Cisco WAAS Mobile Network Design Guide

Software Version 3.4
July 2008

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Text Part Number: OL-17347-01

Contents

List of Tables

List of Figures

About this Document

Intended Audience

This guide provides network architects with best practices for integrating Cisco WAAS Mobile with various distributed network topologies and usage scenarios.

Pre-Requisites

This guide uses several terms that are specific to the Cisco WAAS Mobile administrator interface, WAAS Mobile Manager. Familiarity with WAAS Mobile Manager or access to the *Cisco WAAS Mobile Administration Guide* is a pre-requisite.

Related Documents

In addition to this Network Design Guide, the following documents are also available:

- *Cisco WAAS Mobile Integration Guide*—Provides information required by network engineers as they consider the deployment of the WAAS Mobile server, covering discussion topics such as firewalls, network topology, authentication and accounting.
- *Cisco WAAS Mobile User Guide*—A guide for the WAAS Mobile end user. This complements the on-line help system and provides a reference for offline study.
- *Cisco WAAS Mobile Administration Guide* – Provides system administrators with detailed installation, configuration, monitoring, and support instructions for WAAS Mobile clients and servers.
- *Cisco WAAS Mobile Release Notes*—Release-specific information regarding features added, changed, and removed as well as known issues and issues fixed in the release.

Chapter 1 Introduction to Cisco WAAS Mobile Network Design

Cisco WAAS Mobile is a system for selectively and dynamically optimizing the TCP connections that applications make to content servers. Optimization decisions are made on the end-user machine using all information available at the time the connection is made. Dynamic decisions make WAAS Mobile effective for mobile and static users operating in a complex infrastructure of data centers and branch offices.

WAAS Mobile enterprise application scenarios are broadly defined by:

- the location of end-users, their mobility and the location of VPN concentrators used for remote access
- the location of branch offices and data centers needing acceleration
- the position and quality of data links joining branch offices and data centers

The process of designing a distributed WAAS Mobile installation is as follows and provides the documentation outline for this *Cisco WAAS Mobile Network Design Guide*:

1. Locate client sites and content servers for acceleration.
2. Select WAAS Mobile server sites.
3. Determine network acceleration policy. This determines the relationships between WAAS Mobile clients, WAAS Mobile servers and content servers. It includes all the static and dynamic acceleration and bypass decisions. There may be subsets of clients that require different network acceleration policy.
4. Size WAAS Mobile server sites.
5. Select a load balancing method at multi-server sites.
6. For SSL VPN links that are being optimized with WAAS Mobile, choose whether to optimize application protocols or to optimize transport. This can affect the placement of WAAS Mobile servers and the configuration of WAAS Mobile clients.
7. Choose policies for Internet access. These policies can be required to allow Internet access with split-tunnel and full-tunnel VPNs.
8. Design network addressing of servers in server farms. This is needed for client configuration.
9. Designate one of the WAAS Mobile servers as the controller server and the remaining WAAS Mobile servers as worker servers.

Chapter 2 Select Sites

Site Selection Principles

The aspect of WAAS Mobile that most affects the high-level network design is that each WAAS Mobile client connects to one WAAS Mobile server at a time, the choice being made by the client.

The aim of site selection for WAAS Mobile servers and server farms is that every client can access all the data it needs in such a way that network traffic traverses only:

- LAN-like links; for example, high-speed, very low latency links between offices in the same city
- WAN links accelerated by a WAN optimizing appliance
- remote-access links and WAN links accelerated by WAAS Mobile

The aim of site selection is to optimize whatever remote-access and WAN links are necessary to provide the data access required. In achieving this, the following rules apply:

- WAAS Mobile accelerated traffic should not cross links optimized by a WAN appliance. This tends to push WAAS Mobile servers towards outlying WAN optimizing appliance locations.
- WAAS Mobile content server traffic should, ideally, cross only very low latency links; that is, LAN-like links and WAN links optimized by an appliance. This limits the content servers that can be optimized with a particular WAAS Mobile server.

Site Selection Procedure

The aim of the site selection procedure is to find locations for WAAS Mobile servers that meet the requirements outlined above. The procedure, in brief, is: *The WAAS Mobile server should be placed at the same site as the content server end of each problematic link.* In more detail:

1. Identify problematic client locations and content servers.

   These are the network locations where users report problems accessing the data they need, and the content servers about which they complain. Often these are remote access users in particular geographical or network regions. They may also be in small or large branch offices. Client locations where users do not have problems accessing the data they need are not included in this list. Sometimes only a subset of data services has performance problems. Data services that perform adequately may do so because they are close to the clients, because the applications are WAN-friendly, or because the services are not business-critical.

2. Identify problematic links that terminate at problematic client locations.

   These are the links that are used to access the problematic content servers. They may be wireless remote-access links over the public Internet, secured by software or hardware VPNs. They may also be private WAN links (for example, MPLS or Frame Relay). In practice, it would be unusual for more than one problematic link to terminate at a problematic client location.

3. The WAAS Mobile server should be placed at the end of each problematic link distant from its client location.

With this placement, client access is optimized to every problematic content server crossing only LAN-like links and WAN links optimized by an appliance. Where a content server is reachable from a problematic client location but cannot be accelerated, WAAS Mobile clients at that location must be configured to bypass traffic to it.

Chapter 3 Determine Network Acceleration Policy

Network acceleration policy is used to determine:

- client locations where WAAS Mobile is not active – for example, central LANs
- the WAAS Mobile servers used by clients
- which content servers are accelerated, possibly depending on client location
- which applications (identified by well-known port) are excluded from acceleration

These settings are implemented within client and server configuration and in some cases can depend on the server site to which the client is connected. There are two broad classes of policy:

- Static policy
- Dynamic policy

Static policy takes the form of:

- fixed client subnets on which the WAAS Mobile client is not allowed to connect (WAAS Mobile Manager: Access Control)
- fixed client subnets that determine which WAAS Mobile server to use (WAAS Mobile Manager: Farm Selection/Client IP Map)
- fixed subnets, host names and ports that may or may not be accelerated (WAAS Mobile Manager: Accelerated Networks and Exclusion Lists)
- which existing connections to break and which to preserve when the WAAS Mobile client starts (WAAS Mobile Manager: Auto-reset connection)

Dynamic policy is based on latency, and can be enabled to determine:

- which server site to use, independent of client subnet (WAAS Mobile Manager: Farm Selection/Latency)
- whether to bypass the chosen server, because it is too close to the client (WAAS Mobile Manager: High-speed bypass)
- whether to bypass individual content servers (WAAS Mobile Manager: Latency-based bypass)

The tables below summarize how different aspects of network acceleration policy can be applied.

Table 1. Summary of Static Network Acceleration Policies

Static Policy	Granularity	Usage
Access Control	Per server	Prevent WAAS Mobile client(s) from connecting to the WAAS Mobile server on a LAN. Prevent WAAS Mobile server(s) from accepting connections from distant LANs.
Farm Selection / Client IP Map	Per server site group	Determine which server sites to use, based on client IP map.
Accelerated Networks	Per client distribution	White or black list of content server subnets for acceleration. Use for configuring Internet access for split-tunnel VPNs, configuring transport optimization for SSL VPNs, and avoiding optimization of too-distant content servers
Exclusion Lists – IP addresses	Per client distribution	Black list of bypassed content servers for configuring application optimization for SSL VPNs and for avoiding optimizing miscellaneous servers
Exclusion Lists – ports	Per client distribution	Avoid proxying applications identified by well-known port, independent of originating process. Use for avoiding applications that cannot be proxied and to exclude any application that uses a successful TCP connection to a well-known port to determine whether it is connected to the correct server.
Auto-reset connection	Per client distribution	Use to dynamically break SSL VPN connections.

Table 2. Summary of Dynamic Network Acceleration Policies

Dynamic Policy	Granularity	Usage
Farm Selection / Latency	Per server site group	Choose server farm based on network tests conducted by the client. Where applicable, this is a lower-maintenance alternative to selection by client IP map.
High-speed bypass	Per client distribution	Prevent WAAS Mobile client(s) from connecting to the WAAS Mobile server on a LAN. This is applied *after* other server selection mechanisms.
Latency-based bypass	Per client distribution	Where applicable, a low-maintenance alternative to using Accelerated Networks.

