

# **SANDIA REPORT**

SAND2006-6616

Unlimited Release

Printed November 2006

## **Design and Initial Deployment of the Wireless Local Area Networking Infrastructure at Sandia National Laboratories**

Marc M. Miller, Dallas J. Wiener, Edward L. Witzke, John P. Long,  
Michael J. Hamill, Mark G. Mitchell

Prepared by  
Sandia National Laboratories  
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation,  
a Lockheed Martin Company, for the United States Department of Energy's  
National Nuclear Security Administration under Contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

**NOTICE:** This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from  
U.S. Department of Energy  
Office of Scientific and Technical Information  
P.O. Box 62  
Oak Ridge, TN 37831

Telephone: (865) 576-8401  
Facsimile: (865) 576-5728  
E-Mail: [reports@adonis.osti.gov](mailto:reports@adonis.osti.gov)  
Online ordering: <http://www.osti.gov/bridge>

Available to the public from  
U.S. Department of Commerce  
National Technical Information Service  
5285 Port Royal Rd.  
Springfield, VA 22161

Telephone: (800) 553-6847  
Facsimile: (703) 605-6900  
E-Mail: [orders@ntis.fedworld.gov](mailto:orders@ntis.fedworld.gov)  
Online order: <http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online>



SAND2006-6616  
Unlimited Release  
Printed November 2006

## **Design and Initial Deployment of the Wireless Local Area Networking Infrastructure at Sandia National Laboratories**

Marc M. Miller, Dallas J. Wiener, Edward L. Witzke  
Advanced Networking Integration Department

John P. Long  
Cyber Security Technologies Department

Michael J. Hamill  
Network Systems Design and Implementation Department

Mark G. Mitchell  
Communication and Network Systems Department

Sandia National Laboratories  
P.O. Box 5800  
Albuquerque, NM 87185-0806

### **Abstract**

A major portion of the Wireless Networking Project at Sandia National Laboratories over the last few years has been to examine IEEE 802.11 wireless networking for possible use at Sandia and if practical, introduce this technology. This project team deployed 802.11a, b, and g Wireless Local Area Networking at Sandia. This report examines the basics of wireless networking and captures key results from project tests and experiments. It also records project members' thoughts and designs on wireless LAN architecture and security issues. It documents some of the actions and milestones of this project, including pilot and production deployment of wireless networking equipment, and captures the team's rationale behind some of the decisions made. Finally, the report examines lessons learned, future directions, and conclusions.

## **Acknowledgements**

The authors wish to thank the following people for their contributions to various aspects of this project:

Dennis Bateman

Mike Cahoon

John Eldridge

Rich Gay

Steve Gossage

Barry Hess

Paul Linke

Pat Manke

Scot Marburger

Brian Maxwell

Del Packwood

Alex Quintana

Carlos Quintana

Perry Robertson

Mike Sjulín

Jason Tang

Tom Tarman

Tim Toole

The authors also wish to thank the many wireless networking vendors who provided information and equipment for evaluation.

## Contents

1	Introduction .....	7
2	Wireless Local Area Network Concepts and Topology.....	8
2.1	Smart Access Point Topology.....	8
2.2	Gateway Topology.....	9
2.3	Wireless Switch Topology.....	10
2.4	Centralized Controller Topology .....	11
2.5	Hybrid Topology.....	12
3	Wireless Client Evaluation.....	14
4	Wireless Access Point Evaluation.....	17
4.1	Range Testing .....	17
4.2	Load Testing .....	17
4.3	Product Evaluation.....	25
5	Wireless Local Area Network Architecture .....	28
5.1	Background.....	28
5.2	Wireless Network Usage Policy .....	29
5.3	Wireless Design .....	30
5.4	Throughput Considerations.....	31
6	Wireless LAN Security Approach.....	34
6.1	SNL Wireless Security Design .....	34
6.2	User Access.....	35
6.3	Cyber-Security Threats, Risks, and Mitigations .....	35
6.4	Wireless LAN Security Summary .....	38
7	Rogue Detection and Containment .....	39
8	Pilot Wireless LANs.....	40
8.1	New Mexico Pilots.....	40
8.2	California Pilot.....	43
9	Wireless LAN Equipment Selection and Acquisition.....	45
10	WLAN Deployment.....	48
10.1	NWIS Entries for Aruba Access Points .....	48
10.2	NWIS Entries for Trapeze Access Points .....	49
10.3	Site Survey Methodology .....	50
10.4	Wireless Network Design .....	53
10.5	Installation Issues.....	54
11	Conclusions.....	55
11.1	Lessons Learned.....	55
11.2	Future Directions .....	55
11.3	Summary .....	63
Appendix A	Diffused Infrared Equipment .....	64
Appendix B	Video Compression over Wireless Links .....	66
Appendix C	Bibliography .....	80

