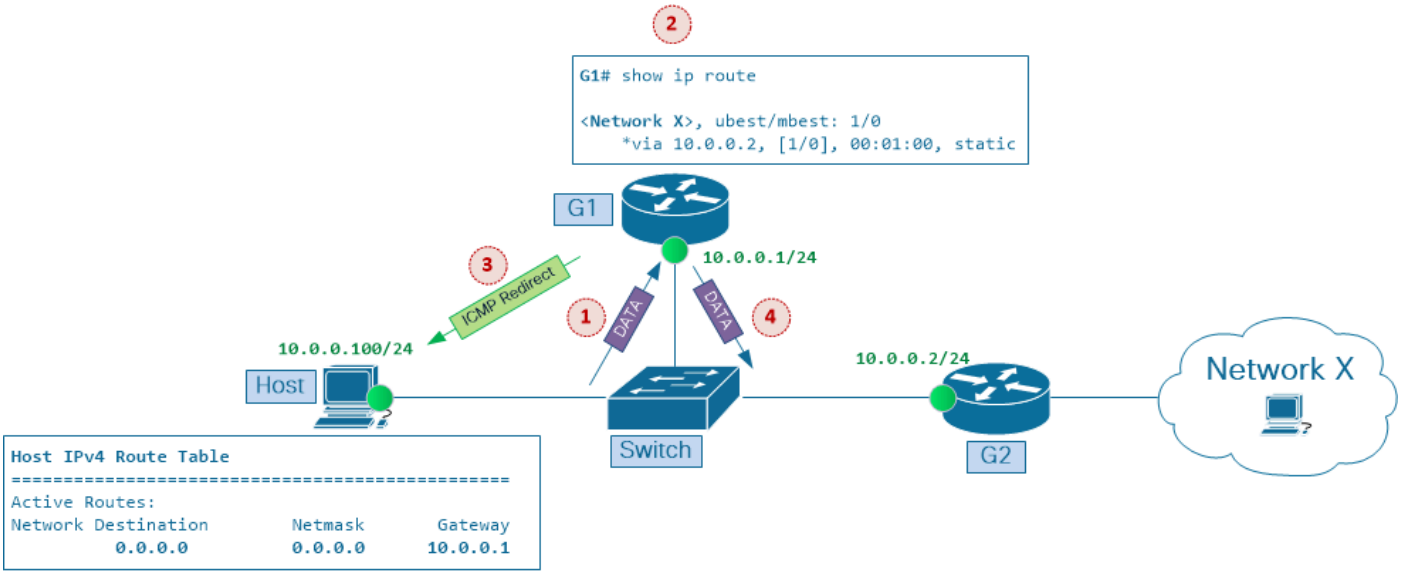

ICMP Redirects in Multi-point Ethernet Networks

Host has IP address 10.0.0.100. The Host routing table has a default route entry that points to router G1's IP address 10.0.0.1 as the default gateway. Router G1 uses router G2's IP address 10.0.0.2 as its next hop when forwarding traffic to destination network X.

This is what happens when Host sends a packet to destination network X:

1. Gateway G1 with IP address 10.0.0.1 receives data packet from host 10.0.0.100 on a network to which it is attached.

2. The gateway, G1, checks its routing table and obtains the IP address 10.0.0.2 of the next gateway, G2, on the route to the data packet destination network, X.

3. If G2 and the host identified by the source address of IP packet are on the same network, ICMP Redirect message is sent to the host. The ICMP Redirect message advises the host to send its traffic for network X directly to gateway G2 as this is a shorter path to the destination.

4. The gateway G1 forwards the original data packet to its destination.

Dependent on Host configuration, it can choose to ignore ICMP Redirect messages that G1 sends to it. However, if Host uses ICMP Redirect messages to adjust its routing cache and starts to send subsequent data packets directly to G2, these benefits are achieved in this scenario

- Optimization of data forwarding path through the network; traffic reaches its destination faster.
- Reduction of network resources utilization, such as bandwidth and router CPU load.