Notes:

1. The terms in the left-hand column are consistent with the WAAS Mobile Manager interface.
2. The term *server site group* refers to a set of server farms (a farm may consist of a single server) configured on a single controller machine.
3. The term *client distribution* refers to the client configurations that are created in the Client Configuration section of WAAS Mobile Manager. All client settings are part of a client distribution identified by its *distribution name*.
4. *Per server* items can be configured differently on individual servers. The policies are specified through the Server Configuration pages of WAAS Mobile Manager and are applied on the server.
5. *Per server site group* items are part of advanced server selection and are made on the controller machine. These items are downloaded whenever a client connects to any server in the site group, and all clients connecting to servers in the site group get the same values. The policies are specified through the Server Configuration pages of WAAS Mobile Manager and are applied on the client.
6. *Per client distribution items* can be configured differently for different client distributions. Different servers can have different configuration details for the same client distribution name. These policy settings allow clients connected to different WAAS Mobile servers to behave differently. The policies are specified through the Client Configuration pages of WAAS Mobile Manager and are applied on the client.

Chapter 4 Select Load Balancing Method

At sites where multiple servers are required, decide whether to use the built-in load balancing capability of WAAS Mobile, or an external load-balancing appliance.

The built-in load-balancing capability makes optimal use of the delta compression capabilities of WAAS Mobile by ensuring that all traffic associated with an active session is sent to the same server for the duration of that session. When using the built-in load balancer, it is recommended that "Random Prioritization" mode be selected. In this mode, the client will first attempt to connect to the server to which it was previously connected. If unable to connect, it will then randomly select another server in the farm.

Chapter 5 Select SSL VPN Interoperation Method

Figure 1 shows a socket interception SSL VPN in schematic form. The arrangement is similar to WAAS Mobile, in that there is interception on the client, and the intercepted traffic is proxied to a server. The difference is that the emphasis is on security. The traffic between SSL VPN client and server is SSL over TCP. The SSL VPN server is hardened and is in the DMZ, so that access to content servers is secure.

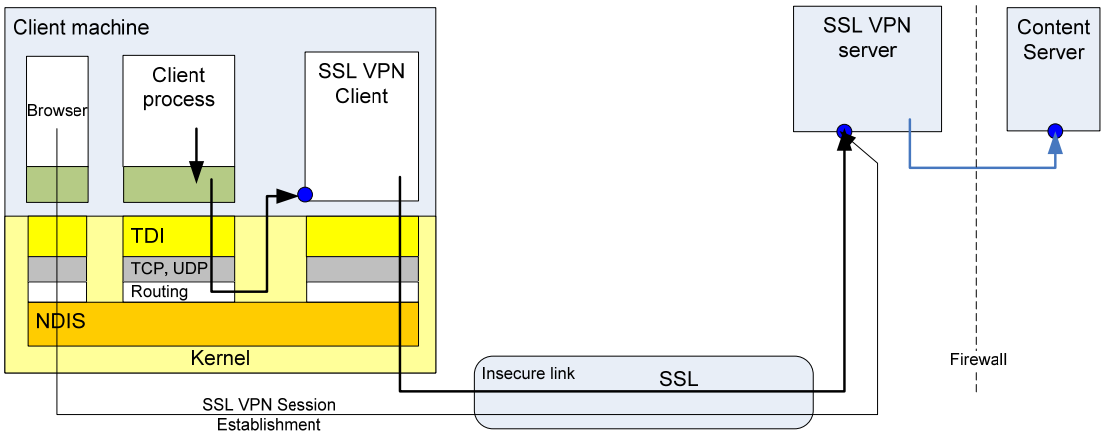

Figure 1. Socket Interception SSL VPN Operation

In the particular case illustrated, the SSL VPN is intercepting socket traffic using a DLL or hook attached to the client process. The TCP connection made by the client process is terminated and its data transferred to the SSL VPN client over another connection made to the loop back adapter on the machine. There are many possible variations.

In the figure, the browser initiates the SSL VPN session that defines the secure connection and traffic is intercepted by the SSL VPN client. Network acceleration policies are applied to WAAS Mobile to enable it to work with the session establishment and the SSL VPN client.

One of the key differences between the way an SSL VPN handles network traffic and the way WAAS Mobile handles traffic is that an SSL VPN is provided to enable access to a secure resource, and access is denied unless an SSL VPN session can be established. In contrast, WAAS Mobile is designed to be transparent. If WAAS Mobile is not running, network access is unimpeded.

Where a remote access client uses both an SSL VPN and WAAS Mobile, a decision must be made whether to use WAAS Mobile to:

- optimize application round-trips, or
- optimize the transport across the remote access link

When optimizing application round-trips, WAAS Mobile intercepts traffic from client processes, while the SSL VPN intercepts the WAAS Mobile optimized traffic and provides transport across the link. When optimizing transport, the SSL VPN intercepts traffic from client processes, while WAAS Mobile intercepts the SSL VPN client traffic and transports it across the link. Each mode of operation has advantages.

For security reasons, application round-trip optimization should be used unless the quality of the link is very poor. Using WAAS Mobile to provide the link transport is less secure because it requires that the WAAS Mobile server be located with the SSL VPN server, usually in the DMZ of a data center, and it is, therefore, accessible from the public Internet.

However, in some cases, application round-trip optimization is not feasible. For example,

- Clientless SSL VPNs use a browser plug-in. The traffic between the browser and its plug-in cannot be intercepted by the acceleration system.
- If the SSL VPN uses a client, it must intercept UDP traffic from the WAAS Mobile client. Not all SSL VPNs can do this. These capabilities have been analyzed for specific VPNs and are available as Deployment Guides or Field Notices.

In such cases transport optimization may be used. Transport optimization is usually available because SSL VPN clients almost always operate in user mode and use sockets-based communication with SSL. The acceleration system can intercept such traffic using the SSL proxy and provide generic compression and transport optimization.

Optimizing Application Round-trips for SSL VPNs

Figure 2 illustrates network traffic flows when optimizing application round-trips for a socket interception SSL VPN. The figure shows:

- SSL VPN session establishment, not proxied
- client process traffic intercepted by the WAAS Mobile client and then the WAAS Mobile client traffic, in its turn, intercepted by the SSL VPN client
- encrypted SSL traffic across the link
- the SSL VPN server passing decrypted traffic to the WAAS Mobile server
- the WAAS Mobile server passing data to the content server

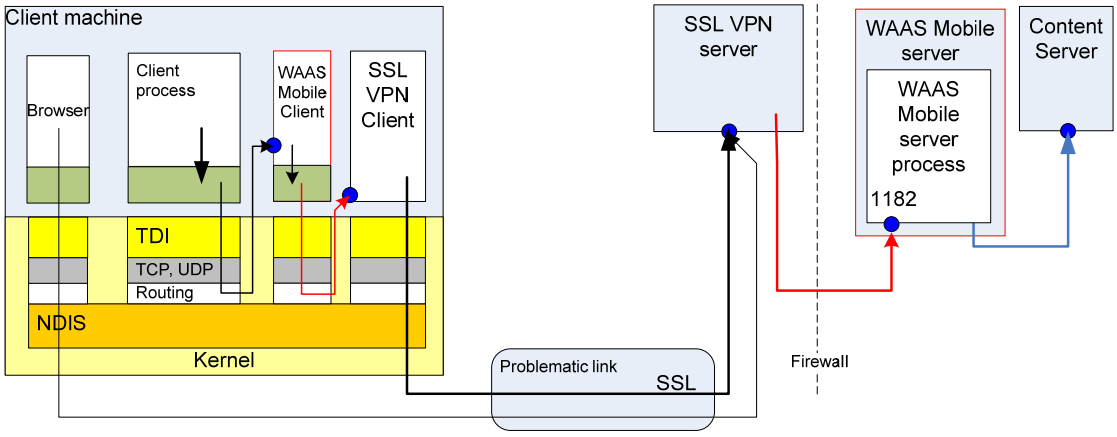

Figure 2. Optimizing Application Round-trips for SSL VPN

In this case the SSL VPN tunnel is placed *between* the WAAS Mobile client and server. The client intercepts the application traffic and, in its turn, is intercepted and proxied by the SSL VPN. Note that the WAAS Mobile server may be in a secure location, as indicated by the location of the firewall. In this arrangement, it is common for the WAAS Mobile client to be unable to contact the WAAS Mobile server until the SSL VPN is established.

