

How To Guide

XRoads Networks

Edge Network Appliance How To Guide:
NetBalancing

Edge Network Appliance How To Guide:
NetBalancing

@2009 XRoads Networks

22642 Lambert St, Suite 403

888-9-XROADS

Table of Contents

NetBalancing Overview

MultiLink Vector Routing

Application Acceleration

Enhanced Session Persistence (ESP)

Application Routing

Best Path Routing (BPR)

Application Redirection

Vector Mappings

NAT Mappings

Server Balancing

Edge Configuration Series

NetBalancing Overview

One of the core functions of the EdgeXOS appliance is its ability to utilize multiple WAN interfaces at the same time, commonly known as link load balancing. The difference between the EdgeXOS appliance and other link load balancers is its ability to provide more granular and detailed control over specific sessions, including the ability to “bond” certain types of session traffic in order to accelerate downloads and dramatically improve performance

XRoads Networks

[Home](#)[Interfaces](#)[Shaping](#)[NetBalancing](#)[Firewall](#)[Site2Site](#)[Tools](#)[Reporting](#)

NetBalancing Modules

These modules provide the network administrator with complete control over their multiple ISP links along with the ability to determine how the additional bandwidth from each link is utilized.

NetBalancing includes several sections, the first section deals with how overall traffic is handled by the appliance. The MultiLink Vector Routing, Application Acceleration, and Enhanced Session Persistence are used to set global preferences. The “Outbound”, “In/Out”, and “Inbound” sections provide for control over specific sessions based on either routing and/or application information. This is where inbound server balancing/failover is configured and where preferences can be setup for routing applications out one link or the other. Next is the “DNS” section which controls the ActiveDNS module (see our specific HowToGuide on ActiveDNS). Finally our latest feature is our simple Server Load Balancing module which provides customers the ability to not only balance traffic between WAN links but also across multiple internal servers as well.

MultiLink Vector Routing

Vector Routing is the technology developed by XRoads Networks to balance traffic between multiple WAN interfaces. The latest generation of this technology incorporates new intelligence which distributes traffic base on administrative weights while still ensuring proper session persistence is maintained.

NetBalancing Selection:

Select Balance Method

IVRA (Intelligent Vector Routing Algorithm) Weighted Spill-Over Round Robin

Select Default Interface (Change On Support Recommendation)

WAN1 WAN2 WAN3 WAN4

App Load Balancing:

Balanced Applications (Change On Support Recommendation)

- Web (HTTP:80/HTTPS:443)
- Email (SMTP:25/POP3:113/IMAP:143)
- Domain (DNS:53)
- File Transfer (FTP:20/21)
- Ping (ICMP)

Enhanced Session Persistence

Enabled Disabled

Strict Control - On Strict Control - Off

Advanced Web Balancing: Enabled Disabled (Turn on advanced web balancing services)

Balance Methods

The default balance method is IVRA which is a proprietary algorithm which automatically determine which WAN link is best for which sessions. It does this by testing each WAN link and automatically tunes the routing so that traffic is balanced based on the weights assigned to each individual interface by the EdgeXOS administrator (see Session Persistence below). Additional options are also available, including Weighted (which only attempts to use the administrative weights and not actual session usage, Spill-Over (which directs more traffic over the secondary links when the primary is full), and Round Robin (which is similar to weighted but is performed in more of a 50/50 manner).

We do not recommend changing the Default Interface or App Load Balancing parameters, these should only be changed if support recommends to do so.

Enhanced Session Persistence

When two or more WAN links are used for routing network traffic, session persistence is paramount. Without excellent session persistence certain applications will not work, others will only work some of the time, and still others will stop working in the middle of a session. Due to the high priority that the EdgeXOS appliance assigns to session persistence it may appear that at times the interfaces are not weighted as configured, this is most likely due to the fact that some session use more bandwidth than others, and while over the long term the WAN interfaces will be balanced as the administrative weights are configured, it may not always appear that way.

Strict Control: Should be used if certain applications are not working correctly or are getting cut off in the middle of sessions. Strict Control provides a higher degree of session persistence but may reduce the amount of balancing that is performed. This is a trade-off but it is generally better to have sessions that work with some balancing than sessions that do not work with equal balancing.

Advanced Web Balancing

This feature can be enabled to spread web-sessions across the WAN links more evenly. Seeing as web traffic is typically over 80% of an organizations traffic, balancing this traffic more effectively means better overall performance.

Application Acceleration

Unique to the EdgeXOS appliance is our Multi-Session Acceleration (MSA) technology. This technology enables the EdgeXOS appliance to significantly and dramatically speed up web-based session downloads.

The screenshot shows the configuration page for Application Acceleration. On the left is a dark blue sidebar with menu items: NetBalancing Selection, Multi-Session, MSA DNS Resolvers, Bypass URLs, and Add URL. The main content area has a top section with a dropdown menu set to 'Application Acceleration'. Below it, 'Multi-Session Acceleration (Advanced Web Bonding Control)' is shown as 'Enabled' with a radio button. Under 'Bonded Content', several file types are listed with checked checkboxes: Adobe / PS Files, Executable / Data Files, Compressed Files, Office Files, Image Files, Media Files, and Other Files. The 'MSA DNS Resolvers' section contains two rows of IP address input fields, each followed by a label '(WAN One ISP DNS Server)' and '(WAN Two ISP DNS Server)'. Below these are 'Update' and 'Save' buttons. The 'Bypass URLs' section features a table with columns 'Select', 'URL Address', and 'Status'. The 'Add URL' section has a text input field with a placeholder '(Enter a URL, like www.google.com)' and 'Add URL' and 'Delete URL' buttons.

Bonded Content

The file types are used to determine which types of content will be bonded during a web session. Some customers may find that certain downloads do not work for some types of content. If that is the case, then the network administrator can simply disable that content type and it should continue to work without any further issues.