Image 2 - Next Hop G2 Installed in Host Routing Cache

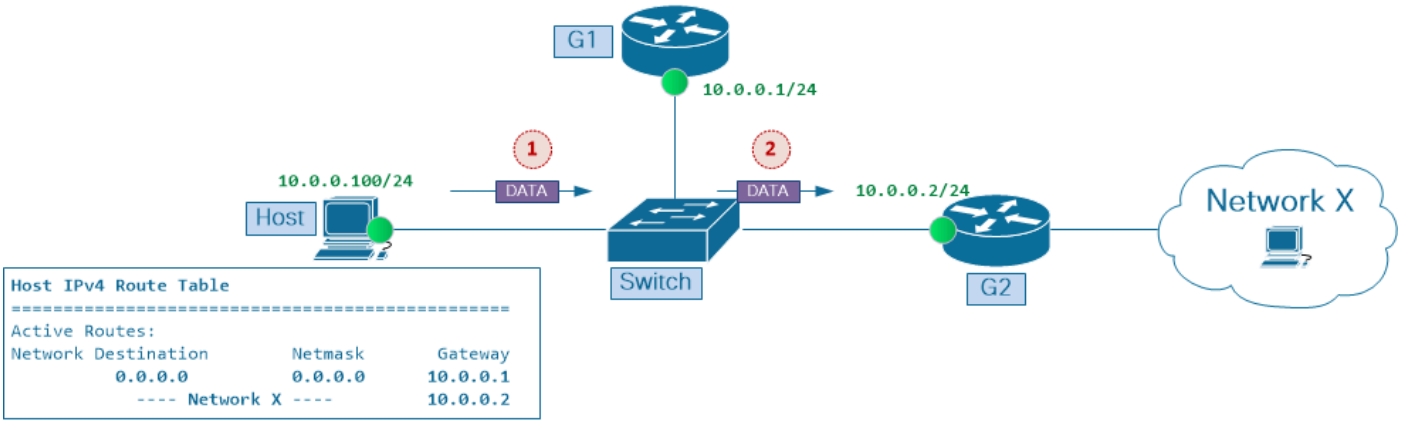

Next Hop G2 Installed in Host Routing Cache

As shown in Image 2, after the Host created route cache entry for Network X with G2 as its next hop, these benefits are seen in the network:

- Bandwidth utilization on the link between Switch and router G1 decreases in both directions.
- CPU utilization on router G1 reduces because traffic flow from Host to Network X does not traverse this node anymore.
- End-to-end network delay between Host and Network X improves.

To understand the importance of ICMP Redirect mechanism, remember that early Internet router implementations relied primarily on CPU resources to process data traffic. Hence, it was desirable to reduce the traffic volume that had to be handled by any single router and also to minimize the number of router hops that a particular traffic flow had to traverse on its way to the destination. At the same time, Layer 2 forwarding (also known as switching) was mainly implemented in customized Application-Specific Integrated Circuits (ASIC), and from forwarding performance perspective was relatively cheap compared to Layer 3 forwarding (also called routing), that, again, was done in general-purpose processors.

Newer ASIC generations can do both Layer 2 and Layer 3 packet forwarding. The Layer 3 table look-up performed in hardware helps reduce performance cost associated with packet handling by the routers.

Furthermore, when Layer 3 forwarding functionality into Layer 2 switches was integrated (which are now called Layer 3 switches), it made packet forwarding operation more efficient. This eliminated the need for one-armed router (also known as router on a stick) design options and avoided limitations associated with such network configurations.

Image 3 builds on the scenario in Image 1. Now Layer 2 and Layer 3 functions, originally provided by two separate nodes, Switch and router G1, are consolidated in a single Layer 3 Switch, such as Nexus 7000 Series platform.

Image 3 - Layer3 Switch Replaces One-armed-router Configuration

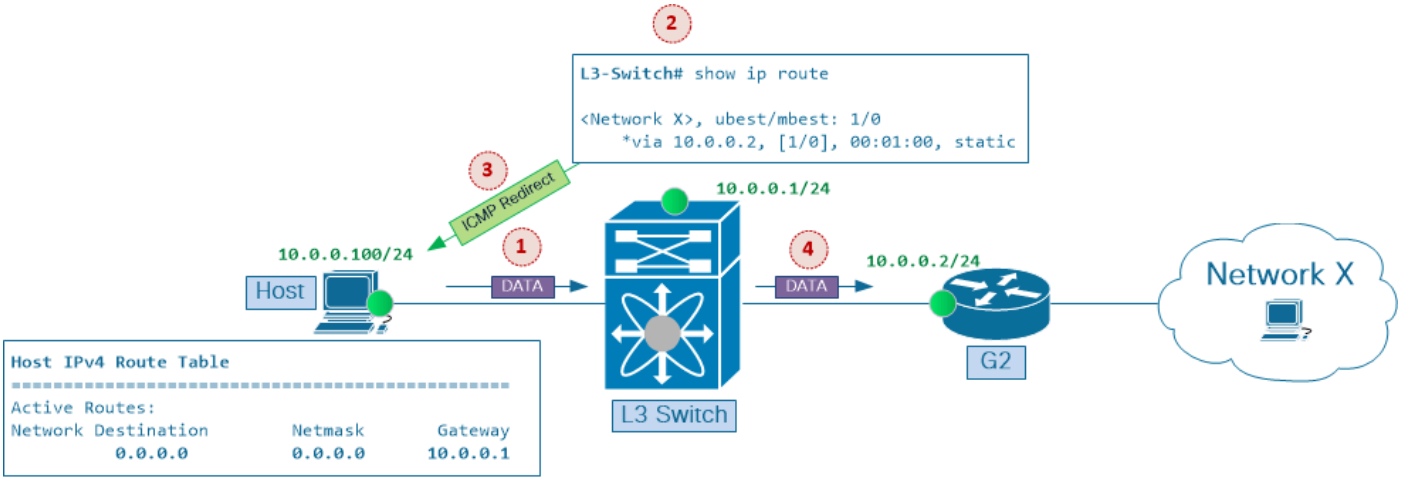

Layer3 Switch Replaces One-armed-router Configuration

This is what happens when Host sends a packet to destination Network X:

1. Gateway L3 Switch with IP address 10.0.0.1 receives data packet from a host 10.0.0.100 on a network to which it is attached.

