

KAK - Keeping High Availability with Forefront TMG's ISP Redundancy Feature

<http://blogs.technet.com/b/isablog/archive/2009/02/16/keeping-high-availability-with-forefront-tmg-s-isp-redundancy-feature.aspx>

Overview

Today, more and more businesses rely on their Internet Service Providers link (or ISP) to handle their outside Internet world communications. Sending emails, browsing the web and any other web related actions are essential business infrastructure services that are only available as long the ISP line is up and running.

Keeping a stable, available and reliable outside Internet connection is one of the critical tasks on every administrator's check list.

Forefront TMG provides a new capability called *ISP redundancy* which basically, enables utilizing not one, but two ISP links for external connectivity, either for traffic load balancing or as a failover backup.

This post is an introduction to the ISP Redundancy feature. It covers the basic setup steps, configuration, monitoring and some caveats and tips and tricks. It assumes you have access to two different ISPs lines and a TMG server hardware that has at least two available NICs.

This article is based on the Forefront TMG Beta2 release; later editions may vary in their User Interface and feature availability.

Important: ISP Redundancy is currently not supported in production environments. It can be used in non-production deployments.

Configuring ISP Redundancy

Once you've passed the initial Forefront TMG setup steps, either by manual configuration or by using the Getting Started Wizard, from Forefront TMG's Management console tree open up the Networking pane and use the ISP Redundancy tab and click *Enable ISP Redundancy* to turn this feature on. Clicking *Enable ISP Redundancy* will open up the configuration wizard. The first configuration step is choosing between two modes of operations:

