

White Paper

Dynamic VPNs Achieving Scalable, Secure Site-to-Site Connectivity

How to cost-effectively replace WAN connections with a more reliable communication infrastructure

Sarah Sorensen
Product Marketing Manager



Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
408 745 2000 or 888 JUNIPER
www.juniper.net

Part Number: 200051-001

Contents

Introduction.....	3
What is an IPSec VPN?.....	4
Why a Standard IPSec VPN is good, but not sufficient	5
Simple, Scalable VPN Connectivity.....	5
Rule-Based VPN Approach	6
Route-Based VPN Approach.....	7
Dynamic Route-Based VPN Approach	8
Scalability Conclusion	8
High Performance.....	9
Throughput.....	10
VPN Tunnel Setup	10
Low Latency	10
Optimized Resource Management	11
Performance Conclusion.....	11
Connection Resiliency	11
Solid State Operation with Component Redundancy.....	12
Device Redundancy- High Availability	12
VPN State Synchronization	12
Redundant Physical Paths	13
High Availability - Full Mesh	14
Dynamic Routing.....	14
VPN Path Monitoring.....	14
Resiliency Conclusion, VPN Can be More resilient than legacy WAN	15
Integrated Security.....	15
Zone-Based Approach Allowing Firewall Control in a Dynamic VPN.....	16
Conclusion: Juniper Provides Scalable, Resilient, High Performance VPN.....	18

Introduction

Private networks carry vital and sensitive business communications between employees, customers, and enterprise partners across widely dispersed geographic areas. All users of these private networks, no matter how remote, expect to be able to access data and resources as if they were located at the same physical site in the enterprise. Thus private networks must be able to provide reliable, “always on” connectivity for all users, while protecting private data and resources.

Originally, enterprises could only rely on service providers to deploy legacy private wide area networks (WANs), using technologies such as point-to-point leased lines, Frame Relay or ATM, to connect users. Organizations were generally pleased by the performance and ease of management of these Layer 2 solutions and, because of limited alternatives, accepted the significant initial and ongoing costs of these solutions. In addition, they were willing to sacrifice flexibility and security to achieve the availability of the network.

Point-to-point leased lines are physical connections between destinations reserved for the use of the sites they are connecting. Service providers provision dedicated links within their network then lay physical lines from their central office to the organization’s premise. As a result, the setup time and costs are high, but, once up, the organization has a dedicated line that they can rely on for their communications. Organizations were generally pleased with the reliability and simplicity of these point-to-point solutions, however, there were shortcomings, namely related to their flexibility and security. The reliance on a physical connection between destinations doesn’t allow changes to be made quickly and can be hard to scale. Also, because these lines are dedicated for the organization’s use only, it is expensive and potentially cost prohibitive to implement redundant links. So, while generally reliable, if a component or line does go down, an enterprise has to wait for the service provider to restore service before network resources become available again. This downtime could cost the company in lost revenue and productivity. In addition, the security risks of these solutions are the same as those that go with any non-encrypted WAN technology, where service providers and unauthorized users can “tap” into the user data streams without the knowledge of the end user.

Packet-switched technologies, such as Frame Relay and ATM, were introduced to provide a lower cost alternative to point-to-point leased-lines. These solutions reduce the cost by sharing the use of the carrier’s transport infrastructure (cloud) with other subscribers, instead of using a dedicated circuit. This allows for a point-to-multipoint configuration, so organizations only need one physical link and interface to reach multiple locations. Since the cost of the service is based on a Service Level Agreement (SLA), in order to achieve a lower price point, the SLA can often be compromised. A lower SLA means enterprises suffer from reduced performance when the cloud is under load. Plus, because the lines are partitioned and the service provider controls the cross-connects of logical circuits, there is the occasional occurrence of the data being delivered to an incorrect recipient, introducing additional security concerns.

The latest evolution in solutions for private networks is IPSec virtual private networks (VPNs). These solutions lower the capital outlay and ongoing costs of legacy solutions by leveraging the cost-effective and ubiquitous connectivity of the Internet to transport private data. VPNs enable enterprises to use the Internet infrastructure to quickly extend the private network across geographically distributed locations and gain the point-to-multipoint advantage. Also, the control of the “WAN” is now in the customer’s hands, not the provider’s. Any changes to the architecture can be made quickly, making it a very scalable solution. Even with all of these benefits, the adoption of VPN technologies has been slowed because organizations want to make sure they can achieve the same simple to manage, reliable communications they received from their leased-lines.


With Juniper Networks, organizations no longer have to make sacrifices. Using the dynamic route-based VPNs in the integrated FW/IPSec VPN product line, organizations can achieve the connectivity, high performance, resiliency and security they need in their private site-to-site networks. With Juniper, organizations are able to achieve the same, if not better, connectivity and reliability than they received from legacy solutions, in a flexible, cost-effective and truly secure site-to-site private network. This paper is designed to provide a framework of the requirements for a resilient, high performance, secure solution and demonstrate how Juniper meets those requirements.