Active Mode FTP Issue

When proxying active mode FTP, the WAAS Mobile client makes a connection to a port listening on the client machine in order to set up the FTP data channel. A socket interception SSL VPN client would intercept this traffic unless specifically configured to avoid it. Commonly, SSL VPNs intercept traffic based on destination network, and only intercept traffic bound for remote machines, so this might not require additional configuration, depending on the specific SSL VPN.

Network Acceleration Policy Settings

A common arrangement is to create an SSL VPN session when the user logs in using a web browser that is directed at a URL on the SSL VPN server. Up to the point at which the SSL VPN tunnel is established, the WAAS Mobile client cannot contact the WAAS Mobile server. The login connection therefore goes directly to the SSL VPN server. Once the connection is established, WAAS Mobile can connect.

The SSL VPN client will periodically verify the status of the session by contacting the SSL VPN server. Such connections *must not be proxied* by WAAS Mobile. If proxied, the SSL VPN would see the request come from the WAAS Mobile server rather than from the client. The difference in network addresses involved would usually cause the original SSL VPN session to be shut down.

To summarize:

1. WAAS Mobile must intercept all applications except the SSL VPN client.
2. The SSL VPN client must only intercept traffic from the WAAS Mobile client.
3. The SSL VPN client must intercept only traffic from the WAAS Mobile client that is directed at the WAAS Mobile server.
4. The SSL VPN session establishment and verification traffic must not be proxied by WAAS Mobile. This can be implemented in all known cases by using network acceleration policy to exclude the SSL VPN server host from being accelerated.

Optimizing Transport for Clientless SSL VPNs

Figure 3 illustrates network traffic flows when optimizing transport for a clientless SSL VPN. The figure shows:

- SSL VPN session establishment, not proxied
- client process SSL traffic intercepted by the WAAS Mobile client, which decrypts the traffic before compressing and encrypting it for transmission across the link
- encrypted WAAS Mobile traffic across the link
- the WAAS Mobile server passing SSL traffic to the SSL VPN server
- the SSL VPN server passing data to the content server

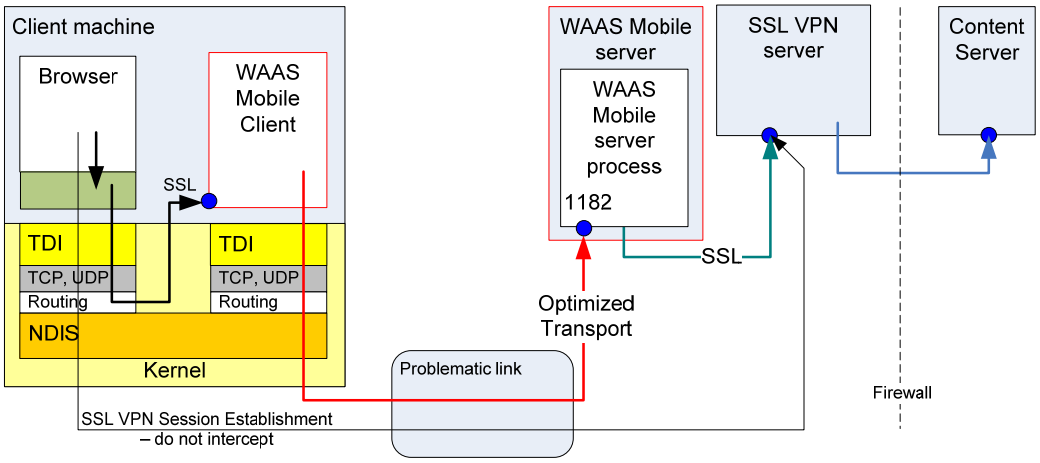

Figure 3. Optimizing transport for a clientless SSL VPN

In this case, the accelerated connection is placed in the only possible location—between the client process and the SSL VPN server. The WAAS Mobile client decrypts the SSL VPN traffic and the WAAS Mobile server encrypts the traffic for delivery to the SSL VPN server. This arrangement requires that the WAAS Mobile server be in the DMZ with the SSL VPN server, as indicated by the location of the firewall. This affects the network addressing of the server and, therefore, both client and server configurations.

Network Acceleration Policy Settings

1. The WAAS Mobile client intercepts traffic from the client process that hosts the SSL VPN application, and from no other process.
2. Only traffic directed to the SSL VPN server (by the client process) is accelerated.
3. If necessary, configure the SSL VPN server to allow traffic to come both directly from the client, and from the WAAS Mobile server.
4. The SSL VPN session establishment and verification traffic must not be proxied by WAAS Mobile. This can be implemented in all known cases by using network acceleration policy to exclude the SSL VPN server host from being accelerated.

Optimizing Transport for Socket Interception SSL VPNs

Figure 4 illustrates network traffic flows when optimizing transport for a socket interception SSL VPN. The figure shows:

- SSL VPN session establishment, not proxied
- client process SSL traffic intercepted by the SSL VPN
- SSL VPN traffic intercepted by the WAAS Mobile client, which decrypts the traffic from the SSL VPN client before compressing and encrypting it for transmission across the link
- encrypted WAAS Mobile traffic across the link
- the WAAS Mobile server passing SSL traffic to the SSL VPN server
- the SSL VPN server passing data to the content server

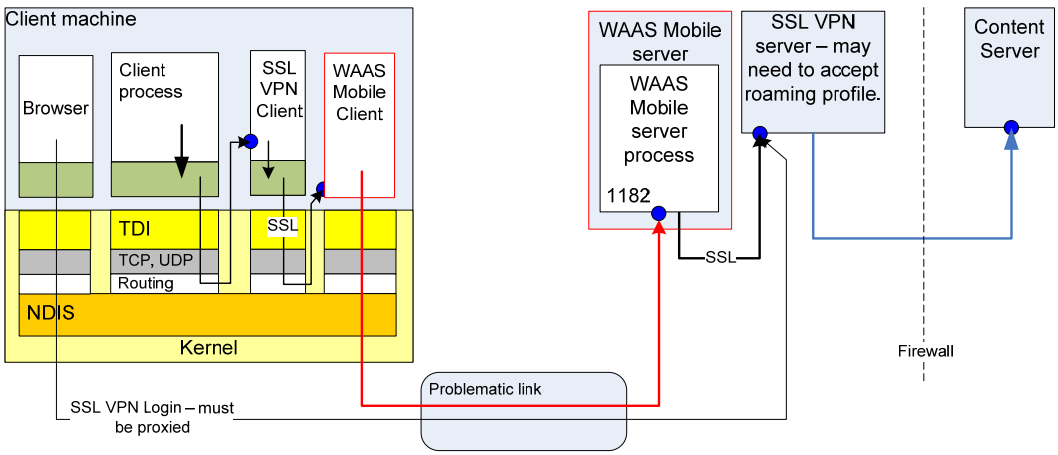

Figure 4. Optimizing transport for a socket-interception SSL VPN

Here, the accelerated connection is placed between the SSL VPN client and server. The WAAS Mobile client decrypts the SSL VPN traffic and the WAAS Mobile server encrypts the traffic for delivery to the SSL VPN server. This arrangement requires that the WAAS Mobile server be in the DMZ with the SSL VPN server, as indicated by the location of the firewall. This affects the network addressing of the server and, therefore, both client and server configurations.

SSL VPN Session Establishment

SSL VPN session establishment is usually accomplished through a web browser. The end-user visits a portal page hosted on the SSL VPN server. The session is periodically refreshed. Among other things, the SSL VPN server checks that the session establishment connection and the data transfer connections from the SSL VPN client come from the same network address. This means:

- WAAS Mobile must proxy both the SSL VPN client and the browser that establishes the session.
- Existing browser connections and SSL VPN client connections must be broken when the WAAS Mobile client starts.

Active Mode FTP

When proxying active mode FTP, an SSL VPN client makes a connection to a port listening on the client machine in order to set up the FTP data channel. The WAAS Mobile client would intercept this traffic unless specifically configured to avoid it.