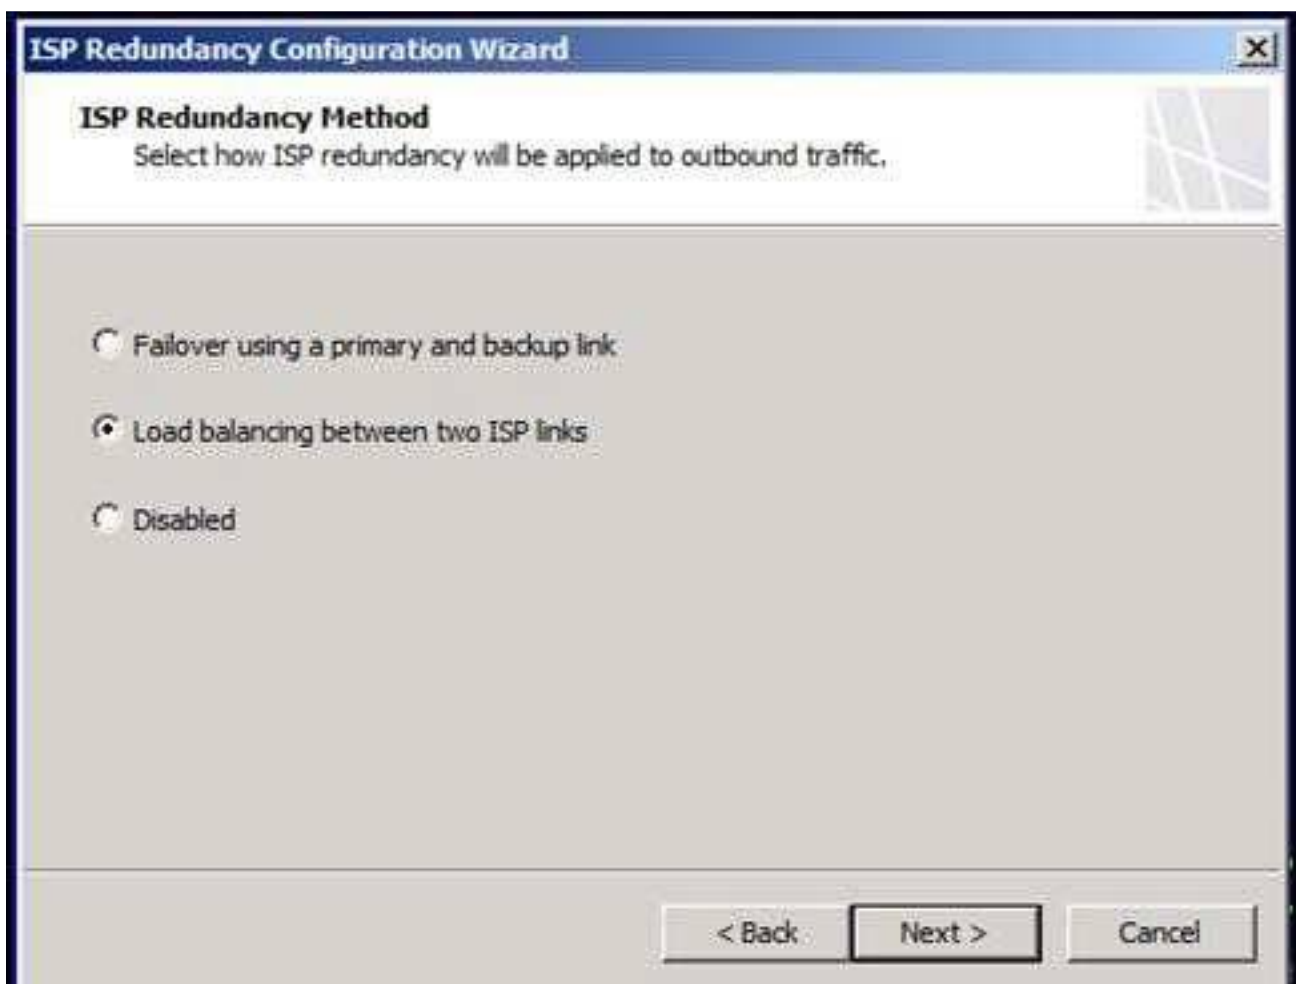


Image 1 – Choosing ISP redundancy method.

Load Balancing – Network connections are distributed between the two active ISP lines. Load factor between the two links can be configured by sliding the percentage rule from one end to the other (See image 2). Distribution levels are determined by the actual number of connections.

Failover – Network connections are routed through the primary ISP Link. The secondary links stays inactive up until the master link connection is broken or disconnected. If the master connection fails the secondary link becomes active by routing the outbound traffic through the second ISP Link. The secondary link will stay active up until the primary link comes back again.