## List of Figures

Figure 1. Smart access point topology.....	9
Figure 2. Gateway topology.....	10
Figure 3. Wireless switch topology. ....	11
Figure 4. Centralized controller topology.....	12
Figure 5. Hybrid topology.....	13
Figure 6. Average Client Throughput on 802.11a Access Points.....	22
Figure 7. Aggregate Client Throughput on 802.11a Access Points.....	22
Figure 8. Average Client Throughput on 802.11b Access Points.....	23
Figure 9. Client Throughput on 802.11b Access Points. ....	23
Figure 10. Average Client Throughput on 802.11g Access Points.....	24
Figure 11. Aggregate Client Throughput on 802.11a Access Points.....	24
Figure 12. WLAN Architecture. ....	30
Figure 13. SNL Wireless LAN Design.....	34
Figure 14. Layout of Access Points in Bldg. 831. ....	41
Figure 15. Security Architecture. ....	42
Figure 16. Access Point Information. ....	48
Figure 17. Access Point Interface and IP Information.....	49
Figure 18. Aruba Site Survey Tool. ....	53
Figure 19. RTP Payload Description. ....	68
Figure 20. Direct Connection Network Configuration. ....	69
Figure 21. Consumer-Grade Router/Switch Network Configuration. ....	69
Figure 22. Wireless Bridge Network Configuration.....	69
Figure 23. File Structure for Global Link Configuration Files.....	70

## List of Tables

Table 1. Throughput of 802.11b Clients from a Wired Network Server.....	14
Table 2. Throughput of 802.11a Clients from a Wired Network Server. ....	15
Table 3. Throughput of 802.11g Clients From a Wired Network Server. ....	15
Table 4. Signal Strengths at Various Distances from the AP. ....	17
Table 5. Trapeze 802.11a Throughput Rates. ....	18
Table 6. Trapeze 802.11b Throughput Rates.....	18
Table 7. Trapeze 802.11g Throughput Rates.....	19
Table 8. Cisco 802.11a Throughput Rates.....	19
Table 9. Cisco 802.11b Throughput Rates.....	19
Table 10. Proxim 802.11a Throughput Rates. ....	19
Table 11. Belkin 802.11g Throughput Rates. ....	20
Table 12. Data from OPNET Model Simulating 2 Mbps LAN.....	65
Table 13. Frame Formats. ....	67
Table 14. Video Quality Comparison. ....	77

# 1 Introduction

A major portion of the Wireless Networking Project at Sandia over the last few years has been to examine IEEE 802.11 wireless networking for possible use at Sandia and if practical, introduce this technology. This has included a study of the technology and existing implementations, development of a wireless network architecture as an extension to the current wired network, evaluation of wireless networking vendors and equipment, implementing pilot networks, procurement of equipment, initial equipment installations, development of procedures for installation and use of wireless networks, and full network rollout. Various other issues, such as efficient radio frequency (RF) channel reuse plans and compressed video services over wireless links were also examined under this project.

This report examines the basics of wireless networking and captures key results from project tests and experiments. It also records project members' thoughts and designs on wireless LAN architecture and security issues. It documents some of the actions and milestones of this project, including pilot and production deployment of wireless networking equipment, and captures the team's rationale behind some of the decisions made. Finally, the report examines lessons learned, future directions, and conclusions.