MSA DNS Resolvers

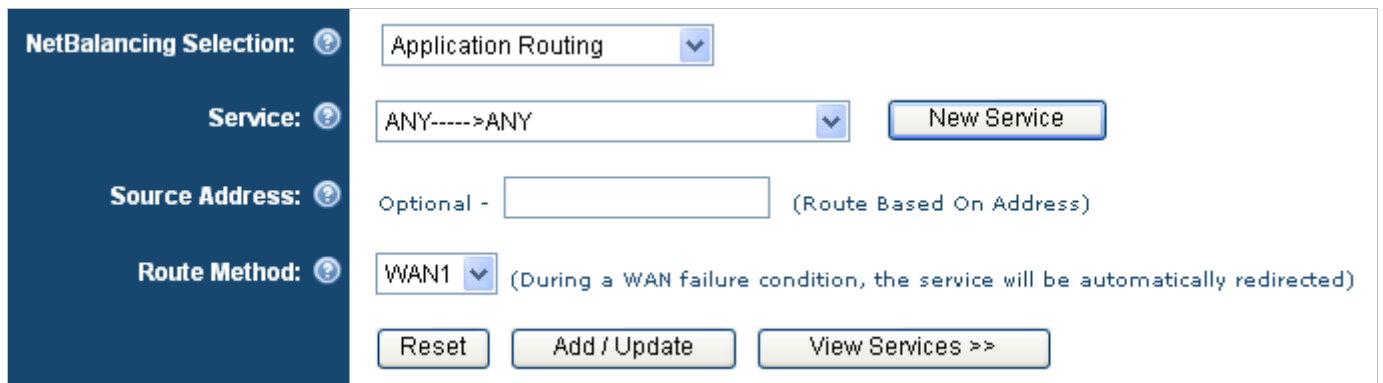
When MSA is enabled the EdgeXOS appliance will perform DNS queries for each session that is initiated in order to determine the best method to download the data. It is critical that the DNS servers provided are highly responsive otherwise it will slow down network connections. These are the same DNS servers which are configured under the LAN Interface section of the appliance. Try to enter DNS servers from two different service providers.

Bypass

There may be times when a certain website is just not working with MSA enabled, if that is the case simply enter the URL that is having an issue and it will bypass the bonding service without having to disable all MSA functions.

Application Routing

In some instances it makes sense to “offload” traffic across one WAN link over another. For example to send all email traffic over the slow DSL link rather than the T1 connection. On the EdgeXOS platform this is done using the application routing.



The screenshot shows the configuration interface for Application Routing. On the left is a dark blue sidebar with the following labels: "NetBalancing Selection: ?" (with a help icon), "Service: ?" (with a help icon), "Source Address: ?" (with a help icon), and "Route Method: ?" (with a help icon). The main configuration area contains: a dropdown menu set to "Application Routing"; a dropdown menu set to "ANY----->ANY" next to a "New Service" button; an "Optional -" label followed by an empty text input field and the text "(Route Based On Address)"; a dropdown menu set to "WAN1" with the text "(During a WAN failure condition, the service will be automatically redirected)"; and three buttons at the bottom: "Reset", "Add / Update", and "View Services >>".

Services

This is a listing of services which have been predefined, the EdgeXOS administrator can add new services if they are not already available. This service would be the application that is going to be forced across a specific WAN link.

Source Address

If a source address is specified than ONLY the address specified is affected by this rule. So it is possible to setup specific routing for a single device. An example would be a server which should only use a specific WAN interface when making a specific outbound connection using the specified application.

Route Method

This is simply the WAN link that should be used when forcing the traffic out. If the link that is being used fails then the rule will automatically be removed and the application traffic will be routed however the global balancing settings are configured

EdgeXOS Best Path Routing™ (BPR)

One of the more advanced features of the EdgeXOS appliance is our Best Path Routing technology. Use this feature to create policy-routes based on network metrics for specific remote networks and/or services. The EdgeXOS BPR is an enhancement to standard policy-based routing techniques in that instead of simply perform standard network tests to determine how routing is performed, it does these tests automatically and intelligently in order to determine the best route when multiple routes are available, in addition the EdgeXOS BPR includes built-in SLA reporting.

| | |
|----------------------------------|--|
| NetBalancing Selection: ? | Best Path Routing (SLAs) ▼ |
| | Use BPR to monitor Service Level Agreements (SLAs) and to ensure the fastest path routing to critical remote locations across the network. |
| Route Description: ? | <input type="text"/> (Network Name) |
| | <input type="text"/> (URL Address - example: www.xyz.com) |
| | NOTE: Must be pingable, and should not be the same as a link Control website . |
| Define Network: ? | OR |
| | <input type="text"/> . <input type="text"/> . <input type="text"/> . <input type="text"/> (Network Address Or Subnet) |
| | SINGLE HOST ▼ (Subnet Mask) |
| Test Node: ? | <input type="text"/> . <input type="text"/> . <input type="text"/> . <input type="text"/> (This is the address that will be pinged) |
| | NOTE: Created automatically if a URL is entered above. |
| Latency: ? | <input type="text"/> 500 ms (Round Trip Time Threshold - Default 80) |
| Packet Loss: ? | <input type="text"/> 50 % (Packet Loss Percentage - Default 3) |
| Jitter: ? | <input type="text"/> 500 ms (Latency Difference Between Tests - Default 50) |
| SLA Reporting: ? | <input type="checkbox"/> (Enable SLA Reports) |
| Route Method: ? | WAN1 ▼ (Select the default WAN interface) |
| | When Threshold Exceeded ▼ (How route selection will be applied) |
| | NOTE: If persistence is an issue for this route, do not select best path. |
| | <input type="button" value="Reset"/> <input type="button" value="Add / Update"/> <input type="button" value="View Routes >>"/> |

Define Network


When setting up any routing policy, the first step is to identify the route that will be manipulated. Here either a specific network/subnet or a URL can be used. Make sure that the test address is associated with the route/URL provided and is pingable using ICMP. If the test node is not pingable the metrics will not work correctly.