Network Acceleration Policy Settings

1. The WAAS Mobile client intercepts traffic from the SSL VPN proxy client (any browser used to establish the SSL VPN session) and from no other process.
2. Use Auto-reset on browser connections and SSL VPN client connections.
3. Only traffic directed to the SSL VPN server host is accelerated.
4. The SSL VPN server should allow traffic to come both directly from the client and from the WAAS Mobile server at different times.

Chapter 6 Choose Policies for Internet Access

This section applies for all VPN access, whether IPsec or SSL-based.

- A full-tunnel VPN directs all network access through the VPN. A split tunnel VPN allows partial network access outside the VPN. For example, a split tunnel VPN may be configured to allow direct access to *local* network resources, or to *all* network resources outside the VPN.
- A full-tunnel VPN forces all Internet access through the VPN and therefore through the enterprise. No special configuration is required to accelerate all Internet hosts.
- A split tunnel VPN may allow access to local content servers directly from client machines. These servers must not be optimized. Use either Latency-based Bypass or an Accelerated Networks blacklist to exclude them.
- A split tunnel VPN may allow general Internet access. In this case, an Accelerated Networks white list is required.

The situation is summarized in Table 3.

Table 3. Network acceleration policies for Internet access with VPNs

Accelerate	**Split Tunnel VPN, local access**	**Split tunnel VPN, Internet access**	**Full Tunnel VPN**
All hosts	Not possible.	Not possible.	No special configuration.
Only enterprise hosts	Latency bypass.	Accelerated Networks white list including enterprise servers.	Accelerated Networks white list including enterprise servers.
Only Internet hosts	Accelerated Networks black list excluding enterprise servers, plus latency bypass.	Not possible.	Accelerated Networks black list excluding enterprise servers.

Chapter 7 Select Controller Machines

A controller machine is not required for a single isolated server site, or multiple isolated server sites where individual clients are able to connect only to one server site.

In other cases, at least one controller machine is required. The number of controller machines is determined by client mobility.

For example, in a national deployment, it is reasonable to expect a traveling work force to be able to use any WAAS Mobile server site. That would require a single controller for all the sites.

In a global deployment, it may be expected that sales forces will stay within their national boundaries, leading to the global deployment being a sequence of national deployments, with a controller per geographical region.

Each controller must be able to contact each of its workers (using TCP on port 1182) in order to receive server farm configuration updates.

Chapter 8 Examples

Single Data Location

Remote Users

Figure 5 shows the simplest scenario, with remote access to a single data center using a VPN concentrator at the data center. For optimization, only one WAAS Mobile server is required, associated with the VPN gateway, as shown in Figure 6.

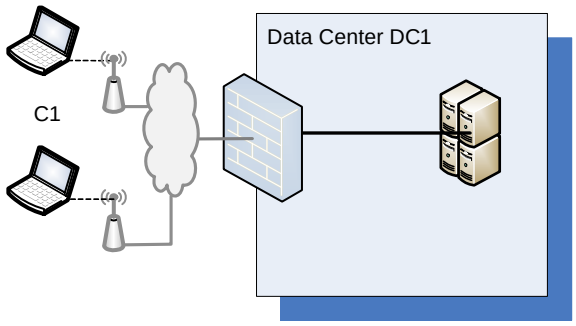

Figure 5. Single Data Location with Remote User Access

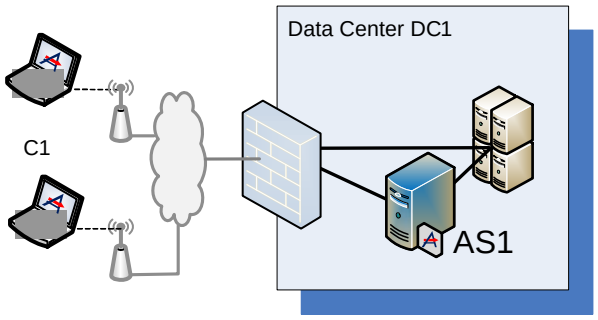

Figure 6. Single Data Location with Remote User Access, optimized

Table 4 shows the server to which the client locations should connect and Table 5 shows that the WAAS Mobile server provides optimized access to the data center. No special policy is required.

Table 4. Server usage

Client	Server
C1	AS1

Table 5. Server access to content

WAAS Mobile servers	Data Center
	DC1
AS1	Yes

Remote Users and Small Branch Office, Single VPN Gateway

Figure 7 shows a small branch office connected over the Internet to a data center. Remote clients connect to the same VPN gateway as the branch office. For optimization, only one WAAS Mobile server is required, associated with the VPN gateway, as shown in Figure 8.

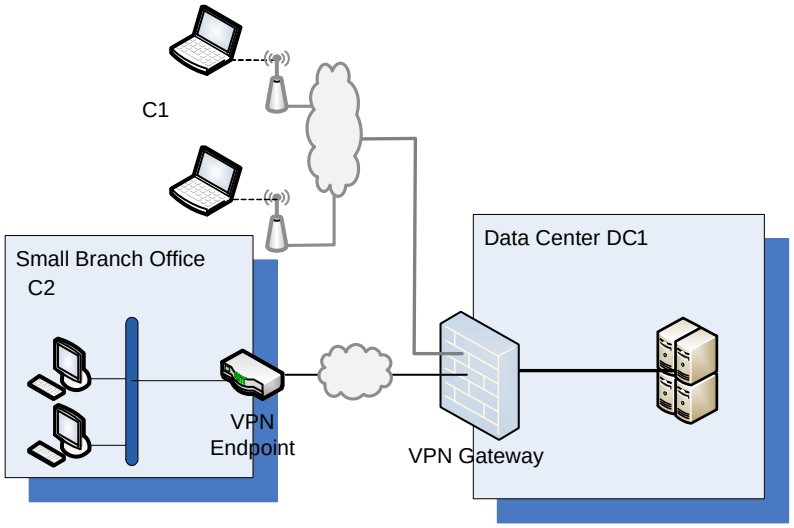

Figure 7. Remote Users and Small Branch Office, Single VPN Gateway

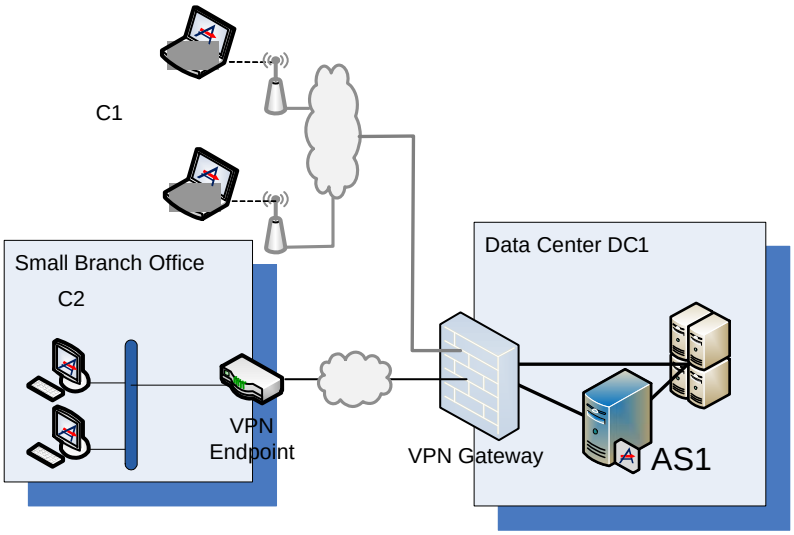

Figure 8. Remote users and small branch office, optimized

Table 6 shows the server to which the client locations should connect. and Table 7 shows that the WAAS Mobile server provides optimized access to the data center. No special policy is required.

Table 6. Server usage

Client	Server
C1	AS1
C2	AS1

Table 7. Server access to content

WAAS Mobile servers	Data Center
	DC1
AS1	Yes

Remote Users and Large Branch Office with WAN Appliance

Figure 9 below shows a large branch office connected by an optimized WAN link to a data center. There is a VPN gateway in the branch office, to which remote clients can connect to get access to resources in the data center. For optimization, one WAAS Mobile server is required, associated with the VPN gateway, as shown in Figure 10.

