OPUS

Extending the Enterprise:

Building Secure Infrastructures Using the Wireless LAN Services Module On Catalyst 6500 Series Switches

prepared for Cisco Systems
October 2004

CONTENTS

ILLUSTRATIONS

Executive Summary

Wireless LANs mean mobility for end-users, but for network managers the technology raises numerous concerns about security, scalability, performance, and manageability.

Cisco Systems is addressing these concerns with its recently introduced Wireless LAN Services Module (WLSM) for Catalyst 6500 series switches. The WLSM brings Cisco's vast set of security and management tools into the wireless LAN arena, seamlessly integrating wireless clients into the existing wired infrastructure. A single WLSM can support 300 access points and 6,000 clients scattered across the enterprise.

Cisco commissioned Opus One, an independent networking consultancy, to assess WLSM performance in the areas of roaming, Non-Stop Forwarding/Stateful Switchover (NSF/SSO) failover, and delay/jitter.

Roaming is a critical measure of mobility and scalability. NSF/SSO failover increases uptime by protecting traffic against the loss of a routing session or a Supervisor Engine. Delay and jitter are critical metrics in assessing performance of any application, especially the voice-over-IP applications now under consideration for many wireless LANs.

Among the key findings of our tests:

- Roaming times are virtually the same regardless of whether the WLSM and access points are on the LAN, or across a WAN

- Roaming times range from 78 to 313 milliseconds, even in the highly unlikely case that more than 100 clients move simultaneously and associate with the same access point

- NSF/SSO reduces downtime caused by Supervisor card failure from many tens of seconds to less than 2 seconds

- Delay or jitter are virtually identical with and without the WLSM in place

This report is organized as follows. We begin by introducing the WLSM and discussing its security and manageability features. Then we move on to discuss test bed configuration, procedures, and results from tests of local and remote roaming, NSF/SSO failover, and delay and jitter.

Introducing the WLSM

The Wireless LAN Services Module (WLSM) is a component of the Cisco Structured Wireless-Aware Network (SWAN) architecture. SWAN extends wireless awareness into key elements of network infrastructure, providing the same level of security, reliability, ease of deployment, management, and scalability for wireless LANs that organizations have come to expect from wired LANs.

The WLSM occupies a single slot in Catalyst 6500 series switches. No separate infrastructure is needed for WLAN management, thus protecting investment in existing equipment.

Further, the WLSM integrates with existing Catalyst 6500 services modules, such as those providing stateful firewall and intrusion detection services (IDS). The WLSM's coexistence with other services modules means security services are extended into the wireless domain, again with no additional infrastructure needed.

In fact, the WLSM brings the entire Cisco security arsenal to wireless networks. In addition to firewall and IDS functions, available security services include access control lists (ACLs) at layers 2, 3, and 4; router ACLs; unicast reverse path forwarding, which blocks malformed and spoofed packets; rate limiters, which prevent denial-of-service (DOS) attacks; and IPSec- and SSL-based virtual private networks (VPNs). Cisco SWAN also adds wireless-specific security mechanisms such as rogue access point detection, reporting, and containment.

The WLSM also allows the creation of wireless "mobility groups," giving network managers the ability to define and individually authenticate different sets of users. For example, a WLSM-equipped Catalyst 6500 might be configured to distinguish between "guest" and "employee" logins, and grant different sets of resource permissions to users in each group.

The WLSM uses Cisco's Wireless LAN Context Control Protocol (WLCCP) to carry authentication messages between access points and the Catalyst 6500. A WLSM-equipped switch acts as an authenticator by learning the location of every associated wireless client node. The WLSM also interoperates with Cisco's existing AAA services.

The switch learns MAC-to-IP bindings of wireless clients either by snooping on DHCP exchanges or by snooping ARP or IP packets from wireless nodes. These two learning mechanisms enable the switch to provide uninterrupted layer-3 connectivity to roaming wireless nodes.

The WLSM maximizes uptime for wireless traffic through the use of Cisco's Non-Stop Forwarding/Stateful Switchover (NSF/SSO) mechanisms. In the event of a Supervisor card failure in the Catalyst hosting the WLSM, NSF/SSO preserves forwarding state during the loss and recovery of routing sessions, keeping disruption to a minimum.

The WLSM also simplifies deployment and management of wireless infrastructure. With a WLSM-equipped Catalyst 6500 series switch, network managers can configure new access points literally “out of the box.” Further, because the WLSM provides a single point of ingress for wireless traffic, it is possible to define and enforce traffic policies for all wireless clients through the use of mobility groups.

Figure 1 below shows the WLSM at work in a typical network deployment. Note that the various elements of Cisco SWAN – including access points, authentication servers, and the WLSM in the Catalyst 6500 – all work together to monitor and manage wireless traffic throughout the enterprise.

Figure 1: WLSM Network Deployment

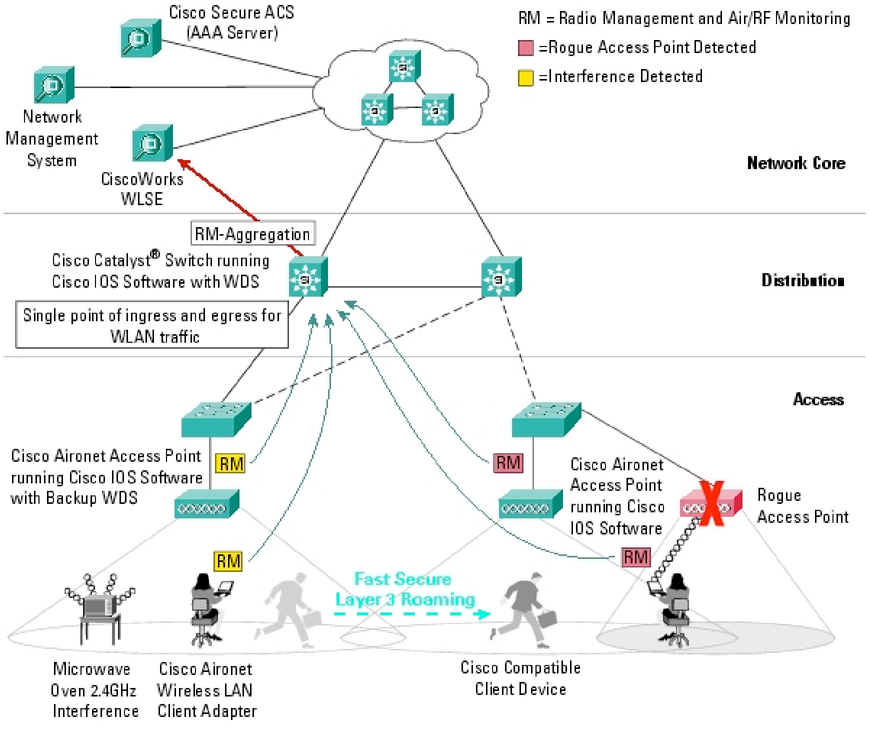
Cisco Secure ACS (AAA Server)
RM = Radio Management and Air/RF Monitoring
🟥 =Rogue Access Point Detected
🟨 =Interference Detected
Network Management System
CiscoWorks WLSE
Network Core
RM-Aggregation
Cisco Catalyst® Switch running Cisco IOS Software with WDS
Single point of ingress and egress for WLAN traffic
Distribution
Access
Cisco Aironet Access Point running Cisco IOS Software with Backup WDS
RM
RM
Cisco Aironet Access Point running Cisco IOS Software
RM
RM
Rogue Access Point
Fast Secure Layer 3 Roaming
Microwave Oven 2.4GHz Interference
Cisco Aironet Wireless LAN Client Adapter
Cisco Compatible Client Device

