

Hybrid Branch Office VPN Architecture (Partial Mesh)

Example configuration files created with — WSM v11.4.1

Revised — 6/28/2011

Use Case

In this configuration example, an organization has multiple locations of different sizes and wants to connect the networks at each site. They have several sites that host shared network resources used by the entire organization. They also have small offices, which need to connect to the shared network resources. The organization wants direct connectivity between the sites that host shared network resources. They also want connections from small offices to converge at a reliable central location. The organization could have resources distributed among multiple sites, or business processes that fit a hybrid VPN architecture.

Note *This configuration example is provided as a guide. Additional configuration settings may be necessary, or more appropriate, for your network environment.*

Solution Overview


In a hybrid VPN configuration, also known as a *partial mesh*, some sites are interconnected directly to each other (a mesh configuration), while other sites connect to a central location (hub and spoke).

Mesh Between Primary Sites

In this solution, the sites that host shared network resources are considered primary sites. VPN connections between the primary sites are configured as a mesh, which means that all these sites have direct VPN connections to each other. This configuration depends heavily on the reliability of the primary sites that host unique network resources, because users must have reliable connectivity to these resources. If the organization adds additional remote locations, it could become necessary to expand the capacity of a primary site. This configuration provides resiliency for primary sites, because a failure at a single site impacts only services directly dependent on it.

Hub and Spoke Between Secondary Sites and a Central Hub


In this solution, the small offices are considered secondary sites. VPN connections to the secondary sites are configured as hub and spoke. One of the primary sites, the Colocation site, is the central hub for VPN connectivity to the small offices. The central hub uses tunnel switching to route VPN traffic between each secondary site and all other sites. The small offices depend heavily on the reliability of the central hub, because it is a possible single point of failure for VPN tunnels from these secondary sites. If the organization adds additional remote locations, it could become necessary to expand the capacity of the central hub.



If most shared network resources are located at a single location, a centralized (hub and spoke) VPN architecture could be a better solution. If shared network resources are primarily dispersed among the remote sites, a decentralized (full mesh) VPN architecture could be a better solution.

How It Works

The XTM device at each primary site establishes a VPN connection to an XTM device at every other primary site. The XTM device at the central hub (Colocation) acts as the primary gateway for VPN tunnels from all secondary sites. The central location receives all data transferred for secondary sites. If the central location receives a connection that is not intended for a resource at the central location, the device at the central location redirects the traffic to the tunnel for the destination. This is known as tunnel switching.



Requirements

Reliable connectivity

While this is a fault-tolerant design, sites that host resources unique to their location should have reliable connectivity appropriate for the resources they host.

A reliable central location

The central location handles the aggregate of all VPN connections. All VPN traffic from the small offices depends on the availability of this site.

Sufficient bandwidth

Switched tunnels require bandwidth at the source, destination, and the central location. As shown in the previous diagram, the Small Office that receives traffic from HQ uses the upstream bandwidth at HQ, the upstream and downstream bandwidth at the Colo, and the downstream bandwidth at the Small Office. Due to encryption and encapsulation overhead, VPN bandwidth is measured at less than line speed.

An XTM device appropriate for each location

XTM device capabilities vary by model. For VPN configurations, you must consider the VPN throughput and tunnel capacity of each model. Network environment, configuration options, and other factors may also help you determine the most appropriate model for each site.

VPN throughput is the amount of data passed over the VPN per second. The central location processes switched traffic twice.