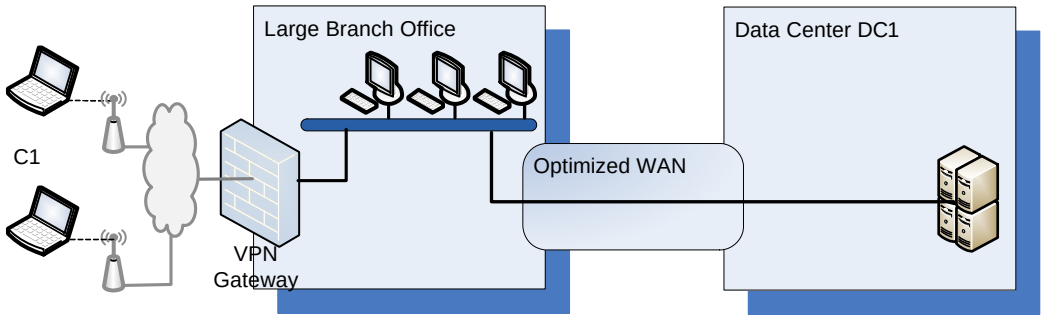

Figure 9. Remote users and large branch office

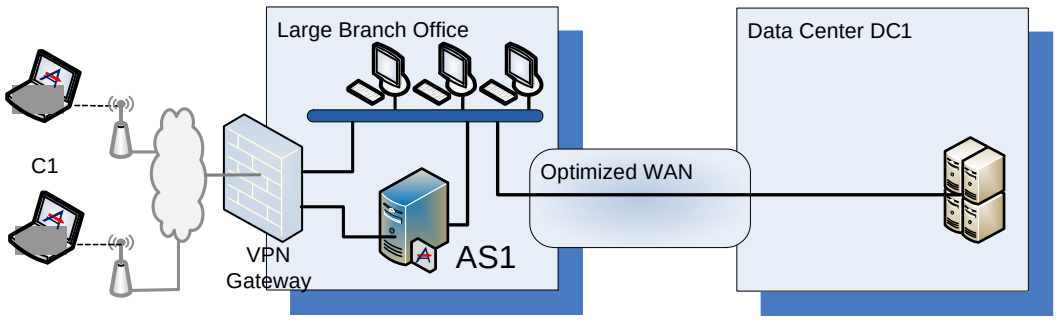

Figure 10. Remote users and large branch office, optimized

Table 8 shows the server to which the client location should connect. The clients in the large branch office with optimized WAN access to the data center do not require WAAS Mobile. Table 9 shows that both WAAS Mobile servers provide optimized access to the data center. No special policy in this area is required.

Table 8. Server usage for remote users at two locations

Client	Server
C1	AS1

Table 9. Server access to content for remote users at two locations

WAAS Mobile servers	Data Center
	DC1
AS1	Yes

Remote Users, Small and Large Branch Offices

Figure 11 below shows one data center connected to both a small and a large branch office. The small branch office and the remote users are connected to the same VPN concentrator at the data center, and both require optimized access. For optimization, two WAAS Mobile servers are required – one associated with each VPN concentrator, as shown in Figure 12.

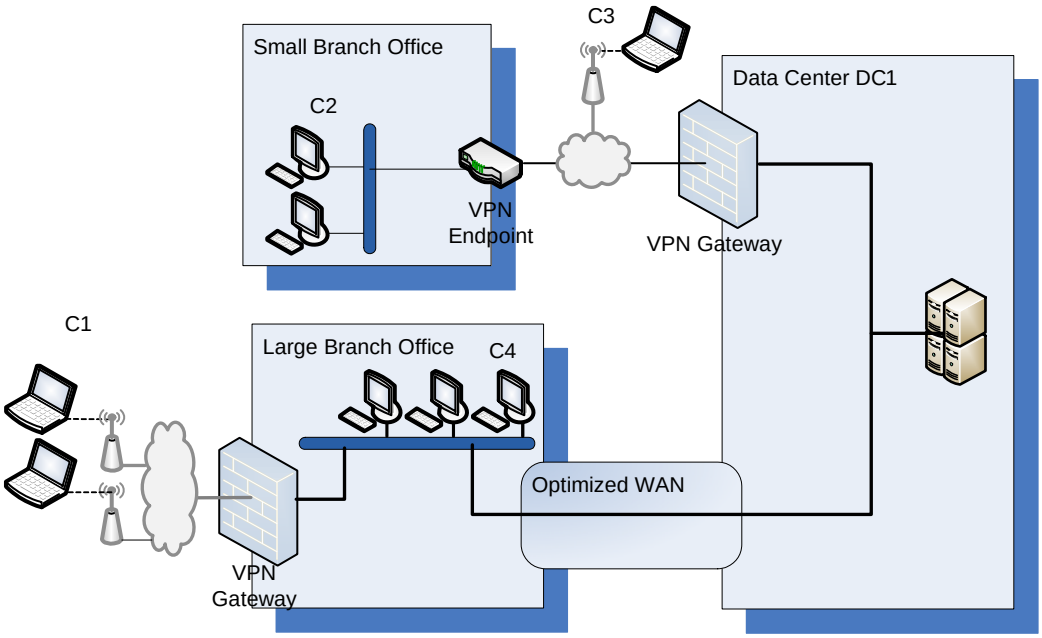

Figure 11. Multiple VPN Gateways

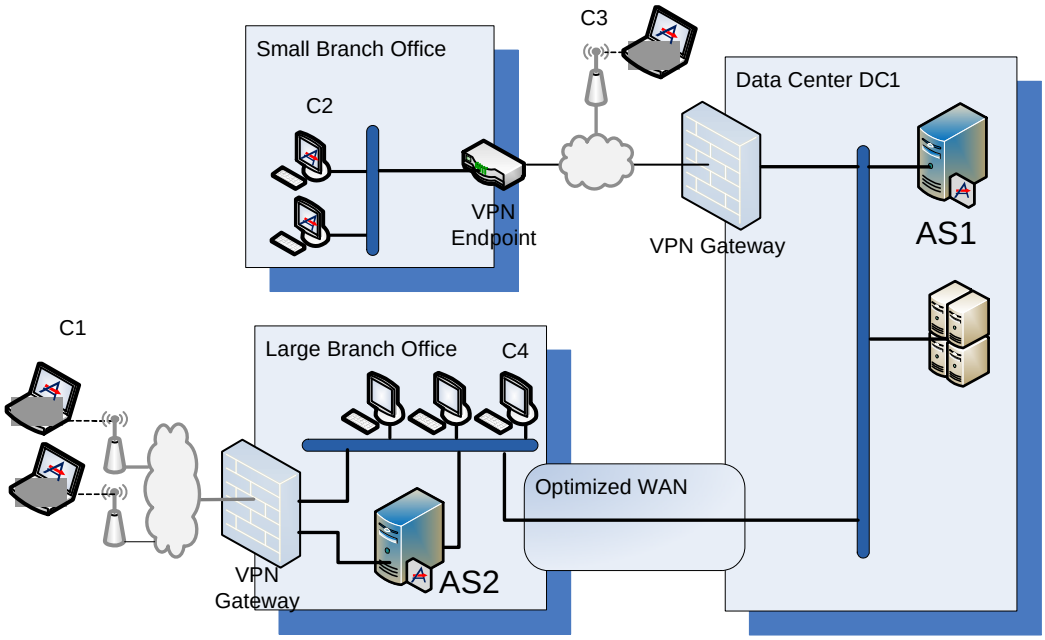

Figure 12. Multiple VPN Gateways, optimized

Table 10 shows the server to which each client location should connect. The clients in the large branch office with optimized WAN access to the data center do not require WAAS Mobile. For the others, the server selection policy can be implemented using wither Farm Selection/Latency or Farm Selection/Client IP Map. Table 11 shows that both WAAS Mobile servers provide optimized access to the data center. No special policy in this area is required.

Table 10. Server usage for remote users at two locations

Client	Server
C1	AS2
C2	AS1
C3	AS1
C4	none

Table 11. Server access to content for remote users at two locations

WAAS Mobile servers	Data Center
	DC1
AS1	Yes
AS2	Yes

Two Data Locations with Optimized Interconnection

Remote Users at one location

Figure 13 below shows two data centers with WAN optimization appliances and one data center supporting remote access through its own VPN concentrator to both data stores. Optimization requires one WAAS Mobile server co-located with the VPN concentrator to optimize client connections through that VPN concentrator to both data locations, as shown in Figure 14.

