A Unified Approach for Securing Wireless, Remote, and Internal LAN Access

Overview   In today's enterprise environment, the proliferation of access methodologies such as Gigabit wired connections, high-speed broadband access, remote access from various locations (kiosk, internet café), and wireless networks, has driven the need for new services. Mobile and remote access to company resources are essential for remote worker productivity and business efficiency. However, larger scale deployments of wireless communication propagate new security and performance demands.

In such an environment, administrators face an ever-growing need to protect critical company resources from attacks. Controlled access to different resources based on user identity/credentials, user groups, and client devices are some of the top security requirements for these environments.

The components of a typical enterprise environment can consist of:

Access Networks

- Wired networks (referred to as *internal LAN*)
- Wireless access
- WAN or public Internet

Resource Networks

- Demilitarized **z**one (DMZ) for services accessible to the Internet
- Remote access gateways that are part of DMZ services
- Data centers with servers and mainframes that host business applications

Access & Control

- AAA infrastructure that consists of authentication and accounting servers

The diversity of an enterprise environment dictates the need to consider multiple aspects when planning for access. Normally, an internal LAN is considered a secure network. Due to its broadcast nature, wireless communications are not considered as secure. Such networks are vulnerable to eavesdropping, rogue access points, and other cracking methods. For remote access, VPN solutions such as dial-up, IPSec VPN, and SSL VPN are commonly used. And, any access to data center devices must be protected and secured. In the data center, access lists are used to prevent unauthorized access, and reverse-proxy servers use authentication mechanisms to provide a higher degree of security for applications.

Challenges   In an enterprise environment, the need for security is constantly evolving. Maintaining individual security methods for each access scenario, be it individual authentication or access control lists (ACLs), simply does not scale well and increases the administration burden, making it prohibitively expensive. There must be a better alternative for securing enterprise access – one that is cost-effective, easy to manage and secure, while addressing performance and scalability requirements.

Basic security requirements consist of:

- Verification of user credentials and services to define user access.

- Client integrity checks that consists of endpoint security verification and of redirecting users to predefined subnets to download compliant anti-virus software, firewalls, operating systems updates, and patches.

- Firewall rules such as granular access control and packet filtering based on protocol, port, and destination.

Very often, the same users access corporate resources from various locations. Therefore, security mechanisms and access policies should be independent of user access methods, such as wireless, internal LAN, and remote access. What is needed is a *unified security policy* with access rules that can guarantee the same user service level agreements regardless of access method. A unified approach is more cost effective, easier to administer and maintain, and less error-prone because all access policies need only be defined once.

This white paper describes the challenges addressed by a unified approach by outlining the technical requirements and characteristics of an internal LAN, wireless, and remote access architectures. F5 uses a *universal access approach* that covers the network layer on up through the application layer – all using a single, unified cost-effective management solution. Case studies show how the F5 approach delivers a unified security policy that addresses the requirements of multiple access methods.

Characteristics of Wireless, Remote & LAN Access

Although a unified approach is very attractive, it must address the different characteristics of internal LAN, wireless, and remote access methods. The following sections describe the characteristics of each access method.

Wireless Access

Wireless networks are designed for high mobility and flexibility. Mobility means that a user can move between various locations (different floors in a building, Internet café, and airport) while staying connected to the network. Flexibility means using various devices (PDAs, Smart Phones, Laptops) to connect to the network.

In a wireless environment, there are typically two connection scenarios:

- A public hotspot access
- Secure access to the intranet

In either case, user authentication is required to gain access. Hotspot authentication routes user traffic to the public Internet, whereas intranet traffic is controlled using VPN encryption methods. Further classification of intranet access can be based on user groups, servers and resources, and security policies that reflect a company-wide policy.

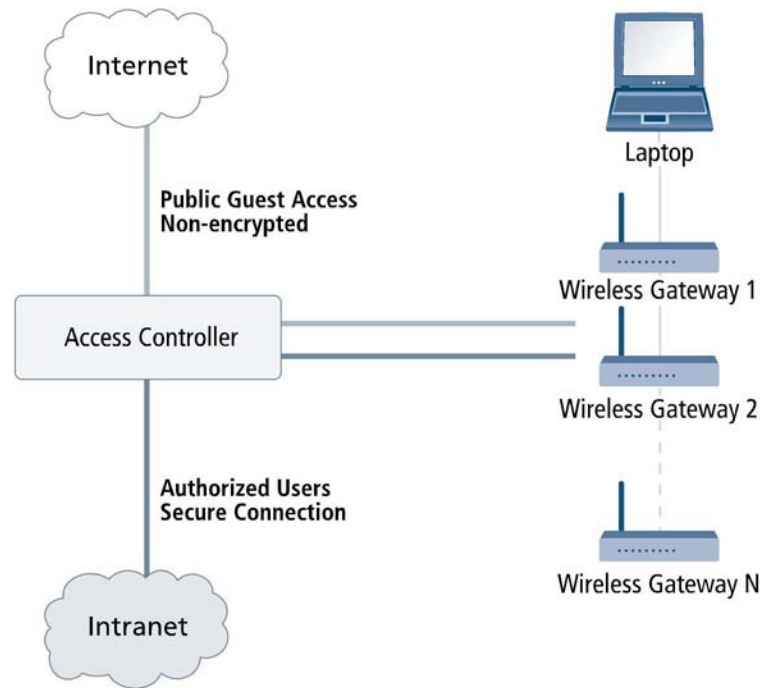Figure 1 shows a typical configuration to accommodate wireless access.

