Cisco WAAS Mobile Integration 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-15418-03

Contents

Table of Figures

About This Document

Intended Audience

This guide is intended for network engineers setting up Cisco WAAS Mobile in various environments and configurations.

Document Outline

The guide will discuss the following topics:

- Network deployment scenarios supported by Cisco WAAS Mobile.
- Configuring the network and firewall for the Cisco WAAS Mobile server.
- Other Cisco WAAS Mobile server features and network configuration options.

Related Documents

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

- *Cisco WAAS Mobile Administration Guide* – Everything needed to set up and administer WAAS Mobile.
- *Cisco WAAS Mobile User Guide* – A user guide for the end user. This complements the on-line help system and provides a reference for offline study.
- *Cisco WAAS Mobile Network Design Guide* – Provides network architects with best practices for integrating WAAS Mobile with various distributed network topologies and usage scenarios.
- *Cisco WAAS Mobile Release Note* – Release-specific information regarding features added, changed, and removed as well as known issues and issues fixed in the release.

Chapter 1 Network Scenarios

This section presents background information and key concepts that will help to clarify the network architectures in which Cisco WAAS Mobile can operate. The following topics are covered in this chapter:

- Cisco WAAS Mobile Enterprise Architecture Overview
- Sample Deployment Scenario (Upstream Proxy)

Cisco WAAS Mobile Enterprise Architecture Overview

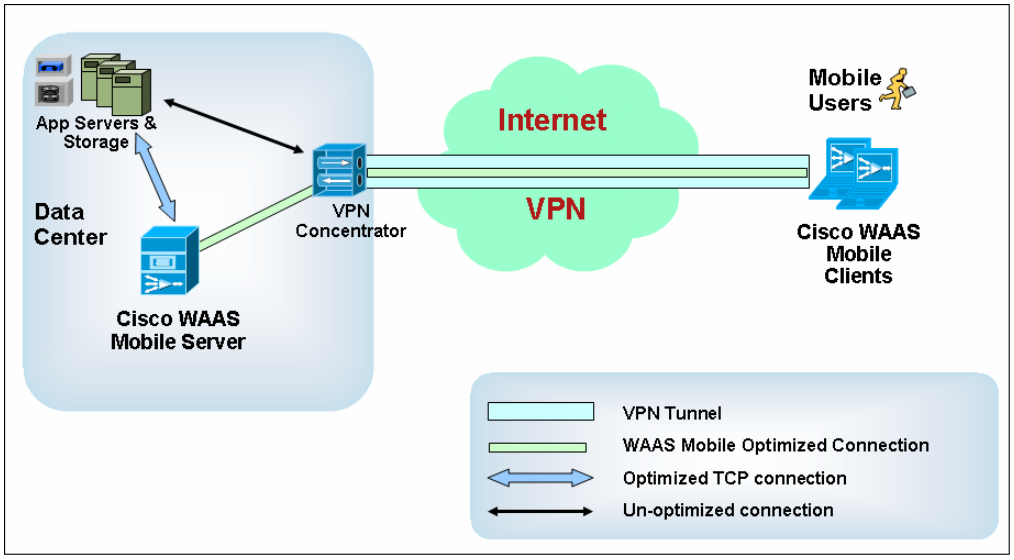

Figure 1 Enterprise Architecture with VPN

The above picture shows a generic enterprise WAAS Mobile deployment with VPN software running on the end user "client" machine. The network connecting the client machine to the enterprise data center has one or more of the following problems: low bandwidth, high-latency, high-packet loss or unreliable connections.

Public (Internet) and Private (Intranet) IP Space

WAAS Mobile can optimize access to both Internet and Intranet content servers. The enterprise administrator has control over which servers in the public and private IP space will be proxied by WAAS Mobile by modifying an "Accelerated Networks" table, which can be managed server-side through the WAAS Mobile Manager administration interface. In addition, the server itself can be deployed in the private IP space as shown in the diagram above or in the public space, allowing optimization without VPN software – this scenario will be covered in the next section.

Optimized TCP Connections

These are TCP connections proxied by WAAS Mobile. On the client machine, these connections originate with the application and terminate at the client. In the data center, the "proxied" connections originate at the WAAS Mobile server and terminate at the content server.

Bypassed TCP Connections

When the client is installed and running, not all TCP connections are sent from user applications through WAAS Mobile. The software provides both deployment-time and runtime options for determining which connections will be proxied and therefore optimized. In addition, the client can be configured to shut itself off automatically in the event that it detects that it is connecting to the WAAS Mobile server over a high-speed connection. In this case, all TCP connections from user applications will be bypassed as indicated by the black arrow in Figure 1.

Accelerated Applications

By default, WAAS Mobile proxies a range of applications including most web browsers, email clients, Windows Explorer for file shares, ftp clients, and thin clients like Citrix and Microsoft Remote Desktop Client (RDC). In addition, any generic application using TCP connections to content servers can be added via its process name. This list of accelerated applications is determined by comparing the name of the process running on the end user's machine to a pre-configured list of "Proxied Processes." TCP connections not in this list will be bypassed.

WAAS Mobile UDP Traffic

The optimized data between the client and server uses ITP, or Intelligent Transport Protocol, in place of TCP. ITP runs on top of UDP bound for port 1182 on the server (WAAS Mobile's "well-known port"). UDP is used for all optimized traffic except for session initiation, which uses TCP and port 1182 on the server.

Sample Deployment Scenario (Upstream Proxy)

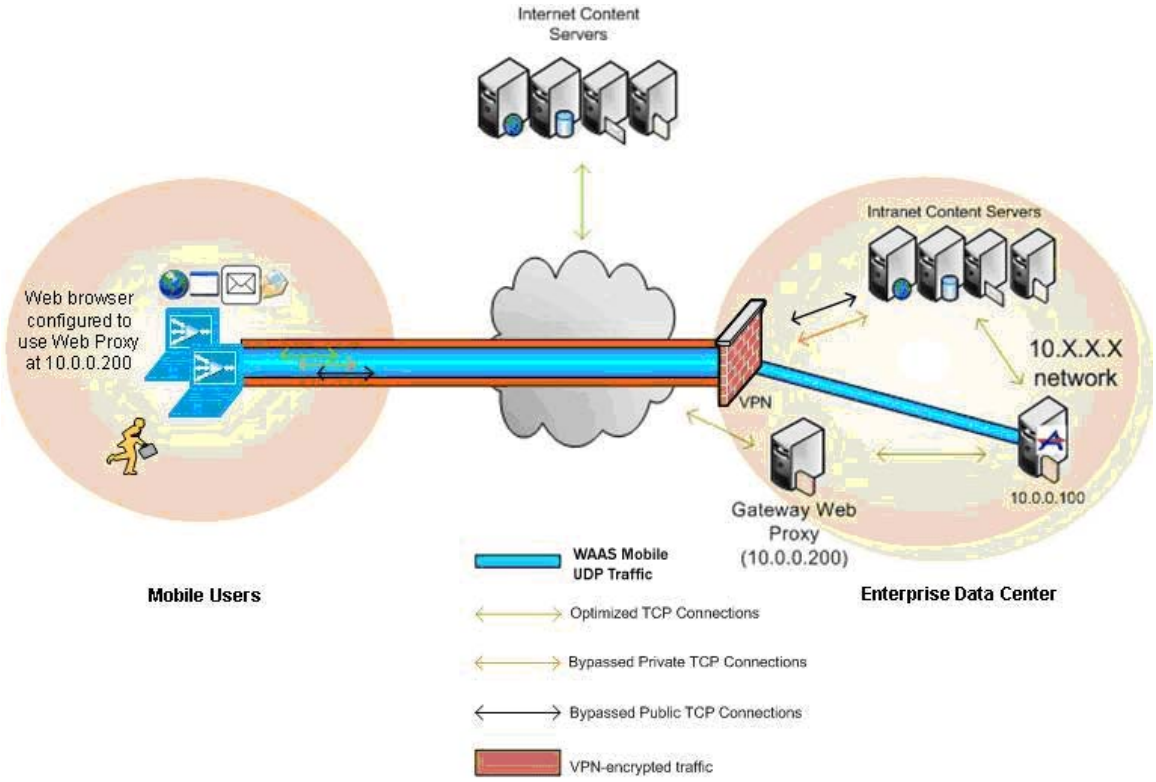

Figure 2 Use of Upstream Proxy

All traffic from mobile users, including both the Internet and intranet bound traffic will pass through the VPN. Requests for public web sites go through a Gateway Web Proxy in the data center. The WAAS Mobile server has a private IP. Public requests through web proxy are optimized.

For security reasons, many enterprises force remote PCs to proxy HTTP requests through a Gateway Proxy. WAAS Mobile sees only an attempt to connect to a web proxy at a private IP address (10.0.0.200 in this example) and it proxies the connection attempt as if it were destined for an Intranet content server. All other TCP connections to public or private content servers from accelerated applications will be optimized normally.

Chapter 2 Ethernet, DNS and Firewall Configuration

This section describes how to configure the Ethernet interfaces on the WAAS Mobile server. The server can use just a single interface for administration and WAAS Mobile data or it can be configured to use an "external" and an "internal" interface. We strongly recommend the latter configuration if the server is deployed to allow access from client machines connecting from the public Internet and not through a VPN. The following topics are covered in this section:

- Ethernet Interfaces
- WAAS Mobile Server Security Note

Ethernet Interfaces

The WAAS Mobile server works best with dual Network Interface Cards. However enterprises can combine the following requirements and use a single interface.

External Interface Requirements

- IP address must be routable/accessible to client machines
- Network mask and default gateway settings depend on enterprise architecture
- DNS server settings required to allow the server to resolve the names of content servers to be accelerated
- Firewall must allow TCP and UDP access to port 1182

Internal Interface Requirements

- IP address accessible to network monitoring station or OAM if SNMP monitoring desired
- Firewall must allow TCP and UDP access to port 161 from the network monitoring station (OAM) if it will be querying for SNMP counters
- Firewall must allow TCP access to port 80 to allow administrators to access WAAS Mobile Manager for server configuration
- Firewall must allow access to port 80 for clients for the system reports on the server feature
- If upstream proxy servers are used along with WAAS Mobile, additional configuration may be required. For example, when deploying Cisco WAAS Mobile with Microsoft ISA server, the Flood Mitigation feature, which is on by default in ISA server, should be disabled. Additionally, the WAAS Mobile server addresses should be added to the ISA server's flood mitigation IP address exclusion list.

WAAS Mobile Server Security Note

Behind a correctly configured firewall, the WAAS Mobile server is a highly secure application server. It has very good protection against Denial of Service (DoS) attacks. The primary DoS risk is from the TCP listener ports 1182, and UDP port 1182 which is the same type of risk associated with a standard web server.

WAAS Mobile encrypts the initial TCP exchange on port 1182 between the client and the server in which the client supplies the user name/password. The control channel uses a public/private key to encrypt the login information in a way that makes the server safe from replay attacks.

Chapter 3 Licensing

WAAS Mobile licensing for enterprises is an Active Users License limiting the number of concurrent active connections.

The Active Users Licensing scheme requires a license for each WAAS Mobile server that limits concurrent active users connected to that server to a specified number of connections. When the number of active users allowed in the license is exceeded, subsequent users connecting will not be accelerated.

> NOTE: All changes in Licensing Configuration require a restart of the WAAS Mobile server.

Configuring WAAS Mobile Server

The Active Users License key is entered on the Licensing page of WAAS Mobile Manager. Entering the license key requires a start/restart of the WAAS Mobile Server process (the WAAS Mobile Manager > Home page).

Figure 3 License Information Entry

As seen in the image above, the WAAS Mobile Manager > Licensing page shows information about the server's license such as the number of users on the license and when the license expires. As soon as the key is entered and the server is running, users can be directed to it.

Chapter 4 Load Balancing and High Availability

Multiple WAAS Mobile servers can be deployed in a server farm to provide high capacity. In general, more than one server is also recommended for production deployments in order to provide high availability and redundancy.

To balance the load across the servers in the farm, as well as to redirect traffic to healthy servers when one is down, WAAS Mobile supports the three options discussed in this chapter:

- DNS Load Balancing
- Load Balancing Appliance
- Advanced Server Selection

DNS Load Balancing

When building a WAAS Mobile client using the WAAS Mobile Manager > Client Configuration, the IP address or host name of the WAAS Mobile server must be specified. If a host name is given, the client will resolve the host name at runtime to obtain the IP address used to connect to the WAAS Mobile server and establish an acceleration session. If the DNS host name has been configured to resolve to multiple IP addresses, DNS load balancing will be used by the client.