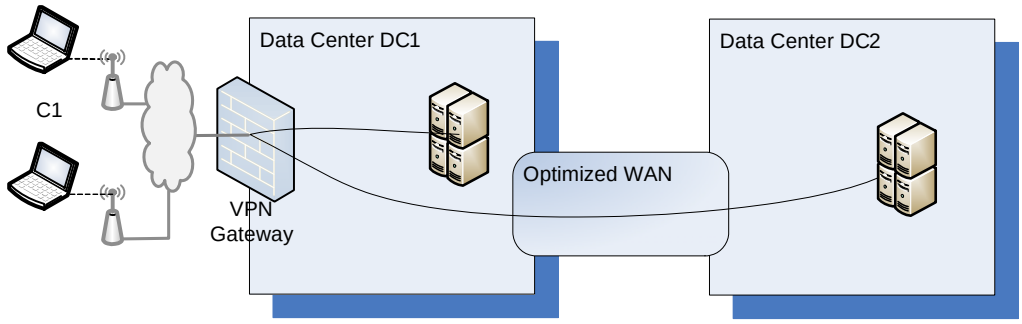

Figure 13. Two data centers with a single VPN concentrator

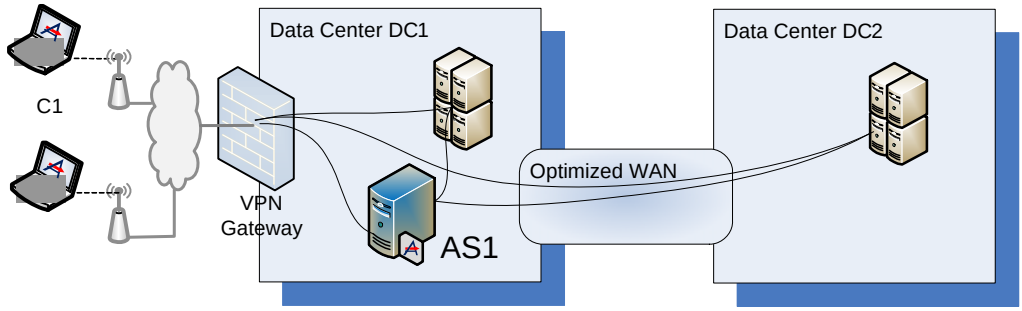

Figure 14. Two data centers with single VPN concentrator

Table 12 shows the server the remote client should connect to. Table 13 shows that the client C1 can obtain optimized access to both the data centers. No special policy in this area is required.

Table 12. Server usage for remote users at two locations

Client	Server
C1	AS1

Table 13. Server access for remote users at two locations

WAAS Mobile servers	Data Centers	
	DC1	DC2
AS1	Yes	Yes

Remote Users at Two Locations

Figure 15 below shows two data centers connected by WAN optimization appliances, each supporting remote access through its own VPN concentrator to both data stores. Site selection places one server in each data center, as shown in Figure 16. Each server optimizes access for the clients that connect to the VPN concentrator in that data center to both data locations.

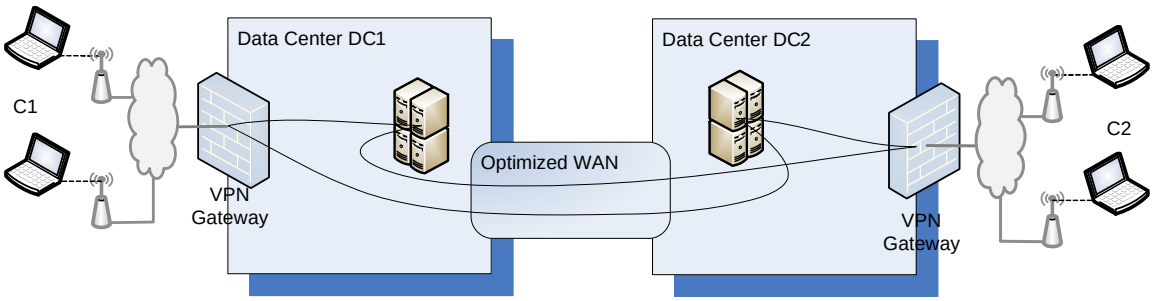

Figure 15. Remote users at two connected locations

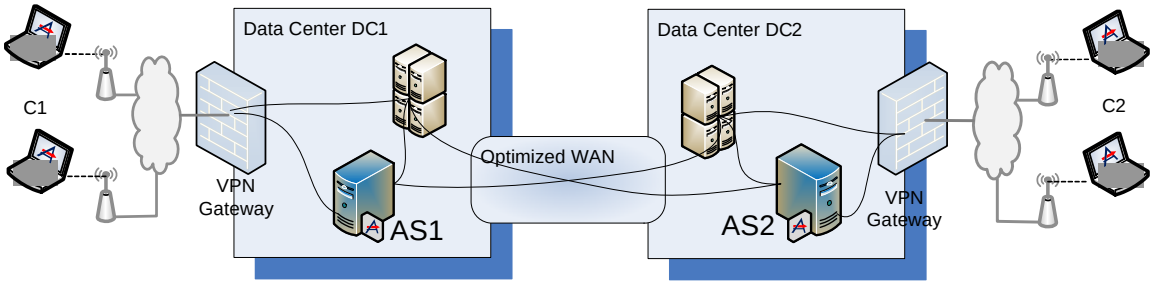

Figure 16. Two data locations with optimized remote access

Table 14 shows which server each remote client should connect to. This can be specified using either of Farm Selection/Latency, Farm Selection/Client IP Map. Table 15 shows that each client can obtain optimized access to both the data centers. No special policy in this area is required.

Table 14. Server usage for remote users at two locations

Client	Server
C1	AS1
C2	AS2

Table 15. Server access for remote users at two locations

WAAS Mobile servers	Data Centers	
	DC1	DC2
AS1	Yes	Yes
AS2	Yes	Yes

Two Data Locations with Unoptimized Interconnection

Remote Users at One Location

Figure 17 shows two WAN-connected data centers with remote users connecting to just one. Static users local to D1 and D2 are not candidates for acceleration and are not shown. Such users would not have WAAS Mobile installed. In this case, optimization requires a single WAAS Mobile server at the data center nearest the remote clients. It can optimize only traffic to data center D1. Traffic to data center D2 must be bypassed. This is shown in the diagram by the direct connection from the VPN concentrator in D1 to the servers in D2.

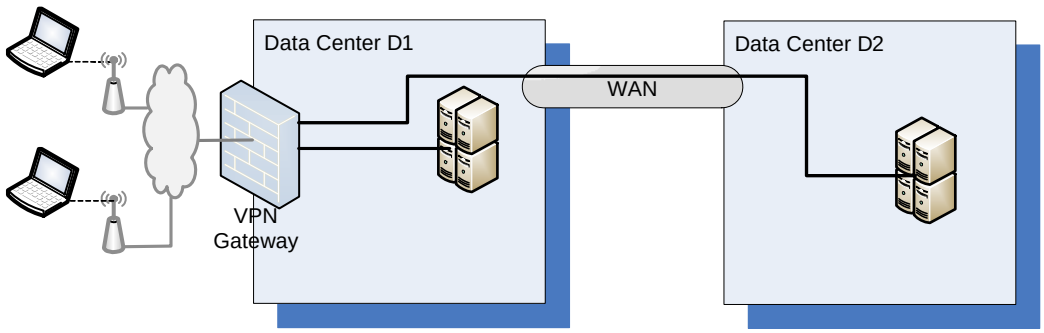

Figure 17. Remote users at one of two data locations

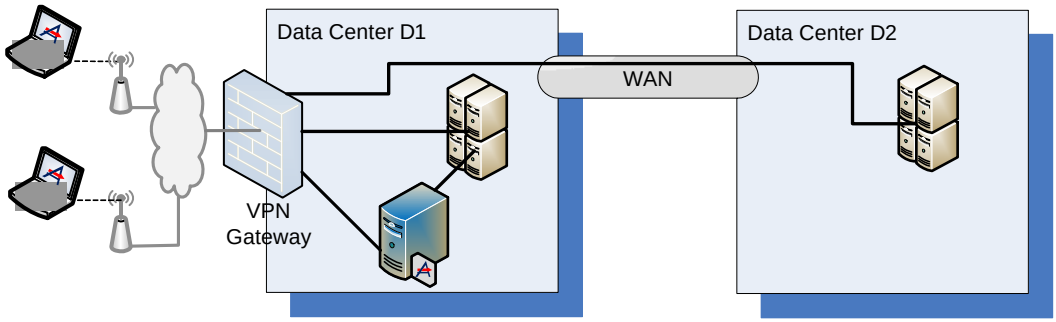

Figure 18. Remote users at one of two data locations, optimized

The required network acceleration policy could be implemented using one of the following:

- Use Latency-based Bypass to exclude the D2 servers from optimization. This requires that the round-trip time across the WAN between D1 and D2 is sufficiently large compared with the round-trip time from the remote clients to D1.
- Use an Accelerated Networks blacklist to exclude the D2 servers.
- Use an Exclusion List based on network addresses.

