Technical Tip

TT-0501201a -- Information -- 01-Feb-2005

Proxy IP Addresses (PIP)

Contents

Contents....1
Introduction....1
Associated Products....1
Description....1
Proxy IPs and Virtual Matrix Architecture....1
Resolution....2
Many-to-one NAT....2
Many-to-one NAT with virtual server load balancing....4
Full NAT with virtual server load balancing....6
References....7

Introduction

This tech tip explains the use of Proxy IP addresses (PIPs) and how to use them in Network Address Translation (NAT) situations on Tigon-based switches running Nortel Web OS (formerly Alteon Web OS) 8 – 10.

It shows the rationale behind their use and includes examples in specific NAT situations, specifically Many-to-one NAT, Many-to-one NAT with virtual server load balancing, and Full NAT with virtual server load balancing.

For information on PIPs in Nortel Application Switch Operating System (formerly Alteon Application Switch OS), please see the appropriate user guide and application guides.

Associated Products

The information in this document is intended to be used with the following product(s) with the indicated software or hardware revisions:

Product Name or Order Number	Revision Information	
	Potentially Affected	*Corrected*
Nortel Web OS (formerly Alteon Web OS) on the Tigon Hardware Platform	8 - 10	n/a

Description

Proxy IPs and Virtual Matrix Architecture

There's been much confusion about PIPs. Most of the confusion has been due to the requirement, starting with WebOS 8.0, that if you put a PIP on one port you must put one on each of the first 8 ports. Let's first review why this is so.

The Virtual Matrix Architecture was introduced in Web OS 8.0. This architecture takes practices used in supercomputing and applies them to the switch. The first 8 ports' processors and memory are clustered together to create one super processor and a large pool of memory. This is done by hashing on certain criteria in packet flows as they ingress the switch and assigning the packet flows to a port and thereby the processors and memory on that port. This means that a packet may ingress port 1 but yet be processed by the ASIC on port 4 even though nothing is plugged into port 4. Port 4 in this situation is called the designated port.

The problem is that no layer 4-7 processing is done until that frame flow gets to the designated port (port 4). This means that even if you had a PIP on port 1 where the frame ingresses it would **not** hit that proxy IP processing and therefore **not** utilize the PIP configured on port 1. It would have to use the PIP on port 4. You can't configure whether or not a port is used in the CPU cluster so therefore you must have a PIP on ports 1-8 since packets could go to any of them when the initial hashing is done to decide the designated port for a frame flow.