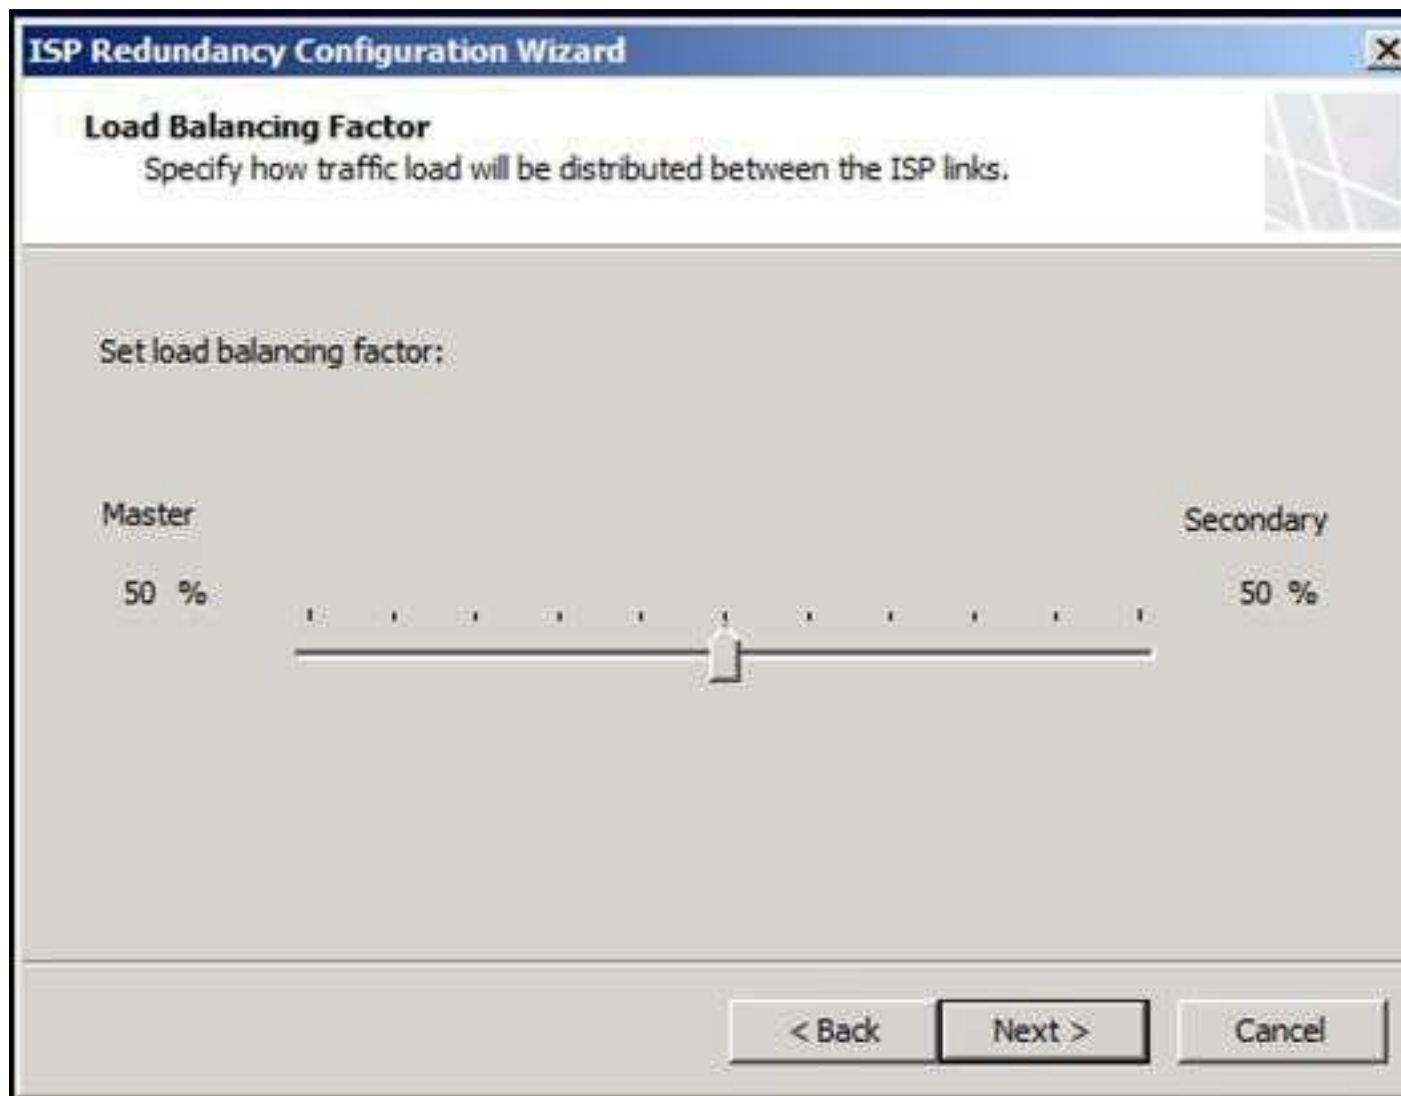


Image 2 – Setting Load Balance Factor for the Load Balance Method

The next steps are similar for both methods. They basically ask for a link name and the specific default gateways and subnet masks of the two ISP links. Within these settings, explicit route destinations can be configured for each ISP link as well. Explicit route definitions, which are only available in load balancing mode, can be useful for defining the DNS servers of both ISPs to be routed explicitly through their respective ISP link. E.g. If we have ISP1 and ISP2 and they have different DNS servers, we'll configure ISP1 Link to explicitly route the DNS ip address from this specific link (by using the *Explicit Route Destinations* buttons). This will make sure Forefront TMG is not querying DNS2 for a name resolution by using ISP1's link. More explicit link traffic control can be set by using another one of Forefront TMG's new Network rules capabilities, see Tips and Tricks section down below for more information.

* Please note that ISP redundancy will function properly if and only if only **one** of the external NICs has its default gateway defined in the NIC properties.