In the sample output from a DNS client cache, the host name myaccelerationserver.com is configured by the enterprise DNS server to be associated with three IP addresses:

```
myaccelerationserver.com
----------------------------------------
Record Name . . . . . : myaccelerationserver.com
Record Type . . . . . : 1
Time To Live  . . . . : 207
Data Length . . . . . : 4
Section . . . . . . . : Answer
A (Host) Record . . . : 10.0.0.101

Record Name . . . . . : myaccelerationserver.com
Record Type . . . . . : 1
Time To Live  . . . . : 207
Data Length . . . . . : 4
Section . . . . . . . : Answer
A (Host) Record . . . : 10.0.0.102

Record Name . . . . . : myaccelerationserver.com
Record Type . . . . . : 1
Time To Live  . . . . : 207
Data Length . . . . . : 4
Section . . . . . . . : Answer
A (Host) Record . . . : 10.0.0.103
```

In this example, 10.0.0.101, 10.0.0.102, and 10.0.0.103 are the IP addresses of the WAAS Mobile servers in the server farm. To load balance across these servers, the WAAS Mobile client will go through the following steps when connecting to the server farm:

- Resolve name to three IPs.
- Select randomly from the three IPs (this balances the load across the three servers in the farm). Assume 10.0.0.102 is selected randomly.
- Try to connect to 10.0.0.102. If the attempt is successful, then use this server for the remainder of the acceleration session.
- If the attempt to connect to 10.0.0.102 fails because the server is temporarily unavailable, then randomly select from the remaining two IPs (this provides high availability). Assume 10.0.0.101 is selected randomly.
- Try to connect to 10.0.0.101.

This simple scheme relies on the intelligence in the WAAS Mobile client to achieve load balancing and high availability equal to that provided by high-end commercial load balancing appliances. This is the recommended scheme for server farms with 10 or fewer servers.

Load Balancing Appliance

Cisco WAAS Mobile is also compatible with standard load balancing appliances such as those used in web server farms. In this scenario, the client is configured to use a single IP address to connect to the server farm. This is the "virtual IP address" of the load balancing appliance. When a new TCP connection attempt is made by the client, the load balancing appliance selects one of the healthy servers in the farm and forwards all IP traffic from that client to the selected server for the duration of the acceleration session.

One advantage of the load balancing appliance over DNS load balancing is that the appliance can be configured to perform frequent health checks of the servers in the server farm. This allows it to be more proactive about detecting failed machines and re-routing traffic accordingly.

> NOTE: When WAAS Mobile's Advanced Server Selection (described below) is not in use, then this is the recommended scheme for very large server farms. WAAS Mobile has been deployed in this way to support server farms of up to 200 servers.

Advanced Server Selection

WAAS Mobile's Advanced Server Selection involves the client making an informed decision on which server to connect to in the enterprise. This feature can be used to provide load balancing and high availability for a single server farm deployed in the enterprise data center or for the more complex scenario where an enterprise has deployed multiple WAAS Mobile server farms in more than one geographic location. In the latter case, the client first selects a server farm (farm selection) and then selects the server within that farm (server selection) to which it will connect. The end result is a dynamic solution that is flexible and adaptable to any enterprise network infrastructure.

To implement a server farm with Cisco WAAS Mobile, a "controller server" is configured with a mapping table defining the server farm(s), containing one or more WAAS Mobile "worker servers." This mapping table provides clients with dynamic instructions on how to select the

appropriate farm and server within that farm regardless of where they connect to the enterprise WAN. Upon startup and periodically thereafter, the controller server will communicate the current mapping table to each worker server across all server farms. Each server will then communicate this information back to each client that is connected to it. The client will attempt to use the same worker server on the next login in order benefit from WAAS Mobile's persistent sessions and/or persistent delta cache.

Configuring Advanced Server Selection

Advanced Server Selection, off by default, is enabled on the WAAS Mobile Manager > Server Configuration > Server Selection page. The controller server communicates this setting and all associated settings to the servers in its farm.

NOTE: All changes in Advanced Server Selection configuration, including enabling/disabling, require a restart of the server.

Server Farm

The Server Farm page is used to configure the servers in a farm and whether there are multiple farms.

Single Server Farm Deployment	If the deployment has only one server farm, enter the IP/Hostname of each server to be in the farm into that field and click Add Server. The farm name field should be left blank, as this single farm will be treated as the default farm which simplifies configuration for the administrator by bypassing any farm selection settings. It is important to note that the controller server's initial self-configuration as a worker server to which clients connect becomes undone once a server is entered in the server farm. If the controller server is to remain a worker server, it must be added to its own server list.

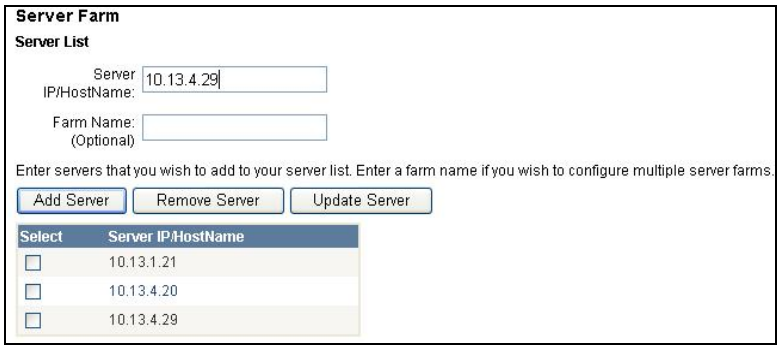

Figure 4 Configure Single Server Farm

Multiple Server Farm Deployment	Multiple server farms are created by entering the IP/Hostname of each server and entering a farm name in the Farm Name field (farm2 and farm3 in the below example, Default Farm will also function as a farm in this configuration).

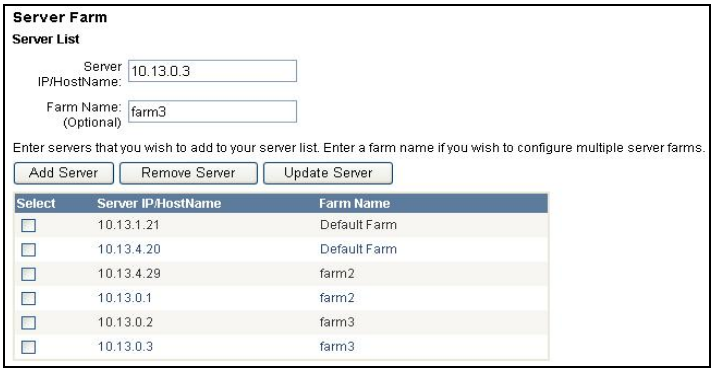

Figure 5 Configure Multiple Server Farms

Entering a farm name automatically disables the simpler default farm configuration used for a single farm deployment and necessitates Server Selection Method and Farm Selection Method configuration.

Server Selection Method

Random and Prioritized are the two Server Selection Method options. Note that Submit must be clicked for a change of this setting to take affect. In all cases, if no other server is available, the client falls back to connecting to the server that was entered when the client distribution was created.

Random Selection	With Random Selection, the client chooses an IP randomly if more than one exists, or if the DNS name resolves to more than one IP. If the first server chosen is not available, then the client will try to connect to another in the farm until all choices are exhausted.
Prioritized Selection	With Prioritized Selection, the client connects to the first server listed in the server farm, and then tries the other servers in the farm in the order listed when the first server is not available.

Farm Selection Method

Farm Selection Method settings only apply in a deployment with multiple server farms. An administrator can choose a Client IP Map or a Latency-based Farm Selection Method. Note that Submit must be clicked for a change of this setting to take affect. Once the client has selected the farm it is to connect to, the server within that farm is chosen based on the Server Selection Method.

Method	Description
Client IP Map	In the example below, any client with IP 10.13.1.x connects to the server farm named farm2 and any client with IP 10.13.4.x connects to farm3.
Latency	If Latency is selected as the Farm Selection Method, the client performs a latency test by pinging all servers in the controller server's server list. The client chooses the server farm with the least latency.

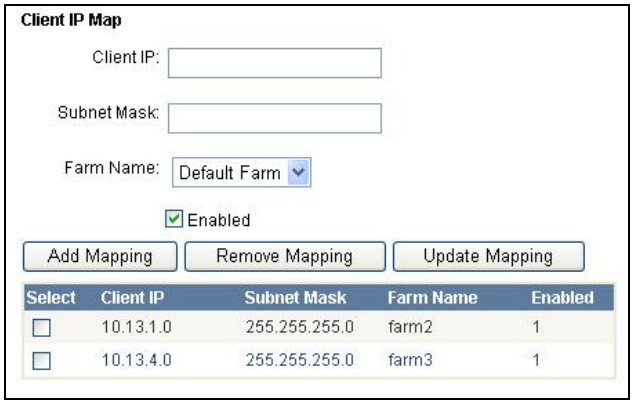

Figure 6 Configure Server Selection by IP Map

Chapter 5 Client IP Mapping Scheme

Using a proxy server like WAAS Mobile makes all content server access appear as though it comes from a single IP address – that of the WAAS Mobile server. This can be problematic for authentication and/or enterprise network traffic analysis. It can also break the load balancing algorithms of application server farms that are based on "sticky" client IPs. In general, it is a commonly accepted trade-off.

If the trade-off is not acceptable, a solution that uses IP Mapping can be used to provide proxy benefits. Cisco WAAS Mobile addresses the problem by assigning a local proxy address for a client from a pool, which in WAAS Mobile is implemented by the *IP Aliasing* feature.

IP Aliasing

IP Aliasing is a feature of WAAS Mobile that allows each user session to be seen by the destination application or content server as a mapped or aliased IP address. This can be accomplished by dynamically assigning a local IP address for each client from a pool of IP aliases or by defining a mapping algorithm that allows the server to determine which IP to present to the destination servers from a given client IP. Once the IP is assigned and created on the network interface, this "local proxy address" will be used for all subsequent communication to destination application servers.

> NOTE: All changes in IP Aliasing configuration require a restart of the WAAS Mobile server.

IP Aliasing from an IP Pool

An IP pool is constructed based on the server settings selected by the administrator via WAAS Mobile Manager. Each assigned IP alias from the pool is created on the network interface and bound for each TCP connection to the destination application server. Therefore, the pool should contain addresses that are routable from the machine to the content server. The assigned IP Alias is reported via the event log and RADIUS accounting, allowing administrators to identify and monitor user traffic. For RADIUS accounting, this information is available in both the start and stop records.

To enable this feature, check the Use IP Aliasing box on the WAAS Mobile Manager > Server Configuration > Aliasing page.

In the Public Network Interface text box, enter the name of the public network interface. You can find this by typing "ipconfig /all" at the command line.

In the Valid Sources text box, enter the IP addresses for clients that you wish to allow to use IP Aliasing. For each entry use the CIDR address format (A.B.C.D/X), where X is the number of bits in the network portion of the address. For multiple entries, use a semicolon (;) to separate each entry. Leave blank to allow all sources, or all clients, to use the IP alias pool.

In the IpPool text box, enter an IP address range and subnet mask to specify a pool of IP aliases to be created on the network interface. To specify more than one range, separate each entry with a semicolon (;).

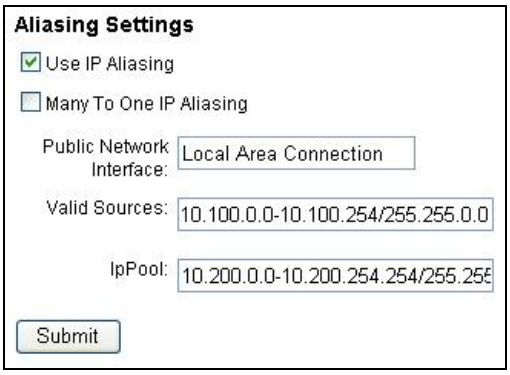

Figure 7 Configure Aliasing Settings

In the example above, if a client with an IP address in the Valid Sources text box logs into the WAAS Mobile server (10.100.10.10 for instance), then it will be assigned an IP from the IpPool range (10.200.50.25 for instance). This assigned IP will be seen at the content server for all accelerated TCP connections from that client.

Many-to-One Mapping

The selection of the local proxy address can also be based on a fixed mapping between client IPs and aliased proxy IPs. One scheme supported by WAAS Mobile is Many-to-One mapping, which involves selecting one server alias for an entire range of client IPs. Specifically, client subnets of the following type X.X.X.0 are mapped to X.X.X.254. The X.X.X.254 IP will be dynamically added to the external interface of the server as the first client from the corresponding subnet logs in. Client IP of X.X.X.254 is not valid in this scheme so network administrators deploying IP Aliasing using Many-to-One mapping will need to allocate IP addresses to client machines accordingly. In addition, the aliased X.X.X.254 address needs to be routable from the server machine to the content servers.