©2005 Nortel Networks Limited. All Rights Reserved

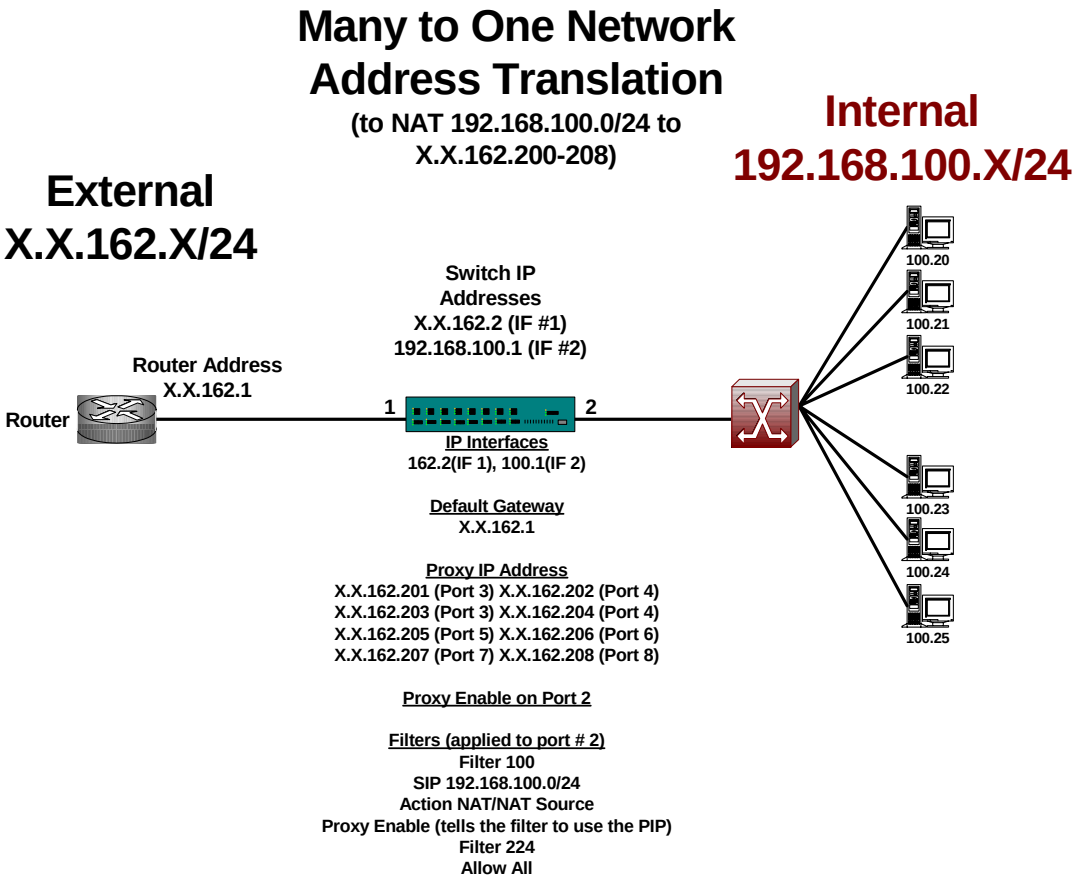
If VMA is disabled then you only have to put a PIP on the ports where you will need them (in this discussion, only on port 1 since that would be the ingress port). However, it is foolish not to take advantage of VMA since it makes our product about 10 times more powerful than it is with VMA disabled. Having said that, this may be prohibitive in certain customer environments because of the use of the additional IP addresses.

Resolution

Many-to-one NAT

Many-to-one NAT is the first situation where a PIP is required. The PIP address is the one address that the many addresses are translated to. Many-to-one NAT is great in that it only allows communication to be established from one direction, which is an excellent security feature.

The following is a visual example of Many-to-one NAT and the configuration:



This configuration would NAT the internal subnet of 192.168.100.0/24 to the PIPs of X.X.162.201-208. Note that this is done through the use of filtering, and PIPs are required on every port since VMA is enabled. Proxy must also be enabled on all ports that have PIPs configured on them.

©2005 Nortel Networks Limited. All Rights Reserved

The line-by-line configuration for this would be:

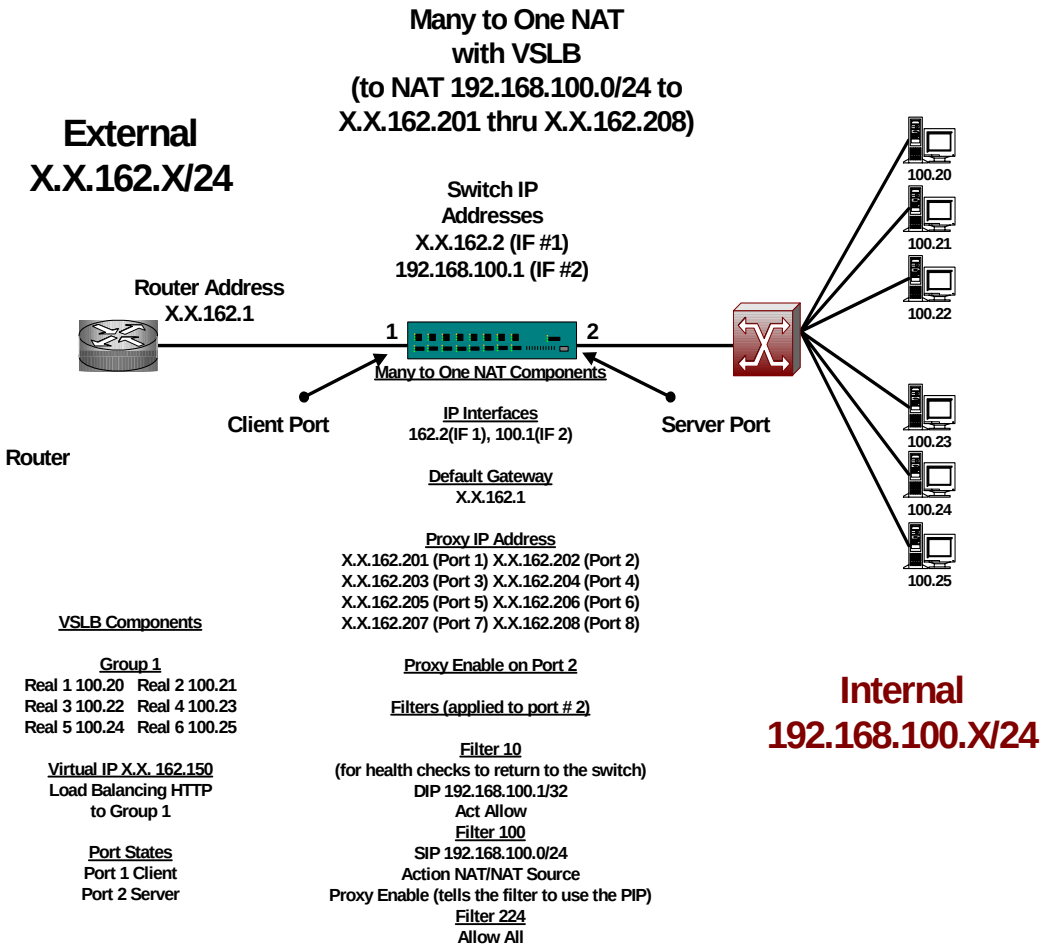
```
/cfg/ip/if 1/mask 255.255.255.0/addr X.X.162.0/ena
/cfg/ip/if 2/addr 192.168.1.0/ena
/cfg/slb/on (SLB must be on for this feature)
/cfg/slb/filt 100/sip 192.168.1.0/smask 255.255.255.0/act nat/nat source/ena
(proxy is enabled by default on the filter so we don't need to enable it)
/cfg/slb/filt 224/ena (default allow all)
/cfg/slb/port 2/filt ena/add 100/add 224
/cfg/slb/port 1/pip X.X.162.201/proxy ena
/cfg/slb/port 2/pip X.X.162.202/proxy ena
/cfg/slb/port 3/pip X.X.162.203/proxy ena
/cfg/slb/port 4/pip X.X.162.204/proxy ena
/cfg/slb/port 5/pip X.X.162.205/proxy ena
/cfg/slb/port 6/pip X.X.162.206/proxy ena
/cfg/slb/port 7/pip X.X.162.207/proxy ena
/cfg/slb/port 8/pip X.X.162.208/proxy ena
```

©2005 Nortel Networks Limited. All Rights Reserved

Many-to-one NAT with virtual server load balancing

Many-to-one NAT is particularly helpful when used in conjunction with virtual server load balancing. The reason is that the load balanced servers can be put on private networks and load balanced by the virtual IP addresses. They can still make outbound requests yet no one can connect to them directly since the NAT filter only works on outbound requests. This is a very nice security feature. Another nice security feature is that PIPs will respond to ARP requests but will not respond to pings. The configuration is simple since you simply configure each application as you normally would.

