Performance Routing with NAT Functionality

Last updated: March 2008

When configuring both Performance Routing (PfR) and Network Address Translation (NAT) functionality on the same router, certain behavior aspects should be taken into consideration. This paper documents PfR behavior when PfR controls the routing for a prefix using static routing.

Introduction

When using both PfR and NAT, there are dynamic optimization aspects that must be taken into account. If PfR decides to optimize a traffic flow or application, then it may introduce some pragmatic consequences which require close collaboration between PfR and NAT. This paper will describe this close collaboration and provide insight on the behavior of PfR in this type of environment.

The pragmatic considerations arise when all of the following conditions occur.

1. Static routing to multiple ISPs (ISP A, ISP B, etc...) from the same router (Router A)
2. Router A is configured as a PfR Border Router
3. NAT is configured on Router A

   ```
   ip nat inside source list <X> interface <INTERFACE_ISP_A> overload
   ```

4. One or more ISPs have Strict Reverse Path Filtering (uRPF) configured

What May Happen in this Situation?

On a site where the above noted considerations are in place, after optimization, applications may fail. The root cause can be found by the ISP usage of Unicast Reverse Path Forwarding (uRPF). Many ISP will have enabled uRPF filtering for security reasons and to avoid sub-optimal traffic flows through their infrastructure. The uRPF will impact the applications because packets may be dropped at the ingress interface of the ISP due to uRPF.

Why are the packets dropped?

When NAT creates a translation slot, it gets an IPv4 global address from the connected ISP, for example from ISP A. If PfR now optimizes the application and changes the best path, it is now through ISP B instead of through ISP A, and ISP B will drop the application packets because the IPv4 source addresses used by the flow is from the ISP A address range and not from the expected ISP B address range.

Solution

The solution is to have PfR and NAT work closely together. This can be done by adding the following configuration:

```
interface virtual-template 1
```

Modify the following command:

“ip nat inside source <X> interface <Interface_ISP_A> overload”

Into

“ip nat inside source <x> interface Virtual-Template 1 overload oer”

How does this fix the problem?

With this configuration, PfR and NAT are cooperating. PfR ensures that new translations get the source IP address from the ISP interface where the packet will be routed. This ensures that uRPF will not drop the packet. PfR forces existing flows to be routed over the interface where the translation was created. If PfR changes the route for a prefix, flows with new translations will be routed to the new interface. However, flows matching existing translations will not be re-routed. Thus no uRPF packet drops.

Are there any negative side effects?

Long lived flows will not be re-routed. The goal of PfR is ensure application performance. If flows are re-routed, uRPF will drop packets and the application will die. Therefore, PfR keeps existing flows routed to avoid the drops. Some translation timers may need to be reduced to prevent new flows from matching existing translations. Ideally, each new flow would create a new translation.

How do I know if my ISP has uRPF configured?

Cisco and the Internet community recommend that all ISPs use uRPF for directly connected customers. uRPF prevents address spoofing since the source address must match the networks that are reachable over the customer to ISP interface. In general, assume uRPF is configured at the ISP. Even if it is not configured, the ISP could configure it an anytime without consulting their customers.

The customer could ask the ISP to ensure uRPF is not configured. However, this could create suspicion between the ISP and customer, because it would allow spoofed addresses to traverse the ISP.

Questions

Why is this issue limited to static routing only?

PfR, NAT, uRPF, and Border Gateway Protocol (BGP) routing to the Internet does not result in uRPF dropped packets.

There are two methods to route to an ISP, BGP and static routing. In both cases, uRPF should be configured by the ISP. However, with BGP, the same customer network is advertised to all ISPs. Since the customer network is reachable from any of the ISPs using a source address from the customer network, it will pass the uRPF check. No packet drops.

Are there any limitations to this solution?

This only applies to circuits terminated on the same router where PfR and NAT are configured. Terminating circuits on multiple routers using NAT is not supported. This solution does not address network configurations where NAT is configured on a non-IOS device where PfR cannot be configured.

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Europe Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: +31 0 800 020 0791
Fax: +31 0 20 357 1100

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA C11-458124-00 2/08