



## DEPLOYING CISCO WAAS MOBILE IN LARGE-SCALE ENTERPRISES

Today, there are already about 450 million mobile workers, and the majority of businesses see the trend to increased mobility growing. Today, it is virtually taken for granted that mobile workers cannot achieve the same levels of productivity when accessing corporate files and applications as do workers back in the office.

Cisco WAAS Mobile is a comprehensive acceleration platform designed specifically to address the IT, user, and business challenges of mobile and home office (SOHO) workers in even the largest organizations. These users' requirements are significantly different and often more challenging than those who work in larger branch offices.

This white paper focuses on the key architectural considerations and product features that provide the scalability, management, and ease-of-deployment required for large, globally distributed enterprises.

### SCALING WAAS MOBILE FOR LARGE ENTERPRISES

When large numbers of Cisco WAAS Mobile clients are distributed on PCs and laptops across a globally-distributed enterprise, there are several factors to consider, including:

- Load Balancing and High Availability
- PC-friendliness
- Manageability
- Architecture Scalability
- Mobility
- Security Supporting text

#### Load Balancing and High Availability

WAAS Mobile's high availability and scalability is based on the following key architectural capabilities:

- continuous health monitoring of server connections with automatic fault recovery, and
- policy-based server farm and server selection.

#### Continuous Health Monitoring with Automatic Fault Recovery

The Cisco WAAS Mobile client continuously monitors the health of the connection to the server, and automatically bypasses acceleration if server connectivity is lost. Because the client makes the decision, users never lose network connectivity if an upstream appliance/server goes down. In the event of a persistent network or server failure, the client automatically selects an alternate server based on server selection policies that are defined and maintained by a Controller Server, as described further below.

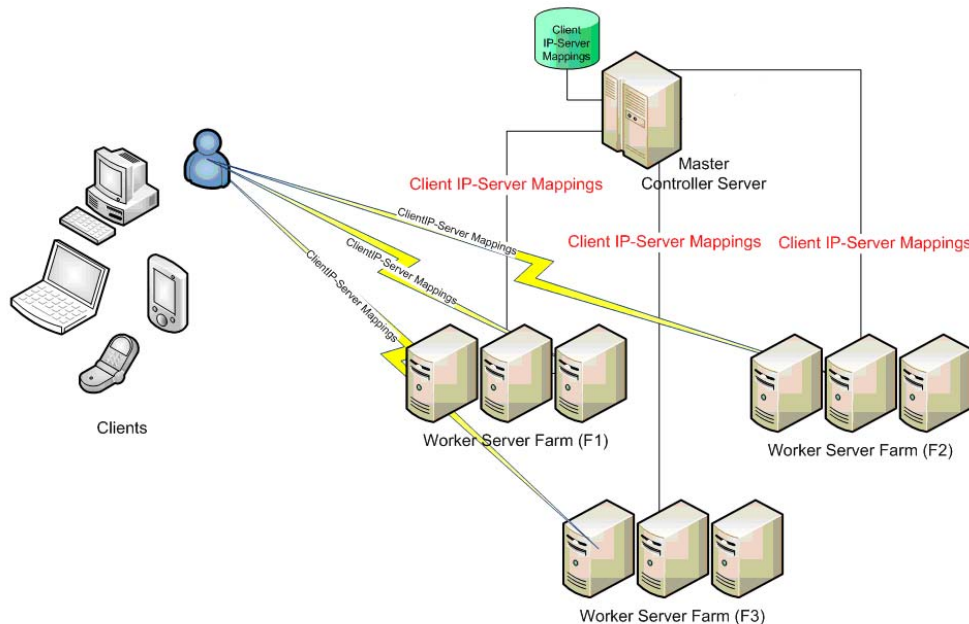
#### Policy-Based Server Farm and Server Selection

Generally, the best acceleration performance is achieved when a client consistently connects to the same server, maximizing commonality between the server and client traffic histories, and resulting in the highest differencing and compression gain.

WAAS Mobile incorporates an integrated load balancing technique for optimizing server selection. WAAS Mobile's Advanced Server Selection, which employs client-side logic, automatically selects the best server farm (in a global deployment) and the best server within the farm. The selection of the best farm can be made based on latency or IP subnet mapping and the selection of the best server within a farm can be based on a prioritized list or random selection.

Advanced Server Selection involves the client making an informed decision as to 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 this case, the client has to determine first which farm to connect to (farm selection) and second which server in that farm to connect to (server selection). The end result is a dynamic solution that is flexible and adaptable to any enterprise network infrastructure.