Should a client visit D1, it would be necessary to avoid using the WAAS Mobile server using either High-speed Bypass, or Access control for clients on the D1 subnet(s) to exclude such clients from acceleration; clients using remote access must get a different subnet from clients within D1.

Should a client visit D2, it would be feasible to use the WAAS Mobile server at D1 for accessing the servers at D1. Either the High-speed Bypass or Access control policies would suffice.

Remote Users at Two Locations

Figure 17 shows two WAN-connected data centers with remote users connecting to both. Static users local to D1 and D2 are not candidates for acceleration and are not shown; these users would not have WAAS Mobile installed. In this example, optimization requires WAAS Mobile servers at each data center, as shown in Figure 20, with each optimizing only traffic to its own data center and bypassing traffic to the other data center. This is shown in the diagram by the direct connection from the VPN concentrator in D1 to the servers in D2.

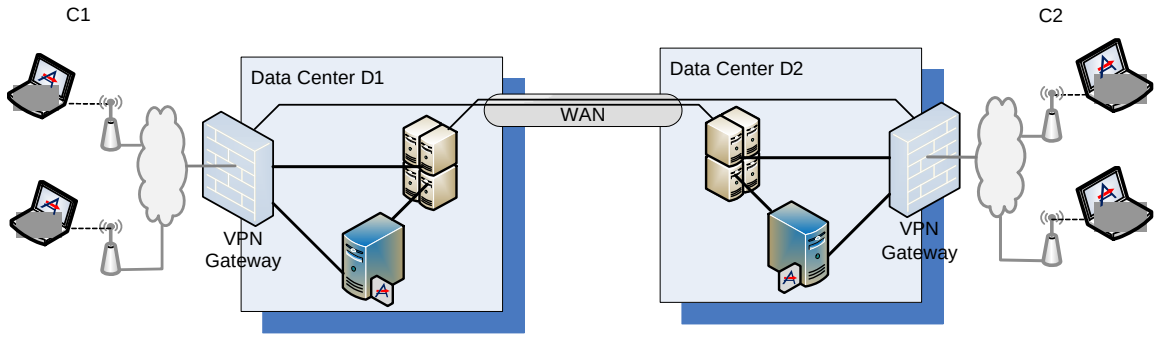

Figure 19. Remote users at two connected data locations

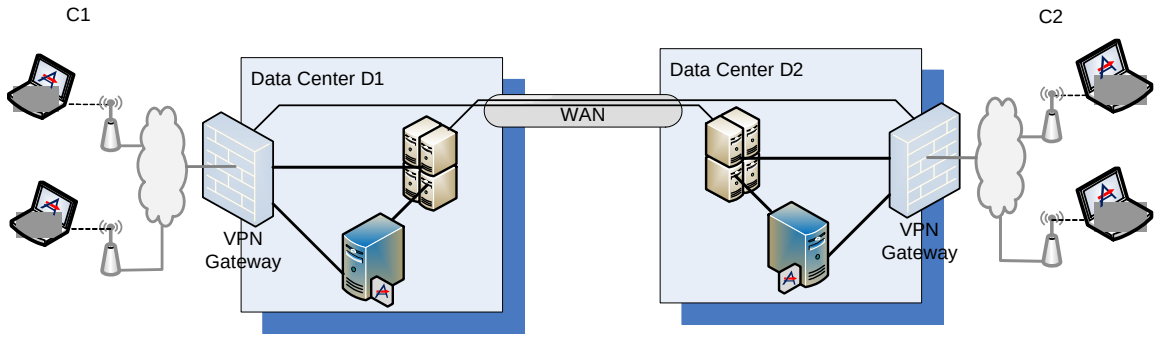

Figure 20. Remote users at two connected data locations, optimized

This is a case with two server sites (farms) in a single server site group. One of the servers can be chosen as a controller. A single client distribution can be used to allow remote users to connect to either D1 or D2, using the correct WAAS Mobile server to accelerate the correct content servers.

Remote Access Only

To use the correct WAAS Mobile server, choose one of the following policies:

- Farm Selection/Latency
- Farm Selection/Client IP Map

To accelerate the correct content servers, use one of the following:

- Use an Accelerated Networks blacklist to exclude the D2 servers
- Use an Exclusion List based on IP addresses

Remote and LAN Users with WAAS Mobile Optimizing the WAN Link

If the remote users also operate on the LAN and wish to use WAAS Mobile to optimize traffic across the WAN link where appropriate, the policy settings are more specific.

Clients in D1 would select the WAAS Mobile server. High-speed bypass would prevent using the chosen server, so there would be no acceleration. Therefore to use the correct WAAS Mobile server, choose Farm Selection/Client IP Map.

To accelerate the correct content servers, use one of the following:

- Use Latency-based Bypass so traffic to each nearby server goes directly.
- Use an Accelerated Networks blacklist applied per WAAS Mobile server.
- Use an Exclusion List based on IP addresses applied per WAAS Mobile server.

Three Data Locations

Figure 21 illustrates multiple data locations connected by a variety of link types. Two of the data locations are connected with optimized WAN links. The third is connected with an unoptimized WAN link.

Figure 22 shows the required locations of WAAS Mobile servers for optimization. As outlined in Site Selection Procedure above, WAAS Mobile servers are located at the data end of each problematic link. The scope of each server is limited by the location of unoptimized WAN links. The figure also indicates the connectivity required.