Parameters

These are the metrics used to determine the status of a particular remote network and if the status meets preconditions needed to change how those remote networks are routed. The defaults are typically fine, the higher the metrics the less sensitive the testing, the lower the more sensitive. To reduce routing changes and threshold alerts increase the metrics.

SLA Reporting

These are graphical reports produced by the EdgeXOS appliance when BPR rules are setup with SLA enabled. These reports are found under the Reporting tab and are very useful in identifying what exactly is happening across each of the WAN links and whether each service provider is maintaining their required service level agreements. The graphical report includes both packet loss and RTT information. The higher the round trip time, the worst that particular path is for routing critical network traffic.



Additional metrics are also collected over time, including jitter information which is helpful when using real-time applications like VoIP and multimedia applications. When the status is green the metrics are within the set parameters.

| Time | Avg | Min | Max | Packet Loss | Jitter | Status |
|---------------------|--------|--------|--------|-------------|--------|--------|
| Current | 325.84 | 323.91 | 327.78 | 0% | 1.933 | ● |
| Fifteen Minutes Ago | 312.20 | 310.72 | 313.68 | 0% | 1.480 | ● |
| One Hour Ago | 193.35 | 193.14 | 193.56 | 100% | 0.210 | ● |
| One Day Ago | 190.53 | 190.37 | 190.70 | 100% | 0.467 | ● |
| One Week Ago | 192.49 | 191.74 | 193.24 | 100% | 0.887 | ● |

Route Method

The methods include “When Threshold Exceeded”, “Best Path Route”, and “Only On Failure”. The first will only change a route when one of the metrics is exceeded, i.e. RTT is set for 100ms and the current path exceeds that amount, the route will be changed to the next lowest path at the time. The second will continuously change the path based on the best performing path, i.e. the path with the lower latency and packet loss. The third will ONLY change the path when an outage occurs of the default path set.

Example: If the WAN1 link is the default, then it will remain on WAN1 unless that link fails or becomes so congested that the path is considered unusable.

NOTE: Using the first two route methods could cause session persistence issues for some applications, it is recommended to use high parameters in both cases.

Application Redirection

This feature is typically used to force outbound traffic to a specific remote server when a network failure occurs, or when a proxy server is being utilized. An example would be when a customer is using an external DNS server, but when an outage occurs must use a DNS server from the other provider as the primary DNS is not available to devices outside of its own network. Another example would be if the EdgeXOS administrator wanted all web traffic to be forced to an external proxy server.

| | |
|----------------------------------|--|
| NetBalancing Selection: ? | Application Redirection ▼ |
| Redirect Description: ? | <input type="text"/> |
| Redirect OnFailover: ? | <input type="radio"/> Redirect Always <input type="radio"/> Redirect On Failover (To apply redirection rule only in case of WAN failover) |
| Redirect Address: ? | <input type="text"/> . <input type="text"/> . <input type="text"/> . <input type="text"/> (Network Address - Must be a 'CIDR' network address) |
| Protocol/Port: ? | TCP ▼ <input type="text"/> (Enter the port you wish to redirect, www = 80) |
| | <input type="button" value="Reset"/> <input type="button" value="Add / Update"/> <input type="button" value="View Redirects >>"/> |

Redirection

This is to determine when redirection will occur, either all of the time as with a remote proxy server, or only when a failover condition exists. A failover condition exists when the primary WAN link, i.e. WAN1 fails. The next step is to enter the IP address of the remote device to which application traffic will be forwarded, i.e. redirected.

Application Definition

This is where a protocol and port need to be identified. The EdgeXOS appliance includes a special protocol by default called VoIP, this is used to perform VoIP proxy redirection typically in the event of a network failover, no port is required in that case.

Advanced Vector-Mappings

This section provides a step-by-step overview of the Vector Map feature. With Vector Map rules any internal address can be forced to send traffic out either of the WAN interfaces.

Vector Map is recommended for any systems accepting inbound traffic whenever WAN load balancing is being used.

NOTE: Never create two Vector Maps with the same "Map Address", as this would cause random routing outages.

| | |
|----------------------------------|---|
| NetBalancing Selection: ? | Vector Mappings |
| Device Name: ? | <input type="text"/> |
| Map Address: ? | <input type="text"/> . <input type="text"/> . <input type="text"/> . <input type="text"/> . <input type="text"/> (Forward Address or Range - Available via the LAN interface) |
| | Optional - <input type="text"/> OR --- Select --- (Enter a source port or port range x:x, if any) |
| Map Interface: ? | WAN1 |
| Apply Order: ? | 1 |
| | <input type="button" value="Reset"/> <input type="button" value="Add / Update"/> <input type="button" value="View Services >>"/> |

The procedure for creating a Vector Map rule is fairly simple. Here is a list of the information required:

- What is the IP address of the internal server?
(must be reachable via the LAN interface or via a static route)
- Which interface will be used to forward this internal server traffic in a load balanced configuration?

NOTE: Load balancing is referred to here not as server load balancing, but WAN load balancing, when multiple WAN interfaces are set to ACTIVE.

The following two examples demonstrate the required configuration when both WAN interfaces will be forwarding traffic to the inbound server(s).

In this case, both WAN interfaces require a Vector Map, one for WAN1 and another for WAN2. For more information on this type of configuration, refer to the HowToGuides.

These examples assume there is a server on the LAN network which has been dual addressed, see Overview Advanced NAT (above) for more. The primary address is 192.168.168.100, and secondary address assigned to the server is 192.168.18.101.

It is further assumed that a One-To-One or One-To-Many NAT rule has been created for WAN1 and WAN2 pointing to 192.168.168.100 and 192.168.18.101 respectively.