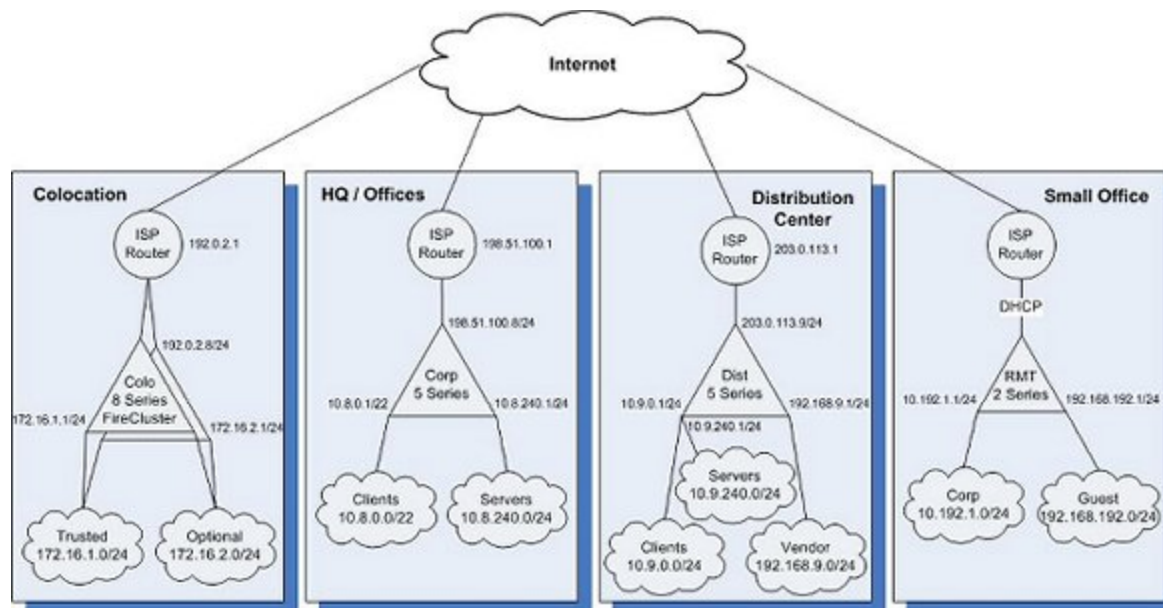
VPN tunnel count is determined by the number of connected networks (as configured in tunnel routes). For offices, this is generally the number of local networks multiplied by the number of remote networks. For the central location, this is the sum total of the tunnel count at all other locations.

For more information about the VPN throughput and branch office VPN tunnel capacity available for each XTM model, refer to the product datasheets: <http://www.watchguard.com/products/resources/datasheets.asp>.

Configuration Example

To illustrate this use case we present an example of an organization that has four locations: a colocation facility (Colo), a corporate office (Corp), a distribution center (Dist), and a small office (RMT). You can also scale up this solution to support additional offices, distribution centers, and small offices.

Topology



The IP addresses for the sites in this configuration:

	Colo	Corp	Dist	RMT
External interface IP address	192.0.2.8/24	198.51.100.8/24	203.0.113.9/24	DHCP
Default gateway IP address	192.0.2.1	198.51.100.1	203.0.113.1	DHCP
Private network allocated to site	172.16.0.0/16	10.8.0.0/16	10.9.0.0/16	10.192.1.0/24
Un-routed network allocated to site	N/A	N/A	192.168.9.0/24	192.168.192.0/24

Example Configuration Files

For your reference, we have included example configuration files with this document. To examine the details of the example configuration files, you can open them with Policy Manager. There are four example configuration files, one for each location in the example.

Configuration Filename	Description
Hybrid-Colo.xml	Central location for the VPNs, the colocation facility
Hybrid-Corp.xml	A corporate office
Hybrid-Dist.xml	A distribution center
Hybrid-RMT.xml	A small office

Configuration Explained

The example configuration files contain branch office gateways and branch office tunnels defined for VPN connections between each site. Each site has three branch office VPN gateways and three branch office VPN tunnels configured.

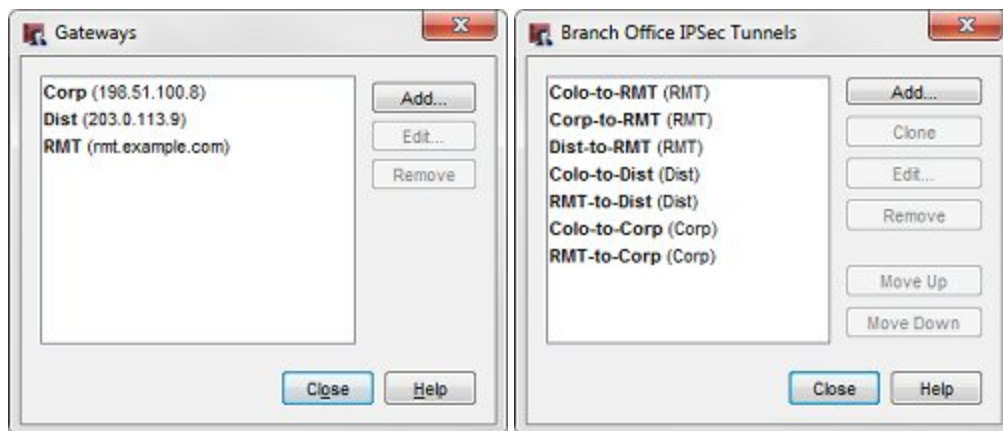
To see the branch office VPN gateways:

1. Start Policy Manager for the XTM device.
2. Select **VPN > Branch Office Gateways**.

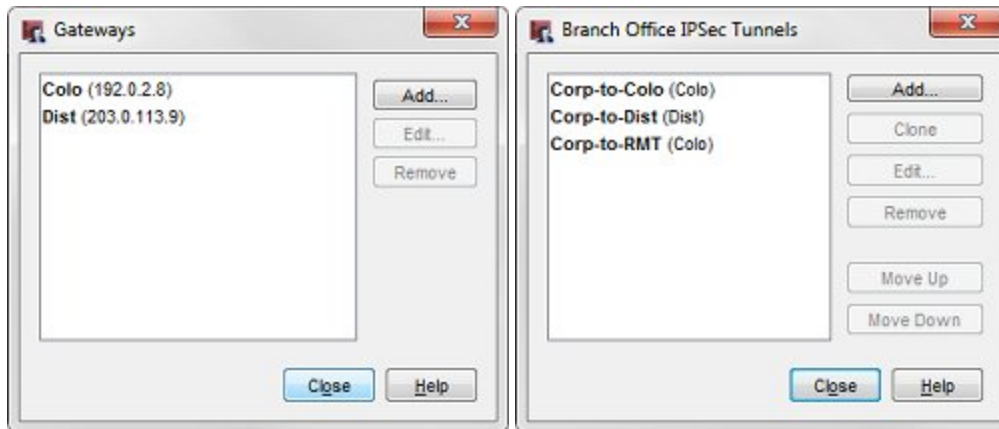
To see the branch office VPN tunnels:

1. Start in Policy Manager for the XTM device.
2. Select **VPN > Branch Office Tunnels**.

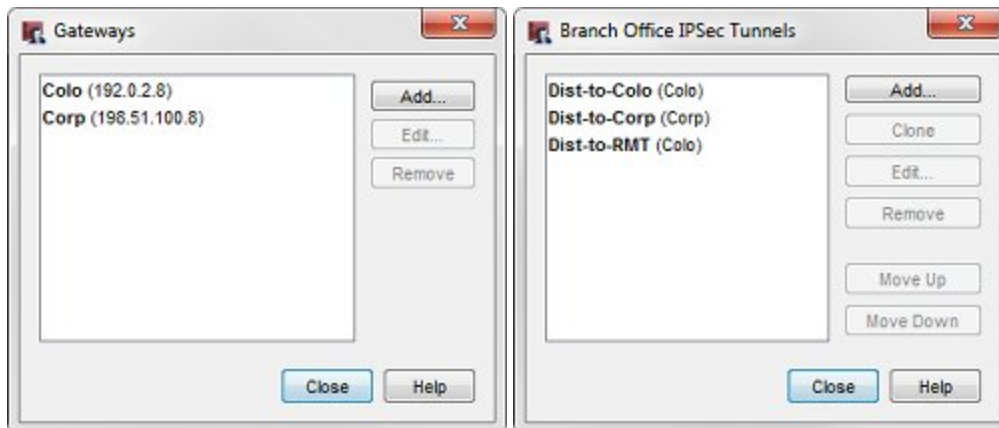
Configuration at the Colocation Site (Colo)



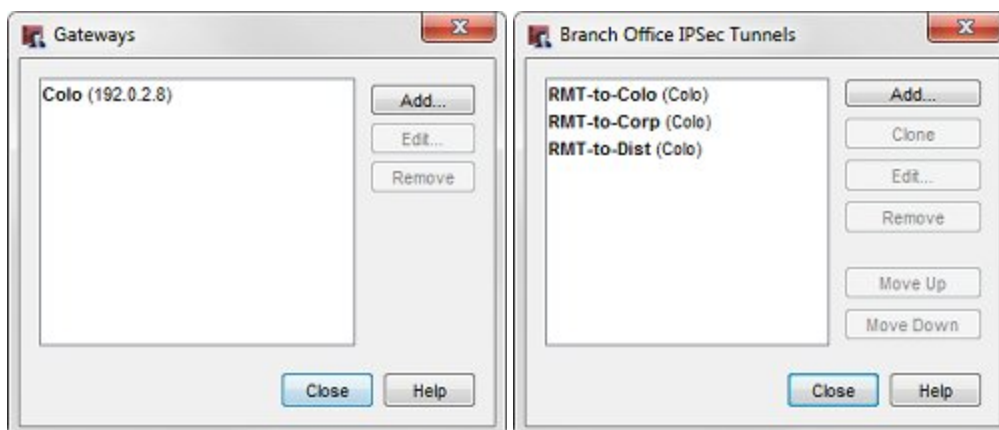
Configuration at the HQ Corporate Network (Corp)