The project team selected 802.11a, with a 54 Mbps signaling rate, operating in the 5.15-5.35 GHz Unlicensed National Information Infrastructure (U-NII) band of the RF spectrum because of its signaling rate and its avoidance of the RF-polluted Industrial, Scientific, and Medical (ISM) bands. (Microwave ovens, Bluetooth devices, some cordless telephones, etc. operate in the 2.4 GHz ISM band, and some cordless telephones operate in the 5.8 GHz ISM band.) The team also decided to support 802.11b (11 Mbps, 2.4-2.48 GHz) because of its use in legacy devices and in handheld devices such as PalmPilots. (Currently 802.11a implementations consume too much power to be practical for most handhelds.) Since most equipment that supported both 802.11 a and b (at the time of our procurement) also supported 802.11g (54 Mbps, 2.4-2.48 GHz), and 802.11g is starting to become available built in to laptop computers, Sandia's wireless network also supports 802.11g.

## 2 Wireless Local Area Network Concepts and Topology

At the point in this project where wireless LAN (WLAN) topologies were being examined, four topologies were being used in enterprise WLAN deployment. They were:

- Smart access point,
- Gateway,
- Wireless switch,
- Centralized controller.

A new hybrid topology that combined the wireless switch and the centralized controller topologies was just coming onto the scene. A brief recap of these topologies will be presented here, although a full description of them can be found in a separate report [25].

### 2.1 Smart Access Point Topology

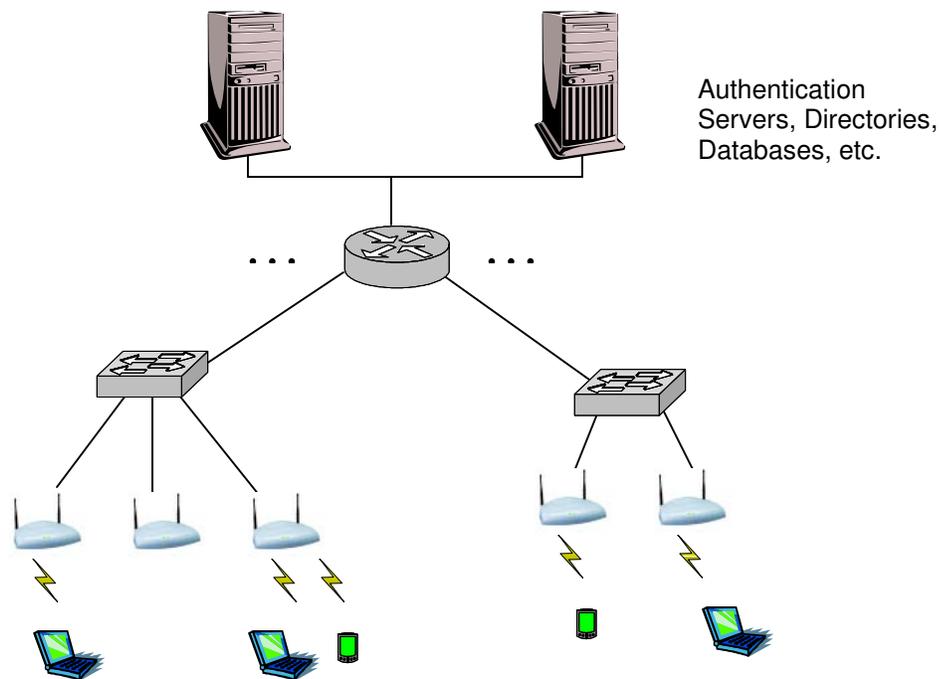
‘Smart’ (also known as “heavyweight” or “fat”) wireless access points (APs) perform not only the wireless 802.11 functions, but also advanced features such as client authentication processing, traffic filtering, AP management and configuration, client disassociation, and event logging. A wireless LAN based on a number of smart access points is said to have a smart access point topology. An example of a smart access point topology is shown in Figure 1.

Smart access points are self sufficient access points that can be placed anywhere in the network and work individually or as part of a network of APs on the same Layer 3 subnet. Smart APs are very flexible because of the many features they offer and the interoperability among different smart access point vendors. For example, a Cisco smart AP could be deployed on the same IP subnet as a Proxim smart AP while still facilitating Layer 2 roaming<sup>1</sup> between the two access points. This interoperability is facilitated by the IEEE 802.11f standard, known as the Inter-Access Point Protocol (IAPP).