Enter the WAN1 configuration information.

Vector Mappings ▾

web

192 . 168 . 168 . 100 - (Forward Address or Range - Available via the LAN interface)

Optional - 80 (Enter a source port or port range x:x, if any)

WAN1 ▾

Reset Add / Update View Services >>

Enter the WAN2 configuration information.

Vector Mappings ▾

web2

192 . 168 . 168 . 101 - (Forward Address or Range - Available via the LAN interface)

Optional - (Enter a source port or port range x:x, if any)

WAN2 ▾

Reset Add / Update View Services >>

The process for actually creating a Vector Map rule is as follows:

- 1 Enter the name for this rule.
- 2 Enter the internal servers IP address for this rule.
- 3 Enter the WAN interface where this traffic will be directed.

Once added the following screen will appear showing the rule that was added.

| Select | Device Name | Address | Interface Map | Port Map |
|--------------------------|-------------|------------|---------------|----------|
| <input type="checkbox"/> | RDP | 10.50.10.2 | WAN1 | ALL |
| <input type="checkbox"/> | VoIP | 10.20.1.20 | WAN1 | ALL |

From this screen the Vector Map rules can be modified or deleted.

VirtualNAT (aka Virtual Server)

This section provides a step-by-step overview of the VirtualNAT feature, generally referred to as a virtual server. XRoads Networks uses the term VirtualNAT over virtual server, because the term virtual server is misleading in that at no time is the Edge device actually performing an server functions.

In reality, VirtualNAT is basically accepting connections for an internal server and passing those connections directly to the internal server without any modifications. Many would define this type of functionality as an application proxy. Again, VirtualNAT is used to avoid confusion with web proxy systems which cache responses, something that VirtualNAT does not perform.



VirtualNAT

Web Server (HTTP/HTTPS)

Create Server Service (Create A New VirtualNAT Service)

. . . (Internal Server Address)

. . . (External Server Address for WAN1)

. . . (External Server Address for WAN2)

. . . (Forward Address - Must be available via an Edge interface)

Reset Add / Update View VirtualNAT Rules >>

NOTE: Currently VirtualNAT is capable of forwarding / load balancing seven applications. In the near future this function will allow for the creation of custom TCP applications for forwarding / load balancing.

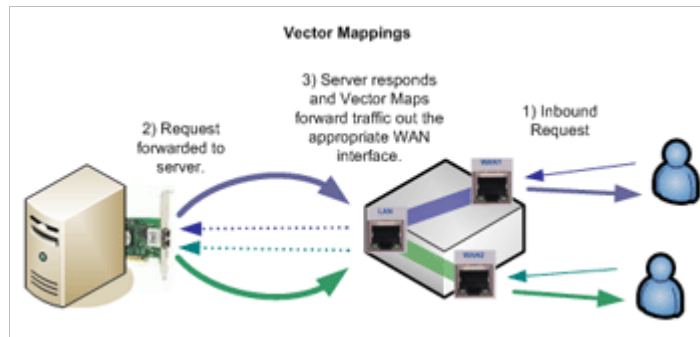
VirtualNAT is far easier to use than creating One-To-Many or other Advanced NAT options because it only requires one step for either redundancy or load balancing and no Vector Maps are required.

VirtualNAT is the recommended method to set up inbound NAT rules for servers.

VirtualNAT drawbacks include: 1) a limited set of applications, 2) TCP only, 3) any logging by internal servers will see all connections coming from the Edge device.

Create A VirtualNAT Rule

The following is an example VirtualNAT rule configuration between two WAN connections and an internal web servers. The following diagram provides the IP addressing information.



This example shows how connections from either WAN₁ or WAN₂ are forwarded to the internal 192.168.168.100 server. Using this method, EdgeDNS entries can be made to create simple inbound load balancing for this server across the two WAN connections.

This configuration also provide complete server redundancy since the EdgeDNS module will only provide address information for UP and ACTIVE WAN interfaces. Should one of the WAN interface fail, the EdgeDNS module will no longer provide that address to external clients.

VirtualNAT

web1

Web Server (HTTP/HTTPS)

Create Server Service (Create A New VirtualNAT Service)

192 . 168 . 168 . 100 (Internal Server Address)

24 . 45 . 65 . 78 (External Server Address for WAN1)

65 . 78 . 45 . 12 (External Server Address for WAN2)

. . . . (Forward Address - Must be available via an Edge Interface)

Reset Add / Update View VirtualNAT Rules >>

The procedure for creating a VirtualNAT rule is fairly simple.

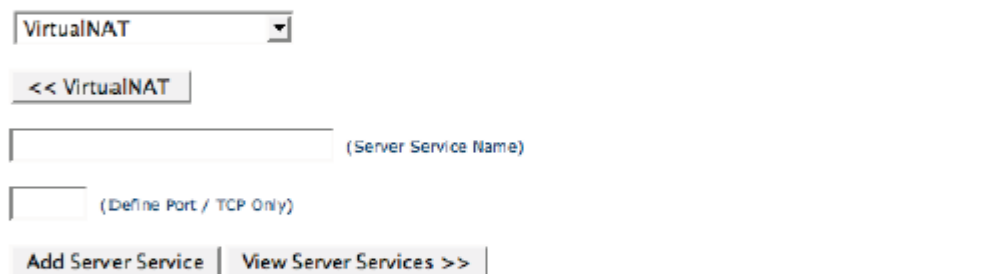
Here is a list of the information required:

- What is the IP address of the internal server? (must be reachable via the LAN interface or via a static route)
- What are the WAN addresses that will be forwarded to this internal server? (should not be the same as the WAN interface addresses)

WAN addresses that are entered into the VirtualNAT configuration will be automatically added to the WAN interface addressing as secondary addresses.