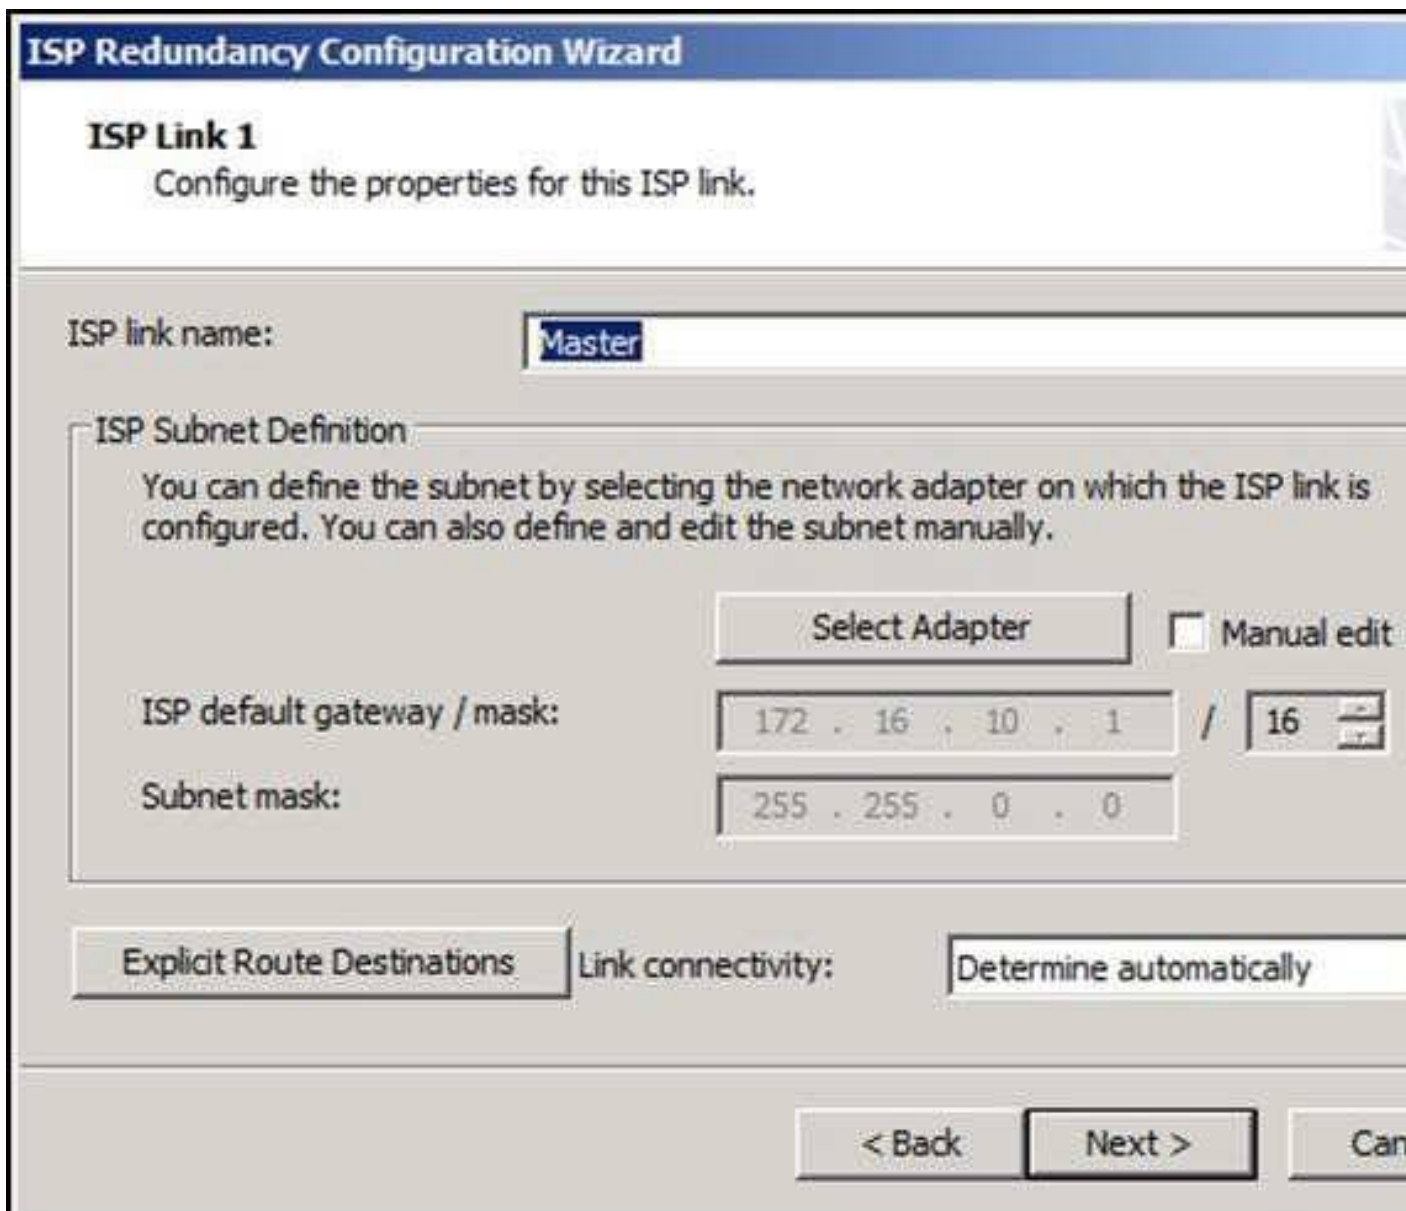


Image 3 – Configuring the 1st link as the master link.