1. Clients and Access Points (AP) send their Radio Management (RM) data to the Cisco AP, switch or router running wireless-aware Cisco IOS Software with Wireless Domain Services (WDS).

2. Cisco AP, switch or router running wireless-aware Cisco IOS Software with WDS uses RM-Aggregation to remove redundant RM data received from the access points and client devices. The WDS device then forwards the aggregated data to the CiscoWorks WLSE.

The WLSM is highly scalable. One WLSM will support up to 300 Cisco Aironet Series access points and 6,000 users when fitted into any existing Catalyst 6500 Series switch equipped with a Supervisor Engine 720.

The access points and clients can be located anywhere in the enterprise; there is no requirement that they be locally attached to the WLSM. In fact, our tests results demonstrate that roaming times are virtually identical whether access points are across the room or thousands of miles away.

It is important to note that the WLSM is not in the Catalyst 6500's data path, and thus does not interfere with data-plane performance. The WLSM facilitates wireless networking by writing Cisco Express Forwarding (CEF) forwarding information base (FIB) entries into the Catalyst 6500's switching ASICs. Once the FIB entries have been written, all traffic is forwarded in hardware and bypasses the WLSM.

The WLSM communicates with access points though industry-standard Generic Routing Encapsulation (GRE) tunnels. This tunneling protocol enables fast roaming even when access points are remotely located.

More information about the WLSM is available at http://www.cisco.com/en/US/products/hw/modules/ps2706/prod_bulletin09186a00802252ba.html.

WLSM Local Roaming

Mobility may be one of the greatest benefits of WLANs, but it also poses some unique technical challenges. How long a delay will notebooks, PDAs, and 802.11 phones experience when moving between coverage areas? How much packet loss occurs during a handoff between access points? Can roaming be monitored and managed from a central location?

To determine the answers to these questions, Cisco asked Opus One to conduct roaming tests using several configurations. Our tests began with baselines involving 32 and 128 clients – scaling up the number of attachments far beyond the levels commonly found in production networks.

We then measured roaming times locally (with clients, access points, and WLSM all on one local network) and remotely (with wide-area network delays separating the WLSM from clients and access points).

To date, obtaining precisely timed measurements of wireless LAN performance has been a challenge because of limitations in existing test instruments. Measurement variance in the hundreds of milliseconds or more is not uncommon.

For this project Opus One and Cisco worked with VeriWave, a startup developing the first commercially available WLAN test instruments offering extremely precise time measurements.

The VeriWave TestPoint uses a custom MAC layer and field-programmable gate arrays (FPGAs) to ensure extremely precise timestamps and packet departure rates. Both mechanisms are critical in obtaining reliable measurements of roaming time (and also failover, delay, and jitter, covered later in this report). VeriWave claims its timestamp precision is measured in nanoseconds, in contrast to the hundreds-of-milliseconds precision of existing test instruments.

Figure 2 below illustrates the test bed configuration for the local roaming tests. Note that there are three Cisco Aironet 1100 access points present: One is attached to a Catalyst 6500 switch equipped with a WLSM, while the other two are attached to a Catalyst 3550 switch. Cisco's SWAN architecture uses industry-standard Generic Routing Encapsulation (GRE) tunnels between the WLSM and Aironet access points.

Figure 2: Local Roaming Test Bed

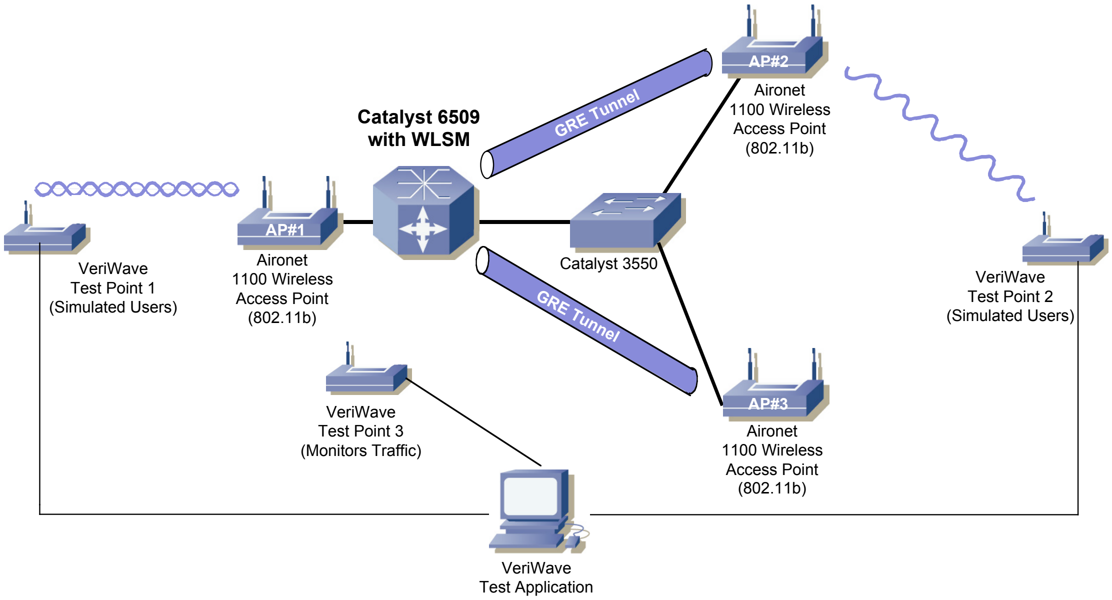

In the local roaming test, VeriWave TestPoint1 associates virtual clients with the access point attached to the Catalyst 6500 (labeled AP 1 in the figure), while TestPoint 2 associates virtual clients with AP 2, which is attached to the Catalyst 3550.

Although we used 802.11b clients for these tests, the WLSM and Cisco Aironet access points support 802.11a and 802.11g standards as well.

The WLSM has a “fast secure roaming” feature that facilitates speedy roaming. Typically, production networks will make use of an external authentication server such as a Cisco Secure Access Control Server (ACS).