What is an IPSec VPN?

IP Security (IPSec) Protocol is an Internet Engineering Task Force (IETF) standard that defines an approach to providing confidential and authenticated information exchange over an untrusted network. The IPSec Protocol specification defines overall packet structure (tunneling), encryption methodology (message confidentiality) and data authentication (message integrity). An IPSec VPN (Virtual Private Network) uses the IPSec protocol and applies the defined encryption and packet-tunneling (i.e. an IP header embedded in another IP header) specification to provide a logical (virtual private) communication overlay to a physical network. IPSec VPNs ensure the data is confidential so, even if packets are tapped or delivered to the wrong destination, only the intended recipient can decrypt it. IPSec authentication also provides message integrity, ensuring the data was not manipulated or altered along the network path.

While a physical network is constructed of network devices interconnected by routers, where the router determines how packets should be forwarded to the ultimate destination, a site-to-site IPSec VPN can be constructed of sites interconnected by VPN gateways, where the VPN gateway determines how VPN (encrypted and authenticated) packets should be forwarded.

VPN tunnels are set up for the exclusive transport of the private data between pairs of gateways. An organization can forward traffic through these tunnels with internal IP addresses because the private data of the overlay network never mingles with the public base network. This enables enterprises to connect remote offices and remote workers without having to reconfigure their IP addresses or pay for the globally-unique IP addresses that are required for Internet transmission.



Why a Standard IPSec VPN is good, but not sufficient

Enterprises are finding IPSec VPN a viable alternative to legacy leased-lines or circuit-switched networks. These solutions leverage the cost-effective and ubiquitous connectivity of the Internet to transport data in a private manner, using the IPSec protocol to encrypt the traffic to provide data confidentiality and data integrity. Up until this point, however, vendors have focused almost solely on the cost and security benefits of IPSec VPNs. While focusing on speeding up encryption to provide reasonable throughput expectations for high speed IPSec VPNs, they have done little to ensure overall connection availability and performance of these solutions. This paper details the new requirements for IPSec VPNs that mark an evolution from just a secure, cost effective option for private networks to one that can offer:

- Scalable Connectivity Through Dynamic Routing - The ability to easily route communications between sites and accommodate network changes and growth with minimal manual intervention.
- High Performance - The ability to set-up and maintain the amount of concurrent sessions required by an organization without adversely affecting the overall performance of the network.
- Network Resiliency - The ability to quickly recover from any type of failure and continue to forward traffic through the VPN tunnels.
- Integrated Security -The ability to provide confidential data transport and data integrity, as well as access control to determine who has access to which resources, in a device that can easily integrate into the network.

Simple, Scalable VPN Connectivity

A VPN exists as a logical overlay to the network transport infrastructure, where first the underlying connectivity needs to be “defined and facilitated” and then the VPN infrastructure needs to be “defined and facilitated.” The way the VPN overlay is implemented will vary from vendor to vendor. A VPN that is difficult to deploy and/or difficult to manage will result in diminishing returns for the VPN, even though there is cost savings over the network infrastructure cost, and will ultimately impede the overall effectiveness of the solution (VPN overlay and the underlying network transport). The VPN must:

- Be easy to manage for all types of network configurations and administratively scale
- Automatically learn and incorporate network topology changes
- Minimize the need for human resources
- Leverage the dynamic nature of the network to increase connectivity

This paper will briefly look at three different approaches and how each approach affects the overall connectivity of the network and how easy it is to set-up and maintain that connectivity. The approaches that will be discussed include:

Rule-based VPNs: work by defining the network topology (IP addresses), then dictating, based on that topology, who can talk to whom over the VPN. While this approach does simplify some VPN deployments, the coupling of the network topology (IP addresses) with the transport of the VPN connection means the network cannot easily accommodate changes or the complex needs of widely distributed networks.

Route-based VPNs: work by separating the physical network from the abstract VPN network to simplify deployment and management. With route-based VPNs, organizations define the VPN overlay links and then define the network routes that will be used for transport. This allows the network route, rather than a policy, determine which traffic goes through the VPN. This provides some flexibility over a rule-based approach, but still requires manual changes to the route tables any time networks are added, deleted or changed.

Dynamic route-based VPNs: work by separating the physical network from the abstract VPN network and enables dynamic routing protocols through the VPN tunnels to completely separate the transport, forwarding decision from the VPN connection, giving enterprises not only simplified deployment and management, but also the flexibility they need to efficiently manage the constant changes inherent in complex networks.

NOTE For more details on each approach, please see Jupiter's Companion Note "How Different Approaches Affect Management and Availability: Comparison of Rule-based, Route-based and Dynamic Route-based VPNs."