Configuration at the Distribution Center (Dist)



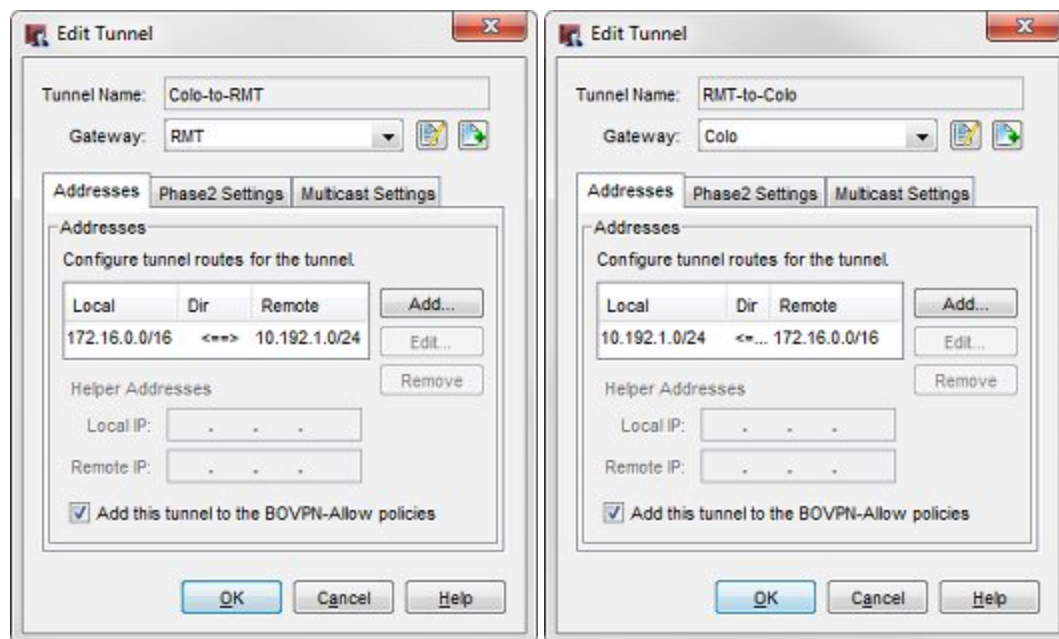
Configuration at the Small Office (RMT)



In the example configuration files, each tunnel is named to represent the local and remote networks it manages. The identifier in parentheses is the gateway used by the tunnel. The Colo has three gateways (one for each of the other primary sites and

one for the secondary site), the primary sites have two gateways (one for each of the other primary sites), and the secondary sites have only one gateway (the Colo).


The tunnel routes have been defined to use the subnets allocated to each site, not the individual networks defined within the site. In this configuration, the small office (RMT) only requires three tunnel routes (not six tunnel routes) to reach the trusted and optional networks at each of the other sites. Any new networks in this allocation established at each site are routed over the existing branch office VPN.



For example, the tunnel routes **Colo-to-RMT** and **RMT-to-Colo** use the subnet IP address 172.16.0.0/16 as the address of the Colo network. This enables these tunnels to handle all traffic between the small office (RMT) network and the Colo trusted (172.16.1.0) and optional (172.16.2.0) networks.

When you look at the tunnel routes, remember that the local-remote pairs are defined relative to the two endpoint networks for the tunnel traffic. In some cases, the local address in a VPN tunnel route is the address of a network at another connected site. For example, in the Colo configuration, the **Corp-to-RMT** tunnel route uses the network IP address of the trusted network at Corp as local, even though it is not physically located at the Colo site.


The subsequent diagram represents all of the local and remote IP addresses of the tunnel routes that are configured between each location.




Tunnel Switching in Action

Now we can use the example configuration to follow the path a packet takes when a user at one location establishes a connection to a resource at a different location over switched tunnels.


A user at the small office (10.192.0.100) tries to connect to a resource at the corporate office (10.8.240.80). The packet first reaches the XTM 2 Series device at the small office. The XTM 2 Series device determines that the destination of the packet is available through the **RMT-to-Corp** tunnel to the Colo gateway.




The XTM 2 Series device sends this packet through the **RMT-to-Corp (Colo)** tunnel.