- 1 Enter the name for this rule and select the appropriate service/application from the drop-down menu.

If a service is required that is not defined, define it using the "Create Server Service" button. Only TCP based services can be defined.



The screenshot shows a configuration form for a VirtualNAT rule. At the top, there is a dropdown menu with "VirtualNAT" selected. Below it is a button labeled "<< VirtualNAT". There are two input fields: the first is labeled "(Server Service Name)" and the second is labeled "(Define Port / TCP Only)". At the bottom of the form are two buttons: "Add Server Service" and "View Server Services >>".

- 2 Enter the internal and/or LAN address for the server
- 3 Enter each of the WAN addresses that will be forwarded to this server

Once added the following screen will appear showing the rule that was added.

From this screen the VirtualNAT rules can be modified or deleted.

| Select | Server Name | Server Address | Service Type | WAN1 Addr | WAN2 Addr | WAN3 Addr |
|-----------------------|-------------|-----------------|-------------------------|-------------|-------------|-----------|
| <input type="radio"/> | web1 | 192.168.168.100 | Web Server (HTTP/HTTPS) | 24.45.85.78 | 85.78.45.12 | |

To add additional VirtualNAT rules, click the << Add VirtualNAT Rule button.

Overview Of Advanced NAT

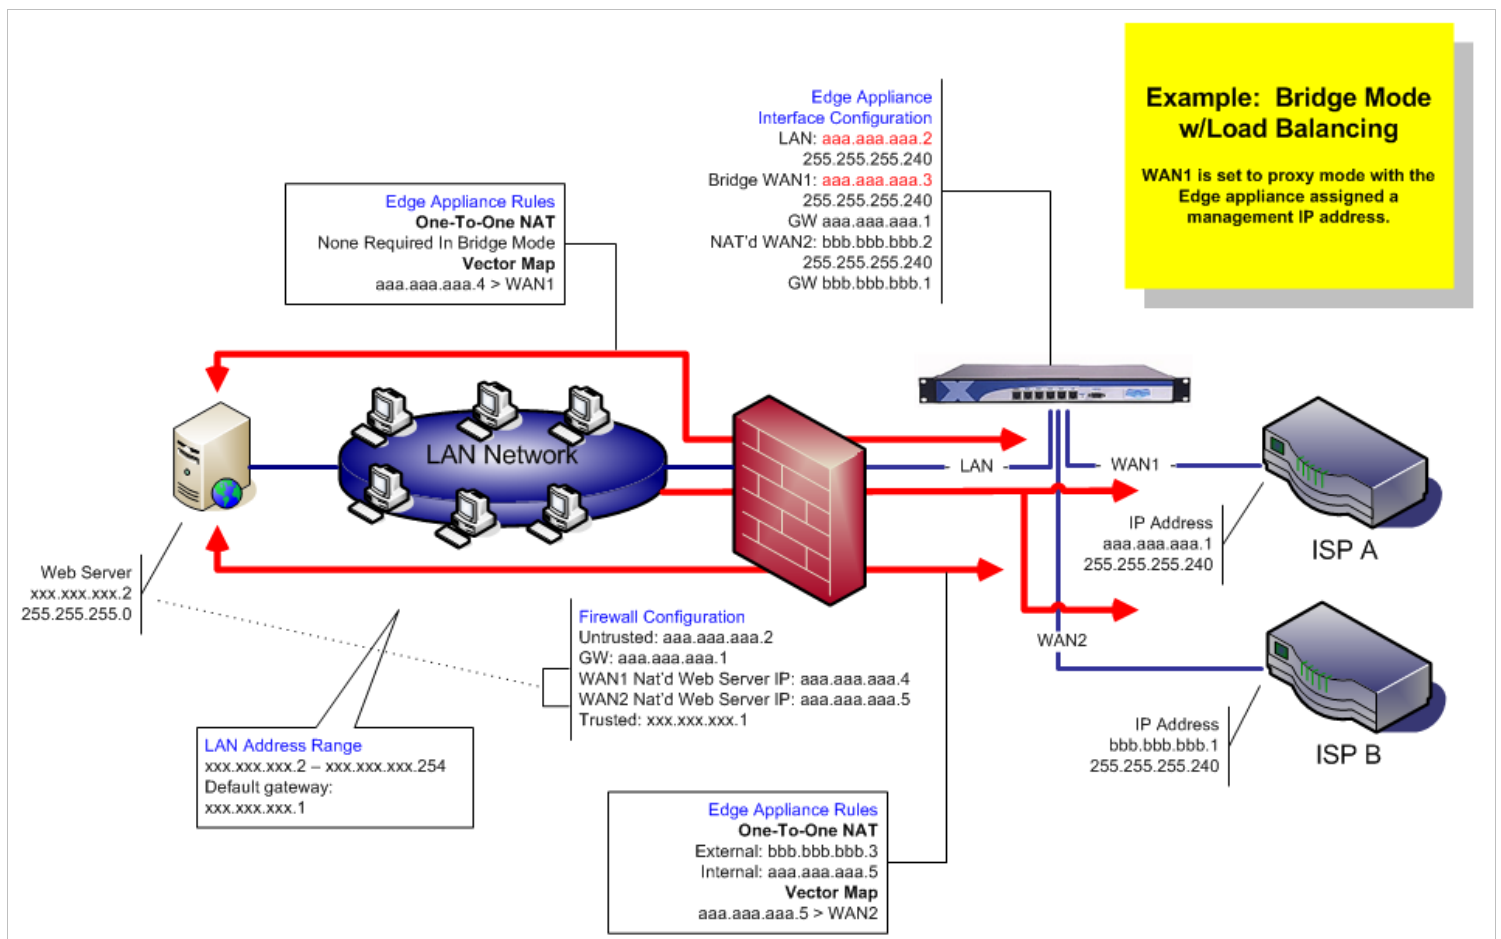
The ADVANCED section of the NetBalancing menu provides a great deal of configurability. The downside to this additional configurability is that the ADVANCED section can be fairly complex to configure.