The following is a visual example:



©2005 Nortel Networks Limited. All Rights Reserved

The line by line configuration for this would be:

```
/cfg/ip/if 1/mask 255.255.255.0/addr X.X.162.0/ena
/cfg/ip/if 2/addr 192.168.1.0/ena
/cfg/slb/real 1/rip 192.168.100.20/ena
/cfg/slb/real 2/rip 192.168.100.21/ena
/cfg/slb/real 3/rip 192.168.100.22/ena
/cfg/slb/real 4/rip 192.168.100.23/ena
/cfg/slb/real 5/rip 192.168.100.24/ena
/cfg/slb/real 6/rip 192.168.100.25/ena
/cfg/slb/group 1/add 1/add 2/add 3/add 4/add 5/add 6
/cfg/slb/virt 1/vip X.X.162.150/ena/service http
/cfg/slb/port 1/clien ena
/cfg/slb/port 2/serv ena
/cfg/slb/on
```

(SLB must be on for this feature)

```
/cfg/slb/filt 10/dip 192.168.100.1/dmask 255.255.255.255/ena
```

(filter 10 is to allow the real server health checks to get back to the switch without being NAT'ed)

```
/cfg/slb/filt 100/sip 192.168.1.0/smask 255.255.255.0/act nat/nat source/ena
```

(proxy is enabled by default on the filter so we don't need to enable it)

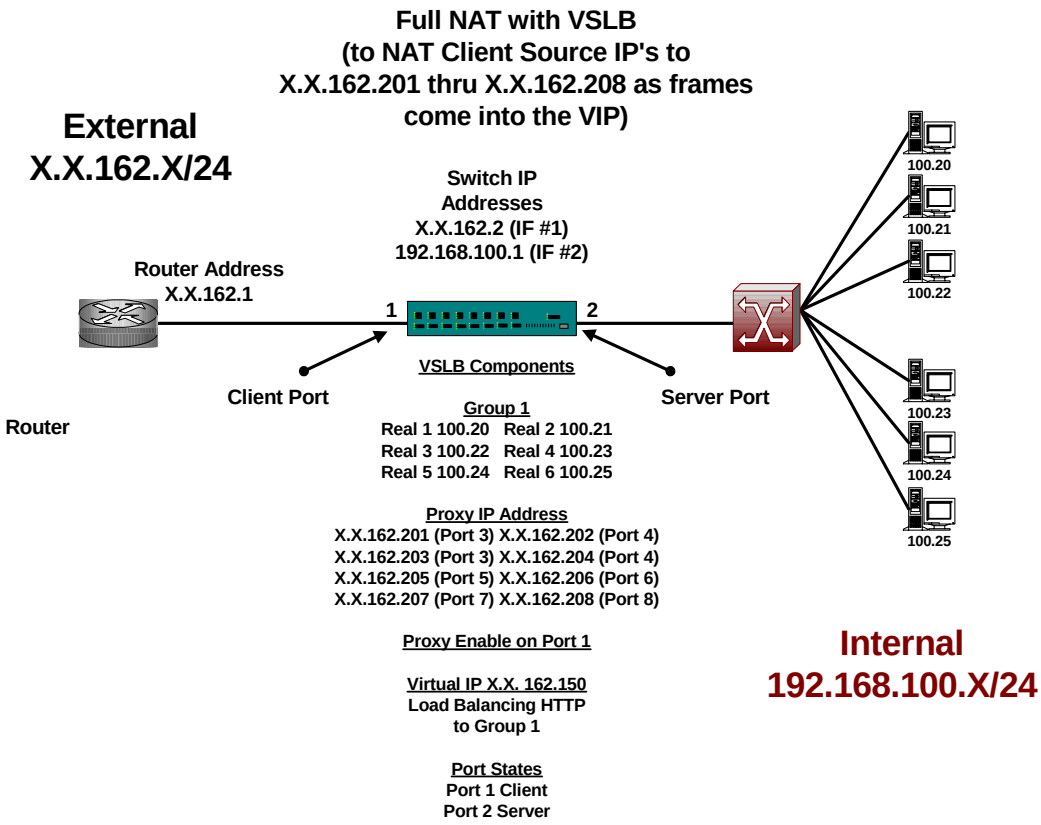
```
/cfg/slb/filt 224/ena (default allow all)
/cfg/slb/port 2/filt ena/add 100/add 224
/cfg/slb/port 1/pip X.X.162.201/proxy ena
/cfg/slb/port 2/pip X.X.162.202/proxy ena
/cfg/slb/port 3/pip X.X.162.203/proxy ena
/cfg/slb/port 4/pip X.X.162.204/proxy ena
/cfg/slb/port 5/pip X.X.162.205/proxy ena
/cfg/slb/port 6/pip X.X.162.206/proxy ena
/cfg/slb/port 7/pip X.X.162.207/proxy ena
/cfg/slb/port 8/pip X.X.162.208/proxy ena
```