As shown in the figure below, WAAS Mobile Controller Server maps client IP subnets to one or more Worker Server farm(s), each with one or more WAAS Mobile Worker Servers. This mapping also provides clients with dynamic instructions on how to select the appropriate farm and server regardless of where they connect to the enterprise WAN.



**Figure 1. Master Controller Server Load Balancing.**

Upon startup and periodically thereafter, the Controller Server communicates the current mapping table to each Worker Server across all server farms. Each Worker Server then communicates this information to each client that connects to it. The client uses a best effort approach to use the same Worker Server on the next login attempt in order to reap the benefits of Persistent Sessions and/or Persistent Delta Cache.

In some scenarios, it may be desirable to employ Cisco Application Content Engine (ACE). Using ACE to load balance WAAS Mobile Servers is suggested for large scale enterprise or service provider data centers where ACE is already deployed. The scenarios in which ACE should be considered for load balancing are where a) the client IP address is fixed (non-mobile users), and b) an intermediate proxy server is not masking the source address. The Cisco WAAS Mobile Advanced Server Selection and Cisco ACE may be used in

conjunction with one another when multiple server farms are deployed; Advanced Server Selection selects the farm and ACE provides load balancing within the farm. (Note that in this scenario, the Controller Server still must communicate with the Worker servers, so a routed path must be enabled between them.)

## **Manageability**

A central management and monitoring facility supports the deployment of WAAS Mobile clients to a large number of PCs. Through the management console, the administrator configures, updates, and monitors clients as well as overall system performance.

Customers with large WAAS Mobile deployments have found the following management capabilities to be essential:

- **Software installation.** Client profiles are packaged as distributions in executable .exe and .cab formats, which can be downloaded using familiar enterprise software management tools or through a web interface.
- **Software upgrades.** Upon login, the client communicates with the server to verify that it is running the latest software release. If not, the client automatically downloads any components that need to be updated without interrupting user activity and without requiring a reboot.
- **Configuration updates.** Upon login, the client checks whether its configuration profile has been modified, and automatically updates its configuration accordingly.
- **Policy-based management.** Administrators may create separate configuration profiles for different user groups, enabling customized configurations based on geographic location, network connectivity, unique application acceleration requirements, or end-user preference.
- **Central monitoring console.** Graphical displays of acceleration and traffic breakdown provide management with the ability to directly measure ROI and administrators with the insight necessary to monitor the health of the acceleration network.

## **PC Friendliness**

The challenge of successfully deploying acceleration software to users around the world, who have a wide range of machines running an even wider range of software and who have very limited IT training and local support resources, should not be underestimated.

The WAAS Mobile client has been designed from the ground-up for optimum operation in the PC environment, and avoids the issues associated with simply porting an appliance design to a PC. As a result, the WAAS Mobile client provides the lowest PC footprint and greatest range of PC hardware, OS, and software interoperability. For example, WAAS Mobile's plug-and-play client requires between 3x and 30x less RAM than competitive products.

WAAS Mobile functions transparently, and provides proven interoperability with user applications. It is compatible with a wide range of browsers, email clients, virus scanners, IPsec as well as SSL VPNs, thin clients (i.e., Citrix, Windows Remote Desktop) and with all recent version Windows including Windows Vista. WAAS Mobile automatically detects other software applications that could potentially conflict, and dynamically adjusts how WAAS Mobile is installed so as to avoid conflicts.

Moreover, since WAAS Mobile intercepts traffic at the socket and transport layers, the client avoids undesirable interactions with PC-based network monitoring utilities (e.g., Wireshark/Ethereal, DUMeter, etc.). Appliance-derived clients are prone to these problems because they intercept traffic at the network driver level. By intercepting at a higher layer of the stack, WAAS Mobile is also able to support acceleration of active FTP traffic as well as HTTPS and clientless SSL VPNs.

When issues do arise and troubleshooting support is desired, a simple one-click diagnostic captures all relevant information, packages it for analysis, and transmits it to the server. Using the system report analyzer tool, technical support engineers can examine the machine environment, statistics between client and server and between server and content server, detailed network connection information, timing and compression of requests and responses, and various performance summaries and analysis graphs.

## Architecture Scalability

The WAAS Mobile server architecture has been designed to smoothly scale to thousands of active users simultaneously accessing content or applications. The architecture has been carefully engineered to maximize server throughput by minimizing consumption of RAM, disk storage, disk seek, and CPU resources. Some of the techniques include:

- Highly scalable storage system
  - Each file or data sequence is only stored once
  - Single instance of a file or data sequence is shared with all users
- Highly efficient memory utilization
  - Only requires 2 MB of server RAM for each simultaneous active download. This is 1 to 2 orders of magnitude more efficient than competitive approaches
  - 1000:1 disk to RAM ratio for search index supports deep histories. Thus, a 1 TB history only requires a 1 GB in-RAM search index. This is 4-5X more efficient than competitive approaches.
- Scalable CPU utilization
  - Multi-threaded architecture makes efficient use of multi-core CPUs
- Optimized disk utilization
  - Because disk throughput is limited by the time it takes to mechanically move read/write heads to new locations on the disk, Cisco WAAS Mobile employs a dynamic disk seek algorithm that optimizes throughput under high load by dynamically trading off acceleration gain versus disk activity, thereby mitigating disk thrashing.