Figure 1: Wireless Access

Remote Access

Remote access enables users to access company resources from practically any location (Internet café, the airport, or a home office). The major requirement is to provide access to several user groups (authorized users, partners, customers), and associate them with the resources that they are allowed access. Access to resources should be possible from any client device (corporate-managed laptops, home PC, kiosks, partner PC, etc.).

Figure 2 shows a typical configuration to accommodate remote access.
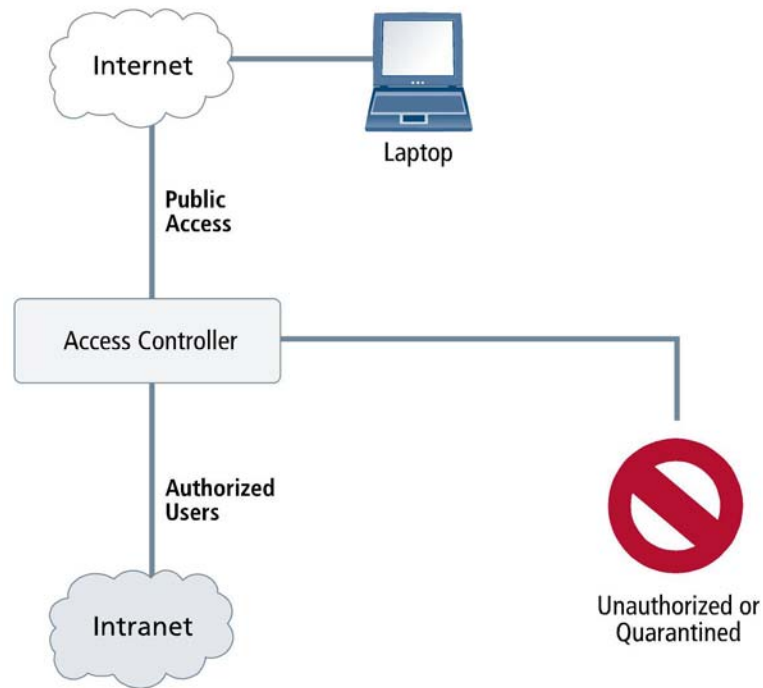
Figure 2: Remote Access

Internal LAN Access

This paper refers to a wired network within an enterprise as an *internal LAN*. Internal LANs are usually considered more secure than wireless and remote access networks; however, opening the internal LAN to any user is a security risk. Many wired networks allow the majority of communications to go unencrypted. Therefore, wired networks are just as susceptible to eavesdropping as wireless networks, especially when an outsider can get a physical connection. Also, internal users may be able to sniff traffic and gain access to other users' email (CEO's email, other authorized users' confidential information).

Figure 3 shows a typical configuration to accommodate internal LAN access.
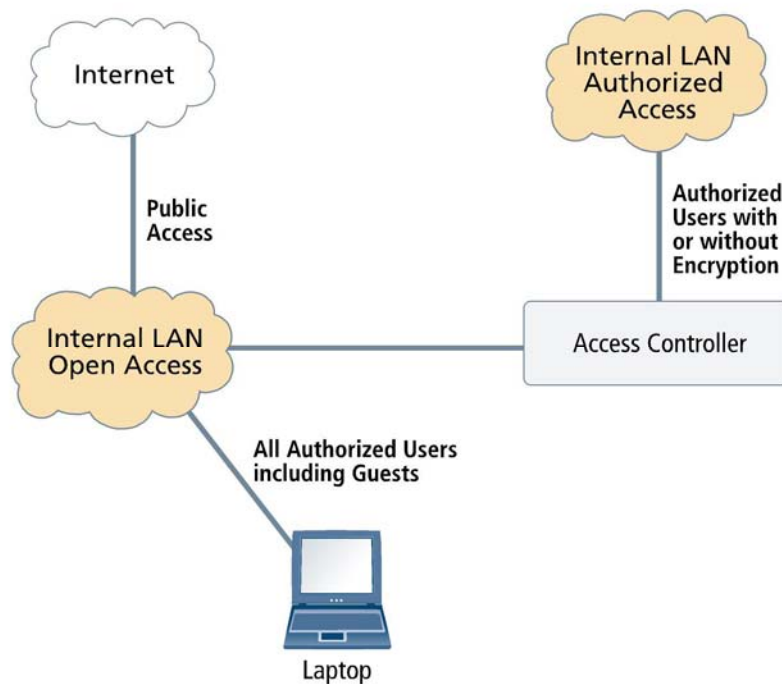
Figure 3: Internal LAN Access