With the WLSM, a wireless client would initially be authenticated via the ACS, but the WLSM would then store the client’s authentication credentials. All reauthentication requests – such as those arising from roaming events – would then be handled by the WLSM, with no external lookups to the ACS required. This greatly reduces roaming time, especially in cases when the ACS is located across a low-speed WAN link.

In our test bed, we pre-authenticated clients on the WLSM without using an ACS. However, since authentication during roaming is performed by the WLSM rather than an external ACS, our authentication times during roaming are equivalent to those on production networks using an ACS.

Once all the necessary 802.11 handshaking is complete, virtual clients associated with AP 1 send UDP/IP packets to clients behind AP 2 at an aggregate rate of 1,000 packets per second. While traffic is being exchanged, we disconnect the power to AP 2, forcing associated clients to "roam" to AP 3, which is attached to the same switch.

The VeriWave TestPoints take three measurements: roaming time, time to data received, and packet loss. *Roaming time* represents the interval from the loss of AP2 to the point where all clients reassociate with AP3. *Time to data received* reflects the interval from the last data packet received through AP 2 to the first data packet received through AP3. This second metric is a more meaningful predictor of application performance, since it reflects the time needed for the client to do actual work. *Packet loss* is simply a count of the number of dropped packets during the roaming event.

To determine WLSM scalability, we repeated the roaming tests with 32 and 128 simulated clients. Both tests represent roaming events that are well above the number of client associations typically found in production settings. For example, Cisco's own scalability estimates for the WLSM assume a maximum of 20 clients associated with each access point. In production networks, a maximum of 12 or fewer associations per access point is common.

Figure 3 below presents results from the local roaming tests. We ran each roaming test five times and averaged the results. Note that the delta between roaming and time to data received is nearly identical, regardless of the number of clients involved.

Figure 3: WLSM Local Roaming Tests

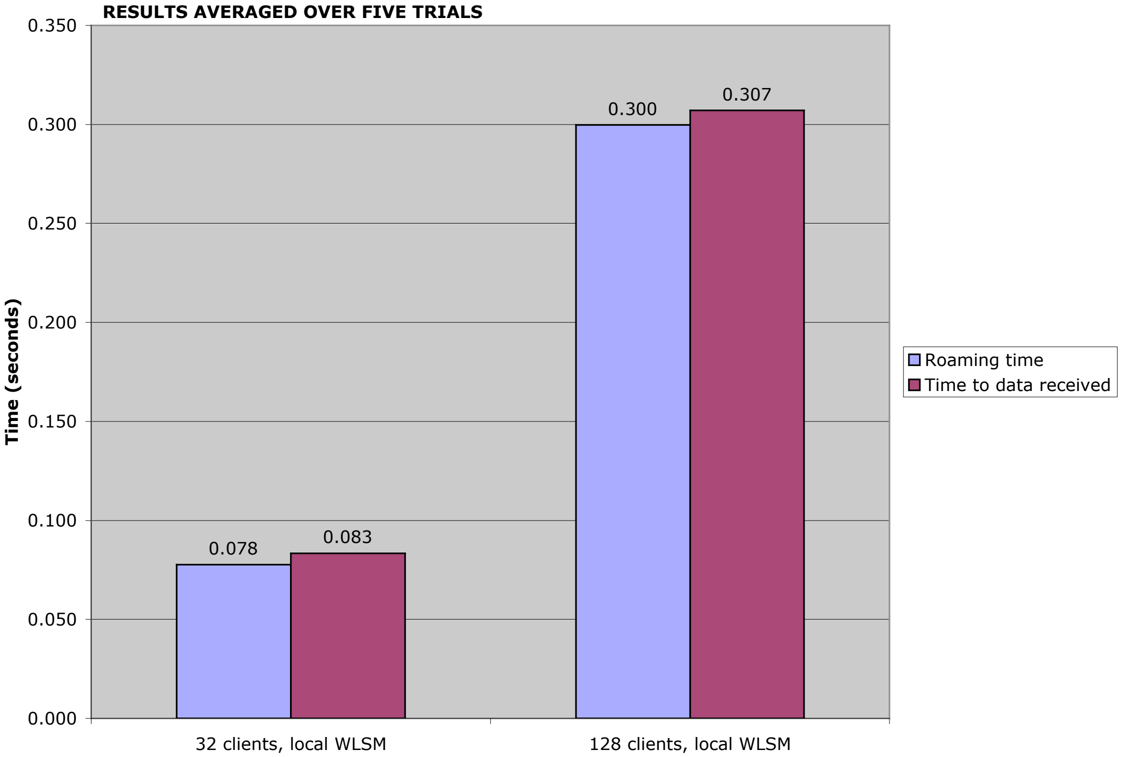

In the packet-loss measurements, the system dropped an average of 214 packets when 32 clients roamed between access points, and an average of 428 packets when 128 clients roamed. Again, the aggregate offered load for both tests was 1,000 pps.

WLSM Remote Roaming

The WLSM manages access points throughout the enterprise, including those access points at remote locations. Because the WLSM is not in the data path between clients and access points, roaming times should be similar regardless of whether the WLSM is remote or local to the access points.

To put that notion to the test, we introduced a large amount of delay – 100 milliseconds, roughly the amount of roundtrip delay between New York and Los Angeles – between the WLSM and some of the access points it managed.

Figure 4 below illustrates the logical test bed for the remote roaming event. In New York, a Catalyst 6500 is equipped with a WLSM and one Aironet 1100 access point. Across a WAN circuit, the WLSM manages two Aironet 1100s attached to a Catalyst 3550 switch in Los Angeles. Over a distance like this, network elements and the speed of light introduce delay of approximately 50 ms in each direction, or 100 ms roundtrip.

