



Hybrid VPN whitepaper


"Optimizing Performance across the Enterprise IP VPN."

Authors

Eric Torres – Vice President for VNO Services

Johann Clamen – Product Manager for VNO Services

www.tatacommunications.com



Tata Communications
"Hybrid VPNs: Optimizing Performance across the Enterprise IP VPN"

Table of contents

Executive Summary	4
Introduction	5
Current “Traditional” Network Services	6
Multiprotocol Label Switching (MPLS)	6
Network Considerations	7
QOS in an MPLS VPN	8
IPSec	9
Network Considerations	9
Quality of Service	13
SSL	13
Network Considerations	14
Security Considerations	16
Hybrid IP VPN	17
Network Considerations	17
Transport Network	18
Edge Design	19
Backup and Recovery	19
Controlling the Network Architecture	20
Conclusion	21
Tata Communications’ Hybrid VPN service	22
Acronym Key	23

Executive Summary

Today's Enterprise networking environment is characterized by international coverage, a single end-to-end network protocol and a wide variety of network topologies, network carriers and access services. The emergence of a so-called "Hybrid Virtual Private Network" (HVPN) is the result of the efforts of organizations to integrate different network components in order to fulfill their need for greater reach, higher performance and lower costs. By matching VPN technology to site requirements, Hybrid VPNs provide a tailored solution with more extensive coverage and lower costs. However, Hybrid VPNs can also carry greater management challenges, making it crucial to develop a comprehensive network strategy and select the right provider to execute. A single-provider Hybrid VPN solution can provide the best combination of cost efficiency and ease of management.

Introduction

Enterprises look to Virtual Private Networks (or VPNs) to provide the same capabilities that private leased lines traditionally delivered, but with cost savings enabled by using a shared network. VPNs create a private network that provides the same level of security assurances as dedicated connectivity. A VPN prevents unauthorized users from penetrating the network, securely carries a company's confidential data over shared or even unsafe (e.g., Internet) networks and ensures that sensitive data will not be intercepted by eavesdroppers.

While most applications do not “remain satisfied” by their network experience, IT professionals are facing new challenges every day to deliver adequate levels of service. Contract management, service level assessment and problem determinations are affected by increasing complexity due to the great number of service elements.

As Hybrid VPNs have come to the forefront, three major concerns have emerged:

1. The plethora of offers, created by national Access Network Providers (of DSL, cable, Ethernet or fiber), with country-specific service levels and contracts often managed by local teams, do not help an organization to optimize its cost/performance ratio. Contract centralization and management have become strategic investment priorities.
2. New networking demands cannot always be satisfied by a simple extension of an existing network solution. The key to an efficient network architecture lies in a clear identification of the requirements and a design that best meets the needs.
3. Once implemented, these solutions should be continuously challenged by a process to ensure that the architecture is maintained and implemented under the best technical and operational controls.

Traditionally, three technically disparate approaches have been used to create a global, IP-based VPN:

Network based – VPNs are created and managed through the service provider network (example: MPLS-based VPN).

CPE based – VPNs are created at the customer edge router; there is no state information about the VPN inside the service provider network (example: IPSec-based VPN).

Application based – VPN or application-specific tunnels are created from IP host to IP host and transparently cross both the customer's equipment and the service provider's infrastructure (example: SSL).

Increasingly, businesses looking to increase productivity and reduce costs have chosen to implement a global solution that leverages a mix of these three VPN strategies – an architecture commonly referred to as a Hybrid VPN. Hybrid VPNs have several crucial advantages, but chief among these is the flexibility to respond to a constantly evolving set of network transport infrastructure requirements and access providers. Hybrid VPNs can also provide solutions to requirements such as Internet access, inter-enterprise communication and connection of switched or mobile users.

Current “Traditional” Network Services

This section will describe the strengths and weaknesses of the three traditional VPN approaches and highlight particular applications and requirements that tend to favor one technology over another.

Multiprotocol Label Switching (MPLS)

Summary

MPLS has become the leading technology for Wide Area Networks. It integrates the advantages of Internet protocol – including addressing, dynamic routing and network meshing – while addressing some of the well-known drawbacks of the Internet with support for private IP addresses, routing using customer-defined routing tables and differentiated transport for various types of data.

Benefits

For reliable performance, MPLS traffic is dynamically routed from any point in the network to any other point, even around trouble spots, heavy loads and congestion. MPLS customers can maintain their own IP addressing plan that is independent from the service provider's other customers. MPLS also provides CoS capabilities and QoS guarantees. By defining a class of service per application to enable appropriate quality of service, MPLS ensures that performance levels map to the traffic's specific characteristics. QoS and performance features are contractually ensured by the backbone provider through SLAs.

Drawbacks and limitations

With a few exceptions, the end-to-end service offered by MPLS providers is generally more expensive than CPE-based VPNs because of the high-quality components required to commit a service level to customers. Furthermore, all network endpoints must be connected to the service provider's network by a high-quality, dedicated link, limiting the feasibility of MPLS-VPN services for individual customers or sites located in hard-to-reach geographies.

