Check Point
SOFTWARE TECHNOLOGIES LTD.

5 July 2018

HOW TO CONFIGURE ISP REDUNDANCY IN NGX R65 - R77.30 VERSIONS

Check Point
SOFTWARE TECHNOLOGIES LTD
INFINITY

Important Information

Latest Software

We recommend that you install the most recent software release to stay up-to-date with the latest functional improvements, stability fixes, security enhancements and protection against new and evolving attacks.

More Information

Visit the Check Point Support Center http://supportcenter.checkpoint.com.

Latest Version of this Document

Download the latest version of this document in PDF format http://supportcontent.checkpoint.com/documentation_download?ID=12314.

To learn more, visit the Check Point Support Center http://supportcenter.checkpoint.com.

Feedback

Check Point is engaged in a continuous effort to improve its documentation.

Please help us by sending your comments mailto:cp_techpub_feedback@checkpoint.com?subject=Feedback on How To Configure ISP Redundancy .

Revision History

Date	Description
05 July 2018	Improved formatting and document layout Added Gaia OS to Supported Operating Systems (on page 5)
15 February 2017	Improved format and layout
29 September 2015	Clarification on Additional VPN Considerations (on page 17)
14 August 2014	• To see some related documents you must log in to the User Center ("Related Documents and Assumed Knowledge" on page 6) • The location of the ISP Redundancy configuration page in SmartDashboard depends on the release ("Initial Configuration" on page 7)
22 April 2012	First release of this document

Contents

How To Configure ISP Redundancy in NGX R65 - R77.30 versions

Objective

This document describes how to configure ISP Redundancy.

Supported Versions

- R77.XX
- R76
- R75.XX
- R71.XX
- R70.XX
- NGX R65

Supported Operating Systems

ISP Redundancy is supported on these operating systems:

- Gaia (only R75.40 - R77.30)
- SecurePlatform (NGX R65 - R77.30)
- Red Hat Linux 7.2 or above (only for NGX R65)
- IPSO 3.8.2 or above (NGX R65 - R77.30)

Note - Consult the Release Notes of your installed version(s) before you configure ISP Redundancy.

Supported Appliances

All Check Point appliances that support Security Gateway.

Before You Start

Related Documents and Assumed Knowledge

Note - To see *Advanced access* documents, you must have a valid Support contract and sign in to the User Center https://usercenter.checkpoint.com. *General access* documents are public.

- sk25129 http://supportcontent.checkpoint.com/solutions?id=sk25129 (*General access*) – Supported platforms for ISP Redundancy
- sk23630 http://supportcontent.checkpoint.com/solutions?id=sk23630 (*Advanced access*) – Advanced configuration options for ISP Redundancy
- sk32225 http://supportcontent.checkpoint.com/solutions?id=sk32225 (*Advanced access*) – Configuring ISP Redundancy so that certain traffic uses specific ISP
- sk31530 http://supportcontent.checkpoint.com/solutions?id=sk31530 (*General access*) – ISP Redundancy Link Interface cannot be created
- sk40958 http://supportcontent.checkpoint.com/solutions?id=sk40958 (*Advanced access*) – How to verify the status of ISP Redundancy links on command line
- sk25152 http://supportcontent.checkpoint.com/solutions?id=sk25152 (*Advanced access*) – Static NAT fails for outgoing connections through gateway with ISP Redundancy in Load Sharing mode
- A general knowledge of Gateway and Security Management Server management, including creating and modifying gateway objects, updating interface topology, and static and hide NAT configuration.
- A working knowledge of DNS is also recommended.

Impact on Environment and Warnings

- None specifically, but care should be taken when one ISP link offers significantly higher bandwidth than the other. If this is the case, the `$FWDIR/bin/cpisp_update` script on the Security Gateway or each Cluster Member should be modified to prevent the low bandwidth ISP from being overwhelmed by normal traffic levels.

Configuring ISP Redundancy

ISP Redundancy ensures reliable outbound Internet connectivity for a single Check Point Security Gateway or Check Point Cluster. It enables connection through redundant ISP connections.

Before you begin to setup up ISP Redundancy, configure the Security Gateway or Cluster object with two external links defined, one for each of the ISPs. This is the simplest deployment option. Other options are discussed later.

Initial Configuration

1. Connect with SmartDashboard to the applicable Security Management Server or Domain Management Server.
2. Open the Security Gateway or Cluster Object.
3. Go to the ISP Redundancy settings based on your version of the SmartDashboard:

SmartDashboard version	Location
R76 and R77.X	**Other** > **ISP Redundancy**
R65, R70, R71 and R75	**Topology > ISP Redundancy**

4. Select **Support ISP Redundancy**.

   Example for R76 and R77.X SmartDashboard:

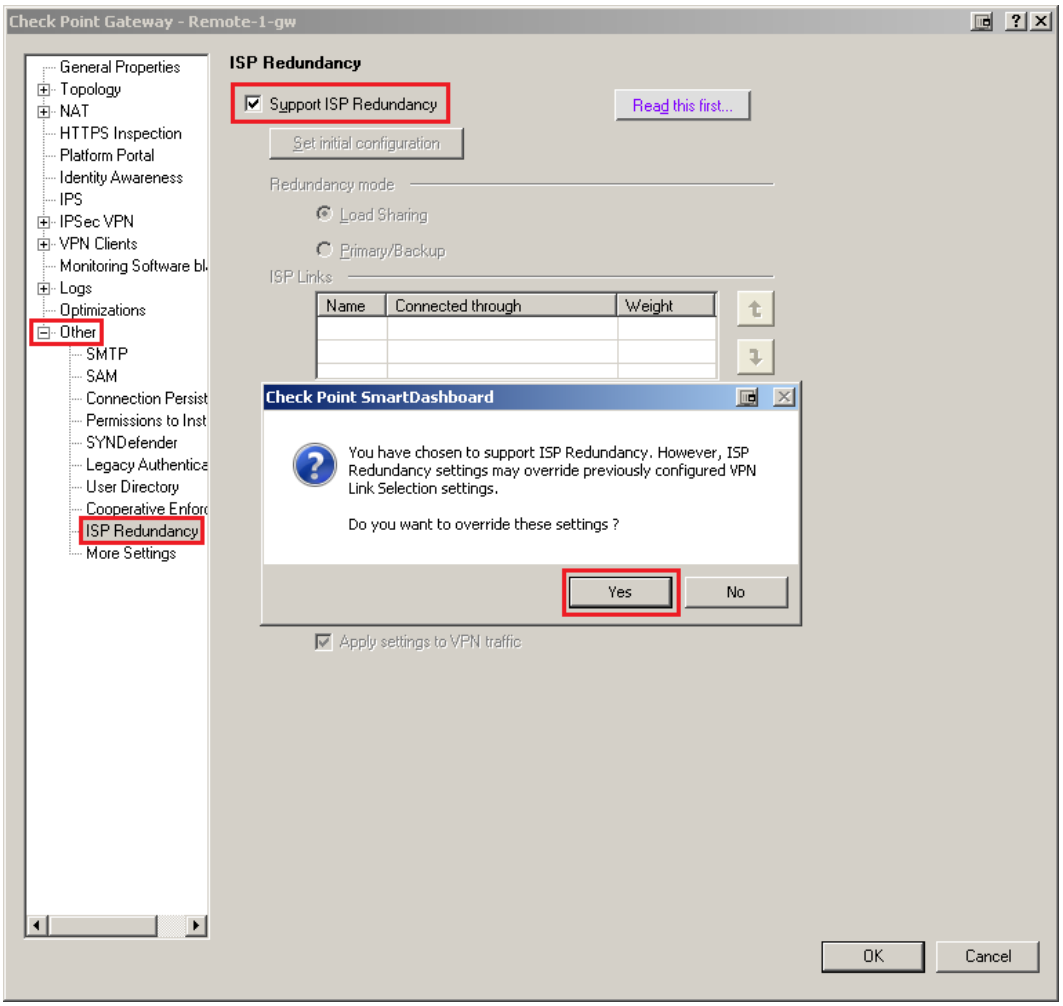

Example for R65, R70, R71 and R75 SmartDashboard:

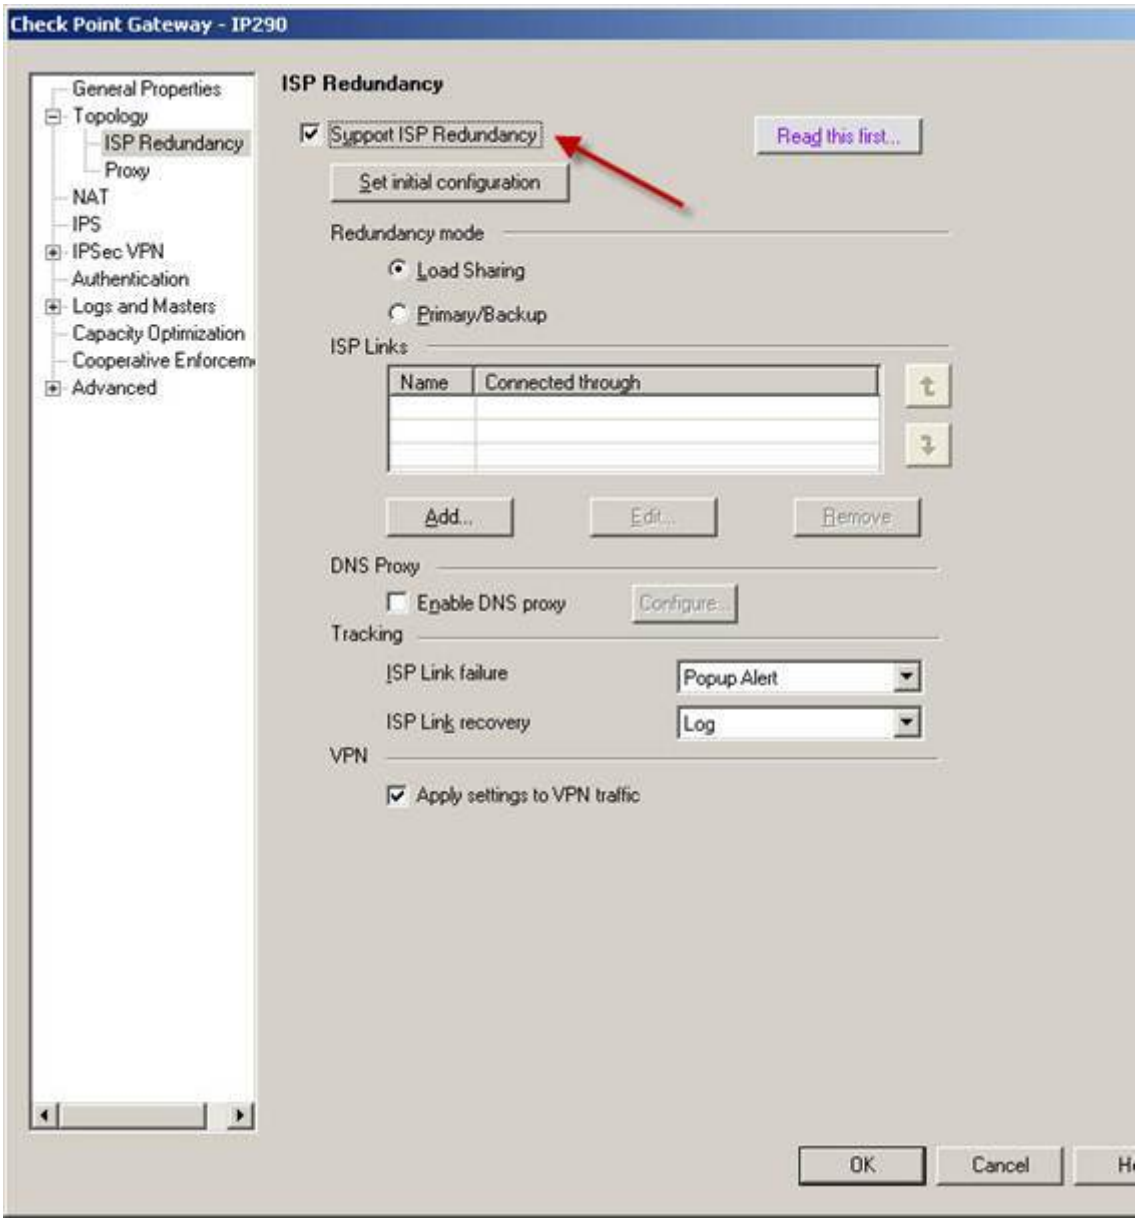

5. Select the ISP Redundancy mode, which controls the behavior of outgoing connections. There are two options available:

   - **Load Sharing**: New outgoing connections are distributed randomly between the ISP links and remain with the assigned link until completion. Should one of the ISP links fail, all new connections are assigned to the link that remains.

     Incoming connections also benefit from the Load Sharing method. In the case of returned packets for a connection initiated outbound, the traffic uses the same ISP link that is used to initiate the connection. Connections that are initiated inbound can use either link to access resources. The Check Point Security Gateway or Cluster can answer DNS requests for internal servers with addresses from both ISPs due to order alternation.

Example for R76 and R77.X SmartDashboard:

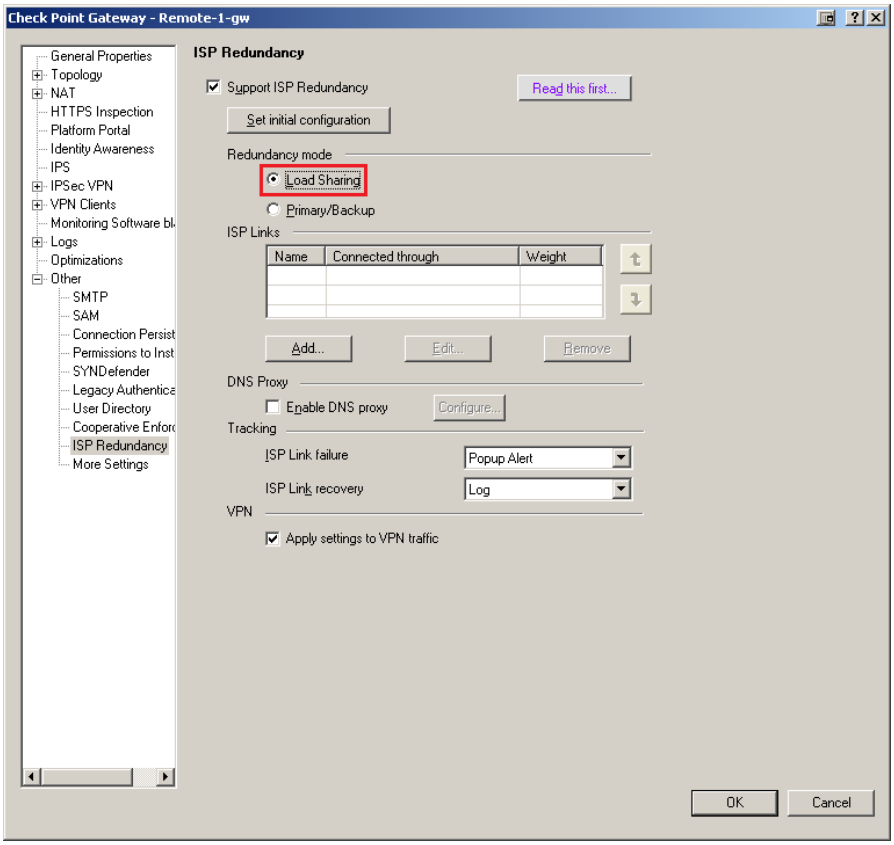

Example for R65, R70, R71 and R75 SmartDashboard:

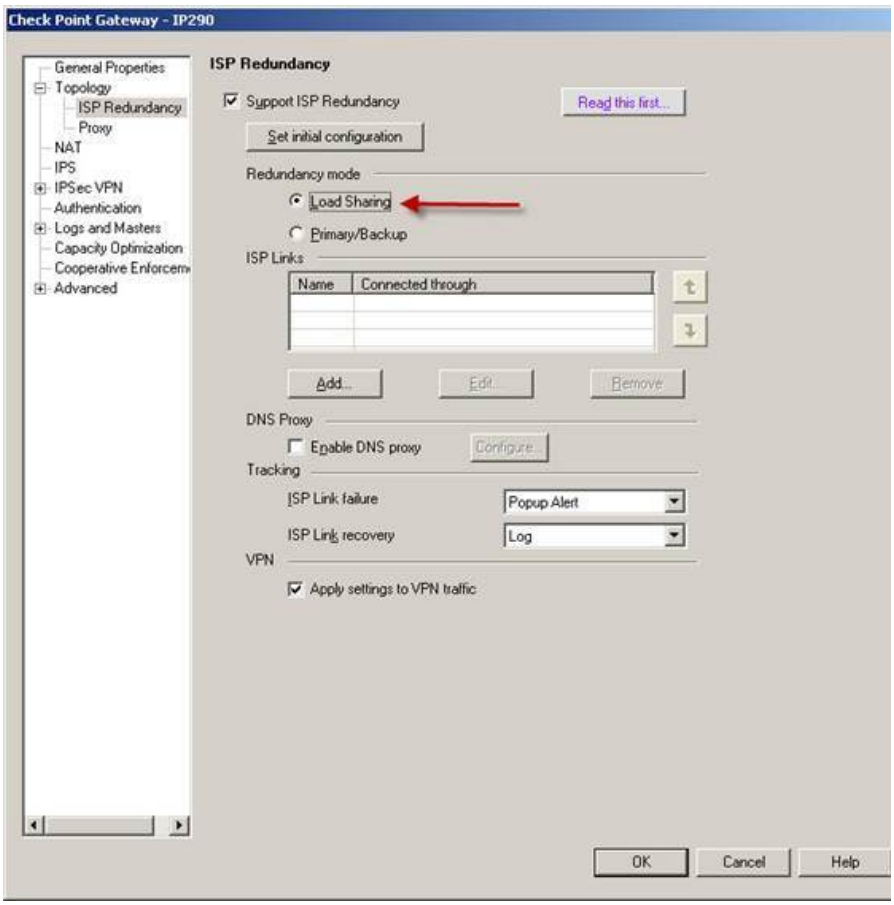

- **Primary/Backup**: New connections use the Primary link as its ISP. In the event of Primary link failure, connections switch to the Backup link, and any new connections use the Backup link as well. Upon recovery of the Primary link, any new outgoing connections begin to use it again while the existing connections on the Backup link continue to use it until completion.

  Example for R76 and R77.X SmartDashboard:

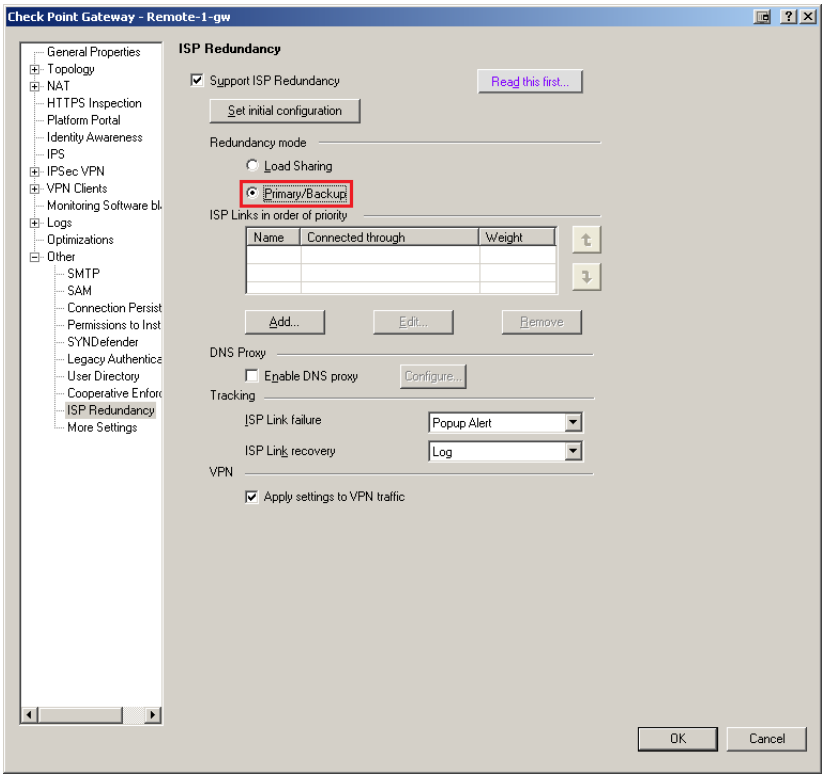

  Example for R65, R70, R71 and R75 SmartDashboard:

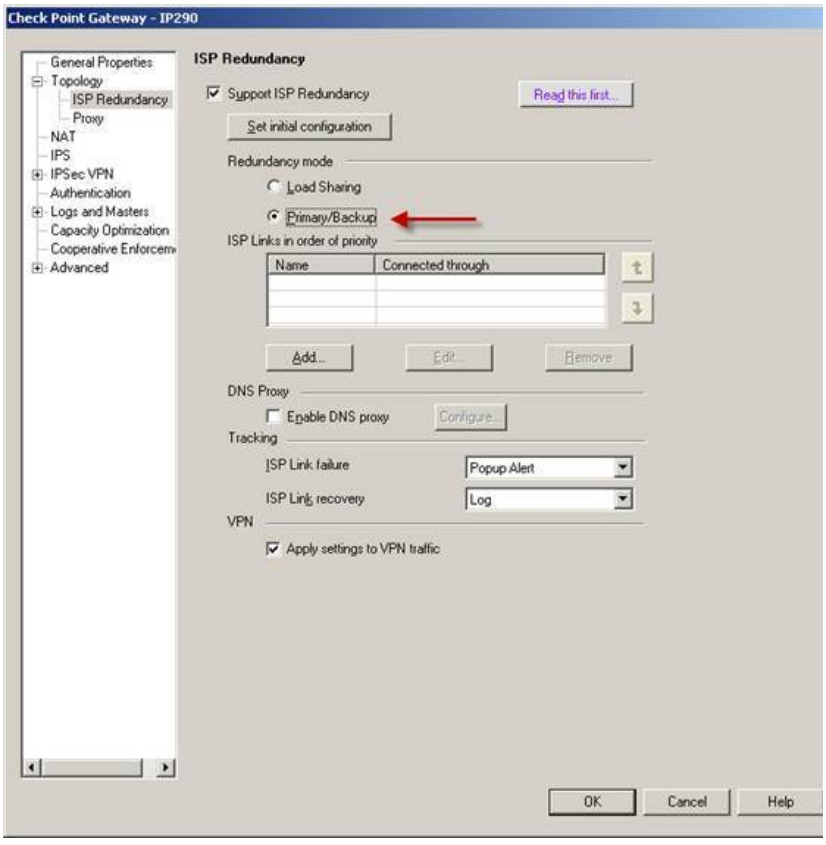

Adding ISP Links

Once the ISP Redundancy mode is selected, add the ISP links.

1. In the **ISP Links in order of priority** section, click **Add**.

   The **ISP Link** window pops up.

   Click the **General** tab.

   a) In the **Name** field, enter the name for the ISP link of your choice. For example: `ISP1`.

   b) From the **Interface** pull down menu, select the external interface that connects to that ISP.

   c) In the **Next Hop IP address** field, enter the IP address of the gateway for the selected ISP. If the interface is a point-to-point interface, leave this field blank.

   Example:

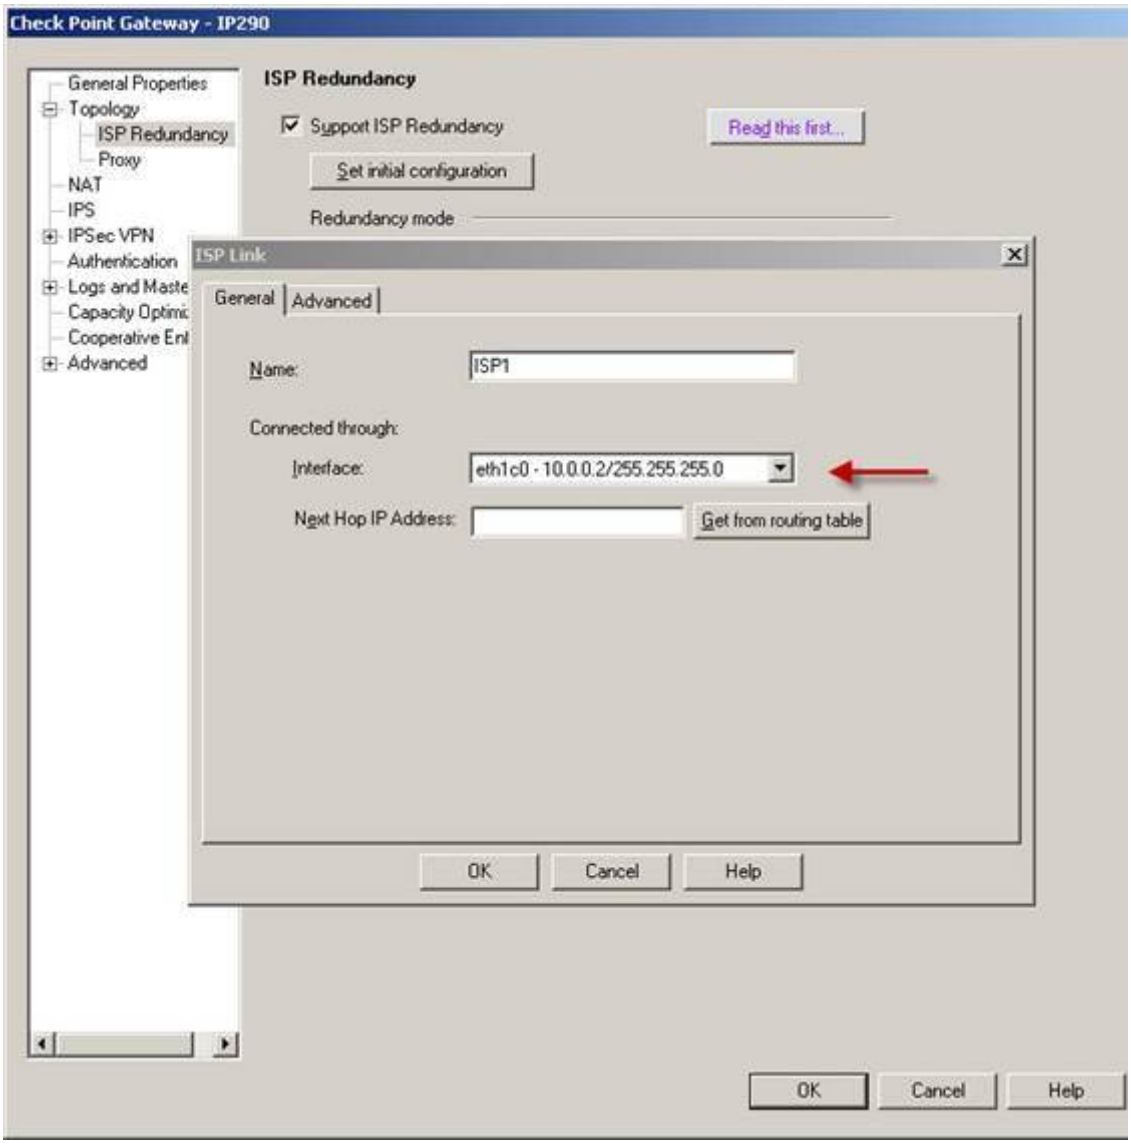

2. Click the **Advanced** tab.

   Here you can define hosts used to perform status checks for this ISP link. Security Gateway sends ICMP Echo Requests to the selected hosts. Failure to respond results in link down status for this ISP. If no hosts are selected, then by default, Security Gateway sends ICMP Echo Requests to the next hop IP address to confirm link status.