Rule-Based VPN Approach

A Rule-Based VPN uses a VPN policy that determines the VPN communication overlay. The network topology is an integral component to the definition of the VPN policy and, therefore, has a negative impact on the resiliency and scalability of the VPN. The reason is that Rule-Based VPNs tie specific traffic and services, or source and destination groups, to one particular IPSec VPN connection, essentially binding that VPN connection to a fixed network route. This means that changes to the network require changes to the policy and on each and every VPN gateway affected by the network change. Often the people responsible for the network are not the same people responsible for its security, so changes to the network, such as adding a server or changing an IP address, may go unknown until someone tries to send traffic and can't because it has not been added on the VPN gateway or in the corresponding policy. In addition, if something happens to the Internet connection, or one of the VPN gateways goes down, the organization has to manually figure out what happened and then manually make a change to get the VPN up and running again. With Rule-Based VPNs, there is a lot of time and effort needed to configure and manage the connectivity, with the effort growing exponentially with the complexity of the network. As a result, rule-based solutions may be fine for small deployments, but don't scale to meet the site-to-site connectivity requirements of large distributed networks.

For example, what happens if something within the Internet goes down or a network connection becomes unavailable? With Rule-Based VPNs, if something happens to a particular connection, the VPN connection also goes down and the enterprise suffers lost connectivity and productivity. The reason is because the network topology is defined within the policy, tying the VPN connection to a static route. For example, traffic from a Tokyo network to a London network gets “routed” to the London VPN (peer), as dictated in the rule. If something happens and that route is not available, the VPN goes down. Even if other network routes are available, the VPN does not know how to use those alternative network routes. For instance, if Tokyo could potentially reach London by going through a New York VPN, it will not be able to do so until the administrator redefines the VPN peer (static route) in the rulebase to reroute the traffic through New York. As a result, site-to-site connectivity is lost until either the problem fixes itself or an administrator figures out what is wrong and makes a change to the rulebase. This makes Rule-Based VPNs very difficult to configure and maintain for large, complex networks.

Route-Based VPN Approach

Route-Based VPNs overcome some of the limitations of Rule-Based VPNs because they separate the physical network from the logical VPN network and allow routing, versus a policy, to determine what gets transported through the VPN tunnel. Route-Based VPNs create logical VPN tunnels between destinations to establish the private network overlay and then the route table determines how the traffic gets there. While this approach simplifies deployment, eliminating the need to define the network topology with firewall rule sets, it requires manual route statements be entered into a route table for each gateway, which can be time consuming and tedious for an administrator to create and maintain. For example, if a network route changes, it needs to be manually updated in the route table for the VPN and its peer gateway. This can pose a lot of work for an administrator of a large distributed network. Another problem is that the administrator responsible for the network may not be the same as the administrator responsible for the VPN, so routing changes may not be made in the VPN in a timely manner; only getting called to the administrators attention when the VPN goes down.

Route-based VPNs do offer a level of flexibility and resiliency over the Rule-Based approach, allowing multiple static routes to be defined between gateways, instead of just one. A “cost” is associated with each route, in relation to the directness of the connection. The “cheapest” route will always be the route of choice for the VPN. A mechanism identifies when a route becomes unavailable, and then removes that route from the routing table, so the VPN will use the next (lowest cost) tunnel to transport the VPN traffic. This will enable the solution to maintain the connectivity of the VPN in the event of a failure, however, an administrator needs to manually ensure that there is always an available route defined in the route table. If a route becomes unavailable and no alternative route has been defined or is available, the site-to-site connectivity will be lost until an administrator is able to manually identify a new viable route and add it to the route table.

Dynamic Route-Based VPN Approach

Juniper introduced the concept of Dynamic Route-Based VPNs, within their integrated firewall/IPSec VPN product line, to the market to meet the requirements of a truly scalable solution. Dynamic Route-Based VPNs separate the physical network from the logical VPN network and allow routing, versus a policy, to automate the transport decisions, just as route-based VPNs. The difference is that Dynamic Route-Based VPNs are able to leverage dynamic routing, which means they can automatically learn the network topology and available routes to maintain the connectivity of the network.

Instead of defining the network topology or manually building route tables, an administrator will enable dynamic routing on each of the gateways. Juniper's Dynamic VPNs treat each VPN link/connection as a dynamic, routable links. This is what enables the solution to run dynamic routing protocols through the VPN tunnels, just as traditional frame relay links did, giving organizations the same ease of use they had with legacy private line and router solutions.

Once dynamic routing is enabled, Juniper's Dynamic VPNs leverage the protocol to automatically learn the network topology, saving organizations the time and resources required to define each and every machine on a network and iterating through the policy every time something is added or changed. New networks are accessible from any tunnel endpoint and are dynamically learned by the other endpoints via dynamic routing protocols. This also reduces the likelihood of mistakes due to human error.