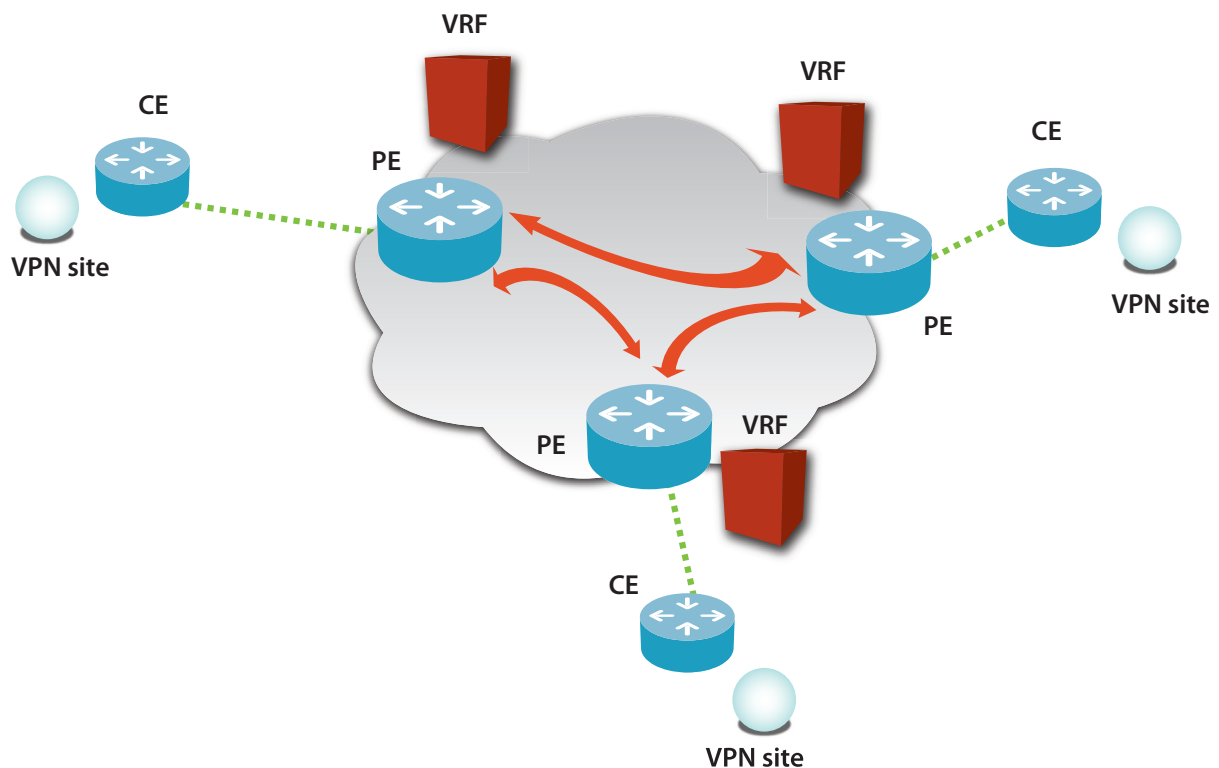
Network Considerations

In an MPLS network, incoming packets are assigned a label (a short, fixed-length identifier) that stores the packet's destination by the Provider Edge (PE) router. The packet label helps switches at each hop to make the decisions necessary to successfully forward the packets through the network.

The Customer Edge (CE) router's role is to route the IP traffic from the customer's site to the operator's MPLS backbone. Each CE builds an internal routing table based on knowledge of the IP networks to which it is linked. MPLS VPN-specific routing tables (VRFs) and CE routing tables are exchanged through the CE/PE link. Since the PEs are permanently communicating with each other across the network, the VRF tables are constantly enhanced, maintained and updated with every known and available route to every part of the VPN. Each CE, therefore, has a dynamic and up-to-date view of the network.

For the customer's dynamic routing protocol, the MPLS VPN appears as an administratively separate transit area (known as an Autonomous System, or AS). A peering process has to be defined to exchange the IP routes, protecting each peer from an uncontrolled routing flood.

Figure 1 – MPLS network



QOS in an MPLS VPN

The MPLS quality of service design pushes all classification work to the edge of the network, so that functions such as classification, policing, congestion avoidance and congestion management are enabled on the egress interface of the CE router. This is a logical approach since the connection to the MPLS cloud is usually the weakest link in a site-to-site connection. MPLS service providers control congestion on the PE egress interface (the output link to the customer's site) by shaping the dataflow according to the bandwidth of the last mile circuit to the customer site.

For a customer-managed CE, the service provider can reassign packet classification to the PE ingress to avoid discrepancies between customer-marked traffic and committed rates per class of service. These classification values then dictate how each IP packet is handled across the network.

IPSec

Summary

IPSec VPNs use Internet protocol (IP) to transmit customer data across various operator transport networks. The confidentiality, integrity and authenticity of the data flow are ensured by authentication and/or encryption protocols; as a result, IPSec is often described as a security feature.

IPSec can be used to create secure VPNs over the Internet, or to add security features in an overlay of an MPLS VPN. When used over the Internet, IPSec provides point-to-point connections.

Benefits

Because IPSec creates a VPN connection through the Internet, it is a widely available and low-cost connectivity option. IPSec enables an extra level of protection through encryption and authentication and allows the customer to maintain its own IP addressing plan through tunnels between sites. IPSec is ideal for connecting isolated company sites as well as enabling temporary connections to a VPN (for an exhibition, as a temporary backup for recovery and so on).

Drawbacks and limitations

Although IPSec VPNs are implemented on the Internet, which is a meshed network, they inherit the complexity of layer 2 VPNs like Frame Relay or ATM. This is due to the point-to-point IPSec tunnels – when a new tunnel is added, customer premise equipment will need to be reconfigured in accordance with the changed logical network topology.

