



Rolling Out New SSL VPN Service

Introduction Typically, service providers offer infrastructure services, such as site-to-site connectivity and data center hosting. In addition to this, they are always exploring new opportunities to offer managed services to corporate customers.

Recent years have seen the growth of several types of Virtual Private Network (VPN) Services. The common characteristic of these services is to virtualize the service entity using the same physical infrastructure as much as possible. In its simplest form, a VPN tries to logically share resources (network infrastructure, security, remote and wireless access), thereby assuring complete privacy to each customer. In a VPN context, the services offered should be tailored individually for each customer. QoS, security and management services are typically desired and they are either provided by the service providers or individually implemented separately for each customer.

Another commonly offered service is the remote access service. Traditionally, the remote access service was built with dial-in solutions. However, recently, with ADSL, cable internet and wireless access technology becoming very popular, IPSec and SSL-VPN based technologies have evolved.

Though different types of remote access have been offered, and several service providers offer remote access as a managed service, it has been predominantly a Customer-Premises Equipment (CPE) type of solution. The reasons for this are as follows:

- Each enterprise has specific security policies; in the traditional remote access service, customer-specific security policies could not be dynamically configured on a single access platform
- A remote access service requires authentication mechanisms (user profiles, etc); enterprises typically have their own authentication (AAA) infrastructure which are used by other services also; therefore, integrating the remote service as a part of an enterprise solution is a straight-forward method
- Layer-2 and Layer-3 separation for each customer VPN

In this white paper, we shall show how several of these service provider drawbacks can be overcome when using F5's FirePass SSL-VPN product and will elaborate on FirePass deployment in various service provider networks.

VPN Access Service providers typically have a single backbone infrastructure which is shared by many enterprise customers. The shared infrastructure can be either a Layer-2 based network (such as Frame Relay, ATM, Leased-Line) or more recently IP-Networks based on the MPLS technology (VPLS, MPLS-VPN).

ATM, Frame Relay and Leased Line networks offer site-to-site connectivity that provide dedicated Layer-2 connectivity.

MPLS-based VPN technologies (IP-VPN and VPLS) have gained wide acceptance and are now well-established in the service provider environment. Carriers exploit MPLS-VPN capabilities to deliver IP-VPN and VPLS services in the backbone. The multi-protocol nature of this technology makes it very interesting from a business point of view. VPNs are enabled by the MPLS technique of designating packets for transmission over explicit routes. The same can be used to construct tunnels through IP backbones and the Internet that shield a company's traffic from that of others.



Along with the deployment of the MPLS-VPN technology, service provider's value-added service offering is a primary source for revenue generation. Remote access service has been a very important requirement for corporate customers as it permits mobile users and remote employees direct access to corporate internal resources. Recent years have experienced the growth of SSL-VPN and the management and deployment costs make this technology very attractive as compared to IPSec-based solutions.

IPSec was initially developed for site-to-site connectivity. With the growth of remote access requirements, IPSec was extended mostly based on proprietary solutions to support remote access clients. The solutions rely on vendor-specific software clients.

This involves additional costs since the tunnel security policies have to be administered. It has also limitations with respect to extranet connectivity and limited support for various types of end user devices (PDAs, Internet-Cafe, etc).

Scalability and maintenance of IPSec solutions also have limitations because of a growing number of VPN connections. For large scale deployments, the administrative costs become cumbersome.

In addition, IPSec is not very well adapted for mobile users, particularly if users move between different locations or want to access dedicated intranet services from Internet cafes, WLAN-Hotspots or a partner's premises.

If an employer wants to login from a home machine, these machines have to be installed with appropriate software clients which are supported by administrators.

IPSec was initially designed for dedicated access to the intranet where client software is pre-installed. In such new scenarios, this can lead to increased configuration and support. Without pre-configuration, it might not be possible to access certain services at all.

IPSec permits full network access to all internal resources within an enterprise. Different service levels (partner and guest access, group access schemes) cannot be implemented easily.

Also, full network access inherently has security risks as the user potentially has access to all internal resources. The IPSec tunnel is established between the client machine and an IPSec aggregation gateway. The IPSec gateway then does not inherently provide security as in the case of reverse-proxy architecture such as SSL-VPN.

NAT and firewall traversal often require complex solutions. In some cases it might not be possible at all to establish IPSec connections if the firewall blocks IPSec traffic. SSL-VPN, on the other hand, relies on SSL technology, widely accepted in the Internet community and widely used in a variety of applications.

SSL-VPN offers a complete, reliable replacement to IPSec remote access with its clientless architecture. The F5 FirePass series offers full network access which is similar to the IPSec solution. An SSL-Tunnel is built between the client machines and the FirePass controller, and this is comparable to IPSec VPN access methods.

In addition to IPSec replacement, SSL-VPN can offer individually restricted access to specific applications by incorporating highly flexible authentication and authorization mechanisms. SSL-VPN permits various customer and application-specific mechanisms and security policies which allow for highly flexible Service Level Agreements (SLA). For most web-based applications, the clientless approach is very attractive as it simplifies administration and lowers maintenance costs.



Services offered by SSL-VPN can be distinguished based on the client authentication mechanisms and different access levels, e.g. defined as partners, guests, employees and administrators. The resources available to each group of users can be flexibly configured.

FirePass also offers webified user interfaces, thus enabling a better user experience when accessing legacy, terminal server or mainframe applications.

In short, SSL-VPN provides the full network access as with IPSec, with additional differentiated services which can be further enriched with the webifier interface to enhance the user experience.

SSL-VPNs have been considered typically as an enterprise solution. Each corporation will define its own access policies and use an SSL-VPN concentrator to allow mobile and remote users to gain access to internal resources. However, if a service provider offers SSL-VPN as a managed service, there are several advantages for an enterprise:

- managed remote access with appropriate SLAs can be offered by ISPs
- cost savings for corporate customers
- quick deployment times

A service provider typically aims at providing a remote access service which is scalable and easily deployable for several customers. Such a service should allow individual customization, permit quick deployment times and be easily manageable. All of these factors are critical in attracting corporate customers to purchase the remote access service from an ISP. An SSL-VPN concentrator is used to provide SSL-VPN as a generic service and if the same concentrator can be used by several customer VPNs and be adapted to the ISP requirements, then it will fit well within an ISP backbone, such as MPLS-architecture.

ISP backbone networks provide the same network infrastructure for several customers and SSL-VPN will provide a value-added service. SSL-VPN can be deployed as a customized service both in a traditional Layer-2 (ATM, Frame Relay) as well as MPLS-based service provider networks. In fact, SSL-VPN could be designed as an extension of the Layer-2 or Layer-3 networks toward the Application OSI Layer.

In a service-provider environment, new challenges arise when an SSL-VPN is used. Several aspects of the remote access service have to be considered when developing such a service:

- network access methods
- resource management
- routing policies
- security policies
- authentication and authorization mechanisms
- administration and management
- high-degree of customization
- user-experience

FirePass SSL-VPN Service Offering

With the FirePass SSL-VPN, the following types of SSL-VPN services can be offered:

Network Based Offering

When SSL-VPN is offered as a Network-based service, a single FirePass or a single cluster of FirePass devices is shared across multiple customers. The deployment scenarios are:

- MPLS-based (IP Core network)
 - Layer-2 VPNs such as VPLS services relying on pseudo-wires
 - Layer-3 VPNs such as MPLS-VPN



- non-MPLS (using Layer-2 Core technologies)
 - ATM backbone
 - Frame Relay backbone
 - Leased-lines

CPE-Based Offering

For CPE-based Offering/Outsourcing, FirePass is deployed at the enterprise site and managed by the Service Provider either completely or partially (such as using outsourcing services). The deployment scenarios are the same as in the former case:

- MPLS-based (IP Core network)
 - Layer-2 VPNs such as VPLS services relying on pseudo-wires
 - Layer-3 VPNs such as MPLS-VPN
- non-MPLS (using Layer-2 Core technologies)
 - ATM backbone
 - Frame Relay backbone
 - Leased-lines

Since the FirePass is installed at an enterprise site, the service provider's core network technology does not affect the SSL-VPN service. FirePass and the SSL-VPN service should be integrated within the enterprise network just like any other DMZ network device (Firewalls, routers, switches, servers, etc).

Service Provider Requirements

Layer-2 networks offer site-to-site connectivity and the overlay routing policies can be individually defined for each customer. MPLS offers a shared Layer-3 (routing and forwarding) infrastructure with individually customizable routing policies. The multiprotocol nature makes it independent of any Layer-2 technology employed. MPLS-based Layer-3 VPNs by default offer a any-to-any communication model. VPLS provides a transparent LAN connectivity between various sites. Multi-protocol Border Gateway Protocol (MP-BGP) is used to distribute VPN-related information [1] and VPLS is based on the RFC Draft [2]. BGP distributes route information across the provider's backbone network so that the provider participates in and manages customer routing information. Carrier's-carrier models provide a further scalable extension to the MPLS-VPNs across multiple service providers. VPLS provides extensions to VLANs which are transported through appropriate tunnels across a MPLS-backbone.

The major advantage of the MPLS-architecture is that a single service provider core can scale for thousands of VPN customers irrelevant of the size of each VPN. MPLS-VPN and VPLS can support both small and very large-scale VPN deployments. Other advantages of MPLS technology lie in the support for end-to-end QoS, traffic engineering, fast reroute, and bandwidth protection. These can be offered as value-added services by the service providers. Any new service for MPLS should take the multi-VPN capability into account. This also means that a service which is developed for a MPLS-VPN network can be potentially scalable for several thousands of customers.

MPLS-VPN and VPLS can be combined with SSL-VPN to provide a very flexible remote access service. Such a service is characterized by its high degree of flexibility and infrastructure reusability. In its simplest form, MPLS-VPN offers security and separate routing and forwarding schemes up to Layer-3 and SSL-VPN provides customized services at the application level.

Enterprises and Service Providers have varying requirements on such an integrated remote access service.



Enterprise Scenario

A typical enterprise solution would require a single or a cluster of FirePass devices which can be setup with the enterprise-wide policies. The policies and access methods are individually defined globally for the entire enterprise. This can be seen as the standard SSL-VPN deployment.

ISP Scenario

An ISP will want to deploy the SSL-VPN as a generic remote access VPN Service which can be reused by several customers. There are two possibilities:

- one SSL-VPN concentrator for each customer
- a multi-Service SSL-VPN concentrator which can be customized for each customer

F5's FirePass is primarily intended as a Remote Access SSL-VPN concentrator and supports SSL-termination and SSL-VPN for enterprise/corporate networks. This is shown in Figure 1.

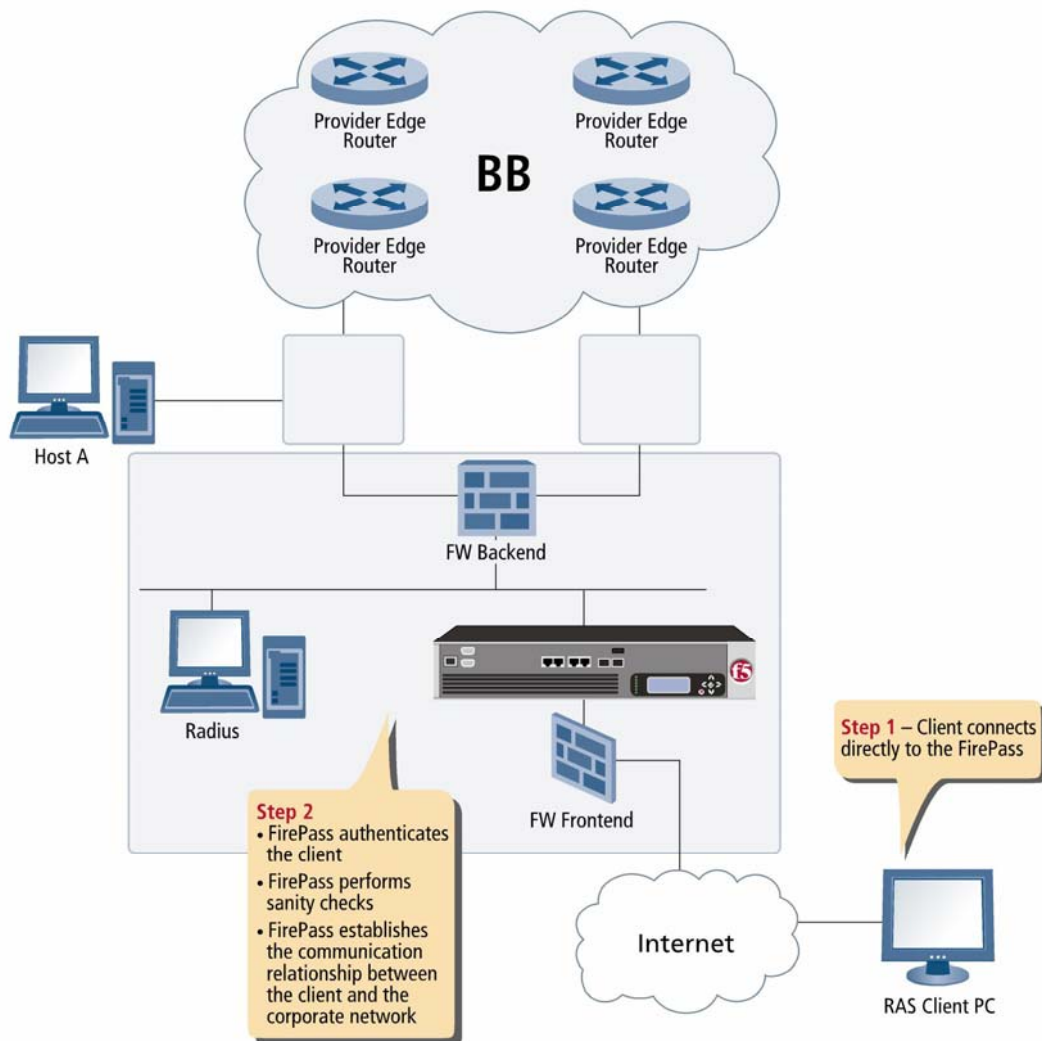


Figure 1: SSL-VPN Concentrator

The disadvantage of this solution is that the above solution may not be able to be deployed by a service provider because the service provider typically may not have access to the VPN client network.



Also, certain SSL-VPN devices may require skilled engineers with a good understanding of networking technologies in order to correctly configure and manage these devices.

On the other hand, if a scalable and generic SSL-VPN solution can be provided with which a service provider can offer SSL-VPN as a service to several customers, then the SSL-VPN could be made available in this capacity.

The resulting scenario is shown in the following. One SSL-VPN FirePass is managed by the ISP and it has the capacity to map various remote access clients to different MPLSVPN networks. Thus, the ISP can use a highly scalable SSL-VPN service (with various levels of Service Level Agreements such as virus scanning, access levels, etc.) which can be defined as standard and/or specialized features either for all customers or individually for each customer.

The return of investment is high for ISP as large-size companies can use the customized services, and a new market segment (small and medium-sized companies) can now subscribe to such services. As a further extension, even individual retail users (private ADSL customers) can now be offered certain sanity checks which were previously not available or only available with additional infrastructure costs.

In the remainder of this document, we shall refer to this generic solution as *virtualized* SSL-VPN as opposed to the standard SSL-VPN.




Figure 2: Virtualized SSL-VPN Service



The second solution is attractive for a service provider as the same infrastructure can be reused for different customers. However, this is only possible if all the enterprise-specific policies can still be supported by the SSL-VPN concentrator.

The service provider solution is shown in Figure 2. Here, FirePass provides a generic platform by virtualizing all the major SSL-VPN functions. This platform offers a virtualized environment where the individual requirements and policies of each customer can be defined separately and then grouped to provide a customized service for each customer. The virtualized FirePass will have to provide the following services:

- Separate interfaces and physically unique physical or logical interfaces for each customer (Layer-2)
- Separate routing with separate routing tables for each customer (Layer-3)
- Group-based policies for each customer VPN
- Group-based resources individually for each VPN

The virtualized SSL-VPN scenario is very interesting for ISPs because not only can a single infrastructure be reused for different customers but also a single point of entry simplifies the management of such a device. Consequently, the service activation and deployment times and the associated costs are reduced accordingly.

The following table compares these two deployment schemes:

Single FirePass Per Customer	One FirePass For Several Customers
Deployment possible by ISPs only if the ISP has access to the internal network of the customer	Deployment can be done by ISPs and the ISP does not necessarily need to have
access to the internal resources of the enterprise	ISPs have to spend separate devices for each enterprise
Same devices for several enterprises	Investment upfront per box CAPEX (Capital Expenditure) is lowered considerably as the solution scales well for several customers
ISP needs to deploy and manage new platforms commonly for various customers	Existing platform can be reused for several customers
Management and roll-out times are not negligible	Fast roll-out times and single management
Individual customization for each customer is possible	Individual customization for each customer is also possible
Individual management capabilities available	Individual management capabilities can be defined and offered with the help of administrative realms (see below for details of administrative realms)
Sanity checks and security policies can be defined individually for each customer	Sanity checks and security policies can be defined for each customer
Each customer needs a separate Entry Point (i.e., IP Address, DNS Entry, etc) to be able to access the FirePass	Single Entry Point for all customers; with the help of Landing URIs (see below for more details) each customer VPN can be distinguished



In the following, we shall present a case study which outlines how the FirePass solution is integrated within a service provider environment.

ISP Requirements

An ISP wants to provide SSL-VPN service to several enterprise customers. Each customer would like to:

- Have separate management capabilities
- Allow and manage their internal/intranet services
- Manage the users and authentication mechanisms
- Specify separate security policies
- Have different routing policies
- Specify and manage the end point policies (such as sanity checks, etc)

Implementation FirePass as an SSL-VPN concentrator offers several valuable features for remote access. We shall now see the various deployment schemes and show how the FirePass can be integrated with:

- **Case 1** : CPE-based offering with various types of remote access services
- **Case 2** : Network-based offering with virtualized SSL-VPN for several customers

We shall also show how the FirePass can be configured for several services in Case 1 and for several customers in Case 2.

Single Entry Point

Single Point of Entry refers to all types of remote access and only one IP Address is needed to access the RAS service.

Case 1

Various services for a CPE-based offering are customized using the Landing URI mechanism. A single domain name and IP Address is required for all services and each Landing URI can represent a specific service.

Case 2

Customization for each individual customer is done using the Landing URI mechanism. Since Landing URI permits a virtual host access method, using the same domain name and only one IP Address is required for all customers, this scales well with a large number of customers and for a large number of services.

In order to allow for a scalable solution, VLAN tagging is used to connect the FirePass with an ISP gateway (e.g. a router or a switch). The ISP gateway will assign each VLAN to a separate customer.

In the case of Layer-2 backbones, either bridging or a distribution switch could be used as shown in Figure 9 and Figure 10. In the case of MPLS-VPN, each VLAN is mapped to a separate Virtual Routing and Forwarding, VRF, (in the case of MPLS-VPN) or a dedicated customer VLAN (virtual circuit identifier) in the case of VPLS networks. FirePass can be either directly connected to a Provider Edge Router which maps the VLAN into a VRF-instance.

Alternatively, a CPE router can be used to terminate the VRFs. In the case of VPLS, the MPLS Edge router can directly map the VLAN to a pseudo-wire or a CPE Switch can map the VLAN to a Trunk Port (e.g. q-in-q technique in conjunction with hierarchical VPLS). Then the VPLS-Edge router can use Q-in-Q technique to map each VLAN to another Provider VLAN and then transport across a MPLS-backbone along a pseudo-wire.



Master Groups

A master group defines the authentication mechanisms that will be used and the resources and services that will be available to all the users belonging to the group. All the policies can be defined separately for a master group.

FirePass offers the possibility to map a Landing URI to a master group.

For CPE-based offering, each service can be assigned a different master group. In Case 2, by defining a separate master group for each customer and associating a unique Landing URI with the master group, flexible VPN access methods can be created. Further separate Landing URIs can also be defined for various services for each customer.



Figure 3: Master Group Definition

Individual Management

FirePass offers Administrative Realms which can be used to create various levels of administrative hierarchies. Users can be declared as Administrators for certain master groups and very granular access privileges can be assigned to them.

In the case of CPE-based management, normally the enterprise customer or the Service Provider will manage the FirePass. In the former case, different administrative realms can be defined in order to authorize various users to carry out specific administrative tasks.

For the Network-based offering, the FirePass device is managed by the Service provider. However, for each customer VPN a separate Administrative realm can be defined. The rights and privileges for the administrators of each customer VPN will depend on the specific customer requirements. For example, if the authentication servers of a customer are located within the customer premises, then the administrator can be allowed to configure or modify the Master Group and authentication server definitions.

Administrative realms can be used as follows:

- realm can administer on a single or multiple master group basis
- administrative realms can be used to define individual master group management
- each VPN (enterprise) will have one or more assigned administrators
- routing table and network resources are configured and administered by the ISP super user and a realm administrator cannot change the network configurations
- all intranet services and policy checks can be individually customized by each group administrator

- external authentication can be used in case the enterprise has its own AAA infrastructure
- visibility and transparency of administrative choices depend on the customer requirements

Thus a realm administrator can be allowed to configure and modify most of the master and resource group functions. The only restriction is with respect to the network configurations. Since the virtualized SSL-VPN offering relies on separate Layer-2 and Layer-3 security which is achieved by separating the VLANs and routing tables, the Administrators of each VPN should not be allowed to access or modify the network configurations or routing table.

VLANS and Routing

Separate routing tables are declared on the FirePass. VLAN can be defined and each routing table can be individually associated with a VLAN. The routing tables define policy-based routing such that all the client traffic (e.g., default route) is forwarded by the FirePass along a specific VLAN interface. The routing table is used as a resource which is assigned to a master group. With these mechanisms, separation of Layer-2 and Layer-3 can be achieved.

The overall view and virtualized FirePass configurations are shown in Figures 5 and 6.

For CPE-based offering, separation of Layer-2 (VLANs) and Layer-3 (routing) may not be necessary. However, for the Network-based offering, each customer network will be separated both at Layer-2 as well as Layer-3. Therefore, separate VLANs and Routing tables are required for each customer.




Figure 4: Overview Virtualized FirePass

Device Management : Configuration : Network Configuration Realm: Full e

Interfaces | **VLAN** | IP Config | Routing | DNS | Hosts | Web Services | Desktop | Misc

IP Configuration

Fields marked by * are required
Fields marked by color are optional

*IP Address/Netmask	*Interface	Broadcast	
172.31.1.200 / 24	eth0		Delete
192.168.10.1 / 24	eth1		Delete
192.168.10.10 / 24	eth1		Delete
10.10.10.1 / 30	VLAN: customer-1		Delete
10.10.20.1 / 30	VLAN: customer-2		Delete
10.10.30.1 / 30	VLAN: customer-3		Delete

Add New IP

*IP Address/Netmask: /

Broadcast IP:

*Interface:

Figure 5: VLAN Separation for Individual Customers

Device Management : Configuration : Network Configuration Realm: Full access Help

Table Number: 254 Table Name: main

*To (IP/Len)	Metric	Interface	*Via (IP)	Src (IP)	MTU	Window (bytes)
0.0.0.0 / 0		eth2	172.31.1.1	<default>		

Table Number: 200 Table Name: customer-1

*To (IP/Len)	Metric	Interface	*Via (IP)	Src (IP)	MTU	Window (bytes)
0.0.0.0 / 0		VLAN: customer-1	10.10.10.2	10.10.10.1		

Table Number: 201 Table Name: customer-2

*To (IP/Len)	Metric	Interface	*Via (IP)	Src (IP)	MTU	Window (bytes)
0.0.0.0 / 0		VLAN: customer-2	10.10.20.2	10.10.20.1		

Table Number: 202 Table Name: customer-3

*To (IP/Len)	Metric	Interface	*Via (IP)	Src (IP)	MTU	Window (bytes)
0.0.0.0 / 0		VLAN: customer-3	10.10.30.2	10.10.30.1		

Figure 6: Separate Routing Table for Each Customer

Resources Mapping

Each master group will be assigned following resources:

- Routing table
- IP Addresses from a IP Address Pool (Figure 7)
- Security Policies, Sanity checks, Firewall Policies

The screenshot shows the 'Network Access Settings' configuration page. On the left is a navigation menu with options like Resources, Device Management, Network Access, Portal Access, and Application Access. The main content area is titled 'Network Access Settings' and contains a table of IP address pools.

Name	IP Address	Mask
Default	192.168.192.0	255.255.255.0
<input type="checkbox"/> customer-1	10.10.100.0	255.255.255.0
<input type="checkbox"/> customer-2	10.10.200.0	255.255.255.0
<input type="checkbox"/> customer-3	10.10.220.0	255.255.255.0

Below the table are 'Update' and 'Delete Selected' buttons. There is also a section for 'Add new IP Address Pool' with fields for Name, IP Address, and Mask, and an 'Add' button. At the bottom, there are checkboxes for 'Use NAPT to access LAN' and 'Use packet filter to access LAN', along with an 'Apply these rules now' button.

Figure 7: IP Address Pools for Each Customer

Individual Authentication Methods

- Separate Authentication either from a shared authentication server or local to a customer

The screenshot shows the 'Master Groups' configuration page. The left navigation menu includes options like User Management, Announce, Endpoint Security, Groups, Master Groups, Resource Groups, Dynamic Group Mapping, Impersonate User, and User Name. The main content area is titled 'Master Groups' and contains a table of authentication mechanisms for different customer groups.

Group Name	Authentication	Resource Groups	Signup Template	Max concurrent sessions	Users	
customer-1	RADIUS	dynamic only	N/A	100	External	Delete
customer-2	Active Directory	dynamic only	No	100	Local	Delete
customer-3	Windows domain	dynamic only	N/A	100	External	Delete
Default	Internal database	0, no dynamic	N/A	100	Local	

A red arrow points to the 'Internal database' authentication method for the 'Default' group. A 'Create new group' button is located at the top right of the table.

Figure 8: Customized Authentication Mechanisms



Individual Service Definitions

On the FirePass, for each master group, different services can be defined. These include:

- Network access
- Web based Email (OWA, iNotes etc.)
- Web based applications
- File Systems
- Application Tunnels (Client/server and web based applications)
- Terminal Services

Individual Security Policies

Similarly, different security policies and sanity checks can be defined at the master group level.

- Sanity Checks which consists of end-point security verification and of redirecting users to predefined subnets to download the compliant anti-virus, firewall, operating systems updates and patches
- Firewall rules such as granular access control, packet filtering based on protocol, port and destination
- Session timeouts

Adding a new Service for a CPE-based Offering

Steps:

1. Define a new VLAN if required
2. Define a new routing table for the service if required OR extend the global routing table to make sure that the intranet devices offering the service are reachable
3. If new resources (IP addresses or a new routing table) are needed for the service, then define a new resource group with IP addresses for the customer
4. If new routing table is required, then associate the resource group with the routing table
5. Define security and routing policies specific to the service
6. Associate the resource group with the routing table
7. Define a new master group and resource group for the new service
8. Define the authentication mechanisms the master table
9. Customize and configure the portal/network/application tunnel services
10. Add a new Landing URI for the service
11. Map the new Landing URI with the Master Group

Adding a new customer VPN to the FirePass

Steps:

1. Define a new VLAN for the customer; a new VLAN is always required in order to separate the Layer-2 for each customer



2. Define a new routing table for the customer
3. Define a new resource group with IP addresses for the customer
4. Associate the resource group with the routing table
5. Define a new master group and resource group for the customer
6. Define the authentication mechanisms
7. Define security and routing policies specific to the customer
8. Customize and configure the portal/network/application tunnel services
9. Add a new Landing URI for the customer
10. Link the new Landing URI with the Master Group

**Integrating
with the
Service
Provider
Backbone**

In the case of a CPE-based offering, the FirePass would normally be placed within the DMZ infrastructure. The deployment of the FirePass devices could be either done by the enterprise customer itself or delegated to the Service Provider. The same is true for the monitoring and management of the FirePass devices.

In the following we shall examine the various deployment schemes for the network-based offering. Service providers offer either a Layer-2 connectivity to interconnect different corporate sites or, more recently, Layer-2 or Layer-3 connectivity is available across an IP-Core network running MPLS in the backbone.

Interconnection with Layer-2 Backbone (ATM, Frame Relay Networks)

The following shows FirePass integration with a Layer-2 backbone network (e.g. ATM, Frame Relay, Leased-Lines).

- One FirePass is used for several customers
- Remote access is offered as an aggregated service
- FirePass is connected to the customer edge router
- FirePass VLANs are mapped to ATM (pvc), Frame Relay DLCIs
- Separation of the traffic happens at Layer-2

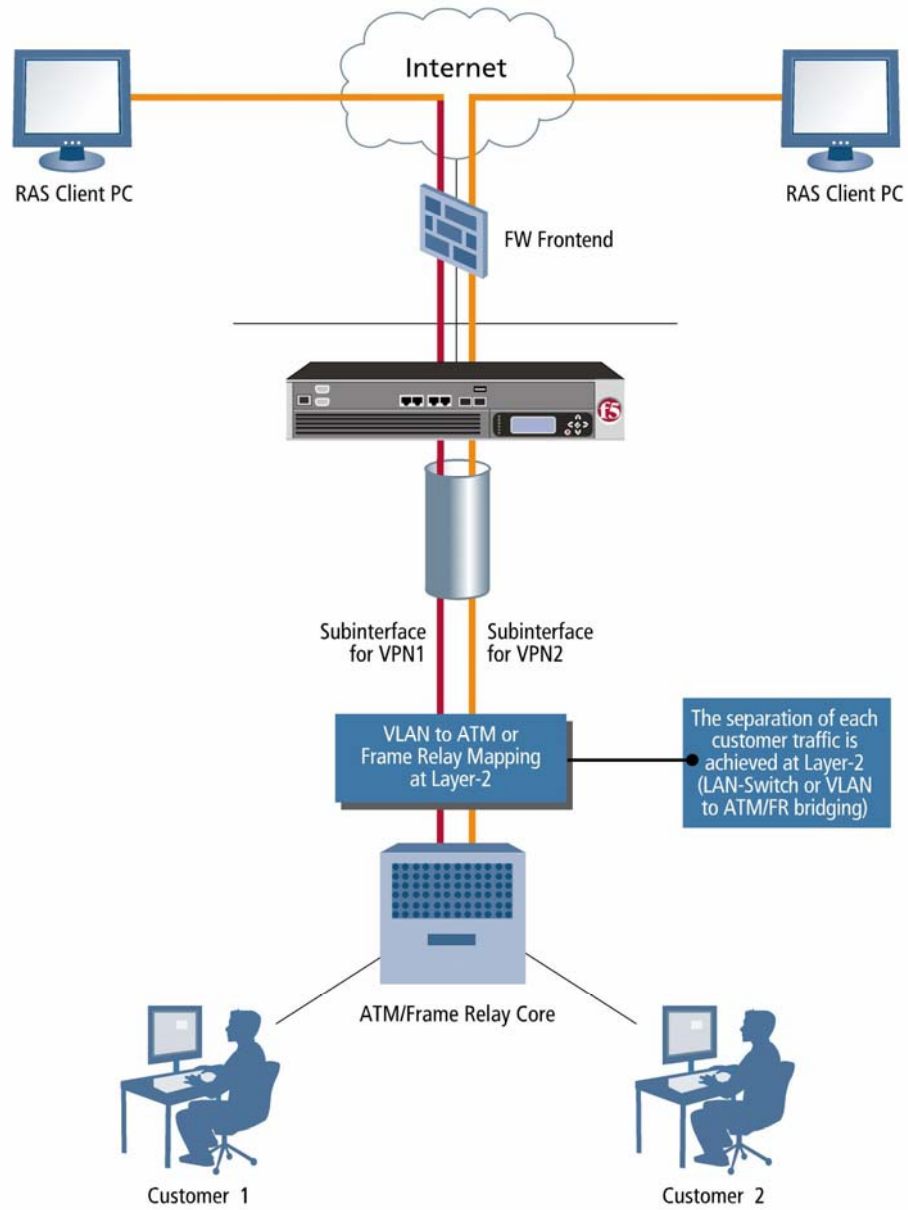


Figure 9: *FirePass Integration with Layer-2 Backbone*

Figure 10 shows one example with all the infrastructure equipment needed.

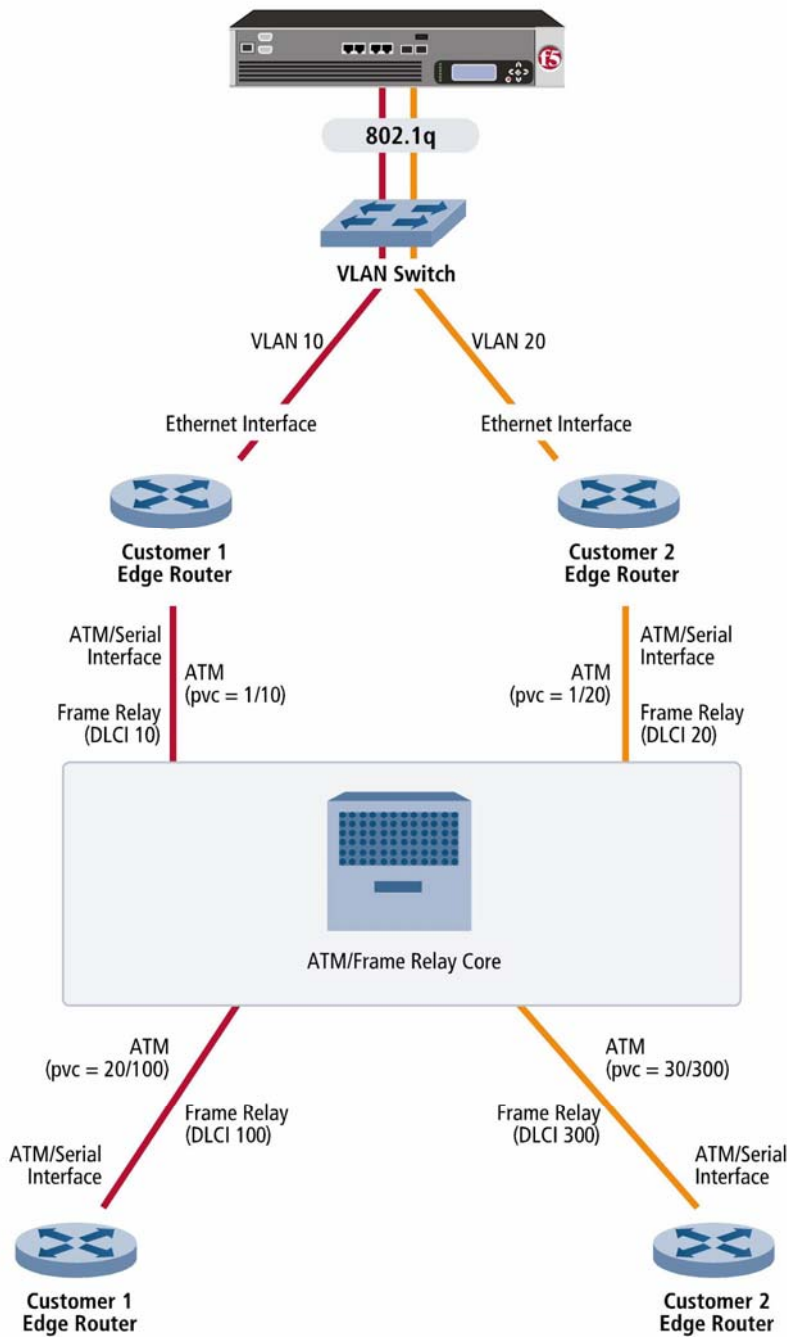


Figure 10: Layer-2 Integration Infrastructure

Interconnection with MPLS-VPLS

VPLS is a Layer-2 VPN connectivity across MPLS-backbones. FirePass offers VLAN separation of the customer traffic and this can be directly mapped on to VPLS Layer-2 service.

- One FirePass is used for several customers

- Remote access is offered as an aggregated service
- FirePass is connected to the provider edge router
- FirePass VLANs directly are terminated on the provider edge router
- Separation of the traffic happens at Layer-2 and this is retained on the provider edge routers
- VPLS traffic is carried over pseudo-wires along the MPLS-backbone

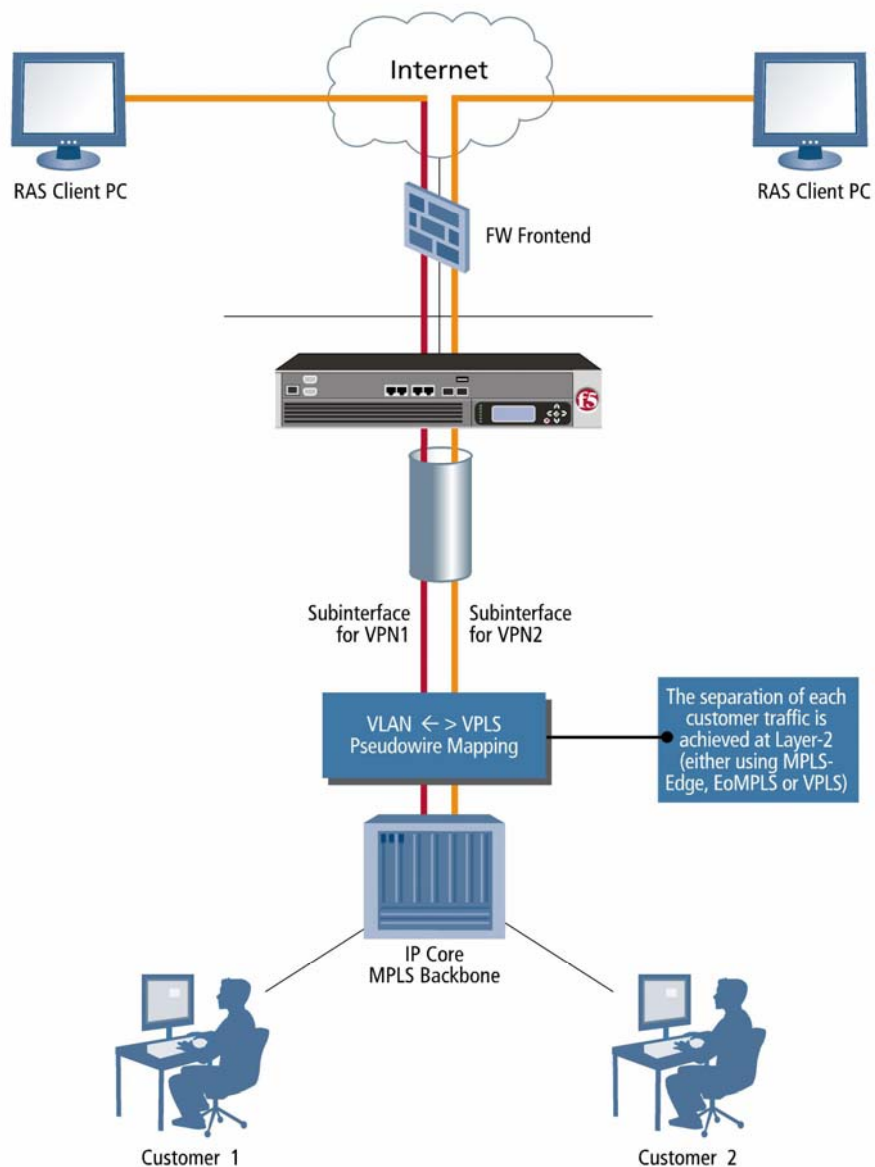


Figure 11: *FirePass in a VPLS Environment*

The infrastructure required for a VPLS implementation with FirePass is shown in Figure 12.

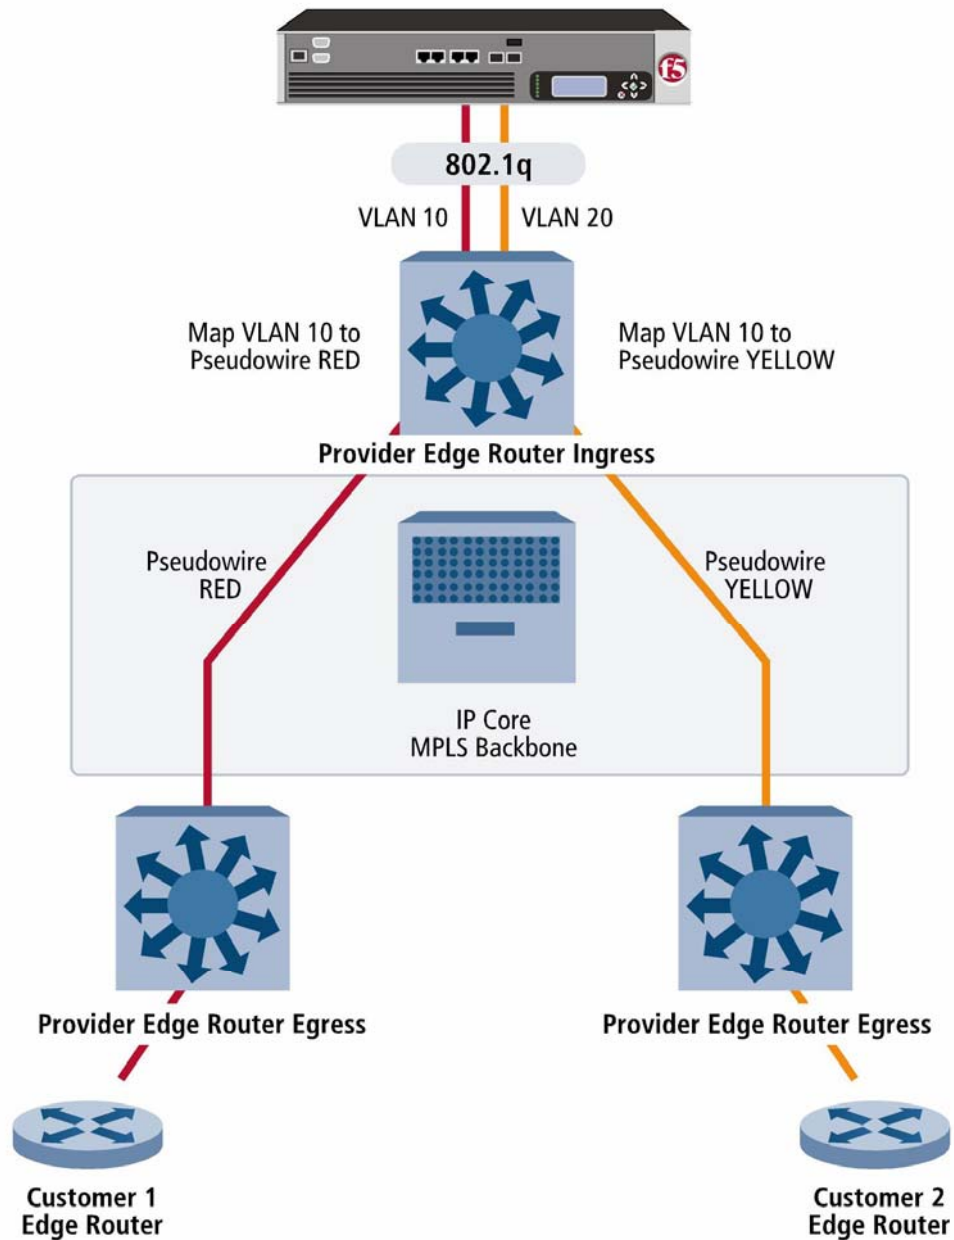


Figure 12: VPLS Integration Infrastructure

Interconnection with MPLS-VPN

MPLS-VPN offers VPN services at Layer-3. It is independent of the Layer-2 technology used. The VLANs from FirePass can be connected to the provider edge core and then terminated at Layer-3. Separate virtual routing and forwarding (vrf-tables) are available for each customer on the provider edge router.

- One FirePass is used for several customers
- Remote access is offered as an aggregated service

- FirePass is connected to the customer edge router
- FirePass VLANs are terminated on vrf at the provider edge router
- Separation of the traffic happens at Layer-3 (vrf routing tables on the provider edge router)
- Routing along the backbone: Direct connection with a PE Router
- Routing can be done in the backbone according to each customer requirement (any-to-any connection up to policy routing can be supported)

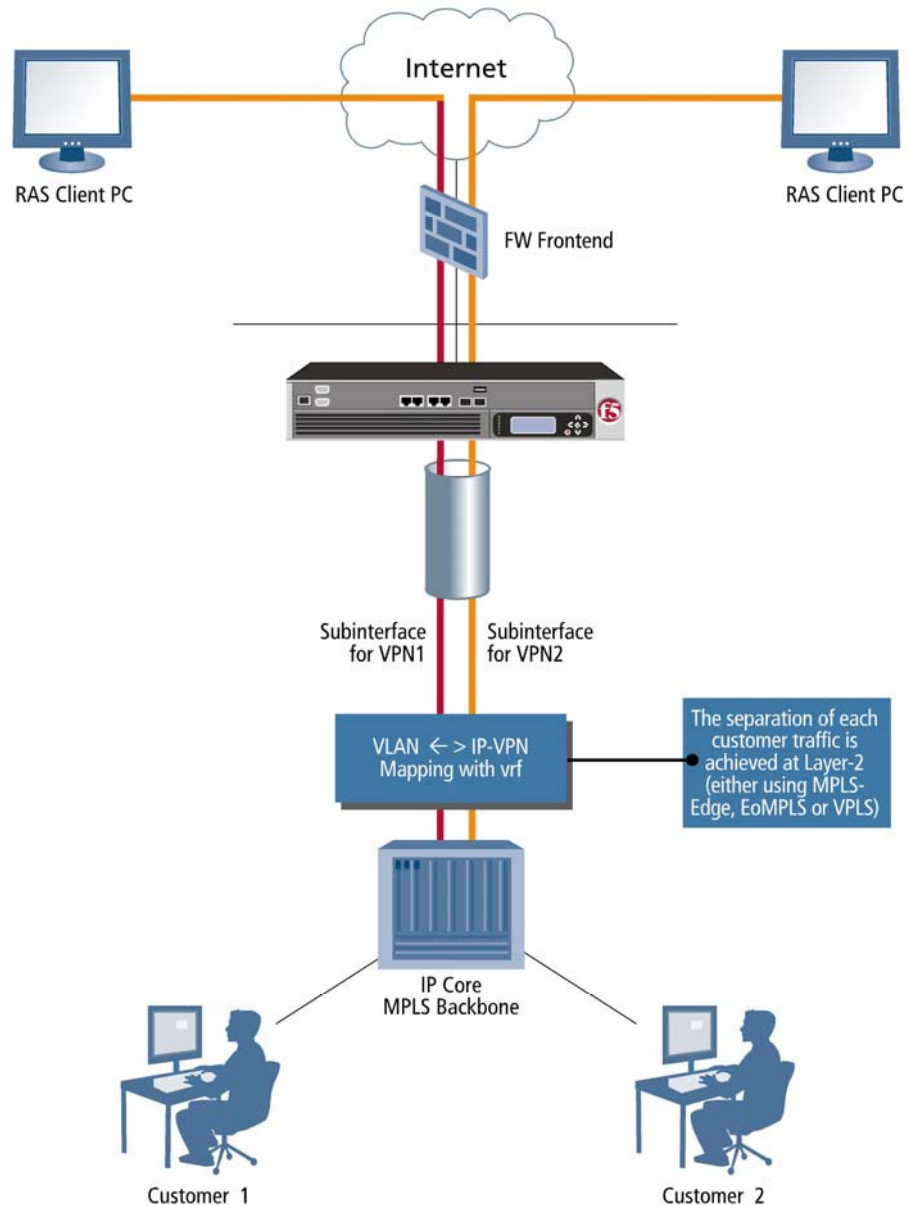


Figure 13: *FirePass Integration with MPLS-VPN*

A possible implementation of FirePass within an MPLS-VPN backbone network is shown in Figure 14.

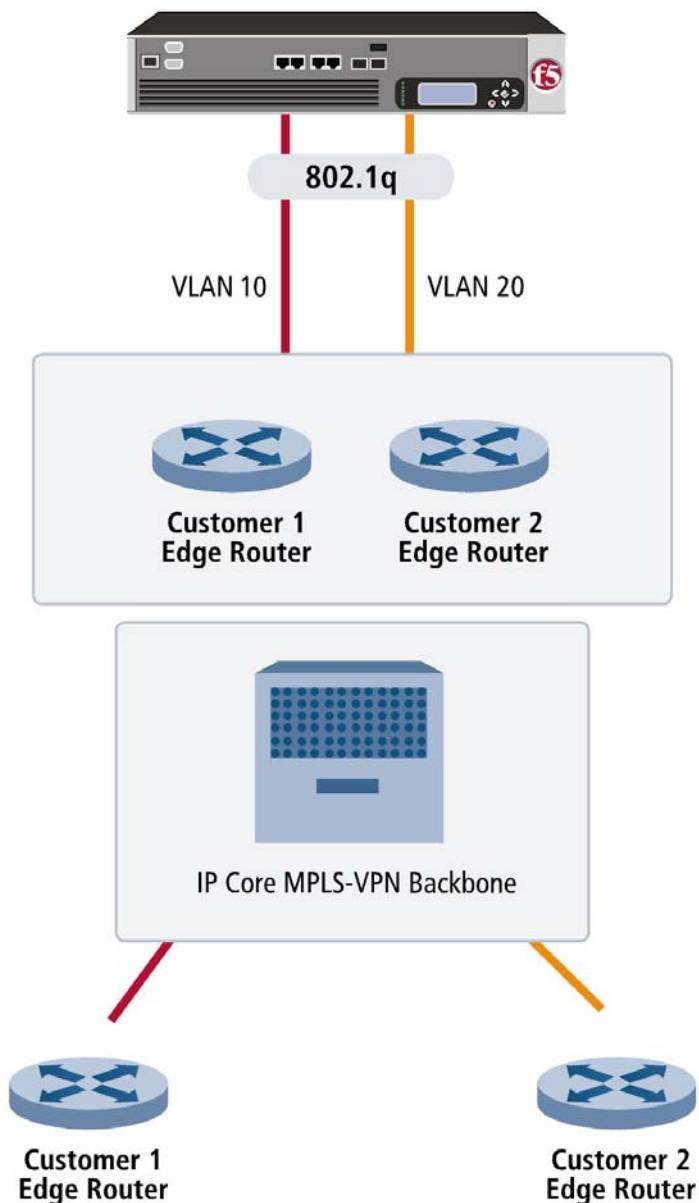


Figure 14: *MPLS-VPN Integration Infrastructure*

Scalability FirePass can be configured in a clustered mode. A cluster can contain up to ten FirePass controllers. Each FirePass can be configured as a redundant fail-over pair. Configuration synchronization and load-sharing are also available for members of a cluster. The bulk encryption performance of each FirePass is 200 Mbps; therefore, a clustered FirePass environment can support a total of 2 Gbps encryption.

A second possibility is to use the F5 BIG-IP traffic management devices to load balance SSL Sessions across multiple FirePass devices, as seen in Figure 15. In this case, all FirePass devices will have identical configurations and a large number of FirePass devices can be used, thus increasing the overall performance of the entire system. The total number of concurrent users and SSL transactions will determine the capacity of the system.




Figure 15: *BIG-IP Load Balancers with FirePass*

FirePass supports a total of 255 routing tables. Therefore, up to 255 customers can be hosted on a FirePass cluster, assuming each customer requires a separate routing table and distinct Layer-3 on the FirePass. Each customer VPN can further contain numerous master groups and several services.

The third scalability factor is the number of concurrent users. FirePass can support up to 2,000 concurrent sessions with latest software release. Thus potentially a cluster can support up to 20,000 concurrent users.

Conclusion FirePass represents an SSL-VPN solution with best in class usability for end-users and administrators. With FirePass, flexible SSL-VPN deployments can be supported. It is best suited both as an enterprise SSL-VPN concentrator as well as for large-scale Service Provider remote access service offerings. The virtualization capabilities make FirePass a very attractive solution for Service Providers. Service Providers can provide a wide-range of SSL-VPN offerings without compromising on flexibility or security. A virtualized FirePass can enable Service Providers to deploy SSL-VPN that supports various backbone architectures and large-scale roll-outs while providing the lowest total cost of ownership as well as comprehensive risk management with granular security controls.

- References**
- [1] BGP/MPLS IP VPNs, draft-ietf-l3vpn-rfc2547bis-03.txt, E. Rosen, Y. Rekhter, October 2004
 - [2] Virtual Private LAN Services over MPLS, draft-ietfppvpn-vpls-ldp-01.txt, Work in progress, November 2003

Author **Dr. Prasad Raja**
Network Consultant
Velera International