In most cases, configuring One-To-Many or One-To-One NAT requires the additional configuration of a Vector Map rule to ensure bi-directional session establishment.

Without a Vector Map rule, inbound connections could be inadvertently routed out a different WAN interface than the one they came in on. In most cases, when this happens the connection will fail and remote access to internal servers will not work properly.

Vector Mappings ensure proper routing of server connections. Vector Mappings are ONLY required for inbound "server" based connections.

The following are examples of different NetBalancing configurations:

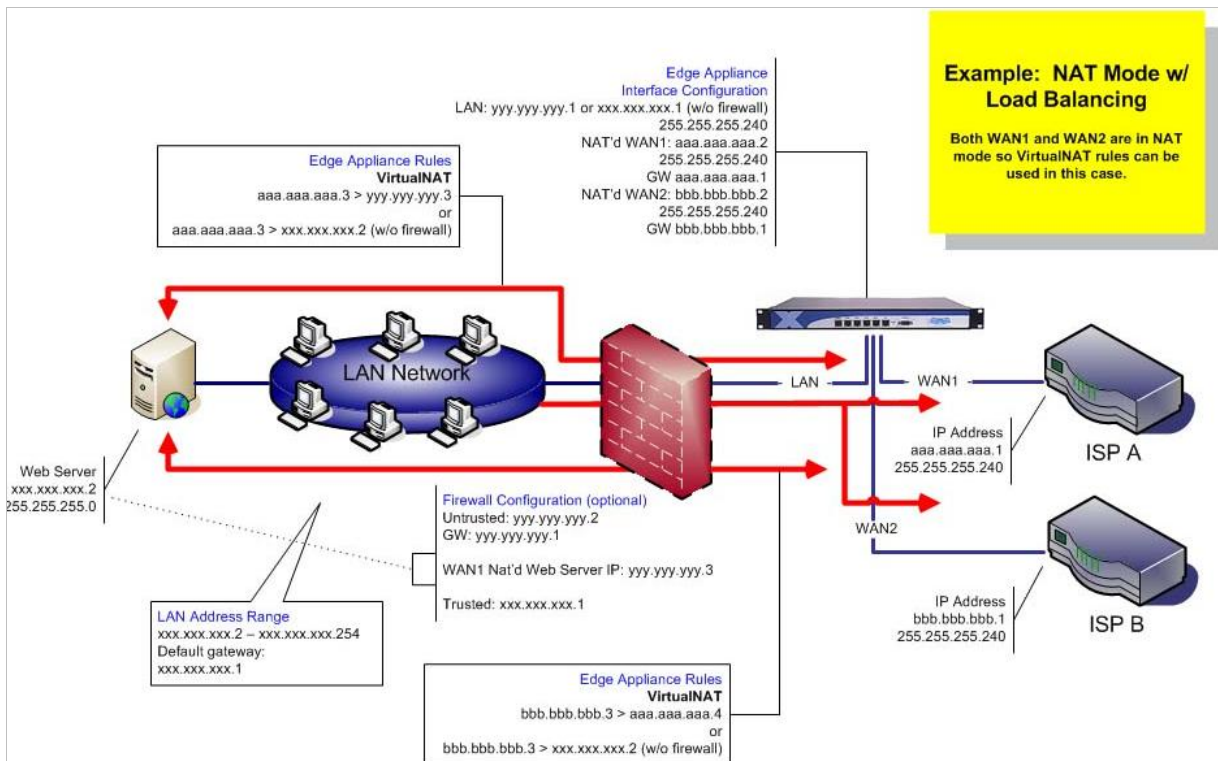


In this example WAN1 is configured in pass-through proxy mode, and WAN2 is NAT'd. Since WAN1 is essentially bridging the WAN1 traffic, no NAT is required for WAN1 addresses they work by default. However WAN2 traffic is NAT'd and thus any server traffic through WAN2 will also need specific NAT rules.

In this case the two WAN connections are also setup to be load balanced, thus both connections will be used at the same time. In general, seeing as this is a web and email server VirtualNAT would be used. However the administrator for this network has determined that they must have logging that reflects the actual end-users IP address, thus VirtualNAT will not work.

The configuration process for is as follows:

- 1 Add a secondary IP address to the web / email server. This secondary address will be used by the WAN2 NAT rules to ensure the server traffic is routed appropriately.
- 2 Create a One-To-One rule (two One-To-Many rules would also work using ports 80) pointing bbb.bbb.bbb.3 -> aaa.aaa.aaa.5 which is the new secondary IP address on the web server that was added in Step 1.
Create a Vector Map directing any traffic from aaa.aaa.aaa.4 to WAN1 and aaa.aaa.aaa.5 to WAN2.
- 3 The bbb.bbb.bbb.3 IP address is automatically added to the WAN2 interface when the One-To-One rule is created.