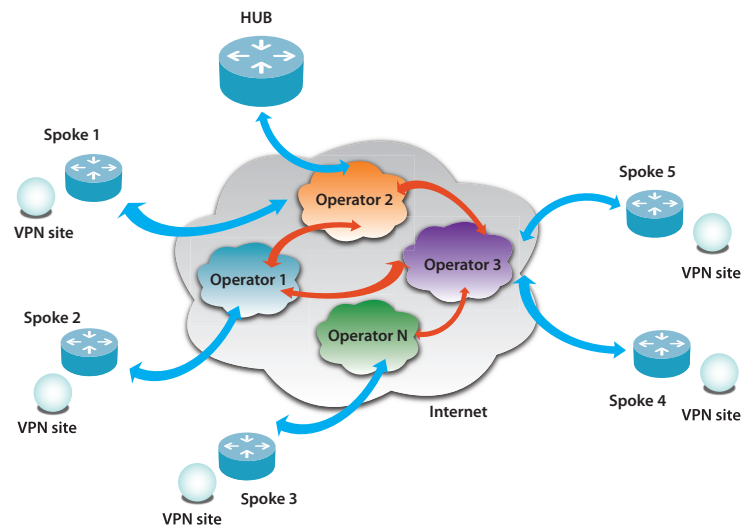
Additionally, packet characteristics may be hidden inside an IP tunnel, preventing the recognition of high-priority flows and making COS usage difficult and limited to the edge of the network.

Specific hardware is often needed in order for the hub to handle encryption/decryption processes in a timely manner, and for isolated workstations, IPSec requires a client software installation.

Network Considerations

IPSec is a layer 3 VPN technology (meaning it operates independent of the applications that may use it) and is usually implemented among routing equipment endpoints. Topologically, most of today's Enterprise IPSec VPNs incorporate hub-and-spoke designs (the simplest consisting of single-circuit, single-spoke connectivity to a hub router at a central facility). The path between two points of the network is determined both by point-to-point tunnels defined between the spokes and the hubs, and by the underlying networks' topologies and interconnections.

Figure 2 – IPSec network, underlying topology



Mobile workers need to install a software client to use IPSec, which then provides full access to network resources and applications and allows them to perform their business tasks remotely. Routing, which is within the public network, is completely transparent to the end user.

Network Topology Considerations


Choosing a network topology is the main difficulty of an IPSec VPN. There are two main basic types of VPN topologies: meshed and hierarchical. In a meshed architecture, each site is connected through an IPSec tunnel to every other site on the network, whereas in a hierarchical architecture, each site (spoke) is connected through an IPSec tunnel to a specific site on the network (the hub).

Another important consideration for an IPSec VPN is the alignment between the logical architecture and the underlying physical infrastructure. Compared to a hub-and-spoke topology with a hub located in London, a direct IPSec tunnel between South Africa and India will improve response times if the Service Provider has a direct physical route east to India.

This means that topologies become very complex. Any change needs to be made with a complete understanding of the global architecture and detailed knowledge of traffic patterns in normal and backup situations. As a result, the hidden costs of a simple action like adding or deleting a site are tremendous.