More importantly, Juniper's Dynamic VPNs are able to automatically survive failures within the network to keep the connection available. If a tunnel is no longer a viable path for a route, or if it is removed from the routing table, a new route will be automatically learned. If there is a path open to get from point A to point B, Juniper's Dynamic VPNs will find it to ensure the communication persists. This makes the ongoing management and maintenance of the VPN possible without much human intervention at all, saving time and resources and providing a scalable solution that achieves the connectivity requirements of large, distributed organizations.

Scalability Conclusion

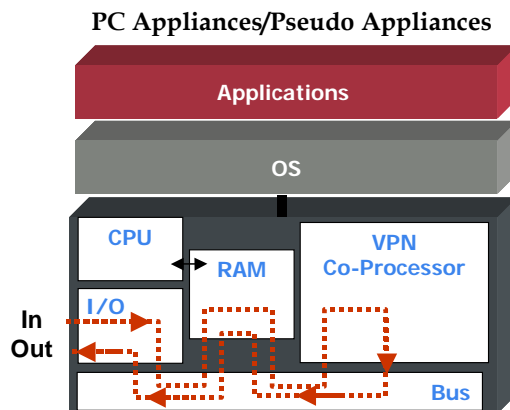
A Dynamic Route-based VPN provides significant management benefits over a Rule-based VPN or a Route-based VPN. We have seen how the topology of the network is integral to the VPN policy in a Rule-based VPN approach and, therefore, creates a significant management burden. When the network changes, the policy needs to be manually changed to reflect the new network. Because of this tight coupling with the network topology, Rule-based VPN can never achieve the level of resiliency described in the Network Resiliency section of this paper. A Route-based VPN decouples the network route from the policy and, therefore, is a slight improvement over the Rule-based VPN approach. However, because Route-based VPNs use static network routes as the foundation for the VPN links, they still require manual intervention, forcing the administrator to manually change the network route in the route table when a site changes. Since the static routes of Route-based VPNs are not compatible with dynamic routing protocols, they simply cannot achieve the level of resiliency described in the Network Resiliency section of this paper. Only a Dynamic Route-based VPN is designed to leverage dynamic routing to minimize the need for human intervention and easily scale to achieve the connectivity required by large site-to-site networks.

High Performance

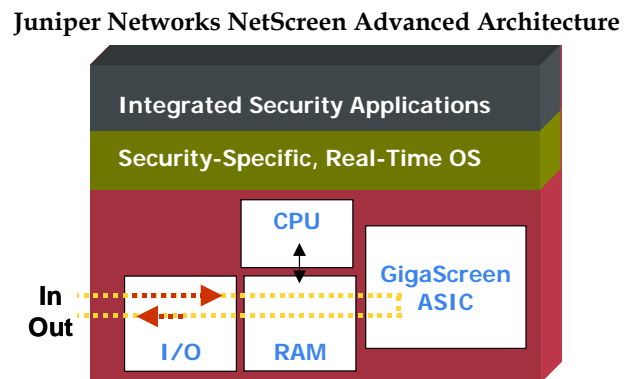
VPN performance is very important when selecting a VPN solution. Performance means more than pure throughput. To get a clear picture of how the VPN is going to be able to perform in real network environments, there are many elements to consider. Considerations are:

- Perform complete VPN functions at maximum speeds to support the number of connections that an organization needs, as well as the overall network throughput
- Offer high performance throughput, regardless of packet size, to meet network requirements
- Quickly establish VPN tunnels when they are needed
- Minimize latency as session loads grow across all applications
- Maximize connection availability by ensuring the solution is optimizing resources and able to maintain existing VPN tunnels, while also adding new ones, to ensure they are ready for use

The first thing to understand is how the solution has been designed to provide VPN functionality. Some solutions are delivered as software on a general-purpose platform or installed on networking devices, such as routers and switches. The design of these solutions can introduce processing latencies and hardware and software interoperability issues that can affect predictability and performance. Juniper’s purpose-built integrated firewall/IPSec VPN product line offers performance advantages over solutions that have cobbled security functionality onto a general-purpose operating system and/or platform because Juniper’s integrated firewall/IPSec VPN devices have hardware and software that have been specifically conceived and engineered for the security functions that they provide.



- General Purpose Processing**
- Data must traverse several unique, non-optimized interfaces
 - Each “API” introduces security risk, interpretation and vendor dependency
 - Processing delay may cause “unpredictable behavior”
 - Difficult to optimize data path



- Security Specific Processing**
- Streamlined, linear packet processing
 - Each processing component is optimized
 - Applications and hardware optimized for security processing and performance

Juniper has built its NetScreen appliances and systems from the ground up to optimize its security functions and provide predictable VPN performance. As a result, Juniper can both accelerate the encryption AND ensure that VPN gateways are able to establish and maintain tunnels, with low latency and high throughput.

NOTE For more detail on an example of how Juniper's purpose-built firewall/IPSec VPN product line improves both performance and functionality delivery, please see the "Juniper Networks NetScreen-5000 System Series Architecture" white paper.