Figure 4: WLSM Remote Roaming Logical Test Bed

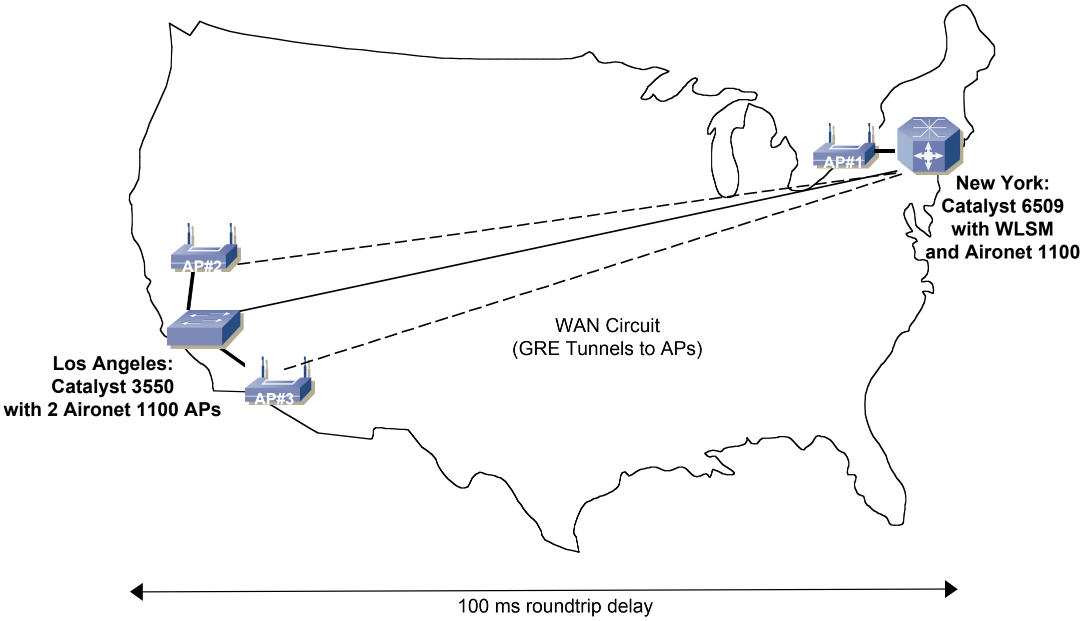

Figure 5 below illustrates the physical test bed for the remote roaming event. The only difference from the local roaming test bed is the addition of a device that introduces delay: The PacketStorm 2600E network emulator from PacketStorm Communications. The PacketStorm device buffers every packet for a user-defined interval.

Figure 5: WLSM Remote Roaming Physical Test Bed

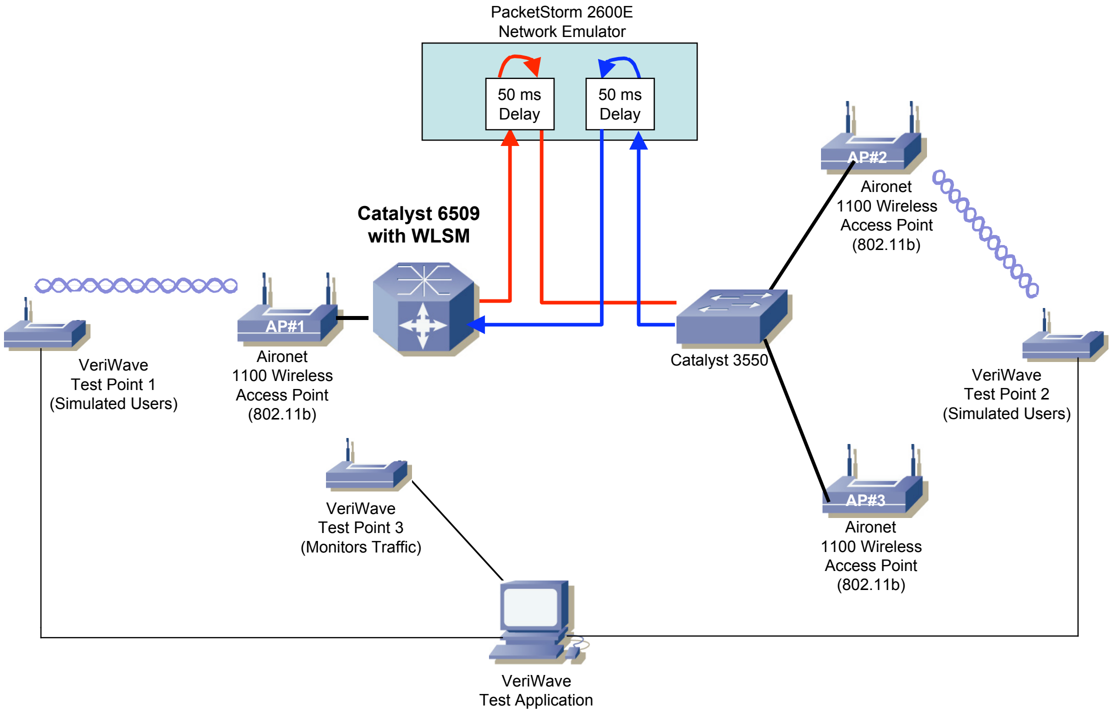

We offered the same traffic in both the local and remote roaming tests – a unidirectional stream of packets at an aggregate rate of 1,000 pps. In the remote case, clients communicated through AP 1 in the "New York" location and AP 2 in the "Los Angeles" location. During the test, we disconnected power to AP 2, forcing the Los Angeles clients to roam to AP 3. As before, we ran each test five times and averaged the results.

Figure 6 below presents remote roaming test results, along with local roaming results for comparison. Note the similarity between the local and remote roaming results. For both the 32- and 128-client test cases, roaming takes virtually the same time regardless of whether the WLSM is local or remote to the access points.

The only significant difference is between the "time to first data" results. However, careful readers will note the difference between the local and remote numbers is about 50 ms – in other words, exactly the amount of delay to simulate distance. We also present the same results with the delay subtracted out, and here the numbers are virtually the same as those in the local test case.