Full mesh networks

Figure 3 – IPSec full mesh network

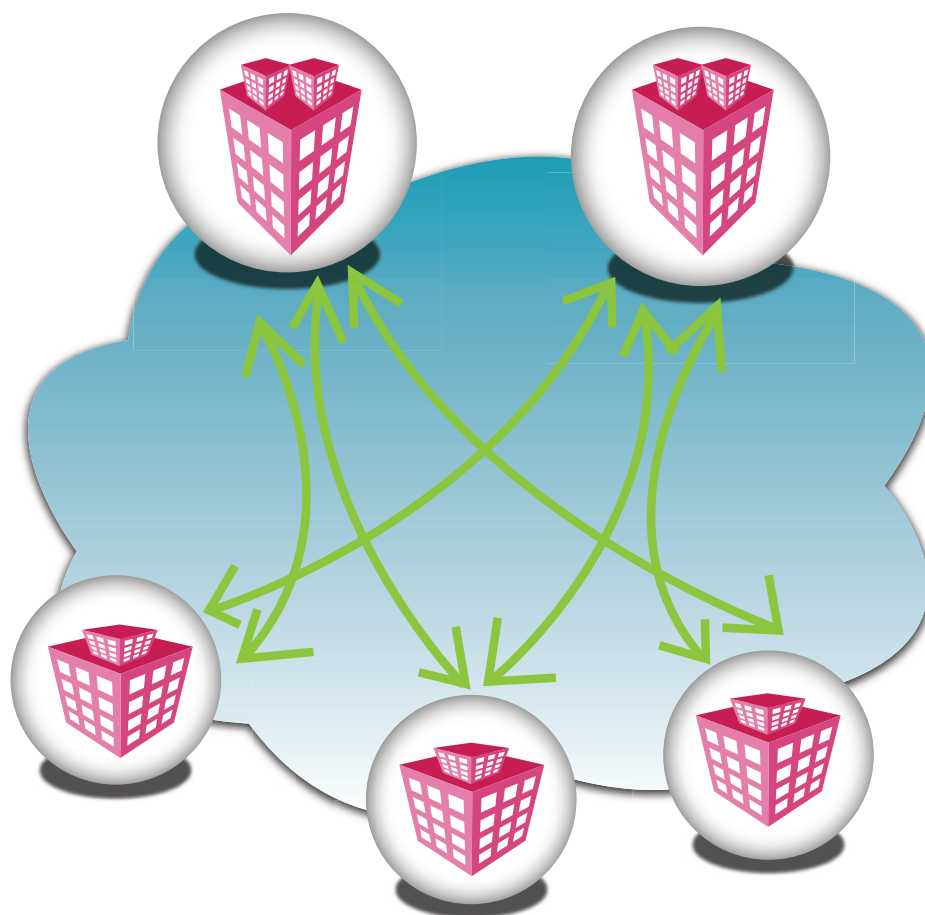


In a full mesh topology, all sites are directly connected to all other sites. At first glance, this arrangement seems to present the ideal solution in terms of efficiency and simplicity. However, the number of tunnels in a full meshed network increases steeply as the number of sites moves beyond a handful. If N is the number of sites, the number of active tunnels will be $(N \times N - 1) / 2$. For example, a 30-site network requires 435 links, and 29 logical connections have to be defined in each customer router. The difficulty of managing such configuration complexity quickly becomes prohibitive.

Additionally, when supporting multicast traffic, a full mesh will cause higher traffic loads on the WAN interfaces of the site supporting the multicast source, since the multicast packets will need to be duplicated for each IPSec tunnel. Network design best practices usually limit full meshed interconnections to four or five sites.

Hierarchical, or hub-and-spoke architecture

Figure 4 – IPSec hub-and-spoke network




In a hierarchical topology, spoke sites are connected to a central hub. While this approach is simpler from an operations standpoint, the delays caused by spoke-to-spoke traffic tromboning through the hub can make it inappropriate for some mission-critical application requirements. Organizations with sites located on several continents can optimize traffic flows more effectively with an any-to-any arrangement or at least by designating a hub in the same continent. To satisfy these requirements, local concentrations can be done in a multilevel hub-and-spoke architecture.

Partially meshed architecture

IPSec VPNs can also use meshed or hierarchical architectures in concert, depending upon the needs of individual sites and the underlying network topology.

Figure 5 – IPSec partially meshed network



Some routing equipment software offers the ability to create dynamic tunnels using IPSec, which significantly reduces the work involved in router configuration. A dynamic tunnel is a special implementation of a manual tunnel that relies on preset options to simplify the process. This technology can support a single-level hub-and-spoke configuration, including a dual tunnel head end redundant design, but also offers the ability to dynamically create temporary spoke-to-spoke connections.


For complex VPNs, dynamic routing offers advantages over static routing in detecting failures, taking advantage of redundancies and avoiding the complexity of maintaining many static routes. Because IPSec tunnels don't support IP routing, IPSec will be run over an encapsulation protocol that supports dynamic routing, like standard IP Generic Routing Encapsulation (GRE) tunnels.

Quality of Service

Because an ISP will prioritize packets equally within an IPsec tunnel, traditional end-to-end QoS is not possible.

However, there are additional service quality concerns that need to be addressed for an IPsec VPN. The internal routing and the peer selection among providers strictly follow Internet rules – packets are routed to the destination network using the closest Network Access Point. But the relationship between different Internet Service Providers is rarely known.

Figure 6 – Closest exit routing between ISP – not always the shortest path



When a VPN spans multiple ISPs, these ISPs should at least over-provision the bandwidth in the core backbone to avoid packet loss or sequence errors in the delivering of packets. This is of great importance for two reasons:

- _ In case of congestion, UDP high-priority packets like voice may be dropped.
- _ Sequence errors can be interpreted by IPsec processing as a “session replay” hacking attempt (anti-replay functionality), and therefore the received packets will be deleted.

SSL

Summary

While SSL is not a networking technology, it has been used by a number of vendors to provide remote-access VPN capabilities by creating application-specific secured tunnels between users and application servers (layer 7 VPN). It has now become an emerging alternative to IPsec for secure remote access for personal computers, PDAs and so forth.

Benefits

SSL is integrated in all leading web browsers, so it doesn't require additional software installation on the end user's workstation. SSL can be used to securely connect mobile users from home, a cyber café or a hotel room. An SSL VPN can be accessed from locations that restrict external access since it operates transparently across NAT, proxy and firewalls (since most firewalls allow SSL traffic).

Drawbacks and limitations

SSL is not designed for site-to-site tunneling. Since SSL VPNs natively support TCP services, such as web (HTTP) or email (PoP3/IMAP/SMTP), this configuration is easy to implement for these services but less trivial for other applications. Security is also a challenge, because while encryption provides confidentiality of the data in transit, it also provides an easy way in for malicious content ("you cannot stop what cannot be seen"). SSL-based VPNs are vulnerable to external unauthorized connections if cookies and session information are not removed after using SSL clients in public areas.

SSL, and its successor TLS, are protocols for managing the security of message transmission on the Internet. Often associated with e-commerce (e.g., links with HTTPS), this protocol enforces strong authentication between web server and browser by leveraging digital signatures and certificates with the highest security levels (class 3 certificates).

SSL gained popularity as companies with a large number of mobile workers looked for a remote access solution that was easily deployable. Present on almost every computer, SSL was the solution, providing remote users with secure access to web or client/server applications and connectivity to internal networks.

Even if SSL is not a networking technology, it has been used by a number of vendors to provide remote-access VPN capabilities by creating secured tunnels for users of the same application (layer 7 VPN). It has now become an emerging alternative to IPSec for secure remote access, but, despite the popularity of SSL VPNs, they are not intended to replace IPSec.