2. The gateway, L3 Switch, checks its routing table and obtains the address 10.0.0.2 of the next gateway, G2, on the route to data packet destination network, X.

3. If G2 and the host identified by the source address of IP packet are on the same network, ICMP Redirect message is sent to the host. The ICMP Redirect message advises the host to send its traffic for Network X directly to gateway G2 as this is a shorter path to the destination.

4. The gateway forwards the original data packet to its destination.

With Layer 3 switches now able to perform both Layer 2 and Layer 3 packet forwarding at ASIC level, it can be concluded that both benefits of ICMP Redirect functionality. First, improvement of delay through the network and second, reduction of network resources utilization are achieved, and there is no more need to have much attention to path optimization techniques in multi-point Ethernet segments.

However, with ICMP Redirect functionality enabled on Layer 3 interfaces, sub-optimal forwarding through multi-point Ethernet segments continues to present potential performance bottlenecks, even though for a different reason, as explained in Nexus Platform Considerations section later in this document.

Note: ICMP Redirects are enabled by default on Layer 3 interfaces in Cisco IOS® and Cisco NX-OS software.

Note: Summary of conditions when ICMP Redirect messages are generated: Layer3 switch generates ICMP Redirect message back to the source of data packet, if data packet is to be forwarded out the Layer 3 interface on which this packet is received.

Sub-Optimal Paths through Ethernet Networks

Interior Gateway Protocols (IGP), such as Open Shortest Path First (OSPF) and Cisco Enhanced Interior Gateway Routing Protocol (EIGRP), are designed to synchronize routing information between routers, and to provide consistent and predictable packet forwarding behaviour on all network nodes that honor such information. For example, with multi-point Ethernet networks, if all Layer 3 nodes on a segment use the same routing information and agree on the same exit point to the destination, sub-optimal forwarding across such networks is rarely the case.

To understand what causes sub-optimal forwarding paths, remember that Layer 3 nodes make packet forwarding decisions independent of each other. That is, packet forwarding decision made by Router B does not depend on packet forwarding decision that was made by Router A. This is one of the key principles to remember when you troubleshoot packet forwarding through IP networks, and is an important one to keep in mind when you investigate sub-optimal forwarding path in multi-point Ethernet networks.

As mentioned earlier, in networks where all routers rely on a single dynamic routing protocol to deliver traffic between end points, sub-optimal forwarding through multi-point Ethernet segments must not happen. However, in real-world networks it is very common to find combination of various packet routing and forwarding mechanisms. Examples of such mechanisms are various IGPs, Static Routing and Policy Based Routing. These features are typically used together to achieve desired traffic forwarding through the network.

While combined use of these mechanisms can help fine tune traffic flow and meet requirements of a particular network design, they overlook side effects that these tools together can cause in multi-point Ethernet networks can result in poor overall network performance.

Static Routing

To illustrate this, consider scenario in Image 4. Router A has static route to Network X with Router B as its next-hop. At the same time Router B uses Router C as its next-hop in static route to Network X.