Throughput

VPN throughput is measured as the end-to-end capability of taking clear text data, encrypting, transmitting and then decrypting the data. This is usually measured in packets or megabits per second. While throughput is only one of the metrics for VPN performance, it has often been used as the only metric. A more critical success factor for solutions is in its ability to process traffic of varying types and packet sizes, under heavy load conditions, without affecting network performance. The Juniper solution is designed to provide consistent and reliable throughput across a wide range of traffic and load conditions, by using purpose-built hardware platforms that optimize all aspects of data and packet processing and encryption/decryption to achieve very fast session ramp rate and enhanced small packet performance.

VPN Tunnel Setup

VPN tunnel setup (sometime called tunnel establishment) is the mechanism and time associated with encryption key negotiation (IKE or manual) and tunnel setup. For every VPN tunnel used, there are potentially several sets of encryption keys that need to be generated and maintained, and each affects performance. Tunnel setup is critical to the success of a central site device that may be terminating hundreds to thousands of VPN tunnels. Juniper's purpose-built firewall/IPSec VPN hardware and ScreenOS combination accelerates IKE (Internet Key Exchange) key generation to improve overall performance to enable the system to quickly establish the VPN security association. Accelerating IKE key generation also provides the less obvious benefit of being able to handle much of the background packet processing, freeing up the CPU to handle the other parts of the VPN tunnel setup. In addition, ScreenOS is designed specifically for these types of tasks, so it doesn't have the general overhead of a general-purpose operating system, resulting in more efficient processing.

Low Latency

VPN latency represents the average delay introduced for a packet to go from point A to point B. Latency is an important performance metric because real-time applications (such as Voice Over IP - VoIP) will not perform well/properly in high latency (long average one way or roundtrip packet traversal). Low latency across packets sizes (from 64-byte to 1518-byte packets) is critical because a system that runs well in one environment/application may not be well suited for another. Juniper's solution uses hardware acceleration to provide minimal impact to overall latency. One significant factor is Juniper's simple packet flow, where packets are quickly processed without unnecessary traversals of PCI busses, which can be a common problem with PC-based platforms using VPN acceleration cards.

Optimized Resource Management

The VPN should be able to optimize processing resources to ensure that it is able to keep existing VPN tunnels running, while simultaneously adding new ones quickly. Maintaining the VPN tunnels is probably the trickiest aspect of VPN performance and where many solutions break down, since they aren't able to process traffic, generate new tunnels and maintain the existing tunnels in a timely manner. This is one of Juniper's strengths, with many customer deployments that contain more than 1,000 nodes in the VPN. The ability to maintain tunnels effectively is due to the same reasons that the Juniper solutions are able to establish tunnels quickly and effectively – the CPU is not as involved in general packet flow, due to the purpose-built hardware's ability to handle the mainline data path and take on a lot of the "heavy lifting," while the custom designed ScreenOS can efficiently handle multiple tasks concurrently.

Performance Conclusion

The solution needs to be able to match the performance requirements of the network. In order to deliver high performance throughput, the solution must quickly perform tunnel set-up, minimize overall latency of packet delivery and ensure that existing tunnels can be maintained, while new ones are added, to ensure VPN availability. Juniper has proven VPN products that are not only able to perform when doing simple VPN throughput measurements, but are also able to perform in the key functions that make a VPN deployment successful and transparent to the user.

Connection Resiliency

Connection resiliency is defined as a solution's ability to provide "always on" availability, by surviving failures at any level. From a VPN perspective, it needs to provide several layers of built in resiliency to deliver a truly fault tolerant solution, which goes beyond simply providing device or network resiliency. Redundancy at the device, network and the VPN level need to be incorporated, because without even just one of the components the VPN's ability to provide a low-latency switch-over is diminished. The following items will be discussed in this section:

- Solid state operation with component redundancy as needed
- Device redundancy that is sharing session and VPN state information
- VPN State synchronization to eliminate need to re-establish VPN sessions during switchover
- Physical path redundancy that provides different connection alternatives
- Dynamic routing support that eliminates need for manual network route definitions
- VPN monitoring to detect VPN connection failure and to trigger a switch-over

Solid State Operation with Component Redundancy

At the lowest level, providing a high degree of connection resiliency means being able to avoid failures as much as possible. The first consideration to achieving the most reliable implementation possible is to select an appliance-based, solid-state solution with the highest Mean Time Between Failure (MTBF) figures. The second consideration is to have redundancy within that appliance for any high stress or mechanical component. These components usually consist of cooling fans and power supplies. The third consideration is that these components should be hot-swappable, so that the system does not need to be turned off before a component can be replaced. The last consideration is to provide sufficient interface densities that enable the solution to offer the physical path/link redundancy that will be discussed shortly. Juniper's purpose-built solutions offer a solid state foundation, with redundant, hot-swappable components and varying interface densities to meet the redundancy needs of different network segments.