©2005 Nortel Networks Limited. All Rights Reserved

Full NAT with virtual server load balancing

Full NAT with virtual server load balancing (VSLB) is the third situation that requires the use of PIP. By design, VSLB NATs the destination IP from the virtual IP address configured on the switch to the address of the real server the switch is load balancing to. By the use of a PIP, we can force the switch to also NAT the source IP. This may be required for delayed binding features in complex topologies as well as when doing one-armed load balancing with delayed binding features. This would force traffic to get back to the correct designated port on the way back out, since it is the owner of the PIP that it substituted for the original SIP which was the client IP address. Unfortunately, the servers are going to think that all traffic is coming from the PIPs when doing this. This is a “no-go” in most customer environments since a lot of the time they log the source IPs of the client requests.

The following is a visual example:



©2005 Nortel Networks Limited. All Rights Reserved

The line-by-line configuration for this would be:

```
/cfg/ip/if 1/mask 255.255.255.0/addr X.X.162.0/ena
/cfg/ip/if 2/addr 192.168.1.0/ena
/cfg/slb/real 1/rip 192.168.100.20/proxy ena/ena
/cfg/slb/real 2/rip 192.168.100.21/ proxy ena/ena
/cfg/slb/real 3/rip 192.168.100.22/ proxy ena/ena
/cfg/slb/real 4/rip 192.168.100.23/ proxy ena/ena
/cfg/slb/real 5/rip 192.168.100.24/ proxy ena/ena
/cfg/slb/real 6/rip 192.168.100.25/ proxy ena/ena
/cfg/slb/group 1/add 1/add 2/add 3/add 4/add 5/add 6
/cfg/slb/virt 1/vip X.X.162.150/ena/service http
/cfg/slb/port 1/clien ena
/cfg/slb/port 2/serv ena
/cfg/slb/on (SLB must be on for this feature)
/cfg/slb/port 1/pip X.X.162.201/proxy ena
/cfg/slb/port 2/pip X.X.162.202
/cfg/slb/port 3/pip X.X.162.203
/cfg/slb/port 4/pip X.X.162.204
/cfg/slb/port 5/pip X.X.162.205
/cfg/slb/port 6/pip X.X.162.206
/cfg/slb/port 7/pip X.X.162.207
/cfg/slb/port 8/pip X.X.162.208
```

References

Always verify your configuration with the appropriate Nortel user guides and command references, available at http://www.nortel.com/cs.

Copyright © 2005 Nortel Networks Limited - All Rights Reserved. Nortel, Nortel Networks, the Nortel logo, and Globemark are trademarks of Nortel Networks Limited.

The information in this document is subject to change without notice. The statements, configurations, technical data, and recommendations in this document are believed to be accurate and reliable, but are presented without express or implied warranty. Users must take full responsibility for their applications of any products specified in this document. The information in this document is proprietary to Nortel Networks Limited.

To access more technical documentation, search our knowledge base, or open a service request online, please visit Nortel Networks Technical Support on the web at: http://www.nortel.com/support

If after following this guide you are still having problems, please ensure you have carried out the steps exactly as in this document. If problems still persist, please contact Nortel Networks Technical Support (contact information is available online at: http://www.nortel.com/cgi-bin/comments/comments.cgi?key=techsupport_cu).