Image 4 - Sub-optimal Path with Static Routing

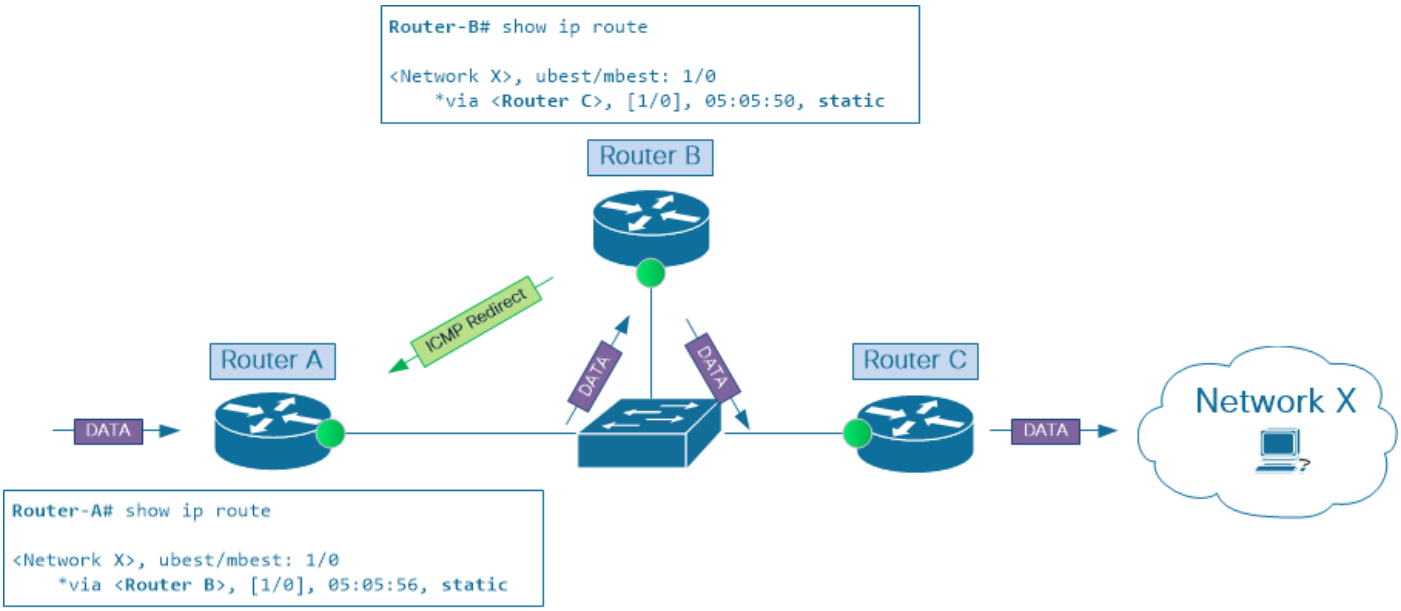

Sub-optimal Path with Static Routing

While traffic enters this network at Router A, leaves it through Router C, and eventually gets delivered to destination Network X, packets have to cross this IP network twice on their way to the destination. This is not efficient use of network resources. Instead, send packets from Router A directly to Router C would achieve the same results, while and consume less network resources.

Note: Even though in this scenario Router A and Router C are used as ingress and egress Layer 3 nodes for this IP network segment, both nodes can be replaced with network appliances (such as Load Balancers or Firewalls) if the latter have routing configuration that results in the same packet forwarding behavior.

Policy-Based Routing

Policy Based Routing (PBR) is another mechanism that can cause sub-optimal path through Ethernet networks. However, unlike Static or Dynamic Routing, PBR does not operate at routing table level. Instead, it programs traffic redirect Access Control List (ACL) directly in switch hardware. As a result, for select traffic flows, packet forwarding look-up on ingress line card bypasses routing information that is obtained via Static or Dynamic Routing.

In Image 4, Routers A and B exchange routing information about destination Network X with one of the dynamic routing protocols. Both agree on Router B is the best next-hop to this network.

However, with a PBR configuration on Router B that overrides routing information received from routing protocol and sets Router C as next-hop to network X, condition to trigger ICMP Redirect function is met and packet gets sent to the CPU of Router B to process it further.

ICMP Redirects on Point-to-Point Links

So far this document referred to Ethernet networks that have three (or more) Layer 3 nodes attached, hence the name multi-point Ethernet networks. However, be aware that ICMP Redirect messages can be generated on point-to-point Ethernet links as well.

Consider scenario on Image 5. Router A uses static default route to send traffic to Router B, while router B

has a static route to network X that points to router A.