In addition to these advantages of the smart access point topology, there are also some disadvantages, such as lack of scalability, the complexity of managing the WLAN, and the relatively high cost of deployment compared to other topologies. Smart access points are designed to be able to function by themselves. They are individually configured and managed. This makes them a perfect choice for small deployments. The smart access point topology is not, however, well suited for large deployments because at some point it becomes too difficult to configure, upgrade, and manage a large number of APs individually.

---

<sup>1</sup> Layer 2 roaming is when a wireless client roams from one access point to another on the same IP subnet.



**Figure 1. Smart access point topology.**

Another disadvantage of the smart access point topology is that it does not permit Layer 3 roaming, which allows users to roam between subnets without dropping sessions. If a user roams from one subnet to another, all sessions are dropped and the user must reassociate, reauthenticate, and obtain a new IP address.

The smart access point topology may also present a security risk because the access points contain configuration data and encryption keys. Network equipment such as switches, routers, and servers also contain configuration data and encryption keys, but physical access to this equipment is usually restricted by placing these devices in a locked room. Access points are usually placed in open areas where it is very difficult to restrict physical access.

## **2.2 Gateway Topology**

Some companies such as Vernier and Fortress Technologies have developed products commonly known as gateway appliances. These devices allow for better security, Layer 3 subnet roaming, and user/group-based security and QoS profiles. An example of a gateway topology WLAN is shown in Figure 2. Gateways are vendor-agnostic, meaning that standard heavyweight access points and networking switches from any vendor will work with them. Gateways allow network administrators to apply QoS and security profiles to individual users or groups of users. Since gateway topology WLANs use heavyweight APs, they suffer from the same management, scalability, and cost issues.

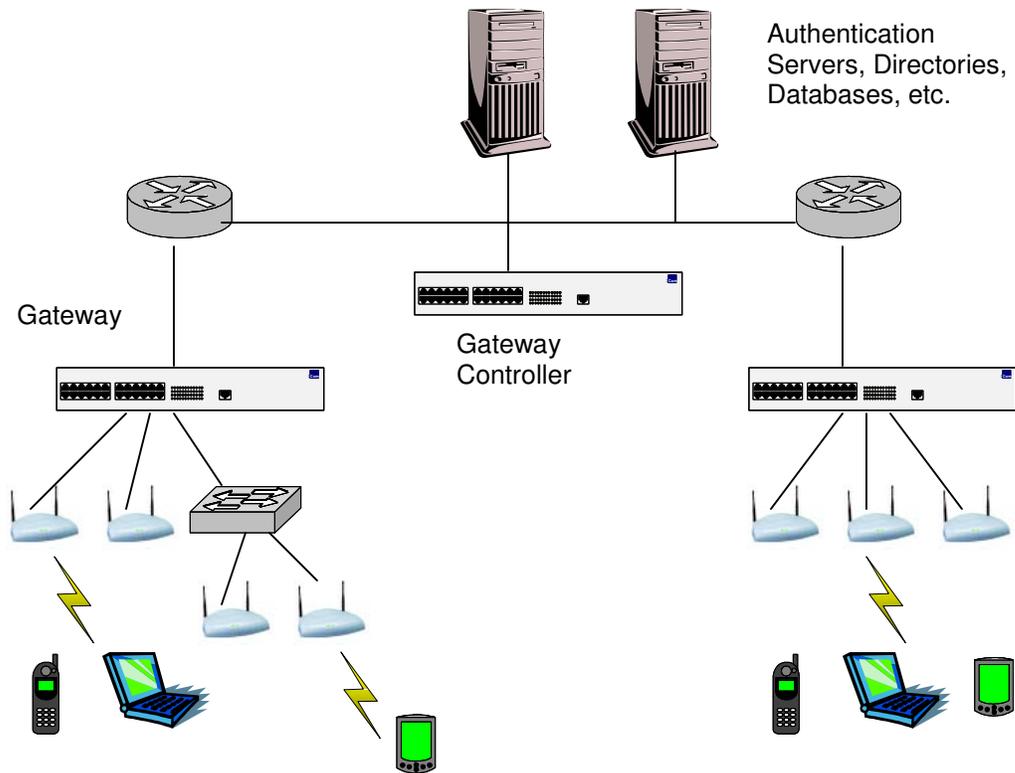