When securing the various types of access methods, you must consider:

- **Susceptibility to Eavesdropping –** Data accessed over a wireless network can be easily sniffed or eavesdropped.

- **Exposure to Intruders –** Intruders can use a rouge access point to gain access to user data as well as sensitive enterprise resources. The WEP security protocol is insecure and can be broken. WPA/WPA2 security is quite new and requires a special administrative skill set to configure and manage.

- **Client Integrity Check –** Client integrity checking and endpoint security are essential to remote access as the risk of exposure to virus and other malware is high. Network administrators need to configure, monitor, and enforce a corporate compliance policy for endpoint security and user credentials, including operating system patch levels, antivirus versions and updates, and firewall versions.

- **Network Segmentation –** Segment various types of access and resource networks with individual policies to access specific resource networks or services within a resource network.

- **End Device Support –** A large range of client devices (from Desktop machines to PDA and Smart Phones) should be allowed access to internal resources without compromising security.

- **Peer-to-Peer Traffic Control** – Although client integrity checks ensure that a client conforms to the enterprise security standards, real-time traffic should be monitored and controlled. Otherwise, peer-to-peer traffic can easily abuse bandwidth, adversely affecting the Quality of Service for Internet access and potentially spreading malware and worms.

- **Robust Connections –** User connections could encounter temporary connection loss or the client IP addresses may change, for example when a wireless client is roaming or the IP address changes with an ADSL connection. Whenever appropriate, mechanisms should be in place to sustain temporary connection loss and re-establish the original session context to ensure application-level transparency.

- **Performance –** Large number of users and bandwidth are typical scalability parameters. Wireless and wired networks potentially require several hundred Mbps or even Gbps throughput capacity.

- **Multiple Geographic Locations –** In an enterprise network that stretches across multiple geographical locations, a decentralized approach might be better than a centralized policy engine even though the policies might be the same for all sites. For example, it might not be appropriate to route all traffic to a central policy engine across a WAN link if local resources have to be accessed.

- **High Availability –** Redundant configurations assure that these security requirements are enforced and controlled so that when a single machine fails, the fail-over machine can take over the operation.

- **Easy Access –** Users should have easy and sustainable access to the network from various locations. From an administrative viewpoint, access methods should not be complex. For example, mobile users need to easily connect to the Internet from wireless hotspots at various locations using a web browser.

- **Manageability –** In the case of wireless networks, managing security policies on each wireless access point is complex, error-prone, and doesn't scale well. For example, endpoint security should be checked before users can access the internal LAN. However, this cannot be done with WEP/WPA/WPA2 implementations.

- **Scalability –** Managing security policies on multiple switches and routers at the port level or using Access Lists is not a scalable model.

  ACLs that restrict VLAN port access to MAC addresses do not scale well or guarantee security. This schema also restricts the mobility of the user.
  Therefore, a centralized security policy management approach can help define a homogenous company-wide policy based on user credentials and roles assigned to each user rather than low-level, error-prone access lists. This reduces administrative and maintenance costs and lowers your total cost of ownership (TCO).

The following table summarizes the access characteristics of each type of network.

| Access	Wireless	Remote	Wired
Susceptibility to Eavesdropping	High risk	Low risk because the traffic is encrypted	High risk
Exposure to Intruders	Open to intruders if deployed without scanners	Intruders can get access only if they can successfully authenticate	Intruders can get access to LAN resources
Client Integrity Check	Not available with WEP/WPA/WPA2	Can be enforced during authentication	Generally not used
Network Segmentation	Wireless access is an insecure network segment Authentication and real-time traffic monitoring mechanisms are needed	Remote access is non-a trusted network Highest security policies are normally needed Variable policies could be used for different services	The wired corporate network can be partitioned into trusted, non-trusted, public and private zones Data centers and special applications require additional authentication

