



Advanced Topics

The following sections provide detailed information about topics that pertain to the BBSM system:

- [VLANs, page 2-1](#)
- [VPN, page 2-5](#)
- [Web Proxies, page 2-7](#)
- [Web Servers, page 2-8](#)

Many features are described in detail in the overview of the configuration chapters. For example, the RADIUS overview and attributes are presented in [Chapter 14, “Configuring RADIUS Billing.”](#)

VLANs

As of BBSM 5.3, the BBSM server supports two VLANs on the internal interface. This feature enables you to put the management traffic and equipment on one VLAN and end-user clients on the other VLAN.

What is a VLAN?

A switched network can be logically segmented into virtual local area networks (VLANs) by functions, project teams, or applications on a physical or geographical basis. For example, a workgroup team can connect all of their workstations and servers to the same VLAN regardless of their physical network connections or intermingling with devices for other teams. VLANs can be reconfigured through software rather than physically unplugging and moving devices or wires.

A VLAN can be thought of as a broadcast domain that exists within a defined set of switches. It consists of end systems, either hosts or network equipment such as bridges and routers, connected by one bridging domain. This bridging domain is supported on various network equipment, such as LAN switches that operate bridging protocols between them with a separate group for each VLAN.

VLANs are created to provide the segmentation services that routers traditionally provide in LAN configurations. Routers in VLAN topologies provide broadcast filtering, security, address summarization, and traffic-flow management. None of the network devices within the defined group bridge any frames, not even broadcast frames, between two VLANs.

Several key issues must be considered when designing and building switched LAN networks.


- LAN segmentation
- Security
- Broadcast control

- Performance
- Network management
- Communication between VLANs

VLANs are extended into the wireless realm by adding IEEE 802.1Q tag awareness to the access point. The access point wirelessly transmits frames destined for WLAN clients on different VLANs on different SSIDs with different Wired Equivalent Privacy (WEP) keys. (WEP is an 802.11 security protocol for wireless network WEP keys.) The only clients that can receive and process packets are those with the correct WEP keys. Conversely, packets coming from a client associated with a certain VLAN are 802.1Q tagged before they are forwarded onto the wired network.

Figure 2-1 shows the difference between traditional physical LAN segmentation and logical VLAN segmentation with wireless devices connected.

Figure 2-1 LAN and VLAN Segmentation with Wireless Components



Incorporating Wireless Devices into VLANs

A wireless LAN (WLAN) is generally deployed in an enterprise campus or branch office for increased efficiency and flexibility. WLANs are one of the most effective methods for connecting to an enterprise network. With Cisco IOS Release 12.01T, you can configure your wireless devices to operate in a VLAN.

The basic wireless components of a VLAN consist of an access point and a set of clients associated with it using wireless technology. The access point is connected physically through a trunk port to the network switch on which the VLAN is configured. The physical connection to the VLAN switch is through the Ethernet port of the access point.

In fundamental terms, the key to configuring an access point to connect to a specific VLAN is to configure an SSID to map to that VLAN. Because VLANs are identified by a VLAN ID, if an SSID on an access point is configured to map to a specific VLAN ID, a connection to the VLAN is established, and clients can access the VLAN through the access point. The VLAN processes data to and from the clients the same way that it processes data to and from wired connections. The fact that the client is wireless has no impact on the VLAN.

The VLAN feature enables users to deploy access points with greater efficiency and flexibility. For example, one access point can handle the specific requirements of multiple users having widely varied network access options and permissions. Without VLAN capability, multiple access points (one for each VLAN) would have to be used to serve user classes based on the access options and permissions that they were assigned.

The following simplified example shows how access points can be used effectively in a VLAN environment on a college campus. In this example, three levels of access are available through VLANs configured on the physical network:

- Student access—Lowest level of access; ability to access school's intranet, obtain class schedules and grades, make appointments, and perform other student-related activities
- Faculty access—Medium level of access; ability to access internal files, read to and write from student databases, access the intranet and Internet, and access internal information such as human resources and payroll information
- Management access—Highest level of access; ability to access all internal drives and files, and perform management activities

In this scenario, a minimum of three VLANs are required—one for each of the above access levels. Because the access point can handle up to 16 service set identifiers (SSIDs), the following basic design can be used. (See [Table 2-1](#).)

Table 2-1 Access Level SSID and VLAN Assignment

Access Level	SSID	VLAN ID
Student	Student	01
Faculty	Faculty	02
Management	Management	03

Using this design, setting up the clients is based on the access level that each user requires. [Figure 2-2](#) shows a typical network diagram using this design.