3. Click **OK**.

4. Repeat Step 1-3 to add the second ISP link.

Configuring DNS Proxy

If behind the Security Gateway there are internal servers that accept inbound traffic, then you need to enable and configure the DNS proxy feature. This lets the Security Gateway intercept DNS queries to your internal servers that arrive at the Security Gateway's external interfaces, and respond to them:

- In the Primary/Backup Redundancy mode, the Security Gateway responds to DNS queries with the IP address of the interface connected to the active ISP link.
- In Load Sharing Redundancy mode:
  - If both ISP links are active, the Security Gateway responds with IP addresses of the interfaces connected to both ISP links.
  - If one ISP link is down, the Security Gateway responds with the IP address of the interface connected to the active ISP link.

Note - For external access to your internal resources to work, each internal server (that is accessible externally) requires an external IP address assigned for each ISP. Static NAT entries must be configured to translate the external IP address to the real IP address of the internal server.

Procedure:

1. In the **DNS Proxy** section, select **Enable DNS proxy**.
2. Click **Configure**.

   The **DNS Proxy** window pops up.

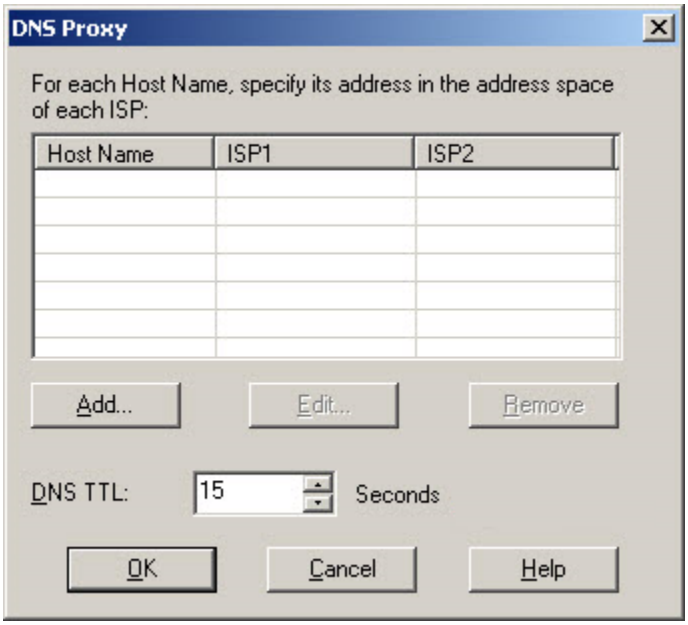

3. Click **Add**.

   The **Host's Addresses** window pops up.

   a) In the **Host Name** field, enter the host name of your internal server. For example: `www.example.com`

   b) In the **Address in ISP 'ISP1'** field, enter the IP address of your internal server for the corresponding ISP link.

   c) In the **Address in ISP 'ISP2'** field, enter the IP address of your internal server for the corresponding ISP link.

   d) Click **OK**.

4. In the **DNS TTL** field, enter the expected time it takes for the Security Gateway to reply. This setting specifies how long the information in the DNS reply may be cached.
5. Click **OK**.

Tracking

In the **Tracking** section, you can select how ISP Redundancy link failures and recoveries are reported. The default is a Popup Alert in the SmartView Monitor for failures and a log entry when the link recovers.

Default options:

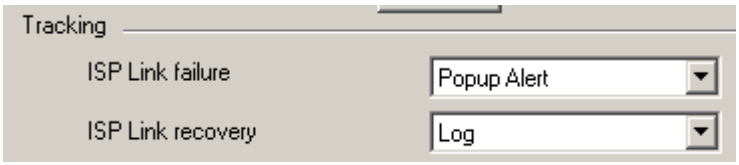

Menu options:

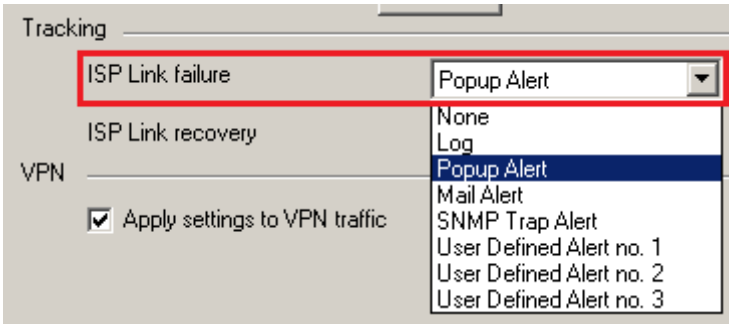

VPN

The selection of **Apply settings to VPN traffic** carries over the configuration on this page to the **IPSec VPN** > **Link Selection** page.

Important - The ISP Redundancy settings override any existing VPN Link Selection setting when this option is enabled.