The XTM 8 Series device at the Colo receives this traffic identified as part of its **Corp-to-RMT (RMT)** tunnel in its local configuration. The local network IP address in this tunnel route in the Colo configuration file is local to the Corp site, not the Colo site.



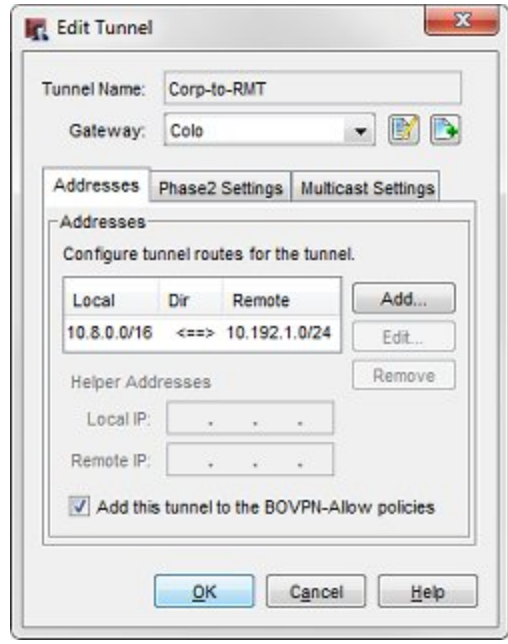
The XTM 8 Series device determines that the destination of the decrypted packet is available through the **RMT-to-Corp (Corp)** tunnel to the Corp gateway.



The XTM 8 Series device at the Colo switches the traffic from the **Corp-to-RMT (RMT)** tunnel to the **RMT-to-Corp (Corp)** tunnel.



The XTM 5 Series device at the corporate office receives this traffic identified as part of its **Corp-to-RMT (Colo)** tunnel, and delivers the decrypted packet to its destination, a server on the corporate office local network.



Conclusion

This configuration example demonstrates how to configure a partial mesh VPN network topology. In this configuration example, each primary site has a direct VPN connection to every other primary site. In the partial mesh topology, one primary site acts as a central hub to perform tunnel switching of VPN traffic between primary sites and secondary sites that directly connect to the central site.

This type of configuration can be a good fit for an organization that has resources distributed among several sites, or business processes that fit a hybrid architecture. The configuration described here can scale up to support additional primary or secondary sites.

This configuration example also shows how to use subnet IP addresses in the tunnel route configuration to reduce the number of tunnels you must configure to connect private networks at each site.

For more information about how to configure branch office VPNs, see the *Fireware XTM WatchGuard System Manager Help*.

About this Configuration Example

This configuration example is provided as a guide. Additional configuration settings could be necessary, or more appropriate, for your network environment.

For complete product documentation, see the *Fireware XTM WatchGuard System Manager Help* or *Fireware XTM Web UI Help* on the WatchGuard web site at: <http://www.watchguard.com/help/documentation/>.

Information in this document is subject to change without notice. Companies, names, and data used in examples herein are fictitious unless otherwise noted. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without the express written permission of WatchGuard Technologies, Inc.

Copyright, Trademark, and Patent Information

Copyright © 1998-2011 WatchGuard Technologies, Inc. All rights reserved. All trademarks or trade names mentioned herein, if any, are the property of their respective owners.

Complete copyright, trademark, patent, and licensing information can be found in the Copyright and Licensing Guide, available online at: <http://www.watchguard.com/help/documentation/>.

About WatchGuard

WatchGuard offers affordable, all-in-one network and content security solutions that provide defense-in-depth and help meet regulatory compliance requirements. The WatchGuard XTM line combines firewall, VPN, GAV, IPS, spam blocking and URL filtering to protect your network from spam, viruses, malware, and intrusions. The new XCS line offers email and web content security combined with data loss prevention. WatchGuard extensible solutions scale to offer right-sized security ranging from small businesses to enterprises with 10,000+ employees. WatchGuard builds simple, reliable, and robust security appliances featuring fast implementation and comprehensive management and reporting tools. Enterprises throughout the world rely on our signature red boxes to maximize security without sacrificing efficiency and productivity.

For more information, please call 206.613.6600 or visit www.watchguard.com.

Address

505 Fifth Avenue South
Suite 500
Seattle, WA 98104

Support

www.watchguard.com/support
U.S. and Canada +877.232.3531
All Other Countries +1.206.521.3575

Sales

U.S. and Canada +1.800.734.9905
All Other Countries +1.206.613.0895