NOTE: Many-to-One mapping requires Advanced Server Selection to be enabled.

Chapter 6 Server Health Checks

The most critical operations task is to determine the health of servers and to remove unhealthy servers from service. Four different types of health tests can be used:

1. Health checks utilized by the load balancers
2. DNS health check
3. Health tests assembled from monitoring application and O/S parameters

If the health check performed by the load balancer fails, no further sessions will be routed to that server. If tests 2 or 3 fail, a server can be programmatically restarted via a network/application monitoring program or removed from service so that it fails the next test by the load balancers.

Load Balancer Health Checks

Most load balancing appliances provide a range of options for performing periodic health tests on all servers. A common method is to use a TCP probe. WAAS Mobile listens for TCP connections for session establishment, so this can be used to determine if at least the listener is operating. If this test fails, the WAAS Mobile service on the machine is no longer operational and should be restarted.

Master server load balancing frequently checks the health of worker servers by connecting and requesting status. This allows it to get a more comprehensive reading of server health and to adjust load balancing accordingly.

DNS Health Check

The WAAS Mobile server can be configured to perform a self-test via a DNS query to verify the status of the computer's Internet connection. If the query fails, WAAS Mobile server will generate an event in the NT error events log. After several failed queries, the server will shut down and server will remain disconnected even when the Internet becomes available again; manual intervention is required to restart the WAAS Mobile server software.

When the server has disconnected due to a failed health check, clients will go into bypass mode; users will be able to browse the Internet and send/receive emails without acceleration until the server computer's Internet connection becomes available and the WAAS Mobile server has been restarted.

The DNS Health Check feature is turned off by default. To enable it, you must add or modify the following key in the server registry:

HKLM\Software\ICT\AcceleNetServer\Network\InternetConnectionCheckMode

Enter a Dword value of 1 to enable the DNS query; enter a value of 0 to disable the query.

It is also possible to customize several other settings related to the DNS health check. You may change the default URL that is used for the query (currently yahoo.com), you may also adjust the wait time between queries, and you may change the number of times the query can fail before the WAAS Mobile server disconnects.

If you wish to customize these settings you may add or modify the following keys:

HKLM\Software\Cisco\WAASMobile\WAAS MobileServer\Network\InternetConnectionCheckAddress

Default: "www.yahoo.com"

Type: STRING

HKLM\Software\Cisco\WAASMobile\WAAS MobileServer\Network\InternetConnectionCheckWaitTime

Default: 0x00000014 (20 seconds)

Type: DWORD

HKLM\Software\Cisco\WAASMobile\WAAS MobileServer\Network\InternetConnectionCheckRetryCount

Default: 3

Type: DWORD

Health Tests from Application and O/S Performance Counters

While the DNS health check can be used as a clear indicator of server health, information from performance counters (SNMP data) may also indicate a problem.

Counters can be used to detect service problems such as excessive CPU utilization. WAAS Mobile should be able to function well at 100% CPU utilization, so that high CPU use alone is not a cause for alarm. However, if other parameters suggest a low workload, then the server may have a thread locked up in an infinite loop. While this condition has not been observed in a production release, it is possible. An alarm can be generated if high CPU usage on a processor occurs when the processing load should be low, as indicated by the following:

- low CPU use on the other CPU of a dual-CPU machine
- low value for the application parameter HttpReq/Sec
- low value for the O/S performance counter for BytesReceived on the internet-facing network card.

If CPU usage is abnormally high, a reboot of the machine may be appropriate.

Server restarts: In the event of an unhandled exception, the server will restart itself. When the server starts up again, an event is written to the event log. Frequency of service restarts is an important metric for monitoring overall service health over time.

Chapter 7 SNMP Counters and Alarms

Introduction

This section describes WAAS Mobile's support for native Windows SNMP alarm generation and access to SNMP counters. In addition, these same values are accessible via Windows NT Events and Windows performance counters respectively.

MIB

The WAAS Mobile server MIB is installed in the WAAS Mobile server software folder. The file name is WAAS MOBILE-SERVER-MIB.TXT. The syntax of the MIB file has been checked using the online MIB checker at http://www.muonics.com/Tools/smicheck.php.

This MIB is also used to document the available Windows Performance Counters and NT Events, which use the same names as the SNMP values.

SNMP Service

Ensure the SNMP service is installed and running. Note that SNMP is not installed on Windows by default. It is located in Control Panel > Add/Remove Programs > Add Windows Components > Management and Monitoring Tools > Details… > Simple Network Management Protocol.

Configure the SNMP service on the Traps tab so that it sends SNMP packets to the management station. The example below shows the OAM management station running at 192.168.1.160.

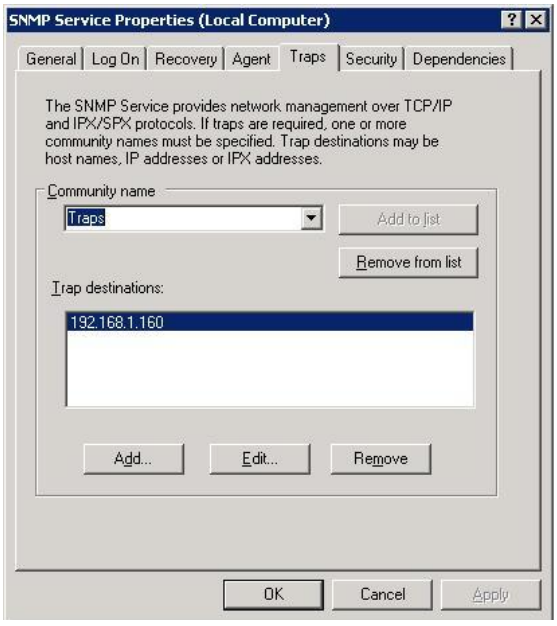

Figure 8 Windows SNMP Service Configuration

Ensure the SNMP service accepts the community name you are planning to use to access the performance variables, for example, **traps**. Read-only access is adequate.

Preparing the SNMP Management Station

To monitor the traps sent by the WAAS Mobile service and to request and display counter information requires an application that can send and receive SNMP data over the network, interpret it according to the WAAS Mobile MIB and display it. Such an application is part of SNMP management station software, which is usually remote from the WAAS Mobile server (in a test environment, it might run on the server machine). There are many such applications available, and they are all managed in a similar fashion. The main configuration aspects, common to all SNMP management station software and relevant to testing traps from WAAS Mobile, are:

- Ensure any needed software is available. For example, on Windows the management station software might rely on the service called "Windows Trap Service" to receive traps (service this is not required by WAAS Mobile – it only provides for receiving traps).
- Import (compile) the WAAS Mobile MIB. This provides the management station with information to interpret the OIDs in the trap data.
- Tell the management station to listen to the WAAS Mobile server. This means to tell it the IP address of that machine and the port it is using to send SNMP traps (via UDP). The port is determined by the SNMP service on the server machine and is usually left at the default (162). You may need to specify the trap community name as well, consistent with the setting used by the Windows SNMP service.
- Tell the management station which "community" to listen to for traps. The default for Windows SNMP traps is "trap", which is set on the Traps tab of the SNMP service. The community is a primitive security mechanism. If you don't listen for the right community, you will not see the trap.
- Tell the management station which community to use for requesting performance counters. Because all WAAS Mobile performance counters are read-only, it is reasonable to use the community **public**.
- Tell the management station whether to use SNMPv1 or SNMPv2.

In addition it is useful to have NetMon or a similar packet capture application available to capture the SNMP packets that the WAAS Mobile server will send to the management station.

Troubleshooting SNMP Notifications ("Traps")

1. Is the WAAS Mobile Server process running? If not, start it and make its start automatic.
2. Are NT Events enabled on the WAAS Mobile server? If not, enable them.
3. Is the SNMP Service running? If not, start it and make its start automatic.
4. Is AccelSnmpXa.dll loaded? If not, check if it is installed and registered properly as described above. If it is correctly installed and registered but not loaded, the DLL is corrupt and the WAAS Mobile server software should be reinstalled.
5. Check SNMP Service Properties Traps tab. Is the community name what you expected? Is the monitoring station address in the list of trap destinations? If either of these conditions is not met, make the appropriate changes.

6. Traps originate on the WAAS Mobile server. To troubleshoot, follow the packet trail using NetMon.
7. Using NetMon on the WAAS Mobile server, capture SNMP packets generated when the server process is restarted, monitoring on the interface on which they are sent. If none are captured, return to step 1 of this troubleshooting guide.
8. If the packets are being sent, inspect their content to see if they are being sent to the expected host and port.
9. On the monitoring station, capture packets on the interface on which they should arrive. If none arrive when they are known to be generated on the server, check the local firewall (if any), network connections and the IP routing arrangements of your network.
10. Check that the server machine IP address corresponds to one of the entities to which the management station is listening, and that the entity is set to use SNMPv1 or SNMPv2.

After all these steps are successfully completed the traps will be displayed.

Troubleshooting Counters

1. Is the WAAS Mobile Server process running? If not, start it and make its start automatic.
2. Are NT performance counters enabled on the WAAS Mobile server? If not, enable them.
3. Is the SNMP Service running? If not, start it and make its start automatic.
4. Is AccelSnmpXa.dll loaded? If not, check it is installed and registered. If it is correctly installed and registered but not loaded, the DLL is corrupt and the WAAS Mobile server software should be reinstalled.
5. Check SNMP Service Properties Security tab. Is the community name what you expect (e.g., public)? Is the monitoring station address in the list of trap destinations? If either of these conditions is not met, make the appropriate changes.
6. Requests for counters originate on the monitoring station. To troubleshoot, follow the packet trail using NetMon.
7. Using NetMon on the monitoring station, capture SNMP packets generated when the management station tries to access the performance variables, monitoring on the interface on which the packets are sent. If none are captured, check that the management station is sending to the correct host.
8. If the packets are being sent, inspect their content to see if they are being sent to the expected host and port.
9. On the WAAS Mobile server, capture packets on the interface on which they should arrive. If none arrive when they are known to be generated on the server, check the local firewall (if any), network connections and the IP routing arrangements of your network.
10. If the packets are arriving, check if any packets are being sent back. It is common for the SNMP service to notify of problems in a response packet. If the returned packets contain "authentication failure" indications, check the Accepted Community Names on the Security tab of the SNMP Service, and check the security name associated with the Context that is being used in by the management station.

After all these steps are successfully completed, the counters will be displayed.

Chapter 8 Persistent Sessions

The Cisco WAAS Mobile "Persistent Sessions" feature maintains acceleration sessions even when web connectivity is lost or when a mobile client switches to a different network such as from Wi-Fi to cellular. When connectivity is restored, the current session is sustained to create a seamless access experience regardless of the changes in the underlying network structure. Downloads and uploads are resumed without loss of data, and no additional log-ins are required.

Persistent Sessions insulates the end-user from problems with RF coverage in wireless networks as well as from problems in poor quality dial-up access. It allows the acceleration system to support advanced wireless network features such as automated Wi-Fi/cellular switchover or hand-offs when roaming through different cellular networks.

Configuration

Persistent Sessions is enabled or disabled for the client via The WAAS Mobile Manager > Client Configuration > Network Settings page.

Session Removal

With the Persistent Sessions feature, the acceleration session between client and server is maintained when network connectivity is lost and until certain conditions are met.

The server always assumes that the most recent session from a client is still active. The server closes a session when one of 3 events occurs:

- The server receives a restart message from the client.
- A request for a new session is received from a client who has an existing session.
- A session remains inactive for an interval longer than a threshold defined in the registry. This is currently set to 1 hour.

The client closes a session when one of 3 events occurs:

- The client receives a restart message from the server.
- A session remains inactive for an interval longer than a threshold defined in the registry. This is currently set to 1 hour.
- A network connection is present but the client has not received any data from the server after a certain amount of time (20 minutes, by default).

IP Independence

In some deployments, clients may not have the same IP when they reconnect or when they roam to a different network. The WAAS Mobile server will recognize the client even if the IP presented to the server has changed. In addition, for deployments of multiple acceleration servers, the WAAS Mobile client-side load balancing feature can be used to ensure that the client will reconnect successfully.

Application Layer Keep-Alives

Many web browsers, email clients, and application servers will terminate a session if they detect an inactive connection. During the time that the client-proxy link is unusable, WAAS Mobile keeps the TCP connections to the client and server applications open for a predetermined period of time. It also sends application layer messages for HTTP and email that prevent shutdown of the application session before service is restored. Other applications will time-out according to their tolerated interval of inactivity.

Chapter 9 Configuring User Authentication

This chapter describes how to configure user authentication using the WAAS Mobile Manager administration interface.