Support for Different End Devices	PDA, laptop and Smart phone Trusted and guest users possible Various operating systems	PDA, laptop, Smart phone, Internet café, home office users, mobile users Various operating systems	Access for desktops, preconfigured machines, laptops Various operating systems
Peer-to-Peer Traffic Control	For non-adhoc networks, depending on the configuration of APs, peer-to-peer can be allowed or all traffic could be routed to a centralized gateway (the universal access controller)	Remote access controller can always inspect peer-to-peer traffic from and to remote users	Peer-to-peer traffic can be inspected only if the traffic first goes across a switch or router
Robust Connections	Mobility and roaming have to be configured across wireless gateways Mobile IP and custom GRE tunnels are often needed	SSL VPN with SSL renegotiation ensures robust connections	Robust by definition
Scalability and Performance	Large number of users Distributed across multiple sites and several hundred Mbps or Gbps throughput needed	A fraction of corporate users will simultaneously use remote access Throughput depends on the corporate Internet bandwidth usually ten or hundred Mbps	Large number of users Distributed across multiple sites and several hundred Mbps or Gbps throughput needed
High-Availability Needed	Yes	Yes	Yes
Easy Access	Guests should have easy access without installing additional software	Internet cafe or PDA access for very easy access mechanisms without any additional software	
Manageability	Integrated policy and user management	Integrated policy and user management	Integrated policy and user management

Multi-Access Overview

Figure 4 shows an overview of the three different user access methods, which includes:

- Remote access and wireless networks are isolated, but with the help of the appropriate mechanisms, their access to the internal LAN is controlled. Remote access controllers are devices that control remote user access before users can access internal resources. Therefore, strong authentication and encryption are essential services for remote access.

- Wireless gateways address issues related to RFID and access point management, and may encrypt wireless traffic. A remote access controller can take care of authentication and VPN encryption before the traffic is forwarded to internal resources.

- In the internal LAN, resources are generally available to users. Security-sensitive applications are placed in one or several protected zones with the appropriate security policies.
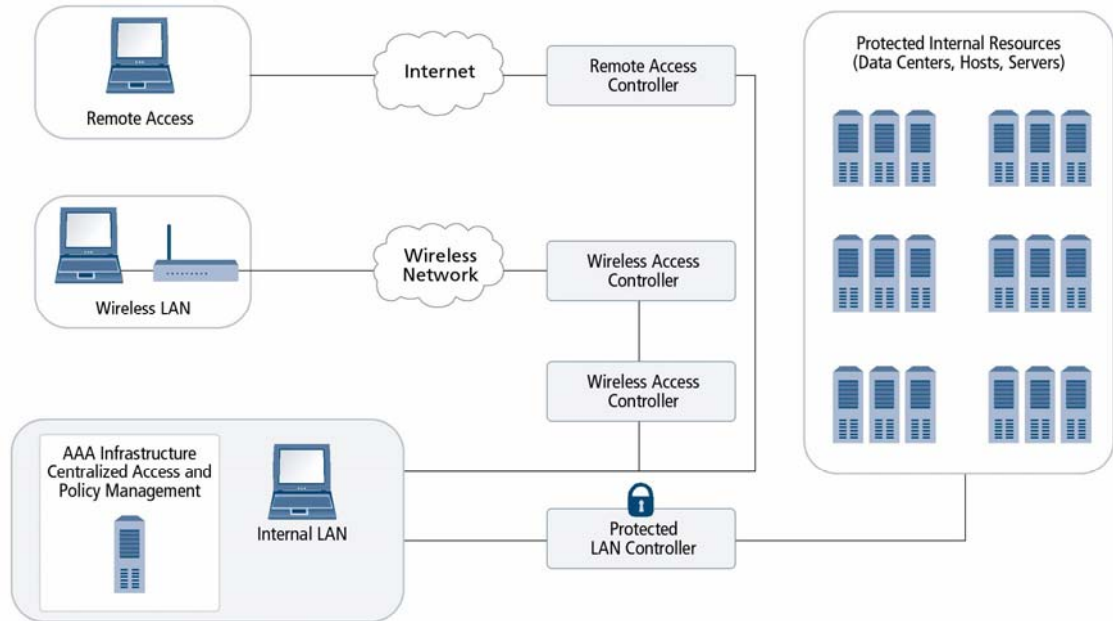
Figure 4: Multi-Access Overview

The same users that normally access the internal LAN during office hours become remote access users when they are outside the office and wireless users when roaming throughout the building away from their desk. Therefore, a single set of access policies applied to the same user would make sense regardless of the type of access method used. Requirements that are common to all access methods include:

- Endpoint security
- Authentication
- Access policy management
- Various service levels, depending on the devices used for accessing (trusted or non-trusted)
- Flexibility and mobility support
- QoS and bandwidth management
- Performance and scalability to support a large number of SSL sessions
- Auditing capabilities