Figure 6: WLSM Local and Remote Roaming Tests

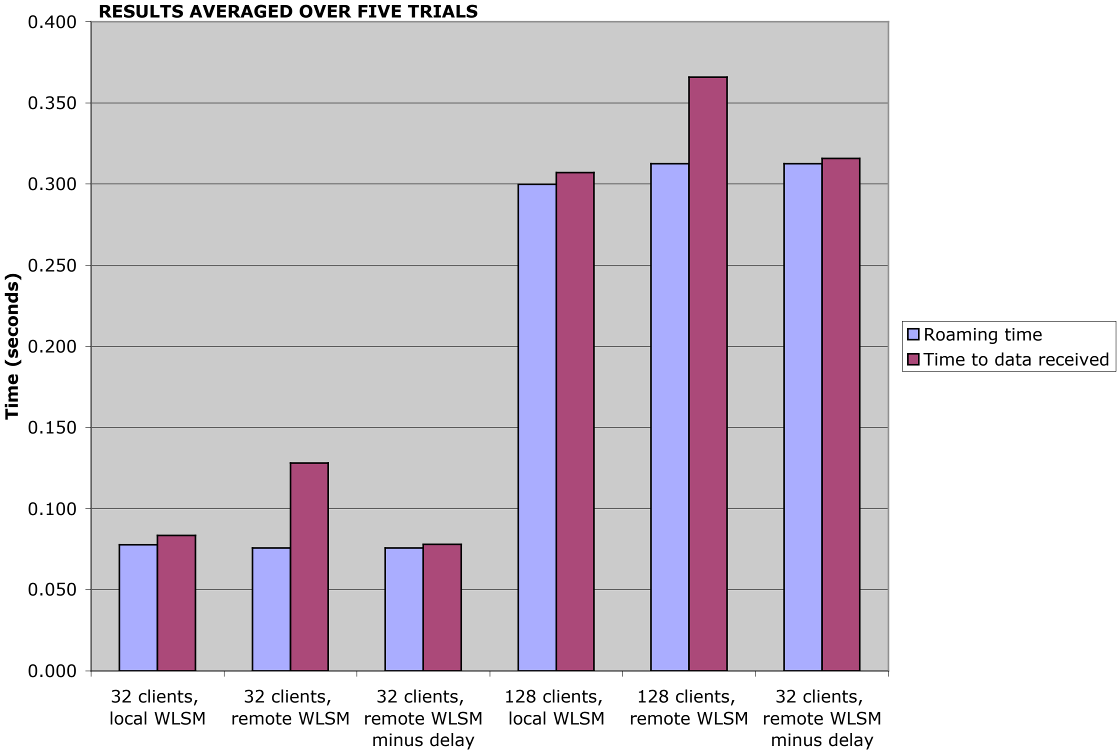

Packet loss was slightly higher in the remote roaming tests than in the local tests, but this again reflects the added delay. On average, loss was 257 packets in the 32-client case, compared with 214 packets dropped in the local roaming test. With 128 clients, average loss rose to 479 packets, compared with 429 packets in the local tests. These increases are strictly a result of added delay along the path and should not be attributed to the WLSM.

The similarity between local and remote roaming demonstrates that the WLSM does not add overhead to normal WLAN operation. Clients like notebooks, PDAs, and 802.11 phones can be managed through a WLSM anywhere in the enterprise with no added performance penalty.

WLSM With NSF/SSO Failover

Ensuring maximum uptime is a concern with any network, and wireless networks are no exception. Cisco extends high availability to the WLSM through its Non-Stop Forwarding and Stateful Switchover (NSF/SSO) mechanisms in the Catalyst 6500.

NSF is a Catalyst 6500 resilience feature that makes use of the graceful restart mechanisms being developed by the IETF. In a nutshell, NSF replicates layer-3 forwarding tables on redundant Supervisor Engine cards, thus preserving forwarding state during the loss and restart of a routing session. If one Supervisor card fails, a redundant card takes over with no effect on routing sessions. This dramatically reduces the time needed to recover from a failure.

Similarly, SSO protects switching state by replicating layer-2 forwarding tables across redundant Supervisor Engine cards.

Cisco asked Opus One to determine failover times for wireless traffic forwarded through a Catalyst 6500 switch equipped with a WLSM and redundant Supervisor Engine 720 cards.

Figure 7 below shows the test bed for the NSF/SSO failover event. As in the roaming tests, virtual clients attached to a pair of Cisco Aironet 1100 access points which, in turn, attached to a Catalyst 6500 switch. The switch housed two Supervisor 720 cards, one in Active mode and the other in Standby mode.

Cisco's NSF works with BGP, OSPF, and IS-IS. We used OSPF in these enterprise-focused tests.

Figure 7: WLSM with NSF/SSO Failover Test Bed

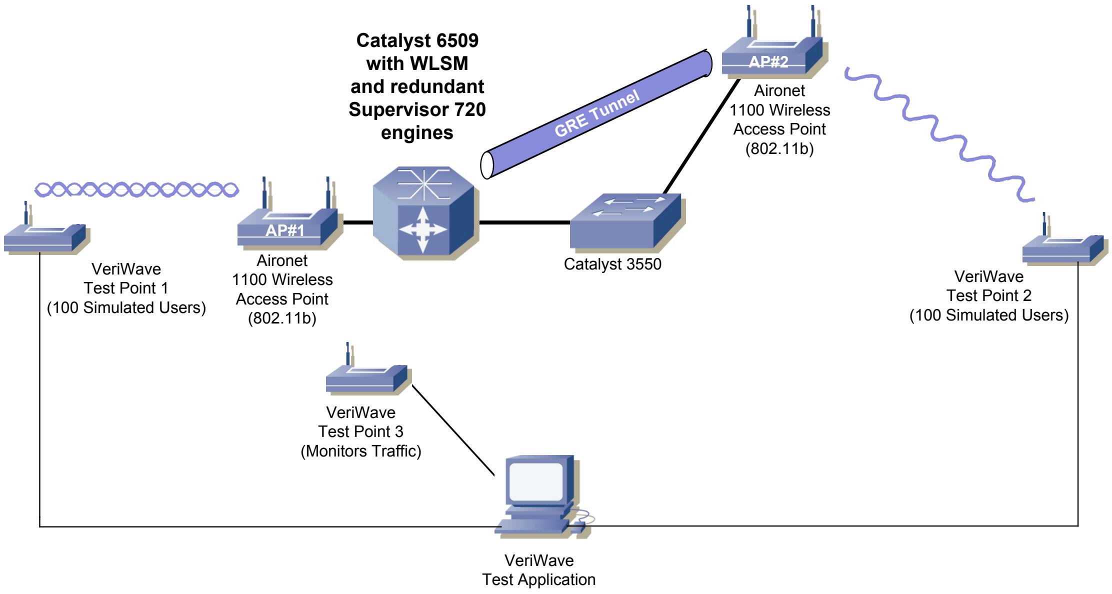

To determine failover time, we configured the VeriWave TestPoints to emulate 200 clients transmitting and receiving data at an aggregate rate of 1,000 packets per second.

At that rate, each dropped packet was equivalent to 1 millisecond of failover time. (We began with baseline tests both with and without the WLSM active to verify that the system dropped zero data packets without a Supervisor failure.)

We then removed the active Supervisor card while offering traffic, thus forcing a failover to the standby Supervisor card, and measured packet loss. We repeated the failover test three times.

We expected some packet loss in these tests. The Supervisor 720 not only establishes and maintains routing protocol and spanning-tree state, but also integrates a 720-Gbit/s switch fabric. Since the loss of an active Supervisor card also means the temporary loss of the switch fabric, some packet loss is inevitable.

Figure 8 below presents results from the NSF/SSO failover tests for the WLSM. The average failover time was 1.776 seconds across three trials. As noted, this loss is a result of the temporary removal of the data path (the switch fabric) in the active Supervisor 720 card.

Without NSF/SSO, the normal routing protocol convergence mechanisms would apply. In the case of OSPF with default timer values, the "router dead" interval does not occur until 40 seconds (the period for the loss of three "hello" messages), followed by additional time to restart the routing session and reconverge the network. This last step may itself last many minutes, depending on the size of the network. With NSF/SSO, failover time is less than 2 seconds.