Authentication Settings

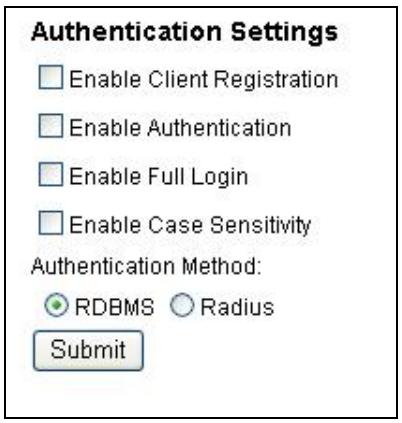

Figure 9 User Authentication Settings

Enable Client Registration

The client registration feature is simply a dialog that accepts input about the user. Client registration information is communicated to the server in the login request packet, and then stored in the server's database.

When client registration is used, the Email address field will be used in the session logging record in place of the user name field of the authentication dialog.

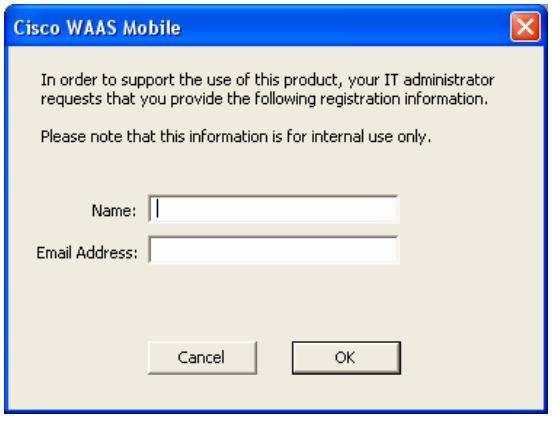

Figure 10 Client Registration Data Entry

Enable Authentication

If this checkbox is checked then authentication using user names is enabled. See below for discussion of User Management.

Enable Full Log-in

This determines whether passwords are also to be used for authentication.

Enable Case Sensitivity

With this checked, user names are case sensitive.

Authentication Method

Select a radio button to enable either RADIUS or RDBMS authentication.

NOTE: All changes in Authentication configuration require a server restart.

User Management

The User Management section allows you to add and remove users from the service.

If the Enable Authentication checkbox is checked, then the WAAS Mobile server is set to authenticate the users that connect to it. Two authentication methods are available:

User name (only) - When Authentication is enabled but Full Log-in is not, the server will authenticate users by user names only. Valid user names are stored in a list on the server via the User Management page shown in Figure 11.

User name and Password (Full Log-in) - When Enable Full Login is checked, the server will authenticate by user name as well as a password, which are entered on the User Management page. It is important to note that if Full Log-in is enabled, a password must be entered for each user – a user will not be able to log in if the password field is left blank.

On the WAAS Mobile Manager > Client Configuration > User Interface page, be sure to select the option Display Authentication UI and, if Full Log-in is enabled, Display Password Entry. The first time the client launches, users will be prompted to enter their user name and password if applicable.

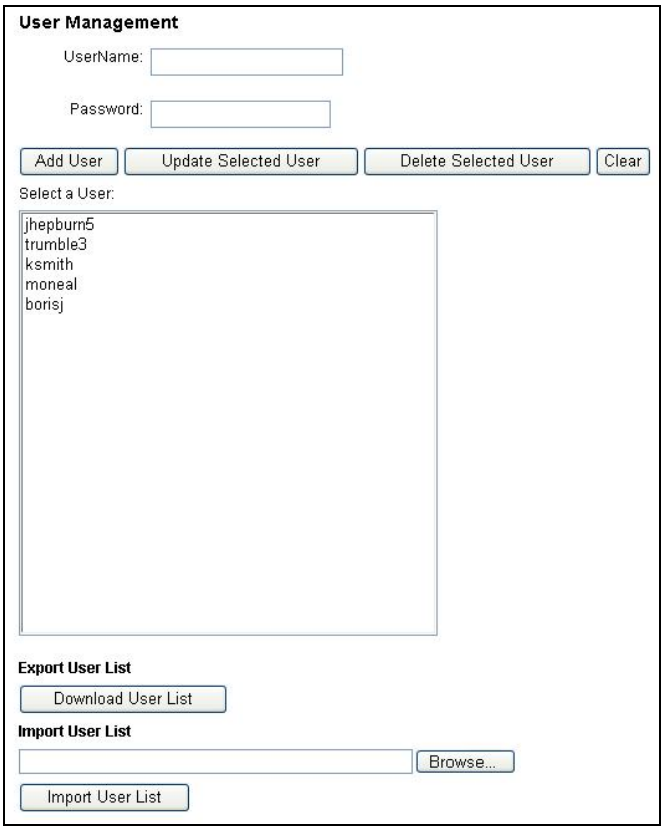

Figure 11 User Management Configuration Form

Adding a User

Type a new user name in the UserName box and click Add User. The name will appear in the user list.

Updating a User Name

Select the user name in the list. Type the new name in the UserName box and click Update Selected User. The change will appear in the user list.

Clearing the User Name Entry

To remove a user's name from the UserName box, click Clear. (This does not delete the user; it simply deselects the name, so that another name can be selected.)

Deleting a User

Select the user name in the list and click Delete Selected User. The name will be removed from the user list.

Exporting a User List

To back up the list of users, or to transfer the list to another server, save the user list to a file by clicking Export to a file. You will be prompted to enter a filename for the user list.

Importing a User List

To import a list of users, find the filename using the Browse button, then click Import this file.

> NOTE: This will overwrite your current user list.

RDBMS Authentication

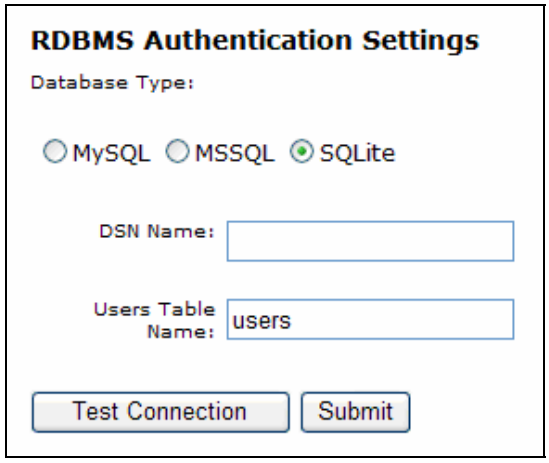

Figure 12 RDBMS User Authentication Configuration

By default, WAAS Mobile uses an internal SQLite database accessed via ODBC, although an alternate back-end database, supporting either MySQL or MS SQL server RDBMS, is also supported. User names and passwords (if applicable) are configured on the User Management page. The Test Connection button confirms that WAAS Mobile Manager can reach the database.

If using one of the alternate back-end databases, make sure the ODBC DSN is a System DSN and that the database name has been created in the chosen SQL database beforehand.

Configuring WAAS Mobile to use RADIUS Authentication

The WAAS Mobile server can be configured to use one or more RADIUS servers as a means of handling authentication requests. This section describes the configuration procedure.

Managing the authentication of many clients in a short time is important for WAAS Mobile. This is because short server outages at peak times result in all clients trying to connect within a few seconds after the server becomes available. This can result in high transient loads on the RADIUS server. Settings are provided for dealing with this by limiting the number of outstanding requests, and accepting client connections that arrive when the limit has been reached. This has the effect of trimming the authentication peaks.

The WAAS Mobile server also detects when a RADIUS server is offline and has the option of allowing users to get acceleration without authentication in this case. Unresponsive RADIUS servers are periodically retried.

Requirements

1. The RADIUS server must be configured so the WAAS Mobile server is a recognized RADIUS client. This requires configuration of at least an IP address and a shared secret on the RADIUS server. A corresponding shared secret must also be configured for each RADIUS server on the WAAS Mobile server (in the registry sub-tree under Options\Auth\RADIUS\Servers). See the following sections for more details.
2. The RADIUS server used in conjunction with WAAS Mobile must support PAP authentication requests. Usually this requires a specific configuration step.
3. On the WAAS Mobile server, it is necessary to choose to send either NAS-Identifier or NAS-IP-Address in the Access-Request packets. The RADIUS server must be configured to be consistent with this choice.
4. The RADIUS server must be able to authenticate the user names-password pairs you supply to the WAAS Mobile client.
5. The RADIUS server should be configured to not terminate strings with a null character. The relevant RFCs require this behavior. However, some RADIUS servers can be configured so that text attributes they send include a terminating null character. This can cause trouble with using Filter-Id.

Troubleshooting

If RADIUS Authentication is enabled on the WAAS Mobile server without any specific configuration for the RADIUS server IP address or anything else, users will be authenticated. This is because the IP address for the RADIUS server defaults to "unknown", which is not a valid IP address. In that case the RADIUS server is assumed to be unavailable and the default action is to authenticate the user. No RADIUS packets are sent in this case.

If a correct IP address is configured for the RADIUS server, but the RADIUS server does not recognize the IP address of the WAAS Mobile server as a RADIUS client, the RADIUS server will not respond. Consequently clients will be authenticated after a timeout period. Subsequent behavior depends on how often and how many clients try to connect. In this case there is likely to be a timeout period before the authentication of all clients. The default period is 25 seconds.

On the other hand, in a production setting, if there are many requests in a short time with no responses from the RADIUS server, it will be ignored after a certain time, and all WAAS Mobile clients will be authenticated immediately. This behavior is configurable. The default is that if at least 30 requests are made in a period of 30 seconds and there is no response, the server is ignored for five minutes. After that, requests are sent to it again.

Sometimes a user name/password dialog appears on the WAAS Mobile client. If the RADIUS server recognizes the WAAS Mobile server, but the Access-Request packets sent to the RADIUS server do not match some criterion on the RADIUS server, it will send back an Access-Reject, and it is this that results in the user name/password dialog. The actual reason for this can be seen in the logs on the RADIUS server (or in the Windows Event log, if it is running on Windows).

Possible reasons are:

- The user name/password is not configured on the RADIUS server.
- The RADIUS Access-Request packets do not conform to some criterion on the RADIUS server. For example, they do not have a recognized NAS-Identifier.

When setting up for initial testing you may find the following situation:

- The RADIUS server logs indicate an incorrect user name/password.
- The test client is authenticated.
- There is no user name/password dialog on the test client.
- A packet trace shows that access-reject packets are being sent.

This indicates a mismatch in the shared secret. A System Report will show that the authenticator in the response packets did not match what was expected. If the WAAS Mobile server is put in production with a mismatched shared secret, it behaves as if it is receiving no responses from the RADIUS server (because the WAAS Mobile server ignores responses with a bad authenticator). Consequently, if there are sufficient client requests happening, the RADIUS server will eventually be ignored and all clients authenticated. For this reason, it is important to ensure that the RADIUS server works properly with a single test client before transitioning to production.

Minimal Configuration Using WAAS Mobile Manager

For a new installation on a machine that has not had WAAS Mobile server installed before, these are the minimal actions needed to activate RADIUS authentication. This simple configuration uses NAS-Identifier to identify the NAS, rather than NAS-IP-Address. It uses a single RADIUS Authentication server and the default RADIUS Authentication port 1812.

1. In WAAS Mobile Manager, choose Server Configuration > Authentication.
2. Check the Enable Authentication box, select the RADIUS Authentication radio button, and click the Submit button.
3. Go to RADIUS Authentication and enter data as shown in Figure 13 below. The NAS Identifier value, WAAS Mobile, must be recognized by the RADIUS server.