Solution

The F5 Solution

Unified Access

The F5 universal access approach is based on F5's FirePass product, Secure Sockets Layer Virtual Private Network (SSL VPN), which is integrated with F5's BIG-IP® Local Traffic Manager. This integration establishes F5 as the only vendor in the industry that can unify and centralize security and access control for remote, wireless LAN (WLAN), and Local Area Network (LAN) users, while scaling to meet the throughput requirements for unified access. The BIG-IP protects and secures access at the network layer (for example, VLANs), offers high-speed encryption capacity, and technology for real-time packet inspection and traffic monitoring. FirePass provides a scalable yet granular access scheme with SSL VPN with enhanced user authentication, access control, endpoint security, policy management, and auditing support.

F5's universal access approach can be either centralized at a single location or distributed across several geographic locations as described in the following deployment scenarios.

F5's universal access approach supports the following:

- Endpoint security
- User authentication
- Network security
- Unified access policy management
- High performance
- Scalability
- Audit and reporting

The following sections describe each capability.

Endpoint Security

Endpoint security encompasses the security measures that should be checked before a client is granted access to a resource. Endpoint security is a proactive way to make sure that only clients who pass all security policies are allowed access. It is a proactive way to prevent any virus or malware from entering the corporate network.

Security measures could be different for different client device types (such as laptop, PDA, Internet café), and consist of checking the client antivirus software and version to validate personal firewall settings and verify client certificates.

One of the strengths of endpoint security is the breadth of access and security control. Access from any device or location to any application can potentially lead to security exposure. With adaptive client security, the endpoint security mechanisms enable an administrator to define different access-levels and to ensure client and resource protection. For example:

- Kiosk users with the cache cleanup feature can access terminal servers, files, intranet, and email. Once the user logs out, the cache location and the cache contents are automatically cleaned up.

- PDA users can only access email and some web applications, but will not be able to access any other servers.

- Laptop users are provided full network access with support for all client / server applications when they pass all endpoint security rules.

F5's universal access approach can dynamically adapt user policy based on the device used. For example, you can define a different set of policies for mobile devices, Kiosk access, laptop policy, and a default policy. Unified endpoint security enables client integrity checking such as checking system registry settings, presence and absence of any specific processes, OS service packs, verifying the client antivirus software version, and different policy decisions in the presence and absence of client certificates.

If a user does not conform to the security policy, you can define a fallback scenario by redirecting the user to a quarantine network where the user can update their client. The quarantine network might provide software and remedial actions such as update the virus program with the correct software version.

User Authentication

F5's universal access approach supports dynamic authentication methods. Based on the authentication result, you can assign users to different groups to specify additional access to specific resources such as access to specific network segments or servers. F5's Pluggable Authentication Module extensions support multiple authentication schemes, including:

- RADIUS
- Active Directory with Kerberos
- Client Certificate LDAP and OCSP
- Basic Authentication

- LDAP/LDAPs
- TACACS+
- RSA Secure ID

This architecture enables you to dynamically define the role of each user based on user credentials.

F5's universal access approach also provides a single sign-on (SSO) facility that authenticates users to Web servers, network resources, and legacy applications without making any modifications to existing applications by automatically passing on user credentials. SSO options include:

- Form-based authentication
- Basic authentication
- NTLM authentication
- Domain authentication

Network Security

One important requirement of F5's universal access approach is the ability to partition the network into various segments to protect and monitor access from one segment to the other.

At the network level, you can use IP addresses, VLANs, MAC addresses, and packet filtering mechanisms to define practically any combination of network security policy based on any network parameter such as originating or destination VLANs, IP addresses, and protocols. You can refine this security with stricter access rules based on authentication results or application responses. With F5's iRules and the Universal Inspection Engine, you can define custom security policies. iRules, based on the TCL programming language, is a simple yet powerful tool for fine grain control, while the Universal Inspection Engine parses the entire payload of user traffic. This enables you to allow, deny, forward, or drop traffic based on IP address, authentication results, or payload. Peer-to-peer traffic passing through the BIG-IP can also be checked for known signatures, redirected to external virus scanners, or dropped completely.

All of these capabilities can be used to implement LAN segmentation, different zones, such as trusted, public, private, protected, etc. Deciding whether authentication and strong VPN encryption is needed can be defined via the BIG-IP that resides on the border of each segment and zone. This configuration becomes the primary point of demarcation for classifying and defining network segments, identifying various traffic types, and monitoring real-time traffic.

In conjunction with SSL VPN access, you can define separate routing tables and VLANs that are associated with each routing table. LAN segments become a protected resource belonging to a user group. Therefore, users from other groups can be denied access to such LAN segments, yielding to network security at Layer 2 and Layer 3 of the OSI model. For example, a user belonging to a public group cannot be switched or routed to VLANs that are part of a private user group.