Or, you can use Switch Discovery Wizard to specify the management and client VLANs when adding switches or access points to BBSM. (CMTSs must be configured using WEBconfig.) Refer to the following documentation about VLANs:

- [Chapter 4, “Configuring Dual VLANs”](#)
- [“Running the Switch Discovery Wizard”](#)

**Note**

Access points are not automatically configured with a default VLAN. If you add or remove a default VLAN from an access point, you must reconfigure your port settings using WEBconfig. Refer to the [“Configuring Access Points” section on page 12-2](#).

VPN

One advantage of the public and private IP addressing feature is that it eliminates the IP Security (IPSec) VPN problems that occur with the use of private IP addressing. All VPN clients are supported. For an overview of public and private addressing, refer to the [“Public and Private IP Addresses \(Multinets\)” section on page 9-2](#).

If the VPN client requires routable addresses, the end user can choose this option from the Connect page. The client has to be configured for DHCP. If the VPN client can work with network address translation (NAT), BBSM can use either routable or nonroutable addresses and the client can be DHCP or static. Cisco VPN 3000 works with NAT.

**Note**

When you finish the session, close the VPN client before disconnecting.

BBSM is VPN Neutral

Technically, BBSM does not support or reject VPN clients. The topology of the in-building network is the parameter that affects the use of VPN clients:

- If the in-building network is deployed with all public IP addresses, the router does not need to perform any NAT or port address translation (PAT) and no problems occur when VPN is used. However, this situation would be unusual and expensive if the operator did not already have a large block of IP addresses.
- More often, the in-building network is configured with private IP addresses, which causes the router to perform PAT to connect various sessions to the Internet. This router PAT activity interferes with certain types of non-Cisco VPN clients. Cisco routers that have the Cisco IOS Release 12.1.4T or later firmware patch installed also support Microsoft PPTP VPNs regardless of PAT activity.

You may experience problems if you use another VPN vendor’s product that uses IPSec with authentication headers (AHs). Although BBSM may support these VPN products, they have not been tested with BBSM.

Public-Private Addressing with VPN

Public-private IP addressing eliminates the following problems that relate to VPN:

- NAT after IPSec

Applying NAT after IPSec encryption to hide IP addresses provides no benefit because the actual IP addresses of the transport devices are hidden by the encryption. Only the public IP addresses of the IPSec peers are visible, and hiding the public IP addresses provides no additional security. NAT is applied after IPSec encapsulation when IP addresses are being conserved. This conservation application is common in hotels, cable or DSL residential deployments, and enterprise networks. In these cases, using NAT after IPSec encryption can interfere with establishing the IPSec tunnel (depending on the type of NAT).

When IPSec uses the authentication header (AH) mode for signature integrity, one-to-one NAT can invalidate the signature checksum. Because the signature checksum is partially derived from the AH packet's IP header contents, the signature checksum becomes invalid when the IP header changes. The packet then appears to have been modified in transit and is discarded when the remote peer receives it.

However, when IPSec incorporates the encapsulating security payload (ESP), devices can send packets successfully over the VPN, even with one-to-one NAT after encapsulation. This scenario is possible because ESP does not use the IP header contents to validate packet integrity. In cases in which many-to-one NAT occurs (also known as PAT), the IP address and the source Internet Key Exchange (IKE) port (normally User Datagram Protocol [UDP] port 500) changes. Some VPN devices do not support IKE requests sourced on ports other than UDP 500, and some devices performing many-to-one NAT do not handle ESP or AH correctly. Remember that ESP and AH are higher layer protocols than IP and do not use ports.

Because many-to-one NAT is common in environments that include remote access clients, a special mechanism called *NAT transparency* is used to overcome these NAT problems. NAT transparency re-encapsulates the IKE and ESP packets into another transport layer protocol, such as UDP or TCP, that address-translating devices can translate correctly. This mechanism also allows the client to bypass access control in the network that allows TCP or UDP but blocks encrypted traffic. This feature does not affect transport security in any way. NAT transparency takes packets that IPSec has already secured and encapsulates them again in TCP or UDP.

The Internet Engineering Task Force (IETF) IPSec working group has not endorsed IPSec over UDP. Although the problems of using IPSec through a device performing many-to-one translations are discussed frequently, no standard approach to solving these problems exists at this time. Therefore, Cisco is implementing support for IPSec over UDP as a short-term solution until a standards-based workaround is available. As a result, IPSec over UDP is compatible only with the Cisco VPN 3000 clients and concentrators.

- Using Cisco VPN 3000 Concentrators and the Cisco VPN 3000 Client with IPSec over UDP

No user intervention is required to perform IPSec over UDP with Cisco VPN 3000 concentrators. An administrator can centrally control when a user should use IPSec or IPSec over UDP through a group-user configuration policy within the Cisco VPN 3000 series concentrator product. If IPSec over UDP is enabled for a particular user, the client and the Cisco VPN 3000 concentrator automatically negotiate using IKE. The setup includes basic information such as the UDP port number and the IPSec-over-UDP requirement. Data will be transmitted successfully through NAT.