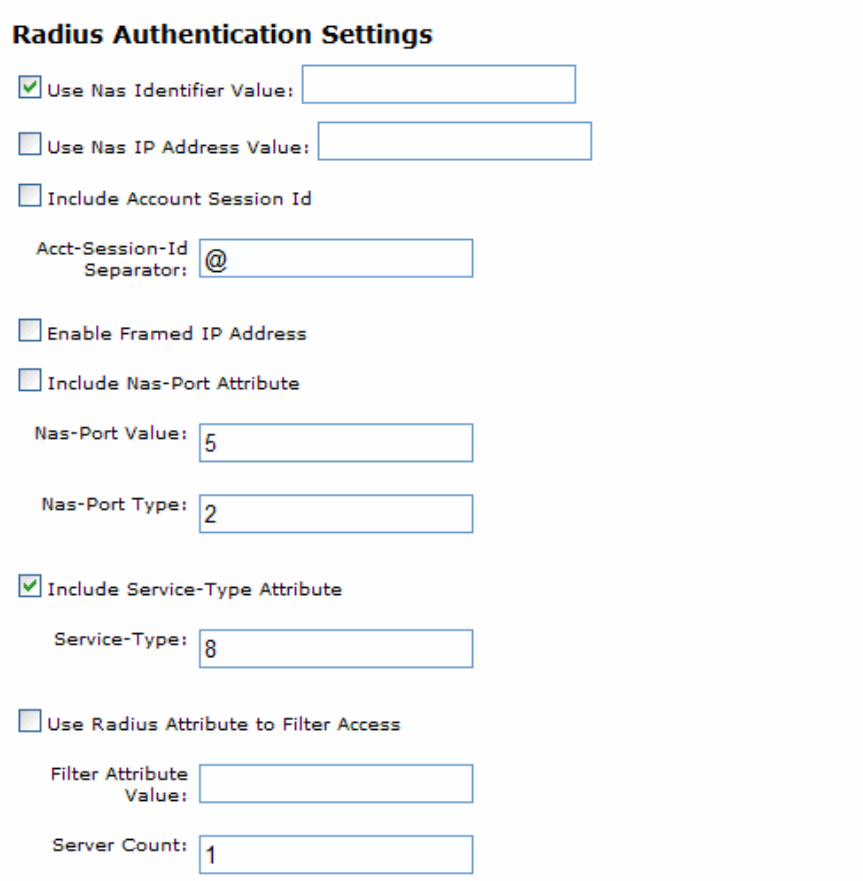

Figure 13 RADIUS User Authentication Settings

Show the RADIUS server details and enter data as in Figure 14 below. The remote address must be the address of the RADIUS server. The shared secret must be associated with the WAAS Mobile NAS Identifier on the RADIUS server.

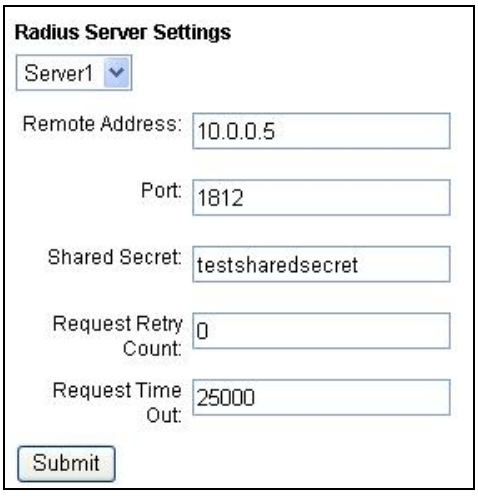

Figure 14 RADIUS Server Data Configuration

At this point, if the RADIUS server is set up consistently with the current simple WAAS Mobile server setup, clients will authenticate using RADIUS, and a failure to authenticate will indicate an incorrect user name/password combination.

Configuring Access-Request Packets

The Access-Request packets always include the user name and password. Other attributes can be optionally included or excluded. The keys and values used to control the attributes are all located under [NASidentifier]\Options\Auth\RADIUS. All values referred to in this subsection are relative to this key.

NAS Identification

Identification of the NAS can use either NAS-Identifier or NAS-IP-Address. The default is NAS-Identifier.

To include NAS-IP-Address, set the value of DWORD UseNasIpAddress to 1 and set the IP address in the string value NasIpAddress in dotted-decimal format.

To avoid including NAS-Identifier, set DWORD UseNasIdentifier to 0.

To use NAS-Identifier but use a value different from the default of WAAS Mobile, set the string value NAS to the desired identifier. There is no limit on the length.

Filter-ID

The Filter-ID attribute is used to allow or deny acceleration to users who have a recognized user name/password. To enable it, create the DWORD value CheckAttribute and set it to 1. This requires that the RADIUS server sends a Filter-ID in its Access-Accept packet, or the user will not be authenticated. The value in the Filter-ID attribute must be "WAAS Mobile". It is case-sensitive. To require a different value, create the key FilterAttribute. Under it create a string value called Value that is set to the required text.

To require the RADIUS server to send a different text attribute (i.e., not Filter-ID) with a specific value, create a DWORD value called Id under the FilterAttribute key and set it to the desired RADIUS attribute value. This is rarely used – Filter-ID is the standard attribute for this use.

Using this attribute requires coordination with the RADIUS server. If the value sent by the RADIUS server does not match the value expected by the WAAS Mobile server, all authentication requests will be rejected.

A user who requests acceleration and for whom the Filter-ID is not sent or is incorrect will have a "WAAS Mobile Alert" user Name / Password dialog box displayed.

Some RADIUS servers can be configured to include a null character at the end of text attributes. This will cause the attribute value to be rejected even though it looks correct in the RADIUS server configuration file. The RADIUS RFCs require that no null be included at the end of text attributes.

NAS-Port-Type

To enable the NAS-Port-Type attribute, create a key NasPortType. Underneath it create a DWORD value called Enable and set its value to 1. This will cause the default value 5,

corresponding to “Virtual” to be sent. To send a value different from the default, create a DWORD value called Value under the NasPortType key. Set its value to the value you require. The standard values from IETF RFC 2865 are:

0 Async

1 Sync

2 ISDN Sync

3 ISDN Async V.120

4 ISDN Async V.110

5 Virtual (the default)

6 PIAFS

7 HDLC Clear Channel

8 X.25

9 X.75

10 G.3 Fax

11 SDSL - Symmetric DSL

12 ADSL-CAP - Asymmetric DSL, Carrierless Amplitude Phase Modulation

13 ADSL-DMT - Asymmetric DSL, Discrete Multi-Tone

14 IDSL - ISDN Digital Subscriber Line

15 Ethernet

16 xDSL - Digital Subscriber Line of unknown type

17 Cable

18 Wireless - Other

19 Wireless - IEEE 802.11

NAS-Port

NAS-Port is not included by default. There are two options for the NAS-Port value:

1. A fixed value is sent.
2. A unique number is sent.

To include the NAS-Port attribute, create a key called NasPort. Under it create a DWORD value called Enable and set it to 1. This will send a unique integer with each request. The integer increases monotonically, even across server restarts. The value is remembered between restarts in the registry value Status\AuthIdBase.

To send a fixed value for NAS-Port, create a DWORD value called Type under the NasPort key. Set its value to 1. Then create a DWORD value under the NasPort key called Value. Set it to the value you wish to send.

To avoid sending the NAS-Port attribute without removing the keys and values you have added, set Enable under the NasPort key to 0.

To switch back from sending a fixed value to sending a unique value, set Type under NasPort to 2.

RADIUS Accounting can be configured to send the same value for the NAS-Port attribute in accounting request packets as the one chosen here. However the configuration is slightly different. Instead of three values (Enable, Type and Value), it uses just two (Mode and Value).

Service-Type

To enable the Service-Type attribute, create a key ServiceType. Underneath it create a DWORD value called Enable and set its value to 1. This will cause the default value 8, corresponding to "Authenticate Only" to be sent. To send a value different from the default, create a DWORD value called Value under the ServiceType key. Set its value to the value you require. The standard values from IETF RFC 2865 are:

1 Login
2 Framed
3 Callback Login
4 Callback Framed
5 Outbound
6 Administrative
7 NAS Prompt
8 Authenticate Only (the default)
9 Callback NAS Prompt
10 Call Check
11 Callback Administrative

Framed-IP-Address

The Framed-IP-Address attribute contains the client IP address. To enable it, create a key called FramedIpAddress. Under it, create a DWORD value called Enable and set its value to 1. To disable Framed-IP-Address without removing these registry settings, set Enable to 0.

Acct-Session-Id

The Acct-Session-Id attribute can be included in the Access-Request packets. It is a text value that can also be included in the accounting request packets. As a result it can be used to help analyze login requests and accounting requests.

To enable Acct-Session-Id, create a key called SessionId. Under it, create a DWORD value called Enable and set its value to 1. The result will be session ID values that look like 11@213, where the first number is the server restart count since installation and the second number is the session ID. The restart count is maintained in the registry value <AcceleNetServerRoot>\Status\RestartCount.

To change the separator used between the session number and the restart count, create a string value called Separator under the Session Id key, and set its value to the desired separator. It can be of any length.

Other NAS Configuration

Dealing with RADIUS Server Failure

By default, if a user is not specifically rejected, they are authenticated. This means that timeouts and other errors are ignored. This is to avoid denying users acceleration service simply because the authentication server is not available.

This behavior is controlled initially by the registry DWORD value TreatNoResponseAsAuth under the key Options\Auth. Its default value is 1. To get finer control over the way timeouts and other errors are handled, set it to 0.

When TreatNoResponseAsAuth is 0, the DWORD registry value ResultIfNoResponse under the key Options\Auth\RADIUS determines what the WAAS Mobile server sends to the WAAS Mobile client when the RADIUS server does not respond to an Access-Request. This lack of response includes the case where the shared secrets on the WAAS Mobile server and the RADIUS server do not match.

The default is to allow access to the WAAS Mobile server (option 0). This can be one of the values in the table below:

Value	Result
0 (default)	User authenticated
1	User not found
2	User bad password
3	Timeout
4	Server busy
5	Server not responding
6	No servers available

Special Settings for Initial Testing

It is recommended that initial testing of the authentication setup be done with a single client and with settings that make it obvious if the RADIUS server is not responding. To do this:

1. Set the DWORD value TreatNoResponseAsAuth under the key Options\Auth to 0.
2. Set the DWORD value ResultIfNoResponse under the key Options\Auth\RADIUS to 3, indicating timeout.

To see these special settings in action, change the shared secret on the WAAS Mobile server to an incorrect value. The result on the test client's Connection Monitor display should be "server busy". This happens because the responses from the RADIUS server are ignored when it is detected that there is a mismatch in the shared secret.

After the RADIUS authentication has been correctly configured these values can be deleted.

Slowing the Request Rate

The following values can be used to reduce the rate at which requests are sent to the RADIUS server. These are only needed in cases where the RADIUS server cannot cope with the peak request rate generated by the WAAS Mobile server. This would occur after an outage of the server, when there were a large number of clients waiting to connect at the same time.

Registry Value Name	**Type**	**Default**
MaxTotalOutstandingRequestCount	REG_DWORD	10000
UseBlockingSendWindow	REG_DWORD	1

The behavior when the maximum outstanding request count is reached depends on the value of UseBlockingSendWindow.

If it is 1, the WAAS Mobile server will wait until a response is received before issuing each new request to the RADIUS server, so that the maximum number of outstanding requests is not exceeded. We refer to this as using a blocking send window.

If UseBlockingSendWindow is 0 the behavior is different. Instead of 'blocking', the request is immediately provided with a response indicating Server Busy. This reduces the load on the RADIUS server as well as quickly authenticating waiting clients.

RADIUS Server Configuration

The minimal procedure described above configured the smallest possible amount of data for a single RADIUS server. This section describes additional configuration details for RADIUS servers, including configuring more than one RADIUS server.

Number of RADIUS Authentication Servers

Up to four RADIUS servers can be used. RADIUS servers are used for requests in a round-robin fashion. Once the server is decided for a request, any retransmissions are made to that server. Set the number of servers in the DWORD registry value *ServerCount* under the Servers key.

Required Values for RADIUS Authentication Servers

For each RADIUS authentication server, at least the following values must be configured under the *Server*<*n*> key, where <*n*> is the number of the server (1, 2, 3 or 4):

RemoteAddress: This registry string value must be the IP address of the RADIUS server. The default value is "unknown", which is not a valid IP address.

SharedSecret: This registry string must match the shared secret set on the RADIUS server. The shared secret set on the RADIUS server can be different for every RADIUS client. The selection of shared secret to compare is made using the IP address of the RADIUS client. So it is necessary to configure a NAS on the RADIUS server and to give it the IP address that the WAAS Mobile server uses to send packets to the RADIUS server. The default value is "*testsharedsecret*".

Optional Values For RADIUS Authentication Servers

RADIUS Listener Port: The port on which the RADIUS server listens may need to be changed from the default (1812). To do this, set the DWORD RemotePort under the *Server*<*n*> key to the desired value. On Windows 2003 Server, the Internet Authentication Service listens for RADIUS authentication packets on both 1812 and 1645 (the originally assigned port). So if Microsoft IAS is being used, this configuration step might not be necessary. However, many sites using RADIUS use 1645, and for them this configuration step is required.

Local Bind Address: Usually the WAAS Mobile server will select an appropriate local address to which to bind for sending RADIUS authentication packets. This bind process determines the source IP address of the packets that are sent. In the case of multi-homed servers and servers with multiple IP addresses per interface it may be desired to specify exactly which address on the machine is used. Do this by setting the registry string LocalAddress under the *Server*<*n*> key to the IP address required.

Authentication Request Retry Parameters

The following parameters control the way authentication requests are retried if no response is received from the RADIUS server. In almost all cases the default values are effective. Their descriptions are provided here for completeness.