F5's universal access approach also provides a series of built-in security features to protect the network from DoS, DDoS, and protocol tampering attacks, including:

- Deny-by-default
- Automatic Defense
- SYN Check
- DoS and Dynamic Reaping
- Connection limits on virtual servers
- Protocol Sanitization
- Packet Filtering
- Resource Cloaking
- Secure Network Translation
- Wireless roaming

- Audit
- Reporting

For information on additional BIG-IP security capabilities, see the white paper at http://www.f5.com/solutions/technology/pdfs/securing_enterprise_wp.pdf.

The FirePass and BIG-IP universal access approach results in a more robust security model based on new services and policies that are easily defined to enhance productivity and lower your TCO.

Unified Access Policy Management

F5's universal access approach uses a single, unified policy to secure access to your network resources at both the network layer and the application layer. This single point of management eases the administrative burden, and coupled with the efficient use of resources, reduces your TCO.

F5 provides different ways to assign user groups and resources based on LAN segmentation, (the source or destination LAN segment where a user resides) or authentication results (for example, user credentials). When authentication is used, you can either statically or dynamically specify group and resource assignment based on responses from authentication servers. For example, if a user authenticates successfully with an Active Directory, the Active Directory can return an attribute that F5 uses to map the user to a specific group.

In addition, FirePass provides the Visual Policy Editor, a unique tool that enables security administrators to graphically define complex security policies to eliminate policy mis-configurations that could result in security holes. The resulting flowchart also gives auditors an easy way to visually audit security policies rather than sorting through complex product configurations.

You can also combine the group management policy with F5's Rate Shaping software add-on module to ensure Quality of Service by reserving bandwidth and prioritizing traffic for critical applications.

Scalability

F5's universal access approach represents a highly scalable model both for the enterprise as well as for Internet Service Providers and Application Service Providers. A single FirePass device can host up to 255 unique URIs so that you can create a unique landing URI for each type of user group. For example, authorized users could land at company.example.com, whereas partners could land at partners.example.com. This virtualization technique can accommodate different user groups while FirePass does the behind-the-scenes mapping to back-end devices.

You can configure F5's universal access approach in different configuration modes:

- Active-passive for increased redundancy
- Clustered for maximum scalability

Clustering up to 10 FirePass devices can guarantee secure remote access for up to 20,000 concurrent users. In a clustered configuration, FirePass eases management by automatically synchronizing FirePass policies. However, you can also manually synchronize configurations, including policy and security rules.

Wireless Roaming

There are two ways wireless users can connect to F5's unified access configuration:

- Users access public resources such as Internet and public printers once they successfully authenticate on an AAA server. The BIG-IP proxies the authentication request and passes

user credentials to the AAA server. Upon successful authentication, the user is granted access to resource networks as defined in the access policies.

- All other access requires a secure connection to the unified access configuration.

Wireless users must authenticate on FirePass to access the internal network. FirePass uses SSL VPN to provide strong authentication and VPN access to comply with the highest security standards.

FirePass SSL connections retain session context with the help of a SSL Session Identifier. SSL renegotiation ensures that the SSL connection and context is re-established even if the IP address changes when roaming or a temporary connection loss occurs. This ensures transparent connections across wireless roaming areas.

Encrypted cookie support maintains the same user authentication context and is useful in conjunction with the single-sign on feature. Therefore, both application session context and user authentication context are retained while wireless users are roaming.

Audit and Reporting

F5's universal access approach provides the following auditing and reporting capabilities:

- Statistics on virtual servers/pools, interfaces, VLANs, Rate Shaping, Packet Filtering, Syslog, and SNMP MIB are logged. Logged information captures user login, session activities, group-level statistics, and a history of user actions. You can also use iRules to log additional information because of F5's deep packet inspection capability.

- A graphical representation of usage statistics, such as CPU load, interface statistics, and user accounting gives you a quick overview of system status.

Deployment Scenarios

The following examples describe how F5's universal access approach manages the security policies for various network topologies. In each case, F5's universal access approach protects your enterprise at the network level as well as the application level, using a highly scalable and flexible architecture. This unified approach delivers a lower TCO for any size deployment.

Example 1: A Campus Network Deployment

A campus network is configured as follows:

- An internal LAN is partitioned into trusted, non-trusted, protected, public, and quarantine zones.

- Access to the public zone does not require any specific authentication or other verification for users accessing from the trusted and non-trusted zones.

- Users from the trusted LAN can access the protected LAN.

- Non-trusted, wireless, and remote access users are authenticated before they can access any protected resources.