Device Redundancy- High Availability

At the next level, device redundancy is needed to reduce the impact of a device failure, eliminating the "single point of device failure," which is often called a high availability configuration. There are two approaches to deploying redundant devices; one is a redundant pair approach, and the other is a clustered approach. In the redundant pair approach, one device is designated and operated as the "master," while the second device in the pair is the "slave." All the master information is replicated to the slave, so the slave can take over if the master device fails. In the cluster approach, any device in the cluster is operating as a peer. All information is replicated between all devices in the cluster, so that any device can be processing traffic. If one device in the cluster fails, then the other devices will take up the additional "load." Redundancy at the device level enables the network to survive a device failure. However, to ensure that connections are not lost and the VPN does not have to be reestablished, these devices need to be sharing state and VPN information.

VPN State Synchronization

Many VPN vendors offer Stateful High Availability, which means that the redundant VPN devices are tracking the state of the communication session, synchronizing the session information to maintain the session if a device fails. VPN State Synchronization is about maintaining and synchronizing the state of the VPN between devices. Performing Stateful HA does not ensure the VPN states are tracked or synchronized – it must be done so explicitly.

If a solution does not provide VPN State synchronization, when a failure occurs, the VPN Security Association (SA) between the two VPN gateways is lost and needs to be re-established. Establishment of VPN SA (generating, negotiating and exchanging encryption keys) is computational and time intensive and can take anywhere between 30 seconds to a minute. As mentioned above, Juniper has an advantage because it can accelerate the computations by using purpose-built hardware. More importantly, however, once established, the SAs should be able to fail-over, as opposed to reconnecting, in the event of a failure. If the solution cannot support the fail-over of the SAs, the overall downtime that results from a failure can exponentially increase based on the number of VPN connections that must simultaneously be re-established. If the failure is at the center of a large VPN network, the resulting downtime could be significant. The solution should be able to

automatically associate the VPN with an available path. When a VPN connection fails, all the sessions and Security Associations associated with that failed VPN connection must be immediately activated on another link to provide maximum VPN resiliency. To provide the highest degree of resiliency and to ensure the lowest possible latency during switch-over, due to a link failure, the solution needs to be able to mirror the SA and automatically associate it with the active connection. Juniper's integrated firewall/IPSec VPNs mirror all VPN security associations, including the timers, keys and certificates, so that, in the event of a failure, VPN tunnels and the sessions that run through them can be re-associated immediately with either the interface or route as quickly as possible.

Redundant Physical Paths

To minimize the impact of a link failure, a solution must support multiple physical connections, or paths, at the same time. The use of parallel physical paths between two sites reduces the dependency on any one mode of transport or network provider and offers persistent connectivity during the failure of one path. There are 3 types of redundant path configurations – dedicated redundant link with HA trigger, dedicated redundant link with on-demand trigger and non-dedicated on-demand trigger (dial backup). Based on cost and resiliency requirements, the organization can determine the best transport alternatives.

- **Dedicated physical redundant path with High Availability fail-over**

The highest degree of path redundancy, where a physical path is connected and the fail-over activation is controlled by a redundancy protocol that determines when and how the link should failover (specifies how sessions and other relevant information is replicated between the devices). When the primary link goes down, a connection is initiated to maintain communications.

- **Dedicated physical redundant path with on-demand fail-over**


An intermediate level of path redundancy, where a physical path is connected and the fail-over activation is controlled by the primary link going down. When the primary link goes down, a connection is established and communication resumes.

- **Non-dedicated physical redundant path with on-demand fail-over**

A minimum level of path redundancy, where the device is connected to a modem or connected when needed, and the fail-over activation is controlled by the primary link going down. When the primary link goes down, a connection is established and communication resumes.

For example, a remote office may be able to live with one-two minutes of communication downtime, but cannot afford to wait for someone to physically come out to reestablish a connection. A cost-effective option may be to use Non-dedicated, On-Demand fail-over where the redundancy is provided using a dial backup/modem/telephone line combination. A large satellite office may require a higher degree of fail-over resiliency, so a pre-established redundant connection over a dedicated DSL might be chosen as the alternate path.

Back-up physical paths enable the solution to persist in the event that a transport mechanism goes down



High Availability – Full Mesh

To ensure that the VPN connection can persist, whether a path or device goes down, the solution needs to combine support for both redundant paths and redundant gateways that are capable of sharing state and VPN information. The ideal high availability configuration is a full mesh, which provides redundancy at every level, so that a connection can survive a single failure anywhere along the path or multiple failures at different levels (i.e. one service provider and one security gateway). Redundant interfaces are one way to achieve full mesh connectivity, enabling the connection to survive an interface or cable failure. A full mesh configuration that incorporates separate ISPs connections, however, can only be achieved using dynamic routing.