Registry Value Name	**Type**	**Default**	**Comment**
RequestTimeout	REG_DWORD	25000	Milliseconds
RequestRetryCount	REG_DWORD	0	
RequestTimeoutGrowthFactor	REG_DWORD	100	%

The main parameters are *RequestTimeout* and *RequestRetryCount*. The default settings mean that just one send attempt is made, and the WAAS Mobile server waits 25 seconds for a response. These settings have been found to be the best ones to allow very high request rates such as happen just after the server comes back from a restart at times of peak demand. In effect they delegate responsibility for retrying requests to the WAAS Mobile client.

The *RequestTimeout* value should be less than the DWORD value ConnectTimeout under the key Network\Session. If the latter value is not present it defaults to 30 seconds.

The RequestTimeoutGrowthFactor is not used unless *RequestRetryCount* > 0. In that case it can be used to increase or decrease the time between successive retries.

RADIUS Server Failure Detection Parameters

The following parameters control the way RADIUS servers that appear not to respond to any requests for a period of time are handled. In almost all cases the default values are effective. Their descriptions are provided here for completeness.

Registry Value Name	Type	Default	Comment
TestInterval	REG_DWORD	30	Seconds
TestRequestCount	REG_DWORD	30	Count of unanswered requests in TestInterval seconds
TestRetryInterval	REG_DWORD	300	Seconds
MaxOutstandingRequestCount	REG_DWORD	256	

The WAAS Mobile server keeps track of the number of authentication requests for which no responses were received. It does this over successive periods each TestInterval seconds long.

If, in one of these periods, there are at least TestRequestCount requests and no responses, the RADIUS server will be assumed to be unresponsive. After this is detected, all requests will be immediately given the response defined in the ResultNoServersAvailable parameter, and that result is modified according to the setting of TreatNoResponseAsAuth described in section Other NAS Configuration, Dealing with RADIUS Server Failure above.

After a time interval of TestRetryInterval seconds (default value five minutes), requests are once again directed to the server.

The WAAS Mobile server also detects some kinds of network problems accessing RADIUS and will periodically attempt to contact the RADIUS server if network problems make it unavailable.

Limiting Request Rate per RADIUS Server

The DWORD value MaxOutstandingRequestCount under the *Server*<*n*> key limits the number of requests that can be simultaneously in progress. Its default value is 256, which is the maximum possible for RADIUS. In almost all cases this parameter can be left at its default value.

Chapter 10 RADIUS Accounting Configuration & Monitoring

The WAAS Mobile server can be configured to use one or more RADIUS servers as a means of recording accounting details for client sessions. This section describes the configuration procedure.

The WAAS Mobile server detects when all RADIUS servers are offline and stores accounting packets internally until a server is again available. This minimizes the likelihood of lost accounting information.

> NOTE: The WAAS Mobile server must be restarted for registry changes to take effect.

Requirements

1. The RADIUS server must be configured so the WAAS Mobile server is a recognized RADIUS client. This requires configuration of at least an IP address and a shared secret on the RADIUS server. A corresponding shared secret must also be configured for each RADIUS server on the WAAS Mobile server (in the registry sub-tree under Options\Accounting\RADIUS\Servers). See the following sections for more details.
2. On the WAAS Mobile server it is necessary to choose to send either NAS-Identifier or NAS-IP-Address in the Acct-Request packets. The RADIUS server must be configured to be consistent with this choice.

Troubleshooting

The first step in troubleshooting is to check the names of registry keys and values you have entered by hand. A misnamed value will be ignored. If you have set a value and you do not see the behavior you expect, first assume there is a misspelled registry entry.

Tools

Troubleshooting can be done using

- The RADIUS server logs
- If the RADIUS server runs on Windows, the Windows Event Log.
- Packet traces using Netmon, Ethereal or something similar on the WAAS Mobile server and the RADIUS server.
- The WAAS Mobile server System Report, if enabled.
- Exported WAAS Mobile server registry subtree.
- RADIUS server product name and version, and configuration files if available.

Operation

When the WAAS Mobile Server service starts, the RADIUS Accounting Client (RAC) reads its configuration from the registry. This includes its NAS parameters, the parameters for each RADIUS server it can distinguish on the network, and the current restart count. It also increments the restart count and writes it back to the registry.

The RAC sends a request to each server, with retries if the timeout is exceeded. If the maximum number of tries is reached without a response, the server is marked as unresponsive and put on a wait list for later retrying. This retry is done with the original parameters. No further requests are sent to the server until it has responded to the test request. If a RADIUS server does not respond to the test request it will be retried later, and so on, indefinitely, using ***ServerRetryInterval*** to determine the timing.

If there is a response to the test request the server is marked as responsive. Subsequent Start, Stop and Off requests are sent to it.

Accounting requests are sent to all responsive servers in a round-robin, using the order specified under the servers registry key.

Each "send" involves sending and waiting for an acknowledgement, possibly multiple times. The number of tries and the timeout is defined per server. The timeout should be coordinated with the round-trip time and server response time. Once the request is acknowledged, the request data is dropped. When the limit on the number of tries is reached the request is considered to have failed for this server. It is returned to the main RAC request list for further retrying, perhaps to another RADIUS server.

Whenever no RADIUS servers are marked as responsive, accounting requests are queued, up to a maximum number specified in ***MaxQueuedRequests***. Requests are dropped, oldest first, to avoid infringing this limit as new requests are added. If detailed logging is enabled in the registry these drops are logged.

If at any time the result of sending an Accounting Request is a network error (for example, ICMP "port unreachable"), the RADIUS server is marked as closed and put on a wait list for later retrying. This retrying involves re-reading the server parameters from the same place in the registry as the parameters originally came from (that is, the same *Server*<*n*> subtree). It uses ***ServerRetryInterval***.

An Acct-Off request is sent when the server is shutting down. The shutdown process waits until it gets a result from the request using the usual retry and timeout parameters for the server it was sent to, and then continues shutting down. The result might be success or failure, meaning no response from the RADIUS server. In either case the server continues shutting down.

Minimal Configuration Using WAAS Mobile Manager

For a new installation on a machine that has not had WAAS Mobile server installed before, these are the minimal actions needed to activate RADIUS accounting. This simple configuration uses NAS-Identifier to identify the NAS, rather than NAS-IP-Address. It uses a single RADIUS Accounting server and the default RADIUS Accounting port 1813.

Go to WAAS Mobile Manager > Server Configuration > Advanced Settings > RADIUS Accounting page and enter values shown below:

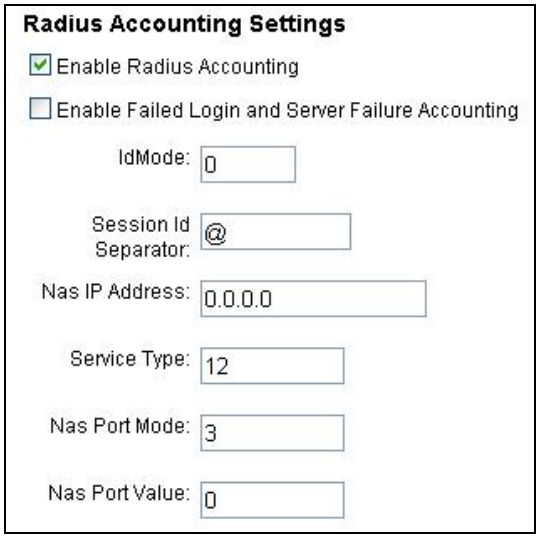

Figure 15 RADIUS Accounting Configuration

On the same page, for the RADIUS server settings, enter values similar to those shown in Figure 16 below.

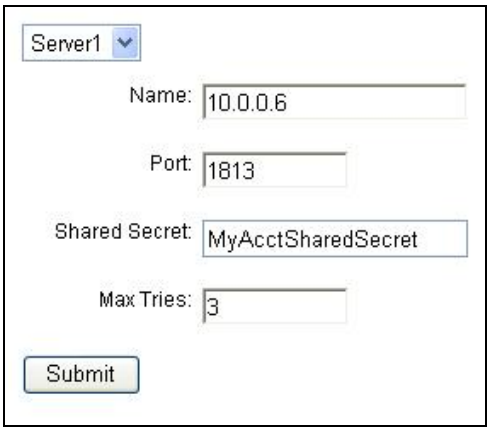

Figure 16 RADIUS Accounting Server Data Entry

Here the "Name" must be the IP address of the RADIUS Accounting server. The shared secret must be associated with the NAS Identifier that is the host name of the WAAS Mobile server (the value of the HOSTNAME environment variable).

Minimal Configuration Using the Registry

For a new installation on a machine that has not had WAAS Mobile server installed before, these are the minimal actions needed to activate RADIUS accounting. This simple configuration uses NAS-Identifier to identify the NAS, rather than NAS-IP-Address. It uses a single RADIUS Accounting server and the default RADIUS Accounting port 1813.

1. Enable RADIUS accounting by setting the DWORD value

   AcceleNetServer\Accounting\Mode = 1

2. Create the keys

   AcceleNetServer\Accounting\RADIUS\Nas

   AcceleNetServer\Accounting\RADIUS\Servers

   AcceleNetServer\Accounting\RADIUS\Servers\Server1

3. Create the string values

   AcceleNetServer\Accounting\RADIUS\Servers\Server1\Name

   AcceleNetServer\ Accounting\RADIUS\Servers\Server1\SharedSecret

Set them to the values configured on the RADIUS server you will be using. Notice that "Name" is the IP address of the RADIUS server, not its DNS name. The shared secret is usually specific for each NAS.

You should also ensure that there is a system environment variable called HOSTNAME and it is set to the name you want this server to be known by.

At this point, if the RADIUS accounting server is set up consistently with the simple WAAS Mobile server setup, RADIUS accounting will be operational. Each time the WAAS Mobile server starts or stops, and each time a session begins or ends, a RADIUS accounting packet will be sent. Note that in this simple configuration the packets contain the NAS-Identifier attribute, and its value is the value of the environment variable HOSTNAME on the machine in question.

Enabling RADIUS accounting (in detail)

Accounting is disabled by setting Accounting\Mode to 0. Setting Accounting\Mode to 1 enables RADIUS accounting. The sessions recorded are only those that are completely initiated. Sessions that were partially constructed but never completed are ignored. Setting Accounting\Mode to 2 enables RADIUS accounting and includes all sessions for which a successful authentication was received. In some cases, the RADIUS server may fail to respond, and this is treated as successful authentication. These sessions are recorded as well.

Configuring Accounting Packets

The keys and values in this section are under Accounting\RADIUS\Nas.

Acct-Status-Type

Every accounting packet contains this and there are four values: On, Off, Start and Stop. These correspond to starting/stopping the WAAS Mobile server, and to starting/stopping each session.

User-Name

This is always sent. if client authentication using user names is in effect. It contains the user name supplied by the WAAS Mobile client, or part of that name. If the client user name contains a "/" character, only the part of the name before that character is used. Otherwise the whole user name is used.

If user names are not in effect ("Active Users" license), the IP address of the client is converted to dotted-decimal form and used.

Acct-Multi-Session-Id

This is sent if there is a user name and it contains a "/". This attribute reports the part of the user name after the "/".

NAS Identifier

The minimal configuration above uses the default NAS identification method, which is to use the NAS Identifier. The minimal configuration did not include any specification of the NAS-Identifier to use. This is because the WAAS Mobile server automatically uses the value of the HOSTNAME environment variable as the NAS-Identifier.

There are two ways to change the value of the NAS-Identifier. To use a different environment variable, set *NasIdEnvName* to the name of the environment variable. To specifically set the NAS-Identifier without using environment variables, set the string value *Identifier* to the desired value.

To avoid including NAS-Identifier, set the DWORD value IdMode to 1. That uses NAS-IP-Address instead, as explained in the next section.

NAS-IP-Address

The alternative to NAS-Identifier is to use NAS-IP-Address. To include NAS-IP-Address instead of NAS-Identifier, set the DWORD value IdMode to 1, and set the value of the IP address in the string value NasIpAddress. Setting the value NasIpAddress also makes the WAAS Mobile server bind to that address, rather than selecting any suitable local address from the ones available on the machine.

To include BOTH NAS-Identifier and NAS-IP-Address, set IdMode to 2.

Framed-IP-Address

This is always sent in Start and Stop packets. It contains the IP address of the client in dotted-decimal notation.

Called-Station-Id

This is always sent in Start and Stop packets. It is the IP address of the WAAS Mobile server as seen by the end-user's machine. It is represented as a string in dotted-decimal notation.

Transformed IP Address

If IP Aliasing is being used, the WAAS Mobile server reports the transformed client IP address in the RADIUS accounting packets. It uses attribute number 114. The attribute is sent as a 4-byte integer in network byte order. When IP aliasing is not used, this attribute is reported as zero.