## Mobility

Mobility creates unique requirements for PC-based acceleration. In particular, traveling users often connect with different IP addresses each time they access the server. Additionally, truly mobile users may transition between networks frequently.

For this reason, WAAS Mobile supports:

- delta cache portability
- persistent sessions.

## Delta Cache Portability

Appliance-derived PC acceleration solutions commonly ‘zero out’ a user’s history each time the PC connects through a different farm, effectively clearing the user’s delta cache every time he or she travels.

Cisco WAAS Mobile, however, is designed for portability, such that the software client immediately begins accelerating content upon connecting to any Cisco WAAS Mobile server. When initiating the connection, the client and server exchange information to determine which portions of the delta cache history held by the client are also available on the server, and will automatically use the history information that they have in common for delta caching. To see the value of this capability, consider the following example in which a user based in North America who generally connects to a North American data center travels to Asia and connects through an Asian data center. The data that is common between the user’s PC and the Asian server will be used while the user is in Asia. When the user returns, the user will resume accessing the North American server and the shared history will once again be used.

## Persistent Sessions

Many web browsers, email clients, and application servers will terminate a session if they detect an inactive connection.

The WAAS Mobile Persistent Sessions feature maintains acceleration sessions even when connectivity is lost or when a mobile client switches to a different network such as from Wi-Fi to cellular. 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. 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.

## Security

PCs used by mobile users and in small offices are exceptionally vulnerable to theft, malicious attack, and traffic interception, and pose a high security risk in any organization. To mitigate these threats, many organizations require that all data be accessed and maintained securely.

Only WAAS Mobile provides a multi-layer, end-to-end solution that enables enterprises to secure data while in motion and at rest. End-to-end security features include:

- Acceleration of secured applications
  - HTTPS acceleration - without distributing private keys
  - Signed SMB file transfer acceleration
- Security for data-in-motion
  - SSL VPN acceleration in all key modes, including clientless, thin client, and full tunnel deployment modes
  - IPsec VPN interoperability
  - Encryption of all WAAS Mobile client/WAAS Mobile server traffic
- Security for data-at-rest
  - Encrypted delta cache store

It should be noted that Cisco WAAS Mobile does not require a predefined open port on the client PC or laptop for incoming connections, which is a clear security risk when using public networks (hotels, airports, etc.).



## **SUMMARY**

Cisco WAAS Mobile is the only acceleration client that can meet the scalability requirements of large enterprises. WAAS Mobile is a comprehensive acceleration platform designed specifically to address the small office and mobile network challenges of even the largest organizations. As outlined in this white paper, WAAS Mobile addresses the full complement of requirements that are essential to a successful enterprise deployment, providing the scalability, security, and ease-of-deployment and management for large, diverse organizations with widely varying PC environments.

**Corporate Headquarters**

Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
www.cisco.com  
Tel: 408 526-4000  
800 553-NETS (6387)  
Fax: 408 526-4100

**European Headquarters**

Cisco Systems International  
BV  
Haarlerbergpark  
Haarlerbergweg 13-19  
1101 CH Amsterdam  
The Netherlands  
www-europe.cisco.com  
Tel: 31 0 20 357 1000  
Fax: 31 0 20 357 1100

**Americas Headquarters**

Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
www.cisco.com  
Tel: 408 526-7660  
Fax: 408 527-0883

**Asia Pacific Headquarters**

Cisco Systems, Inc.  
168 Robinson Road  
#28-01 Capital Tower  
Singapore 068912  
www.cisco.com  
Tel: +65 6317 7777  
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on **the Cisco Web site at [www.cisco.com/go/offices](http://www.cisco.com/go/offices)** .

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus  
Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland  
Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland  
Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden  
Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2008 Cisco Systems, Inc. All rights reserved. CCIP, CCSP, the Cisco *Powered* Network mark, Cisco Unity, Follow Me Browsing, FormShare, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, Fast Step, GigaStack, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MGX, MICA, the Networkers logo, Networking Academy, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, ScriptShare, SlideCast, SMARTnet, StrataView Plus, Stratm, SwitchProbe, TeleRouter, The Fastest Way to Increase Your Internet Quotient, TransPath, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0402R)

Printed in the USA