Image 5 - ICMP Redirects on Point-to-point Links

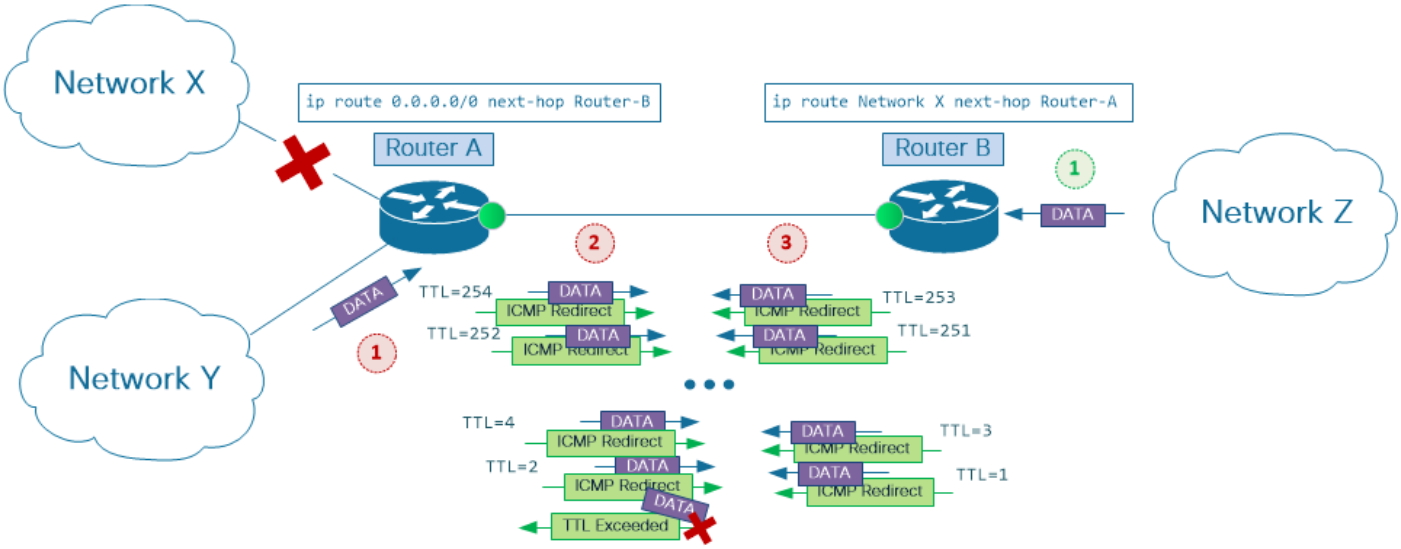

Sub-optimal Path with Static Routing

This design option, also known as single-homed connection, is a popular choice when you connect small user environments to Service Provider networks. Here Router B is a Provider Edge (PE) device, and Router A is a User Edge (CE) device.

Notice that typical CE configuration includes aggregate static route(s) to user IP address blocks that points to Null0 interface. This configuration is a recommended best practice for single-homed CE-PE connectivity option with static routing. However, for the purposes of this example, assume no such configuration is present.

Assume Router A loses connectivity to Network X as shown in the Image. When packets from the user Network Y or remote Network Z try to reach Network X, Routers A and B can bounce the traffic between each other, and decreases the IP Time-To-Live field in every packet until its value reaches 1, at which point further routing of the packet is not possible.

While traffic to Network X bounces back and forth between PE and CE routers, dramatically (and unnecessarily) increases CE-PE link bandwidth utilization. The problem becomes worse if ICMP Redirects are enabled on one or both sides of point-to-point PE-CE connection. In this case, every packet in the flow directed to Network X is processed in CPU on each router multiple times to help generate the ICMP Redirect messages.

Nexus Platform Considerations

When ICMP Redirects are enabled on Layer 3 interface and an incoming data packet uses this interface both to ingress and egress a Layer3 switch, an ICMP Redirect message is generated. While Layer 3 packet forwarding is done in hardware on Cisco Nexus 7000 platform, it is still the responsibility of the switch CPU to construct ICMP Redirect messages. To do this, CPU on Nexus 7000 Supervisor module needs to obtain IP address information of the flow whose path through the network segment can be optimized. This is the reason behind data packet sent by ingress line card to the Supervisor module.

If recipients of ICMP Redirect message ignores it and continues forwarding data traffic to Layer 3 interface of Nexus switch on which ICMP Redirects are enabled, ICMP Redirect generation process is triggered for

each data packet.

At the line card level the process starts in the form of hardware forwarding exception. Exceptions are raised on ASICs when packet forwarding operation cannot be successfully completed by the line card module. In this case, data packet needs to be sent to the Supervisor module for correct packet handling.

Note: The CPU on Supervisor module does not only generate ICMP Redirect messages, it handles many other packet forwarding exceptions, such as IP packets with Time To Live (TTL) value set to 1, or IP packets that need to get fragmented before it is sent to the next hop.

After CPU on the Supervisor module sent ICMP Redirect message to the source, it completes exception handling by forwarding data packet to the next hop through egress line card module.

While Nexus 7000 Supervisor modules use powerful CPU processors that can process large volumes of traffic, the platform is designed to handle most of the data traffic at the line card level without the need to engage the Supervisor CPU processor in packet forwarding process. This allows CPU to focus on its core tasks, and leaves packet forwarding operation to dedicated hardware engines on line cards.

In stable networks, packet forwarding exceptions, if they occur, are expected to happen at a reasonably low rates. With this assumption, they can be handled by Supervisor CPU without significant impact on its performance. On the other hand, with a CPU that deals with packet forwarding exceptions that occur at a very high rate can have a negative effect on overall system stability and responsiveness.

Nexus 7000 platform design provides a number of mechanisms to protect switch CPU from significant amounts of traffic. These mechanisms are implemented at different points in the system. At the line card level, there are hardware rate limiters and Control Plane Policing (CoPP) feature. Both set traffic rate thresholds, which effectively controls the amount of traffic to be forwarded to the Supervisor from each line card module.

These protective mechanisms give preference to the traffic of various control protocols that are critical for network stability and switch manageability, such as OSPF, BGP or SSH, and at the same time they aggressively filter types of traffic that are not critical to control plane functionality of the switch. Most of the data traffic, if forwarded to the CPU as a result of packet forwarding exceptions, is heavily policed by such mechanisms.

While hardware rate limiters and CoPP policing mechanisms provide stability of control plane of the switch and are strongly recommended to be always enabled, they can be one of the main reasons of data packet drops, transfer delays, and overall poor application performance across the network. This is why it is important to understand the paths that traffic flows take through the network and the use of tools to monitor network equipment that can and/or is expected to use ICMP Redirect functionality.