Figure 5 shows the network topology.
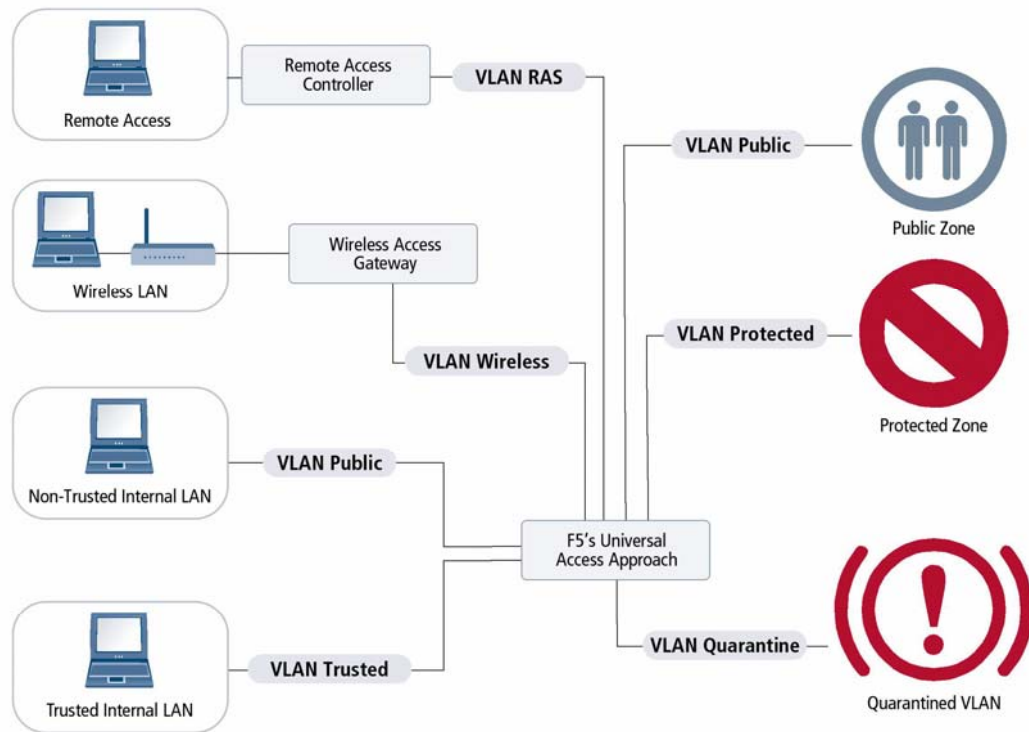
Figure 5: The F5 Universal Access Approach in a Campus Network

The F5 universal access approach lowers your TCO for this type of deployment:

- A central point of demarcation where each zone is separated.

- With the help of separate VLANs, each zone is distinguished. Highly flexible access rules, security policies, authentication schemes, and secure NAT translations can all be defined separately for each VLAN and/or services within each VLAN.

- A single authentication server defines users and their roles to accommodate different access methods.

- You can extend the F5 universal access approach to inspect real-time traffic at multi-gigabit throughput capacity to support Triple Play (data, video, voice), which is often prevalent on campus networks. Deep packet inspection, Universal Inspection Engine, and iRules give you unprecedented control to manage and monitor real-time traffic at a multi-gigabit per second throughput rate.

- Along with SSL offloading and load balancing, the F5 solution supports multi-gigabit encryption throughput and handles several hundreds of thousands of SSL transactions per second.

Example 2: Multi-Site Enterprise Deployment

Consider a large enterprise with several geographically-distributed campus networks. Each campus could represent a small or large campus network with the same requirements described in the previous use case. Wireless and internal LANs should be supported with flexible security and access policies for each access method. Internal LAN segmentation

should also be possible. Remote access service is normally centralized in one or two locations (headquarters).