Figure 8: WLSM Failover With NSF/SSO

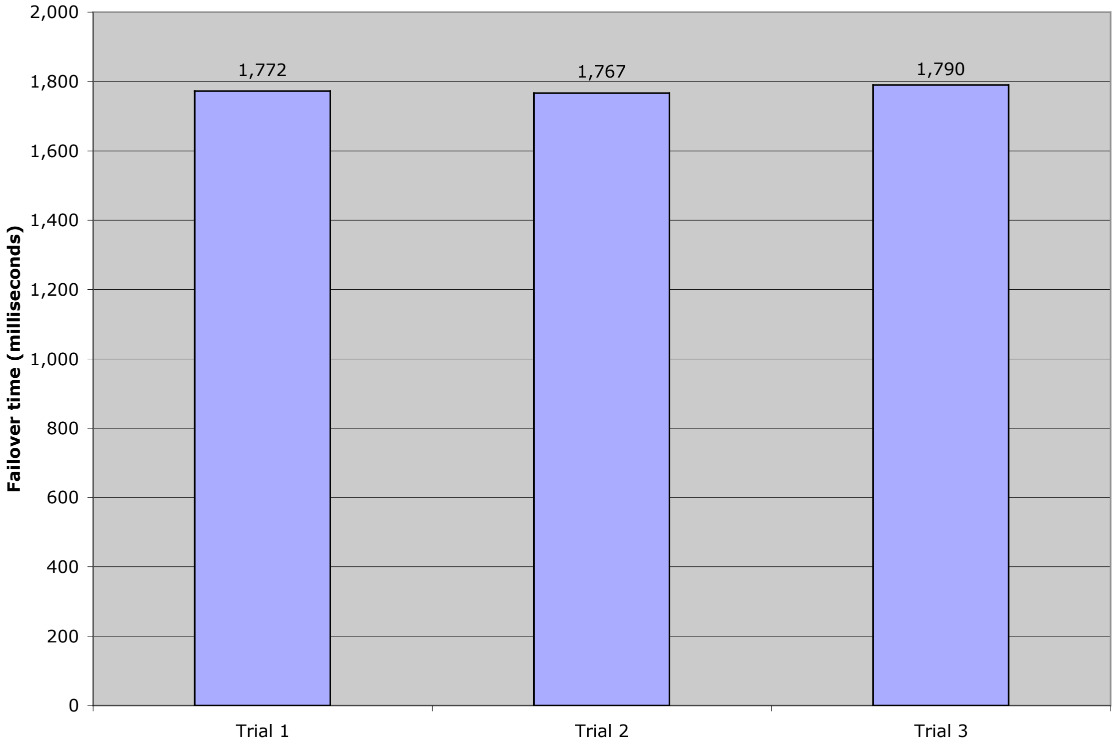

Although we did not use this configuration in our tests, it is possible to achieve zero loss upon failure of a Supervisor card by using line cards equipped with Distributed Forwarding Card (DFC) modules. DFC modules include a local switching engine and switch fabric, obviating the need to go through a central fabric for traffic passing between ports on the same card. Unfortunately, DFC-equipped line cards were not available for inclusion in our tests. Even in this worst-case scenario without DFCs on the line cards, disruption was minimal.

The failover tests demonstrate that loss of a layer-3 routing session introduces only a brief loss in forwarding capabilities.

WLSM Delay and Jitter

Since WLANs may introduce more inherent delay and jitter than wireline networks, these two metrics rank among the most important in assessing WLAN performance. While voice and video applications are the most sensitive to delay and jitter, virtually any application can suffer if a WLAN system introduces excessive or variable amounts of delay. Cisco asked Opus One to validate that the WLSM adds no appreciable increase in delay or jitter.

As noted previously, we used VeriWave TestPoints to obtain delay and jitter measurements. Thanks to VeriWave's highly precise timestamp mechanisms, the numbers presented here are among the most accurate timing measurements of WLAN performance ever published. Where delay or jitter were measured at all in previous tests, the results were usually accurate only to within hundreds of milliseconds. The numbers presented here have timestamp resolution measured in nanoseconds.

To determine the effect (if any) of the WLSM on delay and jitter, we conducted tests both with and without the WLSM in place. In both cases, the VeriWave TestPoints offered a unidirectional stream of UDP/IP packets. For every test, we measured minimum, average, and maximum delay and jitter.

The TestPoint measurements also distinguish between data traffic (what a user would send) and management traffic (what the access points and WLSM send, including 802.11 control traffic such as beacon frames). For both data and management traffic, we ran tests with 64-, 128-, 256-, 512-, 1,024-, and 1,500-byte 802.11 frames[1]. We ran each test five times, and averaged the results of all five trials.

Figure 10 below compares average and maximum delay with and without the WLSM in place.

[1] In this report, 802.11 frame lengths refer to the 802.11 payload, exclusive of layer-2 header and FCS. Note that the 802.11 protocols allow for variable-length headers, depending on the transport method and options in use.

Figure 10: Data Traffic Delay

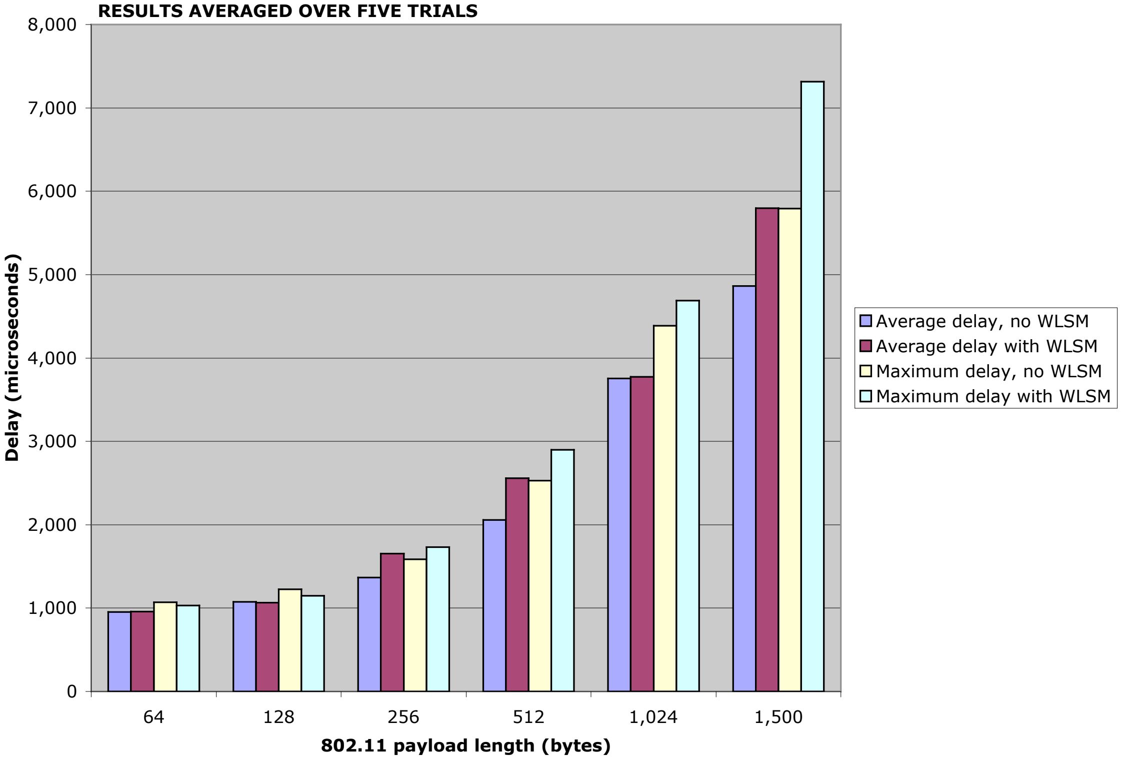

Note that delay is virtually indistinguishable between the two test cases. The very largest difference is about 1.5 milliseconds between the maximum delay measurements with and without the WLSM; in most other cases, the differences are in the tens of microseconds. Given that the threshold where application performance starts to degrade is in the dozens of *milliseconds*, a worst-case difference of 1.5 ms is not meaningful.