Tools to Monitor and Diagnose Traffic

show ip traffic

Both Cisco IOS and Cisco NX-OS software provide a way to check statistics of the traffic that is handled by CPU. This is done with **show ip traffic** command. This command can be used to check receipt and/or

generation of ICMP Redirect messages by Layer 3 switch or router.

```
<#root>

Nexus7000#

show ip traffic | begin ICMP


ICMP Software Processed Traffic Statistics
----------------------------------------
Transmission:
Redirect: 1000, unreachable: 0, echo request: 0, echo reply: 0,

<output omitted for brevity>

 ICMP originate Req: 0, Redirects Originate Req: 1000
 Originate deny - Resource fail: 0, short ip: 0, icmp: 0, others: 0
Reception:
 Redirect: 0, unreachable: 0, echo request: 0, echo reply: 0,

<output omitted for brevity>

Nexus7000#
```

Run show ip traffic command a few times and check whether ICMP Redirect counters increment.

Ethanalyzer

Cisco NX-OS software has a built-in tool to capture traffic flowing to and from switch CPU, known as Ethanalyzer.

Note: For more information on Ethanalyzer, refer to Ethanalyzer on Nexus 7000 Troubleshooting Guide.

Image 6 shows scenario similar to the one on Image 3. Here Network X is replaced by 192.168.0.0/24 network.

Image 6 - Run Ethanalyzer Capture

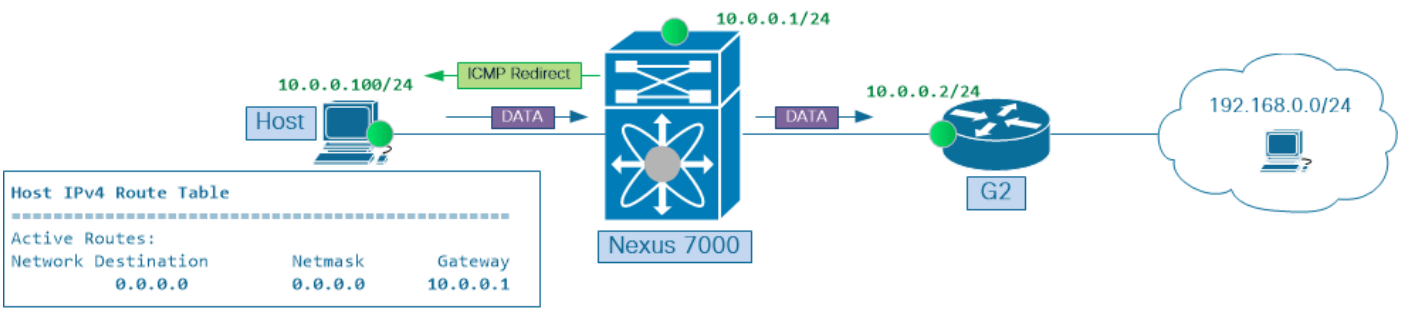

Run Ethanalyzer Capture

The Host 10.0.0.100 sends a continuous stream of ICMP Echo Requests to destination IP address 192.168.0.1. The Host uses Switch Virtual Interface (SVI) 10 of Nexus 7000 switch as its next hop to remote network 192.168.0.0/24. For demonstration purposes, the Host is configured to ignore ICMP Redirect messages.

Use this next command to capture ICMP traffic received and sent by Nexus 7000 CPU:

```
<#root>

Nexus7000#

ethanalyzer local interface inband capture-filter icmp limit-captured-frames 1000


Capturing on inband
2018-09-15 23:45:40.124077   10.0.0.100 -> 192.168.0.1      ICMP Echo (ping) request
2018-09-15 23:45:40.124477     10.0.0.1 -> 10.0.0.100   ICMP Redirect (Redirect for host)
2018-09-15 23:45:40.124533   10.0.0.100 -> 192.168.0.1      ICMP Echo (ping) request
2018-09-15 23:45:40.126344   10.0.0.100 -> 192.168.0.1      ICMP Echo (ping) request
2018-09-15 23:45:40.126607     10.0.0.1 -> 10.0.0.100   ICMP Redirect (Redirect for host)
2018-09-15 23:45:40.126655   10.0.0.100 -> 192.168.0.1      ICMP Echo (ping) request

    2018-09-15 23:45:40.128348   10.0.0.100 -> 192.168.0.1      ICMP Echo (ping) request
    2018-09-15 23:45:40.128611     10.0.0.1 -> 10.0.0.100   ICMP Redirect (Redirect for host)
    2018-09-15 23:45:40.128659   10.0.0.100 -> 192.168.0.1      ICMP Echo (ping) request

2018-09-15 23:45:40.130362   10.0.0.100 -> 192.168.0.1      ICMP Echo (ping) request
2018-09-15 23:45:40.130621     10.0.0.1 -> 10.0.0.100   ICMP Redirect (Redirect for host)
2018-09-15 23:45:40.130669   10.0.0.100 -> 192.168.0.1      ICMP Echo (ping) request
2018-09-15 23:45:40.132392   10.0.0.100 -> 192.168.0.1      ICMP Echo (ping) request
2018-09-15 23:45:40.132652     10.0.0.1 -> 10.0.0.100   ICMP Redirect (Redirect for host)
2018-09-15 23:45:40.132700   10.0.0.100 -> 192.168.0.1      ICMP Echo (ping) request
2018-09-15 23:45:40.134347   10.0.0.100 -> 192.168.0.1      ICMP Echo (ping) request
2018-09-15 23:45:40.134612     10.0.0.1 -> 10.0.0.100   ICMP Redirect (Redirect for host)
2018-09-15 23:45:40.134660   10.0.0.100 -> 192.168.0.1      ICMP Echo (ping) request
2018-09-15 23:45:40.136347   10.0.0.100 -> 192.168.0.1      ICMP Echo (ping) request
2018-09-15 23:45:40.136598     10.0.0.1 -> 10.0.0.100   ICMP Redirect (Redirect for host)
2018-09-15 23:45:40.136645   10.0.0.100 -> 192.168.0.1      ICMP Echo (ping) request
2018-09-15 23:45:40.138351   10.0.0.100 -> 192.168.0.1      ICMP Echo (ping) request
2018-09-15 23:45:40.138607     10.0.0.1 -> 10.0.0.100   ICMP Redirect (Redirect for host)
2018-09-15 23:45:40.138656   10.0.0.100 -> 192.168.0.1      ICMP Echo (ping) request

...
```