Once we've configured the mode of operation and the two ISP links, we can go on and finalize the settings by confirming our ISP redundancy settings presented in the last wizard step.

We may also be required to apply the changes we've just configured and wait for the changes to fully apply. At this point it would be a good idea to check the Web access connectivity by using a client and our Forefront TMG as the gateway or proxy.

Monitoring

ISP redundancy comes with a status dashboard that can be viewed by clicking on Forefront TMG's console tree Networking node and selecting the ISP redundancy tab (see image below)

The ISP Redundancy pane is composed of two main sections.

The ISP redundancy dashboard upper part is a graphical interface for showing the current ISP link's state.

A blue line represents an enabled and active link.

A red line represents an active but disconnected link.

A gray line represents a disabled link.

The lines on the left side represent configuration health and internal issues related to Forefront TMG, where the lines on the right side represent outbound connectivity status.

The lower part of the dashboard is used to show the statistical distribution when *Load Balancing* mode is being used or a simple "which line is currently active" for *failover* mode.

Note that in Load Balancing mode the statistics shown represent actual traffic bandwidth load passing through.

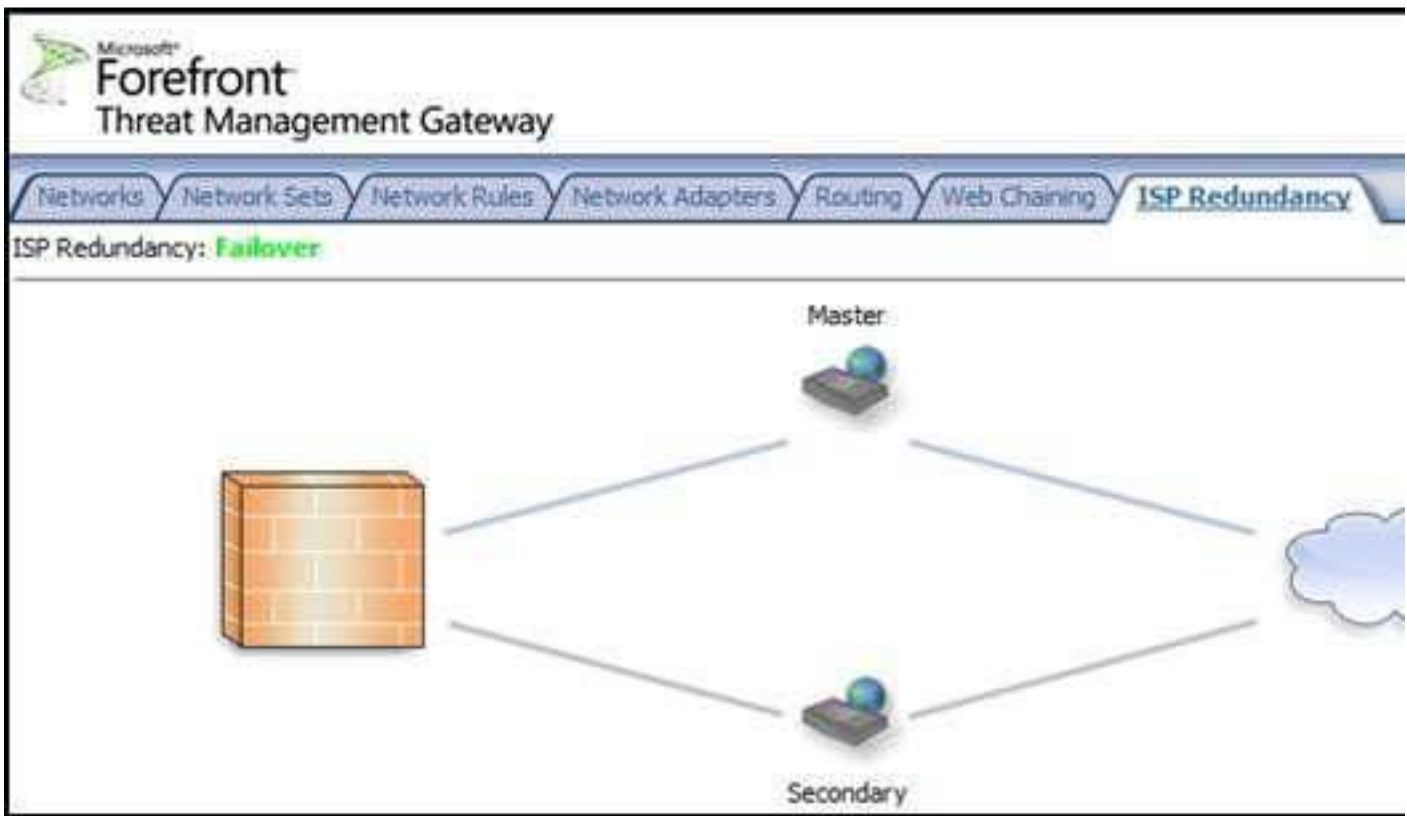


Image 4 –ISP Redundancy status dashboard when working in Failover mode.

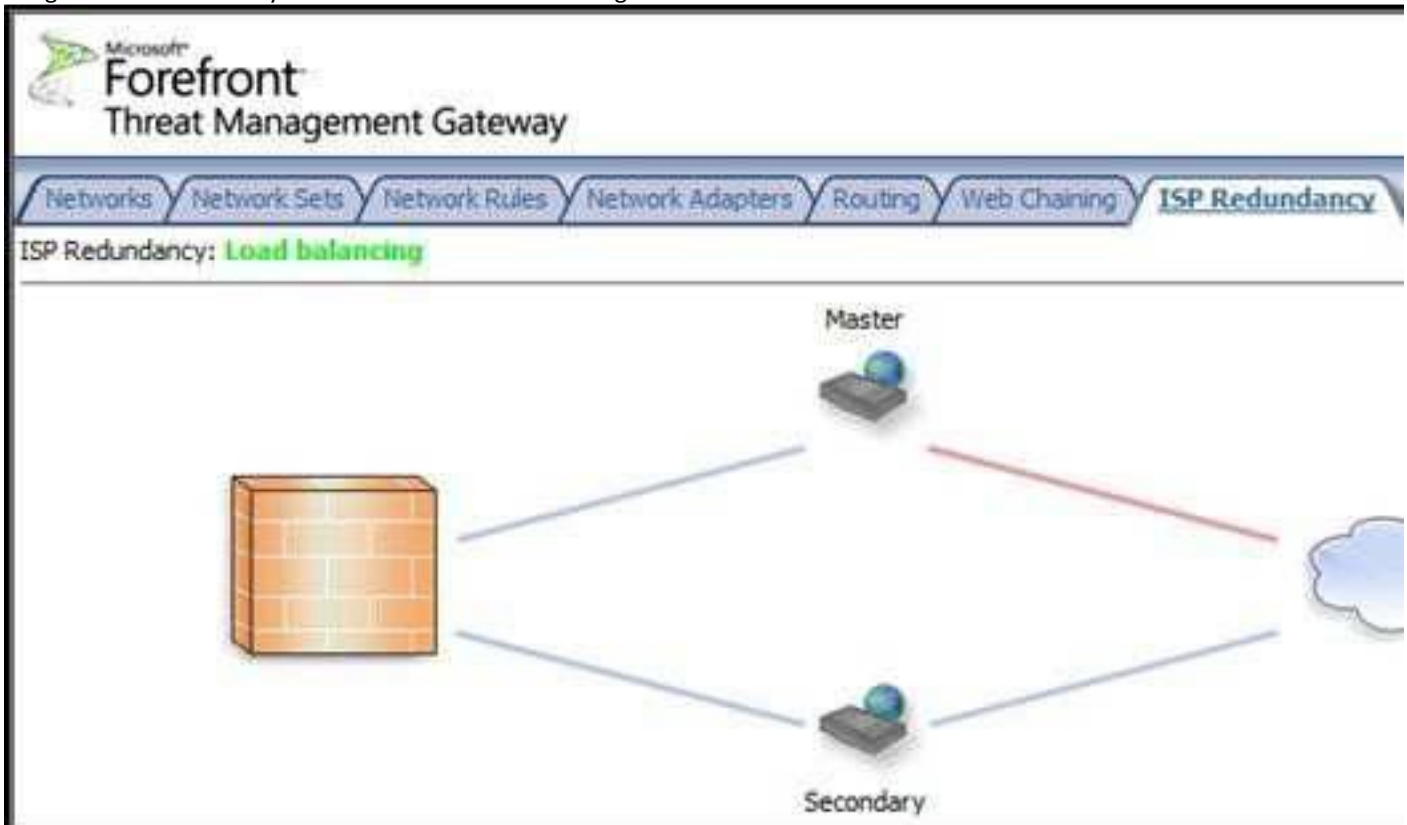


Image 5 – ISP Redundancy status dashboard when working in Load Balancing mode.