Dynamic Routing

In the physical network, dynamic routing can identify when there is a failure in network route, such as an ISP provider losing connectivity, and, if available, choose an alternate path. Supporting dynamic routing through the VPN tunnels will enable the VPN to choose an alternate VPN tunnel when a connection goes down.

Unfortunately, most VPN solutions, both rule-based and static route-based, rely on static routes, which define the route for the VPN traffic between two VPN gateways. For these solutions, if a route goes down, manual intervention is required to change the policy or route table to get around the trouble spots. This results in down time and can end up costing an organization in lost productivity.

Juniper's Dynamic Route-Based VPNs are able to run dynamic routing protocols through the tunnels to automatically make changes to the VPN routes when a failure is identified. The VPN is network aware, which means it can make "best path" forwarding decisions for individual traffic flows. When dynamic routing determines that there is a failure, it will dynamically reroute around that failure to keep the VPN connection flowing. And once a connection is back up, the VPN will learn about it and automatically start using it again, if it is the best path. Dynamic routing support provides a resilient VPN solution that reduces manual intervention.

VPN Path Monitoring

To achieve the fastest VPN fail-over time possible a mechanism is needed to monitor the redundant tunnels to trigger a fail-over in the event a link goes down. Juniper introduced VPN Path monitoring, so regardless whether the solution is using static network routes or dynamic network routes, the fail-over time can be controlled. If the solution uses static network routes, then the VPN path monitor is the only way the solution will know the primary link has failed and the backup link needs to be activated. If the solution uses dynamic network routes, relying on dynamic routing to reroute VPN connections may not deliver adequate fail-over. For example, OSPF and BGP can take up to a minute to learn of a failure. Plus, there are situations where dynamic routing cannot be supported. VPN Path Monitoring provides the ability to set the fail-over interval to occur within seconds.

VPN Path Monitoring will monitor the available tunnels and identify a failure. It does this by periodically sending heartbeat messages through each of the VPN tunnels to a device on the other end of the VPN link. When a heartbeat fails, that VPN link is marked down and an alternate path is selected. To the route table, the interface looks the same as a physical interface would if link connectivity was lost, so the route automatically becomes inactive. If dynamic routing is enabled, dynamic routing algorithms will automatically find the best alternate route. If relying on static routes, the alternate static route (if defined) with the next best metric would be used. If an alternate static route is not available, then the administrator will need to manually add one to the table.

Resiliency Conclusion, VPN Can be More resilient than legacy WAN

A Dynamic Virtual Private Network can be more resilient than a legacy Wide Area Network. Using a combination of all the capabilities described in this section (network resiliency), organizations can take advantage of the appropriate capabilities to create a highly available site-to-site Virtual Private Network that meets their resiliency and cost requirements to replace a legacy WAN.

The VPN device, itself, should have redundancy components. Central and regional sites can have redundant VPN gateways to achieve a high availability configuration, with the devices sharing both session and VPN state information to maintain the VPN in the event of a failure. In addition, the central sites can be setup in high availability configurations (active/passive, active/active, active/active full mesh) to ensure there are multiple VPN connections that can be used to send VPN traffic between sites. The remote offices can have redundant connections using either the same high availability configuration as the central office, or have on-demand dedicated or dial-up connection, either to an alternate ISP or dial-up, so that if something happens to the regional site to which their primary VPN is connected, they can still reach the rest of the network. The solution will use VPN Path monitor to achieve the desired fail-over latency at any of these locations. Dynamic routing protocol support will make all of this possible, enabling the solution to survive failures in the public, physical (Internet) network and in the private, overlay (VPN) network.

Integrated Security

While IPSec VPNs provide confidential connectivity, they do not provide the additional security that an organization needs to ensure access control, authorization, etc. For this, an organization needs to deploy a firewall. Firewalls provide granular security checking, using Stateful Inspection at layers 2 through 4, for access control. Plus, firewalls offer authentication to ensure that users are who they say they are. However, it is important that organizations be able to maintain a conceptual level for their security, so they are not forced to constantly identify specific IP addresses and lose the learning, "self-healing" benefits of dynamic route-based VPNs. In order to facilitate this, the firewall and VPN must be integrated in a way so that the traffic source and destination addresses can be logically separated from the policy itself.

The requirements for integrated, secure networking functionality are:


- Robust firewall and access control capabilities to ensure only authorized users are entering the network and accessing only allowed services on appropriate resources
- Simplicity of Dynamic Route-Based VPNs and a policy-based firewall managed from a single user interface
- Zone-based access control that does not require the use of IP addresses
- Linked VPN and firewall security functionality to ensure no security compromises
- Decoupling access control policies from traffic forwarding decisions

Many vendors simply integrate the VPN policy into the firewall policy, which may adequately protect the private data network and its resources, but does not offer a flexible, easy to manage or scalable solution. This approach dictates that the VPN be rule-based, which means that dynamic routing cannot be used and all the previously discussed management and resiliency constraints associated with a Rule-Based VPN apply.