For data traffic jitter, once again the results with and without the WLSM in place are virtually indistinguishable. As these measurements show, the WLSM introduces no significant added delay or jitter for data traffic.

Figure 11 below shows jitter measurements for data traffic with and without the WLSM in place.

Figure 11: Data Traffic Jitter

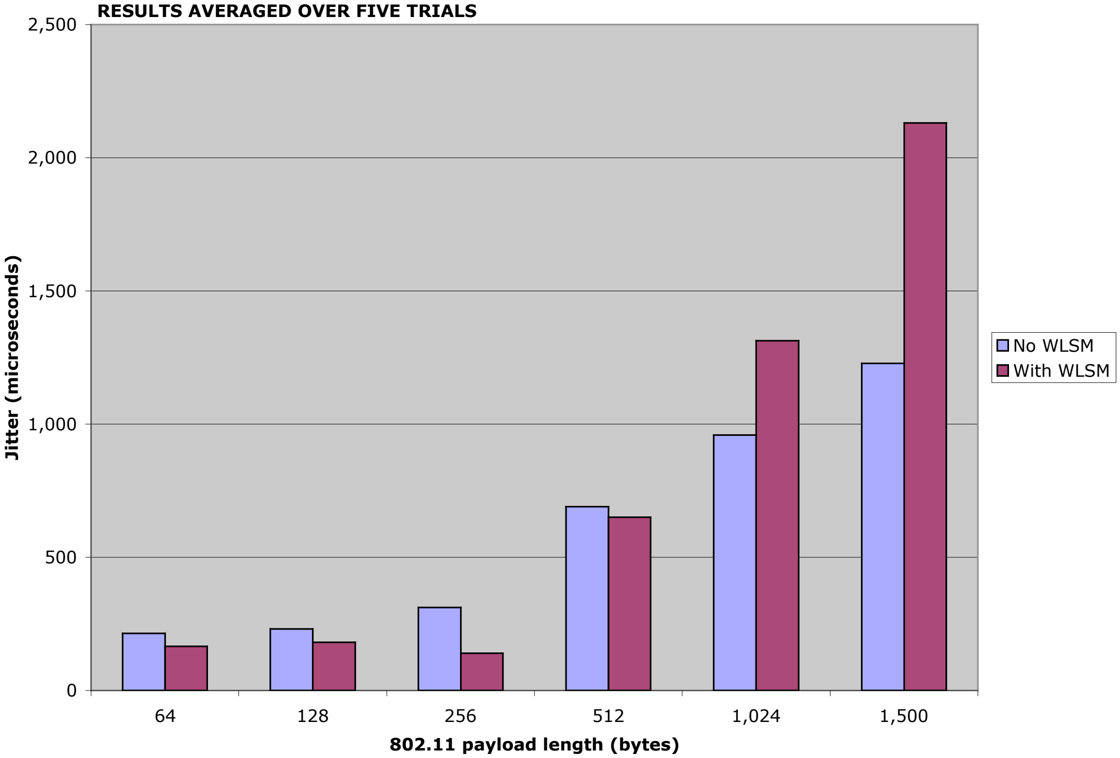

As noted, the VeriWave TestPoints can distinguish between data and management traffic. The TestPoints also measured delay and jitter for management traffic concurrently with the data traffic tests. Once again, the difference between the test cases with and without the WLSM generally was not meaningful.

Figure 12 below presents results from the management traffic delay tests. With the exception of two outliers – maximum delay for 256- and 512-byte management frames – there were virtually no differences between test cases with and without the WLSM. Even these differences may not be meaningful. The vast majority of 802.11 control traffic consists of short frames such as minimum-length beacons. Note that there was virtually no difference between delay for short management frames with and without the WLSM in place.

Figure 12: Management Traffic Delay

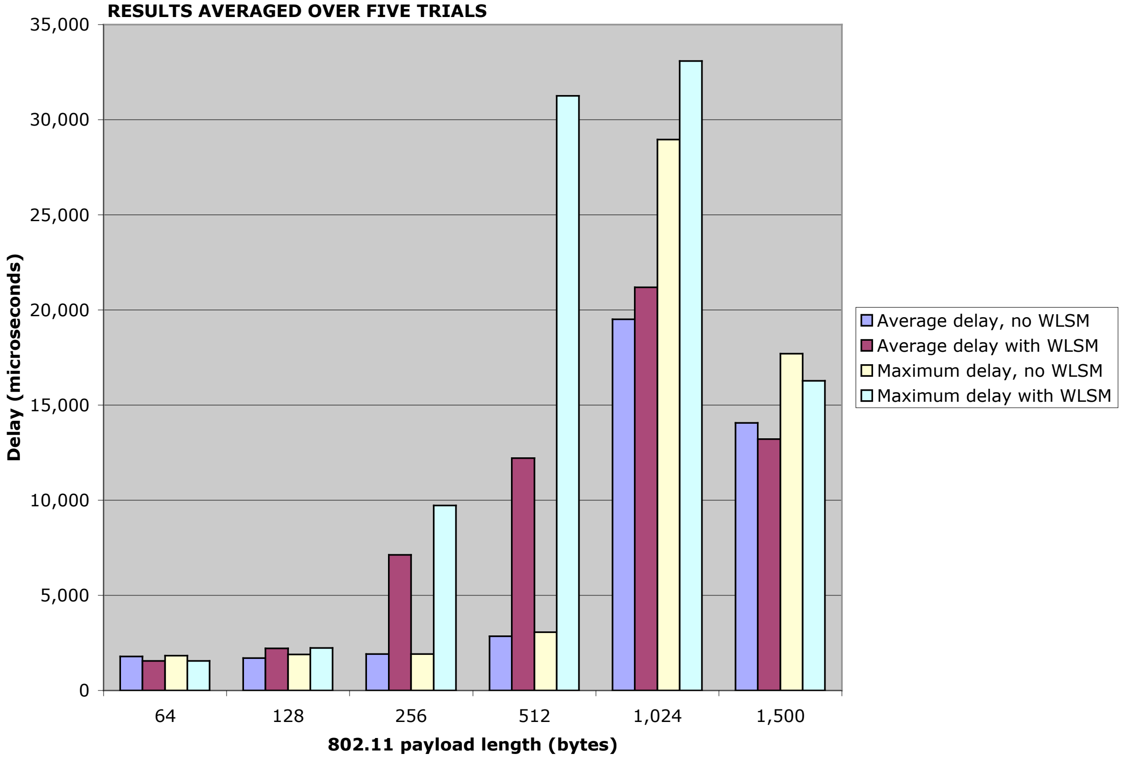

Figure 13 below displays results of jitter tests in management traffic. As with the management delay tests, there is a significant delta between the WLSM and no-WLSM cases with 512-byte frames. However, the same caveat applies here: Most management traffic uses frames far shorter than 512 bytes.