Alerts

ISP redundancy comes with five new types of alerts:

ISP Link address missing – No IP address configured on a network adapter of the Forefront TMG computer in the external network can be associated with the ISP Link.

ISP Link is active – Triggered when an ISP link is active, e.g. traffic can or is passing through.

ISP link is available - Triggered when an ISP link is available.

ISP Link is unavailable - Triggered when an ISP link is unavailable or disconnected.

Both ISP Links are unavailable - Both ISP links are detected to be unavailable

A Forefront TMG admin can utilize these alerts to keep track of the current outbound connectivity health, either by going to the alerts dashboard (from the Forefront TMG management console tree Monitoring node and selecting the alerts tab) or by configuring an email alert action. [More information on monitoring and alerts can be found here.](#)

Tips and Tricks

Diverting traffic to a specific ISP Link by using NAT rules

We saw before that we can define explicit IP addresses to be diverted through a specific link. But there are cases where we are required to divert specific internal network subnets through a specific ISP link.

Forefront TMG introduces new network rules settings that can be used to configure these requirements. For example, if we want a subnet to be routed through a specific link we'll can do the following: Set up a new network rule by clicking on the Forefront TMG console networking node and clicking on the *Create a network rule*. We'll set the source and destination for the network, define it as a NAT and pick the *Use selected IP addresses for each network* option on the NAT Address Selection step.

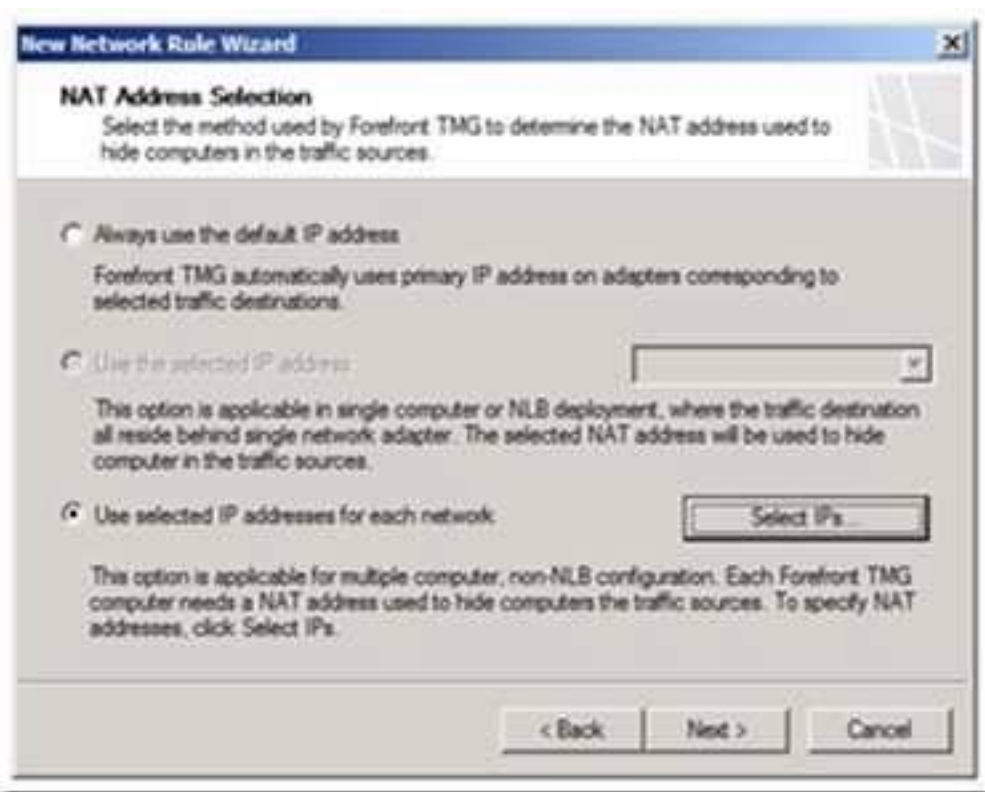


Image 6 – Specifying IPs for each network.