Example:

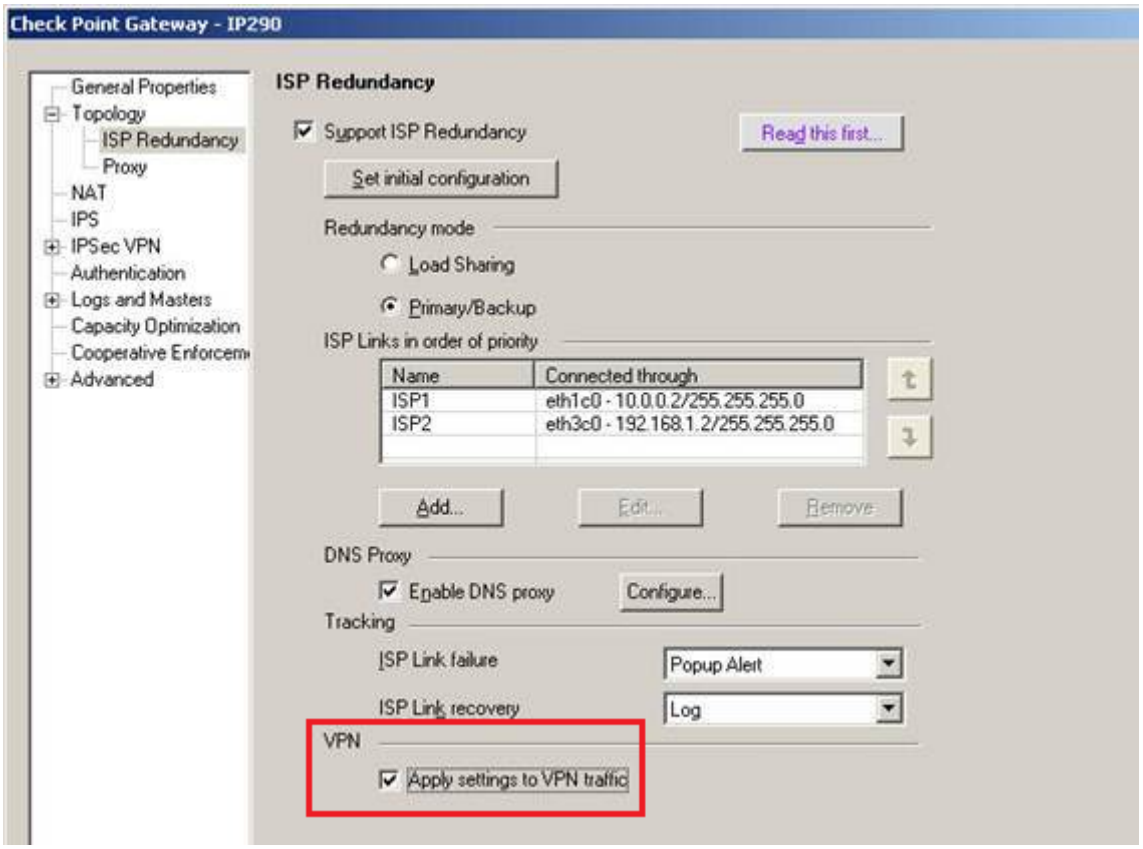

Registering the Domain and Obtaining IP Addresses

The Security Gateway, or a DNS Server behind it, must respond to DNS queries and resolve IP addresses that belong to your internal servers that are publicly accessible. A dedicated DNS server is not required, because the Security Gateway can be configured to intercept the DNS queries ("Configuring DNS Proxy" on page 12).

To Register the Domain and Obtain IP Addresses:

1. Obtain one external IP address from each ISP for the DNS Server, or the Security Gateway that intercepts DNS queries. If external IP addresses are not available, make the DNS Server accessible with Manual NAT.
2. Register your domain with both ISPs.
3. Inform both ISPs of the two IP addresses of the DNS Server that respond to DNS queries for your domain.
4. To allow incoming connections, obtain one external IP address from each ISP for each internal server that is publicly accessible.

Allowing Incoming and Outgoing Connections

1. Define Automatic Hide NAT on network objects that initiate the outgoing connections:

   a) Edit the Network Object that represents your internal network, on which your internal servers are connected.

   b) In the **General** tab of the **Network Properties** window, select **Add Automatic Address Translation Rules**.

   c) Select the **Hide Translation Method** and then the **Hide behind Gateway** option.

   Example:

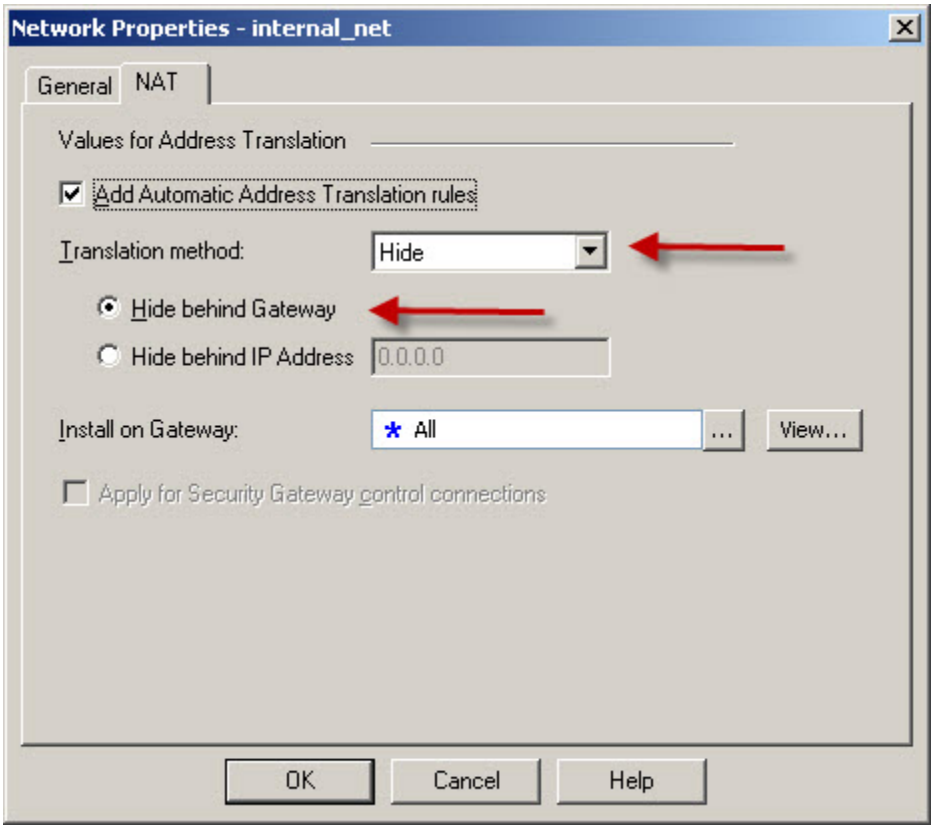

2. To allow incoming connections through both ISP links to your internal servers and the DNS Server, define Manual Static NAT rules.