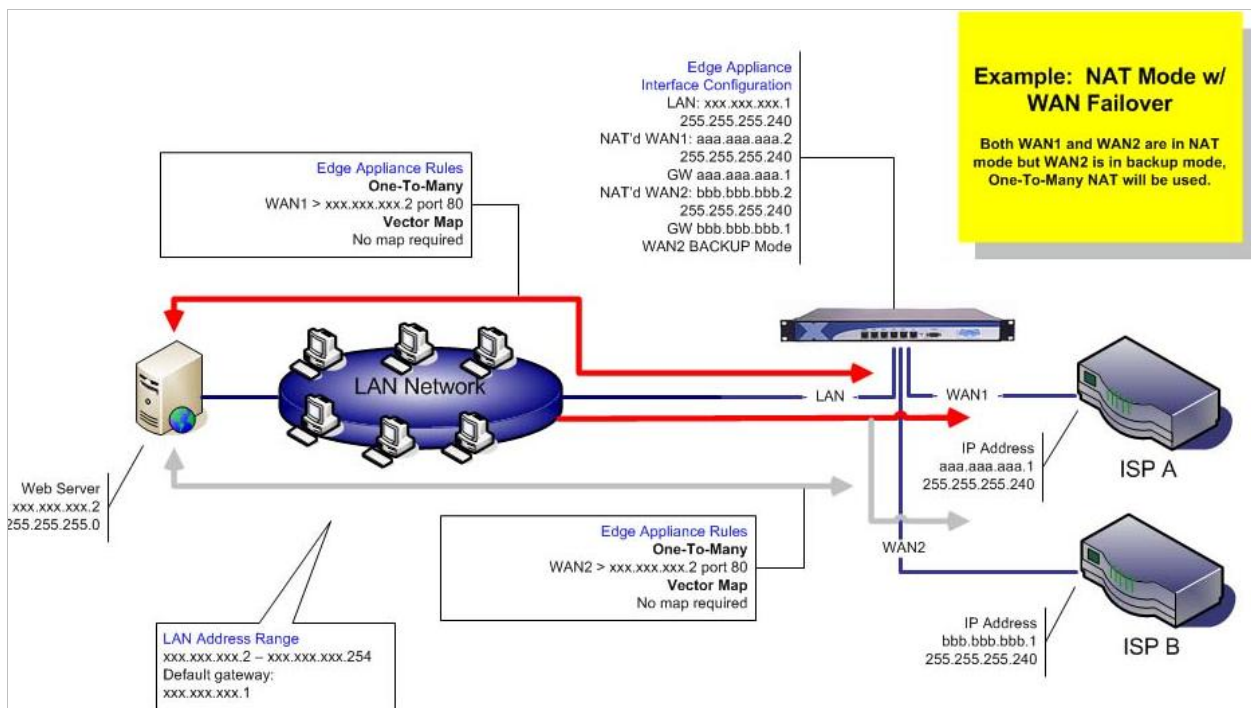
In this case the two WAN connections are also setup to be load balanced, thus both connections will be used at the same time. In general, seeing as both WAN links are setup in NAT mode, VirtualNAT is the best method to use for configuring inbound services.

VirtualNAT sets up a virtual server which proxies the TCP connections to the internal web server. From the web servers' point of view, all connections appear to come from the Edge appliance.

The configuration process for is as follows:

- 1 Setup the WAN1 and WAN2 interfaces. Assign the LAN IP address and test connectivity to the internal firewall and/or server address.
- 2 Create a VirtualNAT rule to the internal server address, either via the NAT'd firewall IP address or the servers LAN address if no firewall is present.
- 3 Bind the address to an external secondary WAN address for WAN1 and WAN2 (these addresses will be added to the WAN interfaces automatically).

No Vector Maps are required for this configuration as the VirtualNAT handles all inbound and outbound connections.



In this example WAN1 is configured in pass-through proxy mode, and WAN2 is NAT'd. Since WAN1 is essentially bridging the WAN1 traffic, no NAT is required for WAN1 addresses they work by default. However WAN2 traffic is NAT'd and thus any server traffic through WAN2 will also need specific NAT rules.

In this case the second WAN connection is setup for failover, thus only the primary connection will be used under normal conditions.

If the Edge appliances VPN module is used in this case, it can be configured to automatically failover from the primary WAN interface to the secondary without any manual re-configuration.

The configuration process for is as follows:

- 1 Setup the WAN1 and WAN2 addressing. The WAN1 addressing will automatically apply to the LAN since we are using proxy mode.
- 2 Create a One-To-Many rule pointing WAN2 -> xxx.xxx.xxx.2 port 80 which will forward traffic destined to the WAN2 interface to the internal server if a failover occurs.
- 3 Create any firewall rules within the Edge appliance for security (see the EdgeWALL HowToGuide for more information).

Advanced One-To-One

This section provides a step-by-step overview of the One-To-One NAT feature. With One-To-One NAT rules any external WAN address can be mapped to an internal LAN address. ALL ports are mapped so there is no need to specify port information.

One-To-One is recommended whenever a large number of ports need to be NAT'd.

NOTE: When a One-To-One NAT rule is created, it automatically adds the WAN address to the external WAN interface as a secondary address to that WAN interface.

One-To-One NAT also provides the ability to load balance traffic across multiple servers. Example: To balance web traffic across multiple internal web servers, specify each server address separated by a comma, this will evenly distribute web traffic across these servers. The Edge device does not perform any specialized tracking so it may not work with some applications.

One-To-One NAT

(Must be different from One-To-Many)

(Check this to forced source NATing when the selected interface is in BACKUP mode)

(Check this to automatically create a reverse Vector Map)

. . . (Must be available via the WAN port selected below)

WAN1

(Forward Address(es) Must be available via the LAN interface)

(Server load balancing can be accomplished by entering the following xxx.xxx.xxx.yyy-xxx.xxx.xxx.zzz)

1

The procedure for creating a One-To-One NAT rule is fairly simple.

Here is a list of the information required:

- What is the IP address of the internal server?
(must be reachable via the LAN interface or via a static route)
- What is the WAN address(es) that will be forwarded to this internal server? (should not be the same as the WAN interface addresses or you will disable remote access management)
- Should this NAT rule have forced source addressing? (Source addressing is not enabled when the link is INACTIVE or in BACKUP mode)
- What order will this rule be applied. The apply order determines when a rule will be applied when multiple rules exist for the same IP address across multiple WAN links.

APPLY ORDER: When using the same LAN (server) address across multiple WAN links, only one WAN link can be used at a time with the same LAN address. However you may want to have failover for this address from one WAN link to the other. This is handled by setting an Apply rule, which determine the priority for this rule when multiple similar rules exist.