- IPSec over UDP through a NAT firewall performing filtering

Although any device that can perform outgoing NAT on UDP packets allows IPSec or UDP to be used, the network administrator needs to ensure that no specific rules exist that block access to the UDP port that the remote network's administrator defines. If the administrator is not blocking the packets, the user can successfully access the remote network in a secure manner.

- NAT before IPSec

When two sites are connected through IPSec, the tunnel does not establish whether or not any of the site network address ranges overlap because the VPN termination devices cannot determine the site to forward the packets to. Using NAT before IPSec overcomes this restriction by translating one set

of the overlapping networks into a unique network address range that does not interfere with establishing the IPsec tunnel. The NAT application is recommended only in this scenario. However, some protocols embed IP addresses in packet data segments. As a general practice, verify that a protocol-aware device carries out the NAT when addresses are translated, not only in the IP header but also in the packet's data segment.

If the packet address was not correctly translated because of embedded addresses before being transmitted, the remote application does not receive the correct IP address embedded in the data segment when it receives the packet and probably does not function properly. Many remote access VPN clients now support a virtual address that the VPN concentrator assigns. Devices at the remote site can connect to the remote access client using this virtual address by one-to-one addressing that translates all transmitted and received packets. If the VPN client does not address-translate packets correctly or if a new application arrives that is not yet supported, the application may not function.

In summary, follow these guidelines:

- Use address ranges at your sites and remote access VPN client virtual address pools that do not overlap with the addresses of other devices that you connect through IPsec. If this is not possible, use NAT in this scenario to allow only for connectivity. (This option is not configured on the BBSM server.)
- Do not hide the public IP addresses of the VPN devices because hiding them does not add security and can cause connectivity problems.
- When you believe that NAT is involved and a remote access client cannot establish a tunnel successfully or send packets over an established tunnel, consider enabling the NAT transparency mode. Do not try to use the NAT transparency mode to resolve connection problems associated with client applications that are not NAT friendly. (The NAT transparency mode option is not configured on the BBSM server. It is not related to the BBSM transparent proxy feature.)

Web Proxies

BBSM can function as a web proxy server in several different ways:

- As shipped, BBSM works as a proxy server for end users whose browsers are configured for a web proxy server. (BBSM uses Microsoft Industry Standard Architecture [ISA] for this function.) In other words, BBSM selectively proxies only the users that require it (those with proxy settings).
- BBSM also has a transparent proxy feature that can be enabled:
 - Enabling transparent proxy forces all users, regardless of browser settings, through the BBSM proxy engine. Refer to [Chapter 9, “Changing the Internal Network IP Address Ranges.”](#)
 - If your router is configured properly, transparent proxy does not need to be enabled for BBSM to operate correctly. Problems occur when the client is using a private IP address and the router is not configured with a static route to the BBSM internal network. In this case, the router cannot route packets to the client and the client cannot browse the Internet.
 - With web proxy enabled globally, you can use Microsoft ISA to monitor the Internet sites that the web browser has visited. (You must turn on ISA logging manually.) The resulting data appears in the ISA log files. This feature is used to gather statistics for marketing purposes and is not recommended for most deployments because using it affects performance.

Refer to [Chapter 9, “Changing the Internal Network IP Address Ranges,”](#) to enable transparent proxy in WEBconfig and log subsequent entries.

**Note**

BBSM has no provision for directing client traffic through a remote proxy server. BBSM supports use of the proxy server only on the BBSM server. However, you can use a Cisco router between BBSM and the Internet to intercept packets using the Web Cache Communication Protocol (WCCP) and then send the requests to a content engine. The content engine can be configured to use the proxy for outbound data. (Cisco IOS supports both versions of WCCP.)

**Caution**

Do not modify the ISA settings to ensure that the BBSM server supports web proxy.

Web Servers

A web server can be hosted on a client machine inside the BBSM network. Although this situation is not common, you can accomplish it by using the following information. Hosting a web server on an internal BBSM network client allows unrestricted access from the external side.

Check the Transparent Proxy check box in WEBconfig and do the following:

- To keep the billing intact, reserve a DHCP address for the client and then set up one-to-one static mapping to that address (if you are using private addresses). (The client is still required to connect to BBSM.) After the client session is enabled, the client can host a web server. The client should not run services such as DHCP and DNS to avoid conflicts with these services running on BBSM.
- You can configure the web server with a BBSM management range IP address if these addresses are not needed for billing. Perform one-to-one NAT of the address in the router (if you are using private addresses). Internet access to the web server is then available.