   If you have only one external IP address from each ISP, and those external IP addresses belong to the Security Gateway, you can allow specific services for specific servers.

In the example below, incoming HTTP connections from both ISPs reach the internal web server 10.0.0.2 (accessible publicly at 192.168.1.2 and 172.16.2.2), and the DNS connections from both ISPs reach the DNS server 10.0.0.3 (accessible publicly at 192.168.1.2 and 172.16.2.2):

Original Packet			Translated Packet		
Source	**Destination**	**Service**	**Source**	**Destination**	**Service**
Any	192.168.1.2	`http`	10.0.0.2	Incoming Web ISP A	`=Original`
Any	172.16.2.2	`http`	10.0.0.2	Incoming Web ISP B	`=Original`
Any	192.168.1.2	`domain_udp`	10.0.0.3	Incoming DNS ISP A	`=Original`
Any	172.16.2.2	`domain_udp`	10.0.0.3	Incoming DNS ISP B	`=Original`

If you have an external IP address from each ISP for each publicly accessible server (in addition to the IP addresses that belong to the Security Gateway), you can allow any service to reach your internal servers, if you give each internal server an internal IP address. In the NAT rulebase, configure:

- Use the external IP addresses in the **Original Packet > Destination** column.
- Use the internal IP address in the **Translated Packet > Destination** column.
- Select **Any** as the **Original Packet > Service** column.

Note - If when you use Manual NAT, automatic ARP does not work for the NATed addresses. On Gaia, SecurePlatform, and Linux Operating Systems, use the `local.arp` file as described in the sk30197 http://supportcontent.checkpoint.com/solutions?id=sk30197. On IPSO Operating System, configure the Proxy ARP.

Alternative Deployment Options

- The simplest and most common ISP Redundancy deployment is with two external interfaces defined, one for each ISP. If only one external interface is available, different subnets can be configured for each ISP. With this deployment, both ISPs are connected to the same interface, but with different next hops, usually through a switch.
- If there is one main ISP and a dial-up connection for the backup, connect one external interface to each ISP. Use the Primary/Backup Redundancy mode to specify that the dial-up connection should only be used, if the Primary link is down.
- In a cluster, each cluster member requires a corresponding external link for each ISP. Follow the cluster configuration guidelines and ensure the physical external interfaces belong to the same subnet as the Cluster VIP address.

ISP Redundancy Script

- The ISP Redundancy script (`$FWDIR/bin/cpisp_update`) on the Security Gateway is executed whenever a Security Gateway starts, or an ISP link state changes. The primary purpose of this script is to change the default route for the Security Gateway according to which ISP links are currently available.
- The ISP Redundancy script can also be used to define advanced options, such as to change link status of a backup dial-up connection when the Primary ISP link becomes active again, or SAM rules to block non-essential traffic when the Primary ISP link is down.

Important - Before you modify this script, make a backup copy. Run the `cpstop` command before you modify the script (in cluster, this can cause a failover), and the `cpstart` command after you save the changes.

Additional VPN Considerations

- The ISP Redundancy script (`$FWDIR/bin/cpisp_update`) must be present on the Security Gateway.
- In a Primary/Backup Redundancy mode configuration, the interface connected to the Primary ISP must be defined as the Primary IP address.
- When ISP Redundancy is configured, the default setting in the **IPSec VPN > Link Selection** page of the Security Gateway object uses the ongoing probing. However, link selection only probes the ISPs configured in the **ISP Redundancy** page. This feature enables connection failover of the VPN tunnel, if connectivity to one of the Security Gateway's interfaces fails.
- The ability of a third party VPN to recognize ISP link failures depends on its configuration. Problems can arise when the third party is unable to recognize traffic from the Secondary ISP link as traffic that originates from the Check Point Security Gateway, or cannot detect link failure and continues to send encrypted traffic to the failed link.

Refer to the **Link Selection** chapter of the *VPN Administration Guide* for more information.

Completing the Procedure

1. Save the changes you made in SmartDashboard.
2. Install the policy.

Verifying the Procedure

ISP Redundancy link status can be manually changed with the `fw isp_link` command.

On the Security Gateway, run:

```
fw isp_link <Name of ISP Link> {up | down}
```

On the Management Server, run:

```
fw isp_link <IP Address of Security Gateway> <Name of ISP Link> {up | down}
```

The command is useful to perform a test, or to disable an ISP link when you know it is not available, but is still reported as active to the Security Gateway. You can also test it if you disconnect the cable from interfaces that connect to ISPs links.

Traffic Behavior

Outbound traffic, such as HTTP traffic, gives the appearance of a seamless transition for clients behind the Security Gateway when the active ISP link fails. With authenticated traffic, such as a VPN connection or FTP transfer, a client behind the Security Gateway suffers a disruption when the active link fails. When the transition to the other ISP link occurs, it changes the IP address that is used to access the remote resources. This, in turn, requires the user to authenticate again from the new IP address.

Index