One-To-One NAT

web (Must be different from One-To-Many)

(Check this to forced source NATing when the selected interface is in BACKUP mode)

24 . 45 . 65 . 78 (Must be available via the WAN port selected below)

WAN1

192.168.168.100 (Forward Address(es) Must be available via the LAN interface)

(Server load balancing can be accomplished by entering the following xxx.xxx.xxx.yyy-xxx.xxx.xxx.zzz)

Reset Add / Update View Services >>

The process for actually creating a One-To-One NAT rule is as follows:

- 1 Enter the name for this rule.
- 2 Check the "force source" button to force source addressing.

NOTE: Source addressing ensures that the address that is NAT'd inbound is also used for outbound sessions.

- 3 Enter the external address and specify the matching WAN interface.
- 4 Enter the LAN address(es) that will be used for forwarding.
- 5 Select the appropriate Apply number, use '1' if this is the only rule.

NOTE: Specifying multiple addresses will force the sessions to be balanced across the IP addresses provided.

Once added the following screen will appear showing the rule that was added.

| Select | Service Name | Interface | External Addr | Internal Addr | Source NAT |
|-----------------------|--------------|-----------|---------------|-----------------|------------|
| <input type="radio"/> | VoIP | WAN1 | 48.53.85.8 | 10.20.1.20 | On |
| <input type="radio"/> | web | WAN2 | 24.45.65.78 | 192.168.168.100 | Off |

From this screen the One-To-One rules can be modified or deleted.

Advanced One-To-Many

This section provides a step-by-step overview of the One-To-Many NAT feature. With One-To-Many NAT rules any external WAN address and port can be mapped to an internal LAN address and port.

One-To-Many is recommended whenever a small number ports needs to be NAT'd.

One-To-Many NAT also provides the ability to specify an address other than the WAN interface address. Additionally, the external port can be different than the forwarded.

Example: To forward port 80 to an internal port 8080, specify the WAN interface, the secondary WAN address, and port 80. Set the forwarding port to 8088.

NOTE: When an optional external address is specified, the One-To-Many NAT rule automatically adds the external address as a secondary address to the WAN interface specified.

The screenshot shows a configuration form for One-To-Many NAT. It includes a dropdown menu for 'One-To-Many NAT', a text input field with a note '(Must be different from One-To-One)', a dropdown menu for 'WAN1', an 'OR' label, a dotted IP address input field with a note '(Select an interface from the drop down or assign a secondary external address)', an 'Optional' text input field with a note '(Only use if different from the forwarding port)', another text input field with a note '(Enter a port number or range - Example: web - 80 or 80:88)', a dropdown menu for 'TCP', and a dotted IP address input field with a note '(Must be available via the LAN interface)'. At the bottom, there are three buttons: 'Reset', 'Add / Update', and 'View Services >>'.

The procedure for creating a One-To-Many NAT rule is fairly simple. Here is a list of the information required:

- What is the IP address of the internal server? (must be reachable via the LAN interface or via a static route)
- What is the port that will be forwarded to this internal server?

One-To-Many NAT

web (Must be different from One-To-One)

WAN1

OR

. . .

 (Select an Interface from the drop down or assign a secondary external address)

Optional - (Only use if different from the forwarding port)

80 (Enter a port number or range - Example: web = 80 or 80:88)

TCP

192 . 168 . 168 . 100 (Must be available via the LAN interface)

The process for actually creating a One-To-Many NAT rule is as follows:

- 1 Enter the name for this rule.
- 2 Select the protocol type TCP,UDP, PPTP, L2TP.

NOTE: PPTP requires two NAT rules, one for protocol PPTP (leave port blank) and another for TCP port 1723.

- 3 Enter the WAN interface OR add the optional external IP address and port.
- 4 Enter the destination port on the internal server.
- 5 Enter the internal servers IP address.

Once added the following screen will appear showing the rule that was added.

| Select | Service Name | Interface | Protocol | Dst Port(s) | Address | Ext Address | Fwd Port |
|--------------------------|--------------|-----------|----------|-------------|-----------------|-------------|----------|
| <input type="checkbox"/> | RDP | WAN1 | TCP | 3389 | 10.50.10.2 | WAN Addr | Same |
| <input type="checkbox"/> | web | WAN2 | TCP | 80 | 192.168.168.100 | WAN Addr | Same |

From this screen the One-To-Many rules can be modified or deleted.

Server Load Balancing

This feature is typically used to force outbound traffic to a specific remote server when a network failure occurs, or when a proxy server is being utilized. An example would be when a customer is using an external DNS server, but when an outage occurs must use a DNS server from the other provider as the primary DNS is not available to devices outside of its own network. Another example would be if the EdgeXOS administrator wanted all web traffic to be forced to an external proxy server.

| | |
|----------------------------------|--|
| NetBalancing Selection: ? | Server Balancing |
| Server Group: ? | <input type="text"/> (Enter the name of this server grouping) |
| | <input type="text"/> (Enter the TCP port to be shared by this group) |
| | <input type="text"/> (IP Address - Server1) |
| | <input type="text"/> (IP Address - Server2) |
| | <input type="text"/> (IP Address - Server3) |
| | <input type="text"/> (IP Address - Server4) |
| Group Information: ? | <input type="text"/> (IP Address - Server5) |

Application Balancing

This is to determine when redirection will occur, either all of the time as with a remote proxy server, or only when a failover condition exists. A failover condition exists when the primary WAN link, i.e. WAN1 fails. The next step is to enter the IP address of the remote device to which application traffic will be forwarded, i.e. redirected.

Inbound Connections

This is to determine when redirection will occur, either all of the time as with a remote proxy server, or only when a failover condition exists. A failover condition exists when the primary WAN link, i.e. WAN1 fails. The next step is to enter the IP address of the remote device to which application traffic will be forwarded, i.e. redirected.