Timestamps in the previous output suggest that three packets highlighted in this example were captured at the same time, 2018-09-15 23:45:40.128. The next is a per-packet breakdown of this packet group

- First packet is the ingress data packet, which in this example is an ICMP Echo Request.

  2018-09-15 23:45:40.128348 10.0.0.100 -> 192.168.0.1 ICMP Echo (ping) request

- Second packet is an ICMP Redirect packet, generated by gateway. This packet is sent back to the host.

  2018-09-15 23:45:40.128611 10.0.0.1 -> 10.0.0.100 ICMP Redirect (Redirect for host)

- Third packet is the data packet captured in egress direction, after it has been routed by the CPU. Though not shown previously, this packet has its IP TTL decremented and checksum re-calculated.

2018-09-15 23:45:40.128659 10.0.0.100 -> 192.168.0.1 ICMP Echo (ping) request

While you navigate through large Ethanalyzer captures that have many packets of different types and flows, it can be difficult to correlate ICMP Redirect messages with the data traffic that corresponds to them.

In these situations, focus on ICMP Redirect messages to retrieve information about sub-optimally forwarded traffic flows. ICMP Redirect messages include the Internet header plus the first 64 bits of the original datagram data. This data is used by the source of the datagram to match the message to the appropriate process.

Use Ethanalyzer packet capture tool with **detail** keyword to display content of ICMP Redirect messages and find IP address information of the data flow which is sub-optimally forwarded.

<#root>

```
Nexus7000#
```

```
ethanalyzer local interface inband capture-filter icmp limit-captured-frames 1000 detail
```

```
...

Frame 2 (70 bytes on wire, 70 bytes captured)
 Arrival Time: Sep 15, 2018 23:54:04.388577000
 [Time delta from previous captured frame: 0.000426000 seconds]
 [Time delta from previous displayed frame: 0.000426000 seconds]
 [Time since reference or first frame: 0.000426000 seconds]
 Frame Number: 2
 Frame Length: 70 bytes
 Capture Length: 70 bytes
 [Frame is marked: False]
 [Protocols in frame: eth:ip:icmp:ip:icmp:data]
Ethernet II, Src: 00:be:75:f2:9e:bf (00:be:75:f2:9e:bf), Dst: 00:0a:00:0a:00:0a (00:0a:00:0a:00:0a)
 Destination: 00:0a:00:0a:00:0a (00:0a:00:0a:00:0a)
 Address: 00:0a:00:0a:00:0a (00:0a:00:0a:00:0a)
 .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
 .... ..0. .... .... .... .... = LG bit: Globally unique address (factory default)
 Source: 00:be:75:f2:9e:bf (00:be:75:f2:9e:bf)
 Address: 00:be:75:f2:9e:bf (00:be:75:f2:9e:bf)
 .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
 .... ..0. .... .... .... .... = LG bit: Globally unique address (factory default)
 Type: IP (0x0800)
```

```
Internet Protocol, Src: 10.0.0.1 (10.0.0.1), Dst: 10.0.0.100 (10.0.0.100)
```

```
 Version: 4
 Header length: 20 bytes
 Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00)
 0000 00.. = Differentiated Services Codepoint: Default (0x00)
 .... ..0. = ECN-Capable Transport (ECT): 0
 .... ...0 = ECN-CE: 0
 Total Length: 56
 Identification: 0xf986 (63878)
 Flags: 0x00
 0.. = Reserved bit: Not Set
 .0. = Don't fragment: Not Set
 ..0 = More fragments: Not Set
 Fragment offset: 0
 Time to live: 255
```

```
 Protocol: ICMP (0x01)
 Header checksum: 0xadd9 [correct]
 [Good: True]
 [Bad : False]
 Source: 10.0.0.1 (10.0.0.1)
 Destination: 10.0.0.100 (10.0.0.100)
```