Service-Type

The Service-Type is sent in all packets. The default value is 12. To send a value different from the default, create a DWORD value called Value under the ServiceType key. Set its value to the value you require. The standard values from IETF RFC 2865 are:

1 Login
2 Framed
3 Callback Login
4 Callback Framed
5 Outbound
6 Administrative
7 NAS Prompt
8 Authenticate Only (the default)
9 Callback NAS Prompt
10 Call Check
11 Callback Administrative

NAS-Port

By default, NAS-Port is not included in accounting packets. To include it, create a registry key *NasPort*. Under it create a DWORD value called Mode. There are then three possibilities:

1. Use a fixed value. Setting Mode to 1 selects this behavior. In that case, create the DWORD value *Value* under the NasPort key and set its value appropriately. RADIUS accounting has a similar, but separate option.
2. Use the session count. This option is provided for historical reasons. To obtain it, set Mode to 2.
3. Use a unique login ID. If this is chosen, the same login ID (an integer) as used for the RADIUS authentication packets will be used in the accounting packets. The ID increases monotonically, even across server restarts. The value is remembered between restarts in the registry value Status\AuthIdBase.

To disable the NAS-Port attribute without removing the keys and values used to control it, set Mode to 0.

Acct-Session-Id

All RADIUS accounting packets from WAAS Mobile server include the Acct-Session-Id. The session ID values that look like 11@213, where the first number is the server restart count since

installation and the second number is the session ID. There is no option to change this order. The restart count is maintained in the registry value Status\RestartCount.

To change the separator used between the session number and the restart count, create a string value called SessionIdSeparator and set its value to the desired separator. It can be of any length. The default separator is "@".

Event-timestamp

This is always sent. It is the time in seconds since January 1, 1970 00:00 UTC at which the event documented by the RADIUS packet occurred on the WAAS Mobile server. The time is from the clock on the WAAS Mobile server, which is not guaranteed to be synchronized with the RADIUS server.

Acct-Delay-Time

This is always sent. It indicates the number of seconds the WAAS Mobile server has been trying to send this packet. An approximation of the time the event occurred on the WAAS Mobile server with respect to the RADIUS server clock can be obtained by subtracting this from the time at which the packet arrived at the RADIUS server.

Attributes For Traffic Volume and Session Length

The attributes listed in the table below are always sent in Stop packets.

Acct-Session-Time	The length of the session in seconds.
Acct-Input-Octets	The uncompressed bytes that would have been sent from the client to the server. This is the size of the data before compression.
Acct-Output-Octets	The uncompressed bytes that would have been sent from the server to the client. This is the size of the data before compression.
Acct-Input-Packets	The compressed bytes that were sent from the client to the server.
Acct-Output-Packets	The compressed bytes that were sent from server to the client.

Acct-Terminate-Cause

Acct-Terminate-Cause is always sent in Stop packets. The values sent can be interpreted using the table below.

The termination codes reported are 32 greater than the values returned from the WAAS Mobile server to avoid colliding with predefined RADIUS values (0-18).

Code	Reported	Name	Description
0	32	NoError	Not used (reserved)
1	33	TransportError	ITP (Transport Protocol) reported that heartbeat stopped, the link is low quality, or the ITP session itself is in a bad state

2	34	SessionError	An error occurred at the compression level for the session
3	35	CtrlSktError	Not used (reserved)
4	36	CtrlSktDead	The TCP connection was reset
5	37	DataSktDead	The UDP connection was reset
6	38	ServerError	Not used (reserved)
7	39	LoginFailure	Login failure
8	40	Timeout	The client did not respond in time to complete the session creation sequence
9	41	CtrlDataInvalid	The client login request was invalid
10	42	CtrlSktHibernate	The client went into hibernation mode

Summary of RADIUS Attribute Options

Registry Value Name	Description		Default
IdMode	This determines how the RADIUS client identifies itself. 0 = Use NasId. 1 = Use NAS-IP-Address. 2 = Use both. If IdMode is 1 or 2, set NasIpAddress.		0
NasIdEnvName	This is the environment variable used to obtain the NasId if the name is not set in the registry. Normally this is HOSTNAME and is set by WAAS Mobile. Usually let this use the default.	String	"HOSTNAME"
Identifier	If set (to something other than the default), this value will be used as the NAS-ID entry in the accounting packets. Usually this is not set, and the HOSTNAME environment value is used.	String	"unknown"
NasIpAddress	The IP address to bind to locally, and to use when IdMode is set to 1 or 2.	String	0.0.0.0
ServiceType	Integer. Sent in all packets.	dword	12

Registry Value Name	Description		Default
NasPort\Mode	0=do not include; 1=use fixed value set below; 2=use session count; 3=use unique authentication id.	dword	3
NasPort\Value	When NasPort\Mode=1, use this value in NAS-Port.	dword	0
SessionIdSeparator	The text to use to construct the session ID, in the form session-id<SEP>restart-count Here the <SEP> is the separator specified in this item.	String	"@"

Other NAS Configuration

Dealing with RADIUS Server Failure

Occasionally, a RADIUS accounting server stops responding to requests for a short time. In that case the WAAS Mobile server keeps the packets that were to have been sent for sending later, when the RADIUS server is available again. The behavior can usually be left at the default. The parameters are documented in the table below.

If there is a network error (usually an ICMP network unreachable packet), the server is put on the "bad server list". The retry goes to the registry to get the parameters the next time it retries. This means that as long as ICMP is enabled you can, for example, change the RADIUS server IP address without stopping the WAAS Mobile server to reconfigure. Change the registry and it will pick it up on the next server retry.

Registry Value Name	Description		Default
ServerRetryInterval	Interval between attempts to send packets to an unresponsive RADIUS accounting server, in seconds.	dword	30
MaxQueuedRequests	This is the maximum number of unsent RADIUS requests that can be queued for sending when all RADIUS accounting servers are offline. Each stored packet is around 100 bytes in size.	dword	100000
MaxFailures	This is used to determine if a server should go on the unresponsive list. This happens whenever the number of failures during the check interval exceeds this number. A server on the unresponsive list is later retried.	dword	3
FailureCheckInterval	This is the interval between checks on the number of failures. This should be set to a period that is large enough to get MaxFailures failures, taking into account the expected logon rate, the server timeout and the number of retries. In seconds.	dword	10

Number of RADIUS Authentication Servers

Up to 10 RADIUS servers can be used. There is no need to enter the number of servers as a value in the registry. The server keys are automatically enumerated to determine the number.

The maximum allowable number of servers is controlled by the value MaxServerCount. The default value is 3, the maximum is 10. This parameter controls the initial search for servers in the registry.

Servers are entered by adding keys called *Server1*, *Server2*, … under the *Servers* key.

Required Values for RADIUS Accounting Servers

For each RADIUS accounting server, at least the following values must be configured under the *Server*<*n*> key, where <*n*> is the number of the server (1, 2, 3…):

Name: This registry string value must be the IP address of the RADIUS server. The default value is "unknown", which is not a valid IP address.

SharedSecret: This registry string must match the shared secret set on the RADIUS server. The shared secret set on the RADIUS server can be different for every RADIUS client. The selection of shared secret to compare is made using the IP address of the RADIUS client. So it is necessary configure a NAS on the RADIUS server and to give it the IP address that the WAAS Mobile server uses to send packets to the RADIUS server. The default value is "*testsharedsecret*".

Optional Values for RADIUS Accounting Servers

RADIUS Listener Port: The port on which the RADIUS server listens may need to be changed from the default (1813). To do this, set the DWORD value Port under the *Server*<*n*> key to the desired value. On Windows 2003 Server, the Internet Authentication Service listens for RADIUS accounting packets on both 1813 and 1646 (the originally assigned port). So if Microsoft IAS is being used, this configuration step might not be necessary. However, many sites using RADIUS use 1646, and for them this configuration step is required.

Local Bind Address: Usually the WAAS Mobile server will select an appropriate local address to which to bind for sending RADIUS authentication packets. This bind process determines the source IP address of the packets that are sent. In the case of multi-homed servers and servers with multiple IP addresses per interface it may be desired to specify exactly which address on the machine is used. Do this by setting the registry string Identifier under the Nas key to the IP address required. The same local bind address will be used for all communication with RADIUS Accounting servers.

Accounting Request Retry Parameters

The following parameters control the way accounting requests are retried in if no response is received from the RADIUS server. In almost all cases the default values are effective. Their descriptions are provided here for completeness.

Name	Type		Default
Timeout	The time to wait for a response to an accounting request, in milliseconds.	dword	5000 (max 60000)
MaxTries	The number of times a request will be sent to the RADIUS server if the server does not respond.	dword	3 (max 10)
TimeoutGrowthFactor	After each attempt, the client can be made to wait longer before the next attempt. This is the factor to use to increase the timeout. It is in %. For example, 150 means that the timeout will be multiplied by 1.5 each time.	dword	200

Additional Server Parameters

The following parameters are provided here for completeness.

Exclude	Use 1 as the value of this to prevent a RADIUS server from being used.	dword	0
MaxSendRate	Used to control the maximum rate at which requests can be sent to this server. In requests per 30 msec interval. Not usually used.	dword	50

Events

It is assumed that the NT events are stored on each server, not centrally. This is important because each WAAS Mobile server will tend to respond identically to problems with the RADIUS servers. This could generate significant network traffic if not properly managed.

There are three kinds of NT events for the WAAS Mobile RADIUS client:

- Informational: start and stop, successful test for RADIUS server availability.
- Warning: one or more RADIUS servers unreachable. An unreachable server will be temporarily removed from the list of servers and will be retried later. Successive retries will use an exponential backoff of the timeout.
- Error: a range of possible error conditions, including missing configuration items, out of memory and other software exceptions, or all RADIUS servers unreachable.

Chapter 11 HTTPS Optimization

WAAS Mobile provides a secure proxy for SSL traffic. This enables SSL traffic to be compressed, just like other non-secure traffic, without compromising security.

In SSL communication, the secure server provides its certificate to the client; the client decides if the certificate represents the server and is trusted. With WAAS Mobile, the secure server's certificate is reissued by the WAAS Mobile server, and it is the reissued certificate that the client compares with expectations. The WAAS Mobile server acts as a certificate authority (CA) to perform the reissuing function. There are two main scenarios:

1. The WAAS Mobile server CA is a root authority (i.e., it is self-signed).
2. The WAAS Mobile server CA is a subordinate authority (i.e., its certificate is issued by another CA).

Use of SSL proxy for client machines is enabled and configured using Client Configuration as described in the WAAS Mobile Administration Guide (see the "HTTP/HTTPS Settings" section of Chapter 4). Server-related configurations, and in particular whether to run the server CA as self-signed or subordinate are controlled through registry settings.

Controlling SSL Proxy Tunnel

The following tables summarize the controls available for SSL proxy. Except where otherwise noted, all registry values are under the key Options\HTTPS.

Enabling

Use of the SSL proxy is controlled with the value DoProxySsl. Set this to 1 on the server to make SSL proxy available to all clients. Set it to 0 on the server to disable SSL proxy for all clients. On the client, DoSslProxy controls use of SSL proxy for the client itself. The client negotiates with the server, requesting SSL proxy. If DoProxySsl is set to 1 on the server, the request is accepted. Otherwise the request is denied and the client does not attempt to proxy SSL traffic. This arrangement allows an administrator to disable SSL proxy for all clients without any client-side changes.

The value of UseSelfSignedCaCert can be used to control whether the trusted CA on the WAAS Mobile server is a root or subordinate CA.

Root CA on the WAAS Mobile server requires the CA certificate to be installed in a trusted store on the client, leading to one or more pop-up dialog boxes. This is the default mode of WAAS Mobile operation and is convenient to use for product evaluation.

Subordinate CA is the appropriate setting for enterprise use. In this case, it is necessary to submit a certificate request to an Enterprise CA to get the CA certificate for each WAAS Mobile server. The section "Certificate Request Controls" shows how to control the production, content and submission of the certificate request. Set UseSelfSignedCaCert to 0 to enable this mode.

The value UseHttpsHostnameList limits the use of SSL proxy to a specific list of servers. It is used on the client. If 1, the value HttpsHostNameList is consulted for a list of hosts for which SSL proxy is allowed. If UseHttpsHostnameList is 1 and the HttpsHostNameList is empty, no SSL proxy is performed.

Name	Type	Default
DoProxySsl	DWORD	1
UseSelfSignedCaCert	DWORD	1

Certificate Request Controls

In the case where the Impersonation CA is subordinate, it is necessary to create a certificate request and submit that to your Enterprise CA. It is necessary to define the data to be sent in the request; the file and folder where the request is created; and how the request will be submitted to the CA. The default settings are appropriate for most situations.