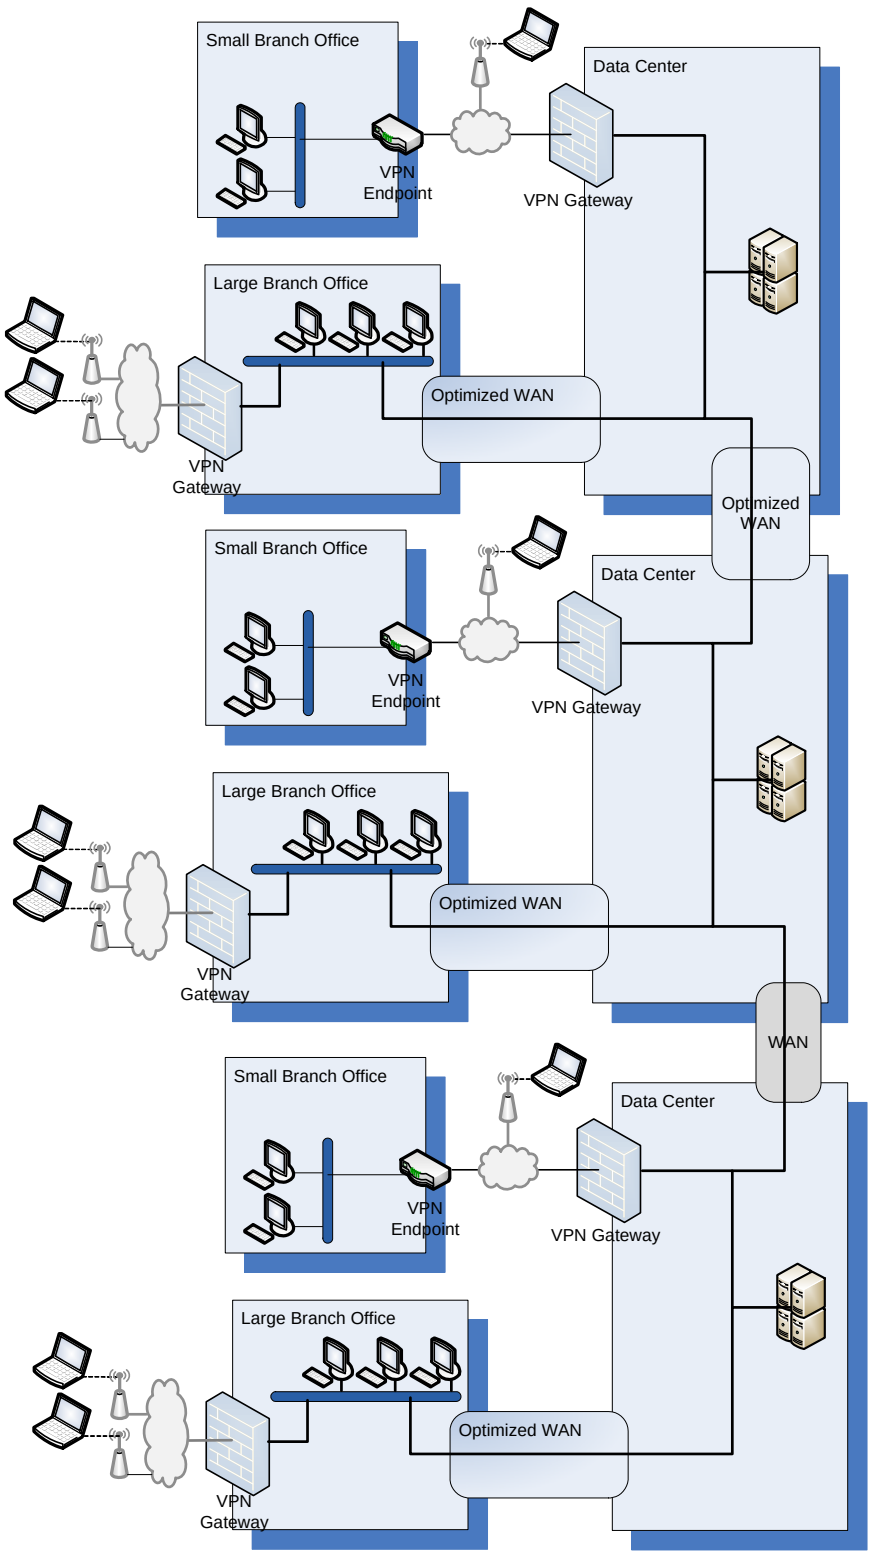

Figure 21. Three data locations

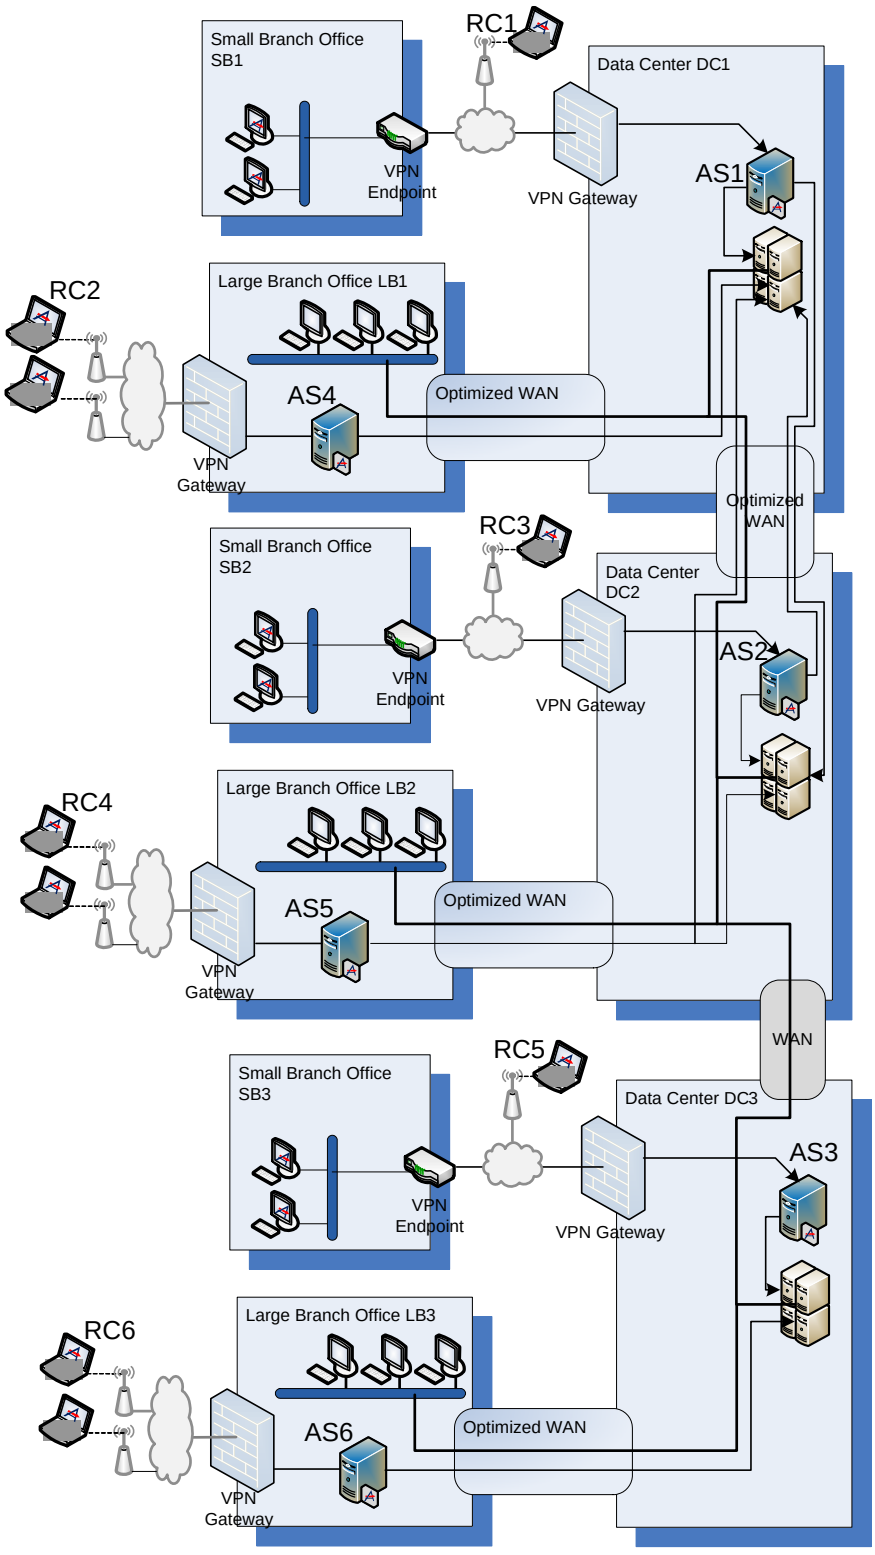

Figure 22. Three data locations, optimized

The use of the WAAS Mobile servers for accelerating content servers in the data centers is shown in Table 16 and the connections of clients to servers is shown in Table 17.

Table 16. Server access for three data center example

WAAS Mobile servers	Data Centers		
	D1	D2	D3
AS1	Yes	Yes	
AS2	Yes	Yes	
AS3			Yes
AS4	Yes	Yes	
AS5	Yes	Yes	
AS6			Yes

Table 17. Client-server connectivity in three data center example

Client	Server
RC1	AS1
RC2	AS4
RC3	AS2
RC4	AS5
RC5	AS3
RC6	AS6
SB1	AS1
SB2	AS2
SB3	AS3

In this case, all the WAAS Mobile servers can be considered part of a single server site group. Selection by Client IP Map will most likely be appropriate, using Table 17 to determine which subnets are associated with each server. The data centers for acceleration by each server can be selected by Latency-based Bypass. If that is inadequate, an Accelerated Networks blacklist can be applied to each WAAS Mobile server.

Cisco WAAS Mobile Network Design Guide