Internet Control Message Protocol

Type: 5 (Redirect)

Code: 1 (Redirect for host)

```
 Checksum: 0xb8e5 [correct]
 Gateway address: 10.0.0.2 (10.0.0.2)
 Internet Protocol, Src: 10.0.0.100 (10.0.0.100), Dst: 192.168.0.1 (192.168.0.1)
 Version: 4
 Header length: 20 bytes
 Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00)
 0000 00.. = Differentiated Services Codepoint: Default (0x00)
 .... ..0. = ECN-Capable Transport (ECT): 0
 .... ...0 = ECN-CE: 0
 Total Length: 84
 Identification: 0xf986 (63878)
 Flags: 0x00
 0.. = Reserved bit: Not Set
 .0. = Don't fragment: Not Set
 ..0 = More fragments: Not Set
 Fragment offset: 0
 Time to live: 254
 Protocol: ICMP (0x01)
 Header checksum: 0xa8ae [correct]
 [Good: True]
 [Bad : False]
 Source: 10.0.0.100 (10.0.0.100)
 Destination: 192.168.0.1 (192.168.0.1)
 Internet Control Message Protocol
 Type: 8 (Echo (ping) request)
 Code: 0 ()
 Checksum: 0x02f9 [incorrect, should be 0xcae1]
Identifier: 0xa01d
Sequence number: 36096 (0x8d00)

...
```

Disable ICMP Redirects

If network design requires traffic flow to be routed out of the same Layer 3 interface on which it entered the switch or router, it is possible to prevent the flow from routing through the CPU if you disable the ICMP Redirect functionality on Layer 3 interface that corresponds to it.

In fact, for most networks it is a good practice to proactively disable ICMP Redirects on all Layer 3 interfaces, both physical, like Ethernet interface, and virtual, like Port-Channel and SVI interfaces. Use the

no ip redirects Cisco NX-OS interface-level command to disable ICMP Redirects on a Layer 3 interface. To verify that ICMP Redirect functionality is disabled:

- Ensure**no ip redirects** command is added to interface configuration.

<#root>

Nexus7000#

show run interface vlan 10

interface Vlan10
  no shutdown

no ip redirects

  ip address 10.0.0.1/24

- Ensure that status of ICMP Redirects on the interface shows disabled.

<#root>

Nexus7000#

show ip interface vlan 10 | include redirects

  IP icmp redirects:

disabled

- Ensure that ICMP Redirect enable/disable flag is set to **0** by Cisco NX-OS software component that pushes interface configuration from switch Supervisor to one of more line cards.

<#root>

Nexus7000#

show system internal eltm info interface vlan 10 | i icmp_redirect

  per_pkt_ls_en = 0,

icmp_redirect = 0

, v4_same_if_check = 0

- Ensure that ICMP Redirect enable/disable flag for a particular Layer 3 interface is set to **0** on one or more line cards.

<#root>

Nexus7000#

```
attach module 7


Attaching to module 7 ...
To exit type 'exit', to abort type '$.'
Last login: Wed Sep 15 23:56:25 UTC 2018 from 127.1.1.1 on pts/0
module-7#

!--- Optionally, jump to non-admin Virtual Device Context (VDC) if verification needs to be done in one
module-7#

vdc 6



module-7#

show system internal iftmc info interface vlan 10 | include icmp_redirect



icmp_redirect : 0x0

 ipv6_redirect : 0x1
```

Summary

ICMP Redirect mechanism, as described in RFC 792, was designed to optimize forwarding path through multi-point network segments. At the start of the Internet, such optimization helped to protect expensive network resources, like link bandwidth and routers' CPU cycles. As network bandwidth became more affordable, and relatively slow CPU-based packet routing evolved into faster Layer 3 packet forwarding in dedicated hardware ASICs, the importance of optimal data transit through multi-point network segments decreased. By default, ICMP Redirect functionality is enabled on every Layer 3 interface. However, its attempts to notify network nodes on multi-point Ethernet segments about optimal forwarding paths are not always understood and acted upon by network personnel. In networks with combined use of various forwarding mechanisms, such as Static Routing, Dynamic Routing and Policy-Based Routing, if you leave ICMP Redirect functionality enabled and do not monitor it properly, this can result in undesirable use of transit node(s) CPU to handle production traffic. This, in turn, can cause significant impact both on production traffic flows and on control plane stability of network infrastructure.

For most networks it is considered a good practice to proactively disable ICMP Redirect functionality on all Layer 3 interfaces in network infrastructure. This helps to prevent scenarios of production data traffic that is handled in CPU of Layer 3 switches and routers when there is a better forwarding path through multi-point network segments.