Data in the Certificate Request

The following data is supplied in the certificate request:

Subject name

A company URL

The public key

Certificate template identification

The Subject Name is a distinguished name that is constructed from SslHostCertIssuerName by appending HostId to the Common Name (CN=) part. HostId as generated is a globally unique identifier across all WAAS Mobile server machines. For example, the subject name will be displayed as

CN = Cisco Inc. 000C0D3F-9490-4AB8-8D68-EEF2292EA562
O = Cisco
C = US

(The CN number is the value of HostId)

The company URL is provided so that end-users can find out information about the company that is running the proxy. This may help them feel more comfortable with the SSL proxy arrangement.

The public key comes from the key set named by KeySetNameCA.

The certificate template identification is for use with Microsoft CAs. Either a version 1 or version 2 certificate template can be specified. The default is a version 1 template with name “SubCA”. This corresponds to the Subordinate Certification Authority template in the Windows 2003 server CA.

Name	Type	Default
SslHostCertIssuerName	REG_SZ	"C=US,O=Intelligent Compression Technologies,CN=ICT Inc."
SslHostCertIssuerNameUntrusted	REG_SZ	"C=US,O=Cisco,CN=Secure SSL

		Server Proxy Authority For Reissuing Untrusted Original Certificates"
HostId	REG_SZ	""
HostIdSelfSigned	REG_SZ	""
CACertUrl	REG_SZ	"http://www.Cisco.com"
CertTemplateVersion	DWORD	1 (can be 0, 1, 2)
CACertType	REG_SZ	"SubCA"
CertTemplateObjId	REG_SZ	"1.3.6.1.4.1.311.10.12.1"
CertTemplateMajorVersion	DWORD	1
CertTemplateHasMinorVersion	DWORD	1 (= "true")
CertTemplateMinorVersion	DWORD	0

Submitting the request

There are two ways to submit the request. One is online, the other is manual.

The request can be submitted electronically if the CA is online and accessible over the network from the WAAS Mobile server. This is a rare situation and applies only to Microsoft CA software. Usually the CA will be set up to accept requests manually. If the CA is online, set the value IsThirdPartyCAOnline to 1, and set the name of the CA in the registry value ThirdPartyCAIdentifier. Use the same format as used in the table below. If the CA is online the process of obtaining a CA certificate for WAAS Mobile server is completely transparent – the request is submitted electronically, the certificate chain is received electronically, and the certificates are installed in the certificate store without administrator intervention.

There are different kinds of certificate request. Most commercial CA software can deal with any of them. The different kinds available are:

1: XECR_PKCS10_V2_0

2: XECR_PKCS7

3: XECR_CMC

4: XECR_PKCS10_V1_5

The default is 1.

The more usual case (the default) is where the certificate request is submitted manually. In this case IsThirdPartyCAOnline is set to 0. The certificate request is created as a file and placed in the location specified by the values CertRequestFileFolder and CertRequestFileName. This happens when the WAAS Mobile server first starts, and also whenever it subsequently starts and does not find its CA certificates in the specified store.

When this happens:

- The SSL proxy capability is switched off on the server, as if DoProxySsl were 0.

- A message is written in the server log file, saying that the certificate request has been made, and documenting where it is located.
- Execution continues otherwise as normal.

The administrator retrieves the certificate request file and submits it to the Enterprise CA. The Enterprise CA issues the certificate as a file, or issues the whole certificate chain as a file. The choice of whether to export certificate or certificate chain depends on the choices made by the end user and the certificate enrollment interface being used.

The certificate file is imported into the certificate store specified by [CertStoreCA, IsMachineCertStoreCA]. This is usually the machine personal store. Then the WAAS Mobile server is restarted. After restarting, check the log file to ensure that the server has restarted successfully.

Name	Type	Default	Client/ Server
CertRequestType	DWORD	1	S
CertRequestFileFolder	REG_SZ	""	S
CertRequestFileName	REG_SZ	"AccelCA.req"	S
IsThirdPartyCAOnline	DWORD	0	S
ThirdPartyCAIdentifier	REG_SZ	"10.0.0.1\CA"	S

Starting WAAS Mobile as Self-Signed CA

In this configuration the WAAS Mobile server acts as a root certification authority.

Leave all HTTPS settings at the default.

No special configuration is required. Simply install the WAAS Mobile server and start it.

Checking that it works

1. Start the WAAS Mobile client. On the first start, a popup dialog box will appear. If the certificate is not accepted, the client will start, but SSL proxy capability will be absent. If the certificate is accepted, SSL proxy capability will be available.
2. Open a browser.
3. Visit a secure site. It should open without any problems (i.e., no pop-up dialog boxes or other warnings relating to security).
4. Double-click the padlock icon. A certificate will be displayed (note: browsers other than Internet Explorer 6 will display padlocks, but displaying the certificate may require a different action).

The issuer name that is displayed should begin with SSL Proxy CA and be followed by a GUID. The GUID should be the same as the registry value Options\HTTPS\HostIdSelfSigned on the WAAS Mobile Server.

If this is not the case, see the section on troubleshooting below.

Starting WAAS Mobile Server as Subordinate CA

In this configuration, the WAAS Mobile server acts as a subordinate certification authority. The WAAS Mobile server will produce a certificate request to submit to your Enterprise CA.

This assumes all settings are left at the default except UseSelfSignedCaCert. This should be set to 0. This configuration is recommended for deployment to production.

Before Server Installation

The procedure in this section is only required once, not for each time WAAS Mobile server is installed.

1. Export the complete certificate chain for the Enterprise CA you will use to create the WAAS Mobile server CA certificate. You can do this using the MMC certificate snap-in on Enterprise CA. Export the certificate chain as a p7b file. (See the Troubleshooting section below for a brief description of the MMC certificates snap-in).
2. Import the chain of certificates into the machine personal store on the WAAS Mobile server machine.
3. Move the root of the certificate chain into the trusted machine store if it is not already there.
4. Check that the Enterprise CA certificate is trusted on the WAAS Mobile server. If it is not, it is likely that one or more of the certificates in the certificate chain are not up to date. Repeat the procedure to this point with up-to-date data.

Server Installation and first start

1. Install the server software.
2. Start the WAAS Mobile server. Check the log file after the start. There will be a message that looks something like this:

   HttpsOptions::InitializeServer: the attempt to initialize the server proxy CA certificate failed, or is not yet complete. The result of the attempt is: ***The certificate request was placed in a following file, which should be submitted to a certificate authority. See the Integration Guide for details. (Filename=C:\Program Files\Cisco\WAASMobile\WAAS MobileServer\AccelCA.req)***

   This indicates where to find the certificate request file.
3. Submit the file (e.g., AccelCA.req) to your Enterprise Certificate Authority (Enterprise CA) to get a certificate file.
4. Import the certificate into the personal machine store on the WAAS Mobile server machine.
5. Restart the WAAS Mobile server. Check the log file. You should see a message like this near the end:

   HttpsOptions::InitializeServer: the attempt to initialize the server proxy CA certificate succeeded. The result of the attempt is: success

6. If you are using a WAAS Mobile server farm, repeat this procedure for each server in the farm.

There is no need to repeat this procedure if the WAAS Mobile server is reinstalled. Only repeat the procedure if there is a note in the WAAS Mobile server log like the one displayed above.

Clients

Your Enterprise CA must be trusted on your client machines. That means that the Root Authority for your Enterprise CA must be in each user's trusted store on each client machine. There are two ways this can be arranged:

- If your Enterprise CA is ultimately certified by a well-known root authority such as Verisign, the certificates for that root authority are already installed in the trusted store as part of the operating system installation. In this case no special action is required.

If your Enterprise CA is self-signed, an administrative action is required to install the certificate for your CA in the trusted store for each WAAS Mobile user on each client machine. This is usually done when the computers and user names are configured.

Checking that it works

1. Start WAAS Mobile client.
2. Open a browser.
3. Visit a secure site. It should open without any problems (i.e., no pop-up dialog boxes).
4. Double-click the padlock icon. A certificate will be displayed. (This is for Internet Explorer 6. Other browsers also display padlocks but displaying the certificate may require a different action).

The issuer name that is displayed should end with a GUID. The GUID should be the same as the registry value Options\HTTPS\HostId on the WAAS Mobile server.

If this is not the case, see the section on troubleshooting.

Troubleshooting

1. Check that DoSslProxy is 1 on both client and server.
2. If UseHttpsHostnameList is 1, check that HttpsHostnameList contains the host you are trying to reach.

MMC Certificates Snap-in

Start the management console using the command "mmc". It should contain a blank window. If it doesn't, make one using the Console > New menu item from the main window. Then use the Console > Add-Remove Snapin menu item, click Add... on the dialog that comes up, select Certificates from the list, click Add, click the Computer Account radio button, click Next..., make sure Local Computer is selected, and click Finish. You might as well also add a view of certificates in the user store as well. Just click Certificates again, click Add, click My User Account, click Finish, click Close, and click OK. You will see a tree view of certificates. Use the context menus to view certificates, export them and import them.

Root WAAS Mobile Server CA

After the client has been run and the certificate popup accepted, check that the WAAS Mobile server CA certificate is in the user's trusted certificate store. If not, send a System Report to the Cisco Technical Assistance Center (TAC) for analysis.

Subordinate WAAS Mobile Server CA

Check the messages relating to the SSL Proxy in the WAAS Mobile server log on startup. If the message following, "The result of the attempt is" is not "certificate request created" or "success," send a System Report to the Cisco Technical Assistance Center (TAC) for analysis.

Check that the WAAS Mobile server CA certificate is present in the personal machine store on the WAAS Mobile server machine. If not, it must be obtained and installed.

Check that the WAAS Mobile server CA certificate is trusted on the WAAS Mobile server. If not, the certificate chain for the Enterprise CA used to issue the WAAS Mobile server CA certificate must be imported into the personal machine store, and the root of the chain must be imported into the trusted machine store.

Run the WAAS Mobile client on a client machine. After it connects successfully, check that the WAAS Mobile server CA certificate is in the user's personal certificate store on the client machine after the WAAS Mobile client has been run. If not, send a System Report to the Cisco Technical Assistance Center (TAC) for analysis.

Check that the WAAS Mobile server CA certificate in the user's personal certificate store on the client machine is trusted. If not, send a System Report to the Cisco Technical Assistance Center (TAC) for analysis.

Inspecting Reissued Host Certificates

In rare cases it may be necessary to inspect the reissued web server certificates. To do this, set PersistHostCertificates to 1 on the client and/or server in order to write the host certificates processed into the personal store. Be careful if you choose to do this on a production server, as many certificates will be written to disk.

Name	Type	Default	Client/ Server
PersistHostCertificates	DWORD	0	CS

Popup (or other alert) On Client

On occasion you will visit a secure web site and the browser will present you with a popup dialog box. The dialog box presents information about the reissued web server certificate that has been created by the WAAS Mobile server from the original web server certificate. There are usually three things to look at:

- Is the reissued certificate trusted?
- Is the reissued certificate within its date range?
- Is the name on the reissued certificate correct?

If the reissued certificate is not trusted, it means the original certificate was not trusted on the WAAS Mobile server. The usual cause of this is that the root of the certificate chain for the web server certificate is not in the trusted machine store on the WAAS Mobile server machine. Web server certificates are almost always issued by a globally recognized CA that is pre-installed on all major operating systems. Lack of trust for a certificate from a public web server strongly suggests that the original certificate should be viewed with suspicion.

The situation is entirely different when the web server is inside your own enterprise. In that case the likely cause for this error is that the root of the web server certificate has not been imported into the trusted machine store on the WAAS Mobile server machine. This will almost certainly be the case in three common scenarios:

- The web server certificate is self-signed. In this case, import the web server certificate itself into the trusted machine store on the WAAS Mobile server machine.
- The web server certificate is signed by an enterprise CA that is self-signed. In this case, import the CA certificate into the trusted machine store on the WAAS Mobile server machine.
- The web server certificate is issued using a two-level enterprise CA, in which the root certificate is self-signed. In this case, import the root CA certificate into the trusted machine store on the WAAS Mobile server machine.

On Windows operating systems, use the certificates MMC snap-in to import, export and view certificates.

Another cause of lack of trust is date range problems. The WAAS Mobile server issues an untrusted certificate in this case. If everything else is OK and the date range is wrong, it is common practice to accept the certificate.

If the name in the certificate (more precisely, the Common Name part of the Subject Name of the certificate) does not match the name of the web server requested, you have to seriously question whether the certificate should be accepted. On the other hand, you can make this happen by using the IP address to contact the web server instead of using a domain name. Seeing as IP addresses are not usually included in web server certificates, this will generate a warning that the certificate name does not match.

Cisco WAAS Mobile Integration Guide