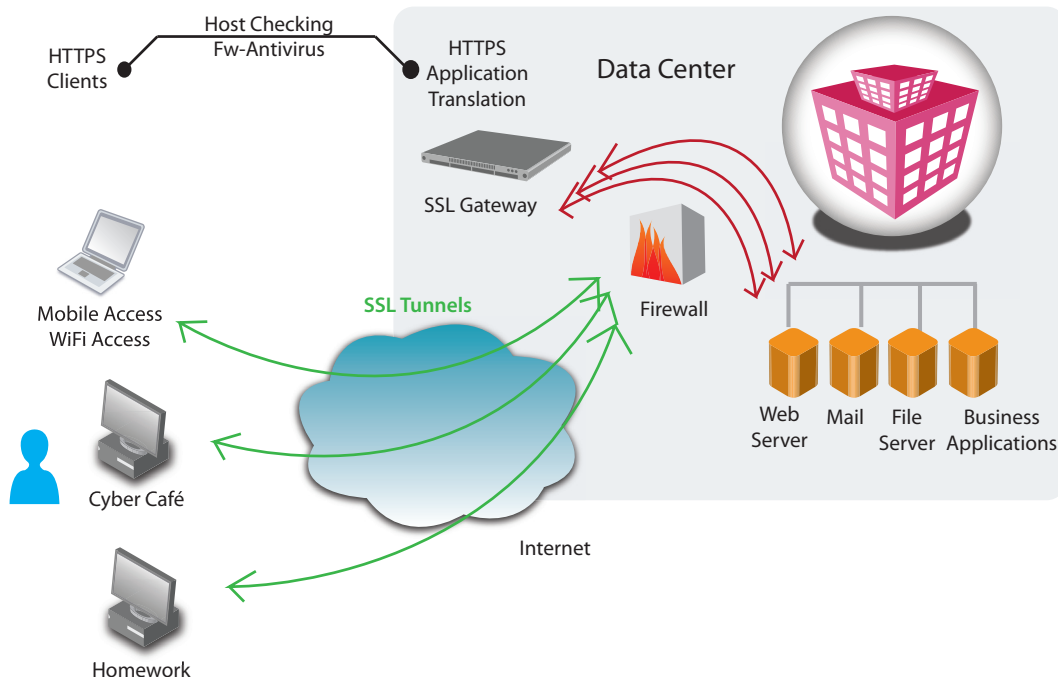
Network Considerations

SSL VPN does not generally require the installation of client software applications (these VPNs are often referred to as "clientless"), since SSL is an integral part of every browser. Most of today's computers are therefore already equipped with the technology required to connect mobile users to an internal network.

It is possible, therefore, to use the web browser to allow a remote endpoint to tunnel client/server applications (users authenticate to a web portal, typically an SSL VPN Gateway, and simply download a small plug-in that will take client/server traffic and tunnel it over SSL).

However, IT infrastructure for SSL-based VPNs must include application proxies (SSL must be aware of each connection or application session). Additionally, a server with enough memory and processing power is required, since this SSL VPN gateway is the endpoint for the secure connection and provides various features such as authentication, encryption and integrity protection or access control.

Figure 7 – Possible SSL VPN architecture



The main types of SSL VPNs are SSL portal VPNs and SSL tunnel VPNs. They are both used to provide remote users access to multiple services controlled and administered by an SSL VPN gateway (usually a firewall, a router with SSL VPN capabilities or a single-purpose SSL VPN-enabled hardware system).

SSL portal VPNs (or clientless): Users access, after identification, a web site with an SSL connection to reach multiple network services channeled through the web page. The web page acts as a portal to other services (e.g., links to other web servers, web-based email, file sharing, applications that run on protected servers or any other services that can be channeled through a web page) but is not a link among applications. To the user, an SSL portal VPN is a web site with more service options available.

SSL tunnel VPNs: An SSL connection actually creates an application-dedicated tunnel between the remote user and the local network. An application installed locally on the user computer ("heavy client") or a web-based application handling active content such as Java or Flash ("light client," downloaded transparently) establishes the connection between the client and the remote server through the proxy using a handshake. The proxy between the remote user and the local network has no access to the encrypted data transferred. The user will then securely access the requested network services. To the user, an SSL tunnel VPN may appear quite different from a typical web site because of the use of the tunneling plug-in or application.

It is also possible to create a non-application-specific tunnel – an IP tunnel that connects the user directly to the office LAN, analogous to an IPSec client.

Security considerations

Since web browsers are widely available, SSL can be used to connect mobile users from home, a cyber café or a hotel room.

An SSL VPN can be accessed from locations that restrict external access since it operates transparently across NAT, proxy and firewalls (most firewalls allow SSL traffic).

For mobile users, the SSL Gateway will generally protect enterprise resources by implementing the following security features:

- _ User authentication
- _ Host checking
- _ Security enforcement (such as a firewall, or up-to-date antivirus as per enterprise policies)

The SSL Gateway also handles application and resources access rights based on the remote user and computer and can apply differentiated access levels for employees and partners, depending on whether the computer being used is a company computer.

Most of the time, security best practices require the use of specific HTTPS filtering (SSL scanning) software in the SSL VPN client, because:

The growth of encrypted transactions is one of the major security concerns for corporate networks; a significant portion of the traffic going through the corporate gateway is encrypted and thus not sufficiently scanned for viruses and cannot be checked for security-policy compliance.

While encryption provides confidentiality of the data in transit, it also provides an easy way in for malicious content (“you cannot stop what cannot be seen”).

Most of the time, users themselves take the responsibility for accepting third-party security certificates, even though they do not usually have the necessary knowledge to make this decision.

SSL-based VPNs are vulnerable to denial-of-service attacks mounted against their TCP connections, and to external unauthorized connections if cookies and session information are not removed after using SSL clients in public areas.

Hybrid IP VPN

Summary

Hybrid VPN is the capability of a single service provider to combine the range-of-connectivity option to address the customer's network challenges globally with a cost-effective and centrally managed solution.

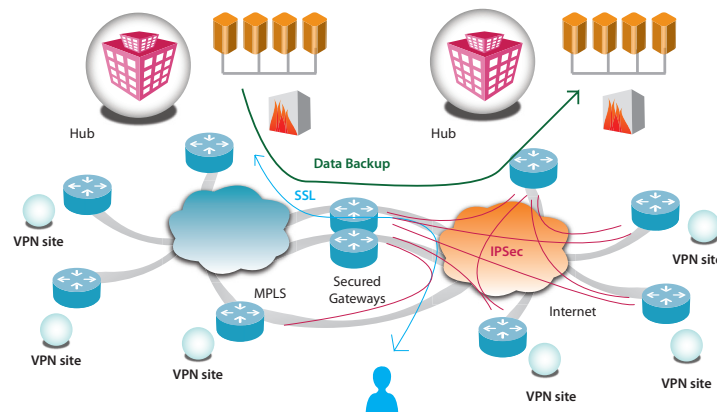
Hybrid VPNs allow customers to mix MPLS and IPsec VPNs on a per-site basis. This increases the network provider's reach, permits a tailored solution for each site and provides less expensive connectivity for groups of sites.

In order to connect provider- and CPE-based VPNs, one of several secure gateways is deployed between the MPLS backbone and the public IP network. These gateways support IPsec tunnel termination by authenticating and mapping the IPsec tunnel to the customer MPLS VPN, and ensure isolation between the public Internet and the Virtual Private Network.

Hybrid VPNs require significant technical resources to ensure that the global solution takes into account all constraints on connectivity, transport, VPN technology choice and design, sizing and application performance requirements.

Since no provider can cover all sites, it is important that a Hybrid VPN be able to seamlessly merge multiple provider access into a single global network, so as to optimize performance by choosing the right technology, reduce costs and increase resiliency by sourcing several providers or technologies per site.

Figure 8 – View of a Hybrid VPN



Network Considerations

For a Hybrid VPN, a simple star topology will often be inappropriate for application or user performance needs. The network designer must consider transit-delay-sensitive applications and choose appropriate access technologies, while also selecting an appropriate design for the transport cloud.

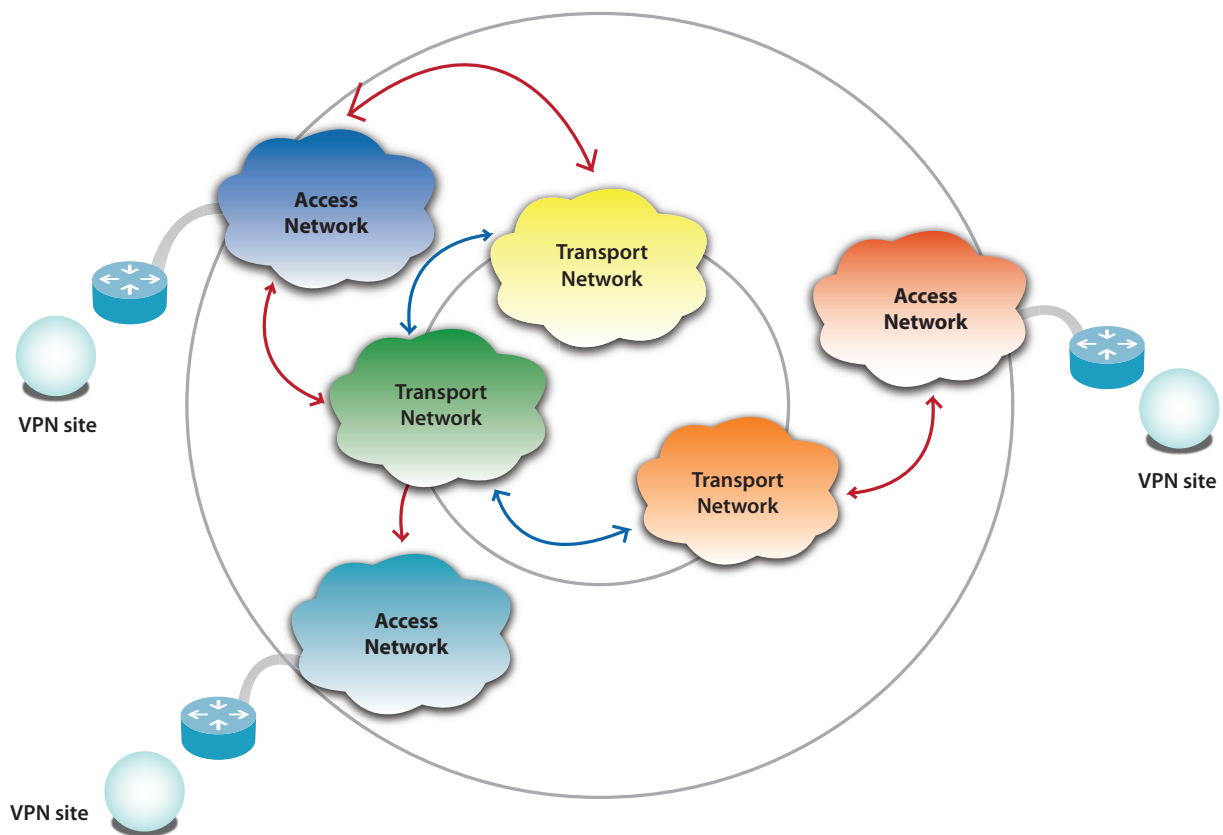
Transport Network

A Hybrid VPN solution should integrate the best transport options to meet performance objectives and costs. The available options can be the decisive factor in choosing a private or public solution. For each region, multiple providers and/or access types should be considered to obtain the best performance, and also to provide network resiliency and diversity as required. The following is a non-exhaustive list of decision elements to design the transport network:

- _ **Transit delay** - Company requirements versus Network Service Provider SLAs
- _ **Traffic map** - Dependent on geography, data centers locations and application flows
- _ **Transit networks** - Network Service Provider cores and peering points in each geography
- _ **Number of IPSec circuits** - Topologies should be chosen after considering traffic flows, content location, delay constraints, complexity of the solution
- _ **Hosting capabilities** - Network Service Provider ability to host gateways in order to optimize the circuit-based network

The resulting transport network can be more complex than the usual cloud that describes the Enterprise Network, because it uses several access and transit networks.

Figure 9 – The transport network (inside view)



Edge design

Decision criteria for the access links should include:

- _ **Service access** - Availability of local MPLS or Internet access
- _ **Site criticality** - Required redundancy level
- _ **Volume** - Throughput required at peak traffic hours
- _ **Distance** - To the nearest PoP
- _ **Last mile access** - Leased line, DSL, Cable, Ethernet, Wimax
- _ **Transit delay** - Provider SLA pass-through and packet queuing

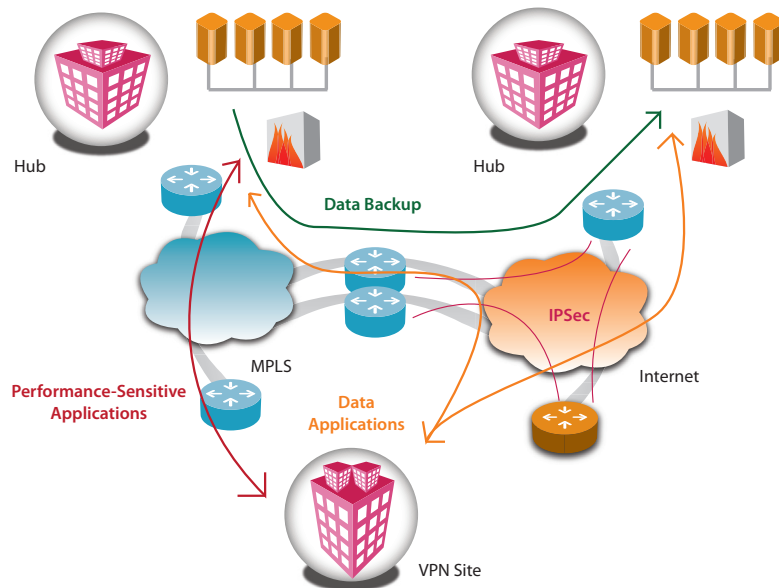
Backup and recovery

An IPSec VPN may be used as a backup solution for a primary MPLS connection. This configuration offers a dual (or single) CE, dual access links and dual transport networks to the customer's site.

This option can be used as an active or standby backup, and can also offload best-effort data through an inexpensive broadband connection. Load-balancing capabilities can be realized by using policy-based routing, or performance routing that takes into account the measured performance of each link.

Furthermore, this option can offer secured local access to the Internet.

Figure 10 – IPSec as a low-cost secondary link



A disaster recovery site already connected to the Internet can be easily connected to the transport cloud. Any site's IPSec router configuration can then be restored at a geographically diverse backup site. This router will learn the VPN routes via a dynamic routing protocol, giving access to the customer's whole Hybrid VPN.

Controlling the Network Architecture

Summary

Appropriate control of the network is critical to ensure application performance and user satisfaction in the long term.

The initial network architecture should be complemented by a set of rules to control future evolution. These rules should clarify the different requirements, constraints and solutions that have been met with the initial design, as well as the principles to be used to create global standards.

A set of processes and tools should be defined to track network usage, network and application performance, service availability, fault management, and managed changes required by new applications or increased network usage.

To deliver the full cost advantages of a Hybrid VPN solution, a service provider should also provide a single contact for all management and service assurance. A globally managed solution will ensure that no hidden costs arise for network management tasks.

In the network computing environment, a decision at any layer has the potential to disrupt the quality of service delivery and create dissatisfied users. A "control architecture" process should be put in place in order to standardize and manage the infrastructure as it evolves.

The first activity in network design should be to define a network architecture and to put in place a process to keep this architecture under control, rather than let opportunities, local initiatives or reactive problem-solving dilute the original objectives.

A network reference (an accurate snapshot of the network in constant evolution) should then be created and maintained to ensure knowledge of:

- _ **Architecture components** - Transport topology, edge design, and a complete list of flows and application characteristics
- _ **Administrative matters** - Service Providers, contracts and SLAs
- _ **End-to-end performance elements** - Response time, network availability, service availability (this list may be completed by the customer, an integrator or a third party)

A single function should analyze network indicators, be proactive in problem determination, and constantly optimize the network to reduce costs and enhance end-user satisfaction. This includes, but is not limited to, Network Service Provider performance elements such as round-trip delays, link utilization, packet loss and network downtime.

Trouble ticket reporting can be used to evaluate operational efficiency using measures such as number of incidents by severity, mean time to repair by incident severity and end-user satisfaction ratings to assess the subjective part of the service.

Proper control of the network architecture ensures a repeatable, scalable approach for anticipating the evolution of the network and provides the best framework for achieving ongoing cost management.

Conclusion

Hybrid VPN solutions intelligently combine a wide range of technology, sourcing and design options that can be adapted to the different functional, quality and economic requirements faced by customers today. Overcoming the management challenges associated with a multi-technology, multi-access network requires the right partner who can determine the exact needs of a customer and simplify network management by being the single source supplier, integrator and provider of end-to-end management.

Tata Communications' Hybrid VPN Service

Tata Communications' Hybrid VPN service

Through its genuine Virtual Network Operator (VNO) capability, Tata Communications brings together the procurement reach, technical skills, processes and tools required to provide a Hybrid VPN solution that incorporates the benefits of both the MPLS and IPsec worlds.

Dealing with a single partner to provide an optimized Hybrid VPN permits operational simplicity and reduces total cost of ownership. A single contact covers all communication needs, including ordering, implementation, service management, troubleshooting, change management and global reporting.

Tata Communications' Hybrid VPN uses a mix of MPLS and IPsec VPNs to provide the optimum service on a per-site basis. By tailoring connectivity options on a per-site basis, Tata Communications provides the most efficient and cost-effective overall solution.

Tata Communications' VNO service extends its reach beyond its own worldwide MPLS backbone to enable service in more than 140 countries. Third-party access is fully integrated into the managed network solution. Using local providers directly in a number of countries reduces the overall solution cost.

The ability to use multiple local providers in the same location is also important and enables complete national coverage as well as fully redundant and resilient solutions where required. Finally, integrated management of third-party providers, and continuous benchmarking on price, performance and customer satisfaction, delivers the best long-term solution to Tata Communications' customers.

MPLS is preferred for business-critical applications that are sensitive to network congestion whenever an MPLS PoP is close enough to minimize international long hauling. MPLS can also be provided through a third party using network interconnections or procurement of local layer 2 access.

IPsec will be used when data services are the primary need, applications are not very sensitive to loss and jitter, the site is in a hard-to-reach country or low cost is the driver. IPsec can also be used to back up MPLS access or to carry the best effort traffic at a site while the critical traffic remains on the primary MPLS access. The IPsec network design is implemented by Tata Communications' experienced network design team to provide the best topology and routing solution.

Additionally, SSL can be added to address the needs of mobile workers.

Acronym Key

ADSL	Asymmetric Digital Subscriber Line
ATM	Asynchronous Transfer Mode
CE	Customer Edge
CoS	Class of Service
CPE	Customer Premises Equipment
DSL	Digital Subscriber Line
DSCP	Differentiated Services Code Point
FTP	File Transfer Protocol
HTTPS	Hypertext Transfer Protocol over Secure Socket Layer
HTTP	Hypertext Transfer Protocol
IP	Internet Protocol
IPSec	IP Security
ISP	Internet Service Provider
LAN	Local Area Network
MPLS	Multiprotocol Label Switching
NAT	Network Address Translation
OSPF	Open Shortest Path First
PE	Provide Edge
PoP	Point of Presence
QoS	Quality of Service (packet prioritization)
SLA	Service Level Agreement
SSL	Secure Socket Layer
TLS	Transport Layer Security
TCP	Transmission Control Protocol
UDP	User Datagram Protocol
VPN	Virtual Private Network
VRF	Virtual Routing and Forwarding
WAN	Wide Area Network

Tata Communications Limited along with its global subsidiaries (Tata Communications) is a leading global provider of the new world of communications. The company leverages its Tata Global Network, vertical intelligence and leadership in emerging markets, to deliver value-driven, globally managed solutions to the Fortune 1000 and mid-sized enterprises, service providers and consumers.

The Tata Communications portfolio includes transmission, IP, converged voice, mobility, managed network connectivity, hosted data center, communications solutions and business transformation services to global and Indian enterprises & service providers as well as, broadband and content services to Indian consumers. The Tata Global Network encompasses one of the most advanced and largest submarine cable networks, a Tier-1 IP network, connectivity to more than 200 countries across 300 PoPs and more than one million square feet data center space. Tata Communications serves its customers from its offices in 80 cities in 40 countries worldwide. Tata Communications has a strategic investment in South African operator Neotel, providing the company with a strong anchor to build an African footprint.

The number one global international wholesale voice operator and number one provider of International Long Distance, Enterprise Data and Internet Services in India, the company was named "Best Wholesale Carrier" at the World Communications Awards in 2006 and was named the "Best Pan-Asian Wholesale Provider" at the 2007 Capacity Magazine Global Wholesale Telecommunications Awards for the second consecutive year.

Becoming the leading integrated provider to drive and deliver a new world of communications, Tata Communications became the unified global brand for VSNL, Tata Communications, Teleglobe, Tata Indicom Enterprise Business Unit and CIPRIS on February 13, 2008.

Tata Communications Ltd. is a part of the \$62.5 billion Tata Group; it is listed on the Bombay Stock Exchange and the National Stock Exchange of India and its ADRs are listed on the New York Stock Exchange (NYSE: TCL).

www.tatacommunications.com

Tata Communications
"Hybrid VPNs: Optimizing Performance across the Enterprise IP VPN"