This option enables defining the specific ISP Link IPs by pressing the **Select IPs** button

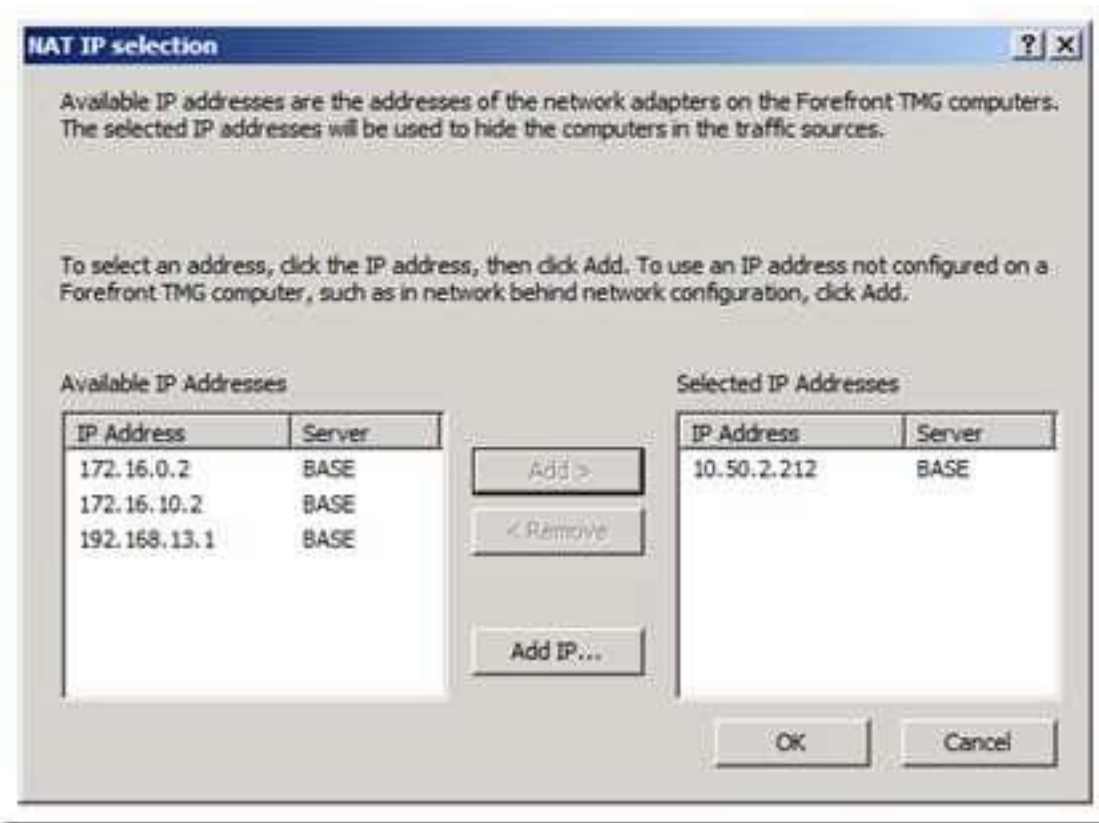


Image 7 – Selecting specific IPs for a specific network.

And hitting ok should do the trick. That way we make sure that the specific internal network that was defined in the source step of this rule, would only go out using the 10.50.2.212 IP.

We can also consider setting the load balance to 90%/10% in favor of the strong ISP link so that regular traffic would mostly go through the strongest link.

Caveats and Troubleshooting

There are some things that are either not ready in this beta version or require specific configuration to work properly that you should be aware of when using the ISP redundancy feature.

- One and only one of the Ethernet cards being used by the ISP feature can have their gateway properties defined by the NIC properties. The other card gateway would be defined from the ISP Redundancy configuration properties step.
- Network rules override ISP redundancy settings. For example, when using *Load Balancing* mode, if a network rule routes specific internal subnet through ISP1, and there is a problem with ISP1's connectivity, that specific sub network **would not** be routed to ISP2.
- The ISP Redundancy configuration steps also have the option of simulating a broken link or forcing Forefront TMG to mark it as active. This can be useful for simulating a broken connection to test the system's reactions or to make sure problematic links are still being defined as active.
- Use the logging tab (under Forefront TMG's management console tree Logs and Reports node) for network activities or ISP redundancies own monitoring dashboard to understand the system's current state.
- Post your issues and questions to relevant TAP or beta distribution list.

This article goal was to provide a basic overview of the ISP Redundancy feature. For any issues, question, suggestion don't hesitate to send your feedback to the TAP or Beta distribution list.

Roiy Zysman

TMG's End to End Scenarios Team.