Figure 13: Management Traffic Delay

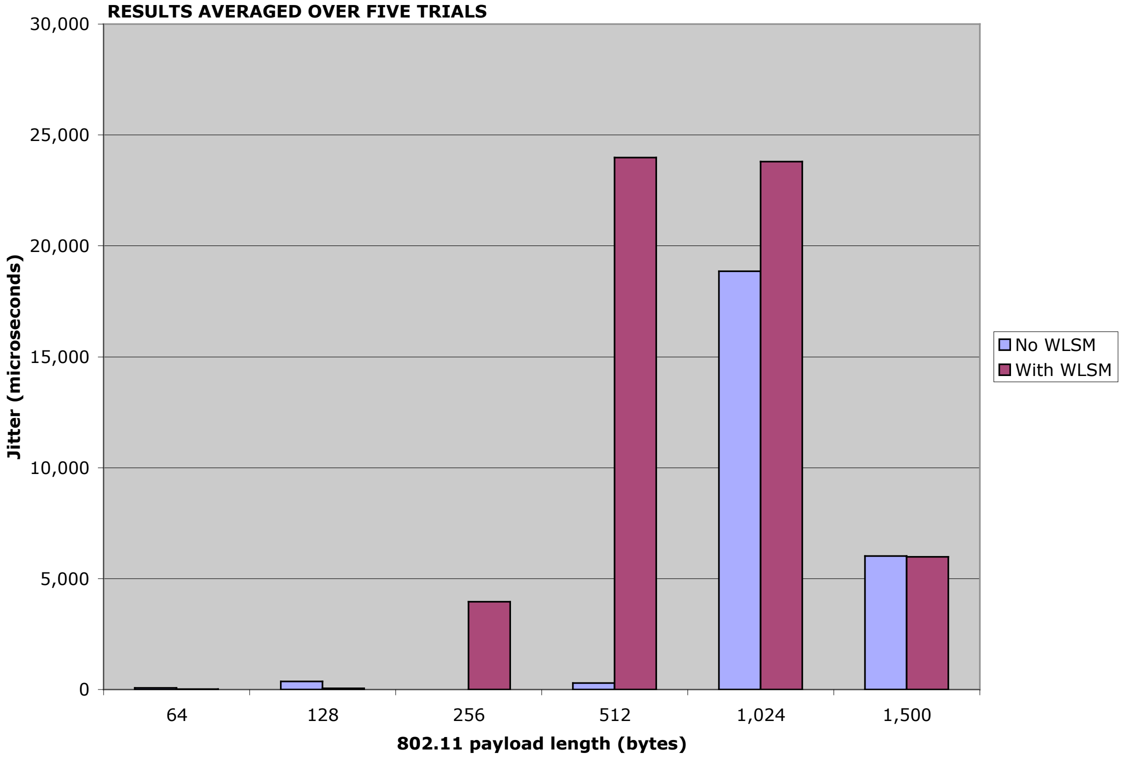

Table 1 below summarizes results from all delay and jitter tests in tabular form. Again, note that all results are averaged across five trials.

Table 1: Delay and Jitter With and Without WLSM

Test case	Data average delay (usec)	Data maximum delay (usec)	Data jitter	Mgmt. average delay (usec)	Mgmt. maximum delay (usec)	Mgmt. jitter (usec)
No WLSM, 64-byte frames	955	1,071	214	1,790	1,825	70
With WLSM, 64-byte frames	958	1,032	165	1,553	1,559	12
No WLSM, 128-byte frames	1,077	1,227	231	1,696	1,901	357
With WLSM, 128-byte frames	1,064	1,147	180	2,207	2,236	49
No WLSM, 256-byte frames	1,368	1,586	312	1,907	1,907	0
With WLSM, 256-byte frames	1,656	1,733	140	7,128	9,732	3,955
No WLSM, 512-byte frames	2,058	2,529	689	2,850	3,073	300
With WLSM, 512-byte frames	2,557	2,899	651	12,219	31,256	23,985
No WLSM, 1,024-byte frames	3,757	4,388	959	19,515	28,947	18,864
With WLSM, 1,024-byte frames	3,772	4,686	1,313	21,195	33,093	23,796
No WLSM, 1,500-byte frames	4,865	5,792	1,228	14,060	17,709	6,021
With WLSM, 1,500-byte frames	5,798	7,315	2,131	13,214	16,275	5,975

Conclusion

These tests demonstrate that Cisco's WLSM extends connectivity to wireless clients with no compromise in scalability, performance, or manageability. With a WLSM-equipped Catalyst 6500 series switch, it is possible to deploy, secure, and manage an entire enterprise's wireless clients – all with no performance overhead.

As our test results show, roaming times are virtually identical throughout the enterprise, regardless of whether the WLSM is in the same location as the access points. Cisco's NSF/SSO mechanisms help ensure uptime during the loss of a routing session or even an entire Supervisor card. And the WLSM adds no delay or jitter to wireless traffic, a critical concern for voice-over-IP traffic.

Taken together, these results demonstrate that it is possible to extend all the advantages of a wired infrastructure to wireless clients. With the WLSM, enterprise network managers can build large-scale wireless infrastructure while maintaining wireline levels of performance, manageability, and security.

Acknowledgements

Opus One gratefully acknowledges the support of VeriWave, a supplier of highly reliable wireless test equipment. In addition to supplying its TestPoint test system, VeriWave also developed custom scripts for this project. Special thanks to VeriWave CTO Thomas Alexander and director of applications Carl Brown for their assistance in configuration, test execution, and results analysis.

TESTED BY veriwave™

About Opus One®

Opus One® is a consulting and information technology firm based in Tucson, AZ. Founded in 1989, Opus One's corporate goal is to help our clients make the best use of information technology. We focus on efficient and effective solutions in the areas of data networking, electronic mail, and security. For more information, see http://opus1.com or contact us at:

Opus One
1404 East Lind Road
Tucson, AZ 85719
+1-520-324-0494