This approach does simplify the integration of the two functionalities, since the access control through the VPN is specified in the same rule. But this simplification makes it difficult to accommodate varied requirements for large, complex networks. When the rule says “Encrypt,” it implies “Permit,” with “Authenticate” optional. This approach doesn’t give organizations the ability to send the same type of traffic to the same destination through one of two tunnels. Because an organization needs to manage the VPN within the firewall rulebase, it introduces unnecessary constraints that end up dictating how the device can be deployed and costing organizations time and money in ongoing management. Removing the access control function from the actual VPN definition allows greater flexibility in policy enforcement and enables the network to be more dynamic. There is a better way.

Zone-Based Approach Allowing Firewall Control in a Dynamic VPN

Juniper has introduced a zone-based approach to network security to enable organizations to apply the necessary security to VPN traffic, without losing the benefits of dynamic routing. A security zone represents a collection of one or more network segments to which inbound and outbound traffic policies can be applied. Examples of typical zones would be trusted, untrusted and DMZ zones, however, custom zones can be created to meet the unique needs of an organization. Through the use of policies, the traffic flow from zone to zone can be controlled, defining the kinds of traffic that is permitted to pass from specified sources to specified destinations at scheduled times. This means that there is no need for additional devices to be added to enforce service controls. Traffic of different VPNs can be brought into different custom zones, allowing the organization to set different policies for different VPNs.



At the broadest level, an organization can create a zone that allows all kinds of traffic from any source in a zone to any destination in all other zones without restriction. The narrowest application would be a policy that allows only one kind of traffic between a specific host in one zone and another specified host in another zone during a specific period of time. The key is to have a solution that enables the organization to control the level of security that is applied to the traffic that arrives or exits from a particular VPN tunnel or set of tunnels.

This model enables the device to separate the network into areas or zones that are protected from each other. The zones can encompass one or more physical or logical interfaces, including VPN tunnel interfaces. This means that IP addresses are not required by the firewall to differentiate the traffic, so the solution does not lose the efficiencies of using dynamic routing to automatically learn the topology with route-based VPNs. Instead the firewall uses the interfaces that are defined in each zone to determine what to do with the traffic.

A security policy is created for each zone, and the firewall applies the policies between pairs of security zones to control the type of traffic that is permitted or denied passage between the zones. When a new interface is added to a security zone, the policies that are in effect for that zone are automatically enforced on the new interface. This makes the integration of new sites into the VPN easy and quick, since all the organization needs to do is add the interface for the new site to an existing security zone.

Juniper also decouples the security policies that control what type of traffic goes in or out of a zone from traffic forwarding specifications. This allows a security policy to be dynamically associated with a VPN configuration. For example, a security policy can be defined that allows all traffic from the "HQ" zone to the "Remote Sites" zone. The routing table in the solution is able to determine which of several candidate tunnel interfaces is to be used to reach a given remote site. The tunnel interface status is reflected in the routing table, which permits automatic VPN link selection, while maintaining consistent policy enforcement. In this way, organizations benefit from the simplicity of both a dynamic route-based VPN and a policy-based firewall and the integrated security for data privacy, access control and authentication.

Conclusion: Juniper Provides Scalable, Resilient, High Performance VPN

Juniper is the only vendor that provides a scalable, resilient, high performance Virtual Private Network that has integrated Firewall capabilities. The Juniper firewall/IPSec VPN product line leverages dynamic routing to provide a solution that is simple to deploy and maintain. After defining the gateways and enabling dynamic routing, the solution automatically learns the network topology and identifies the best route for the VPN tunnel. When a new network is added, the dynamic VPN automatically learns its topology and reconfigures all appropriate existing gateways, so that the organization spends minimal time and resources worrying about making sure that everything is configured correctly

Juniper's Dynamic VPNs provide the performance and system level resiliency that organizations require for their site-to-site networks. Juniper is the first to offer redundancy at the path, device and VPN tunnel level, with state and VPN sync, full-mesh deployments and its unique VPN monitoring, to ensure connectivity persists in the event of a failure.

Juniper built its integrated firewall/IPSec VPN product line from the ground up to tightly integrate VPN, firewall and networking functionality, so that the security device can interact with complex networking environments and facilitate the communication and data flow it is designed to protect. Juniper's unique security zone-based policy application enables organizations to add the simplicity of a policy-based approach to its dynamic VPN, while retaining the network fluidity of dynamic routing. Juniper also decouples the security policies from traffic forwarding specifications to enable the policy to be dynamically associated with a VPN configuration.

The Juniper Dynamic VPN solution provides a powerful communication infrastructure that is secure, easy to manage and cost-effective. Take the first step to replacing your traditional private network solutions today. To learn more about the Juniper VPN solution, go to <http://www.Juniper.com>.