Figure 6 shows the configuration of several geographically-distributed campus networks.
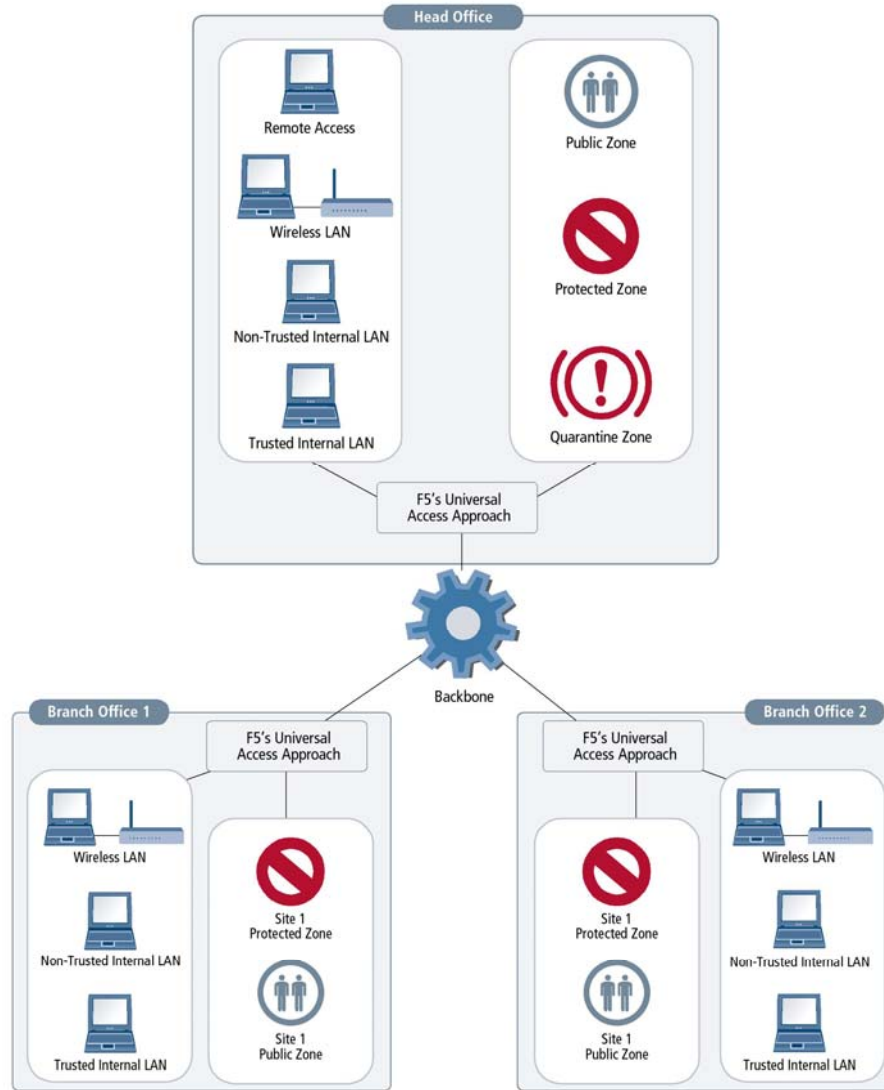
Figure 6: The F5 Solution in an Enterprise Network

The F5 universal access approach lowers your TCO for this type of deployment:

▪ The F5 solution resides at the head office for all access domains, including remote access.

▪ Each branch office has a dedicated F5 universal access approach to partition the campus area into various zones.

▪ The head office contains the definitions of user's access rights and user groups that specify security policies and access rules.

▪ With clustering, F5 configurations can be automatically synchronized.

▪ Branch offices can have protected and public LAN partitioning.

■ Resources to the head office can be accessed from branch offices.

Conclusion

The F5 integration of the BIG-IP and FirePass products is the only solution that provides a unified approach for securing and managing infrastructure and access for internal LAN, wireless, and remote access methods. The F5 universal access approach benefits organizations by providing:

■ A single security policy base for wireless, remote, and internal LAN access.

■ Unified authentication method for all enterprise users and guest users, regardless of access method.

■ A single administration method that results in highly cost-effective management.

■ Secure access to all enterprise critical resources.

■ SSL VPN with encryption (AES, 3DES) and Transparent Proxy support for a large number of users, requiring multi-Gbps throughput.

■ All SSL VPN capabilities for selective service, enabling access to pre-defined applications for specific user groups, portal access, and Webifyers.

■ Enhanced security against network, DoS, DDoS, and protocol tampering attacks with DoS Reaping, Rate Limiting, Packet Filtering, Deep Packet Inspection, and Secure Network Translation capabilities.

■ Unified and visual access policy management, auditing, and monitoring to lower your TCO.

■ High availability and fast delivery for applications.

■ Considerable savings in Capex/Opex expenditures.

■ Lower TCO for both campus and geographically-distributed environments.

■ Full virtualization capability to reduce Capex/Opex for managed service deployments.

About F5

F5 Networks is the global leader in Application Delivery Networking. F5 provides solutions that make applications secure, fast and available for everyone, helping organizations get the most out of their investment. By adding intelligence and manageability into the network to offload applications, F5 optimizes applications and allows them to work faster and consume fewer resources. F5's extensible architecture intelligently integrates application optimization, protection for the application and the network, and delivers application reliability - all on one universal platform. Over 10,000 organizations and service providers worldwide trust F5 to keep their applications running. The company is headquartered in Seattle, Washington with offices worldwide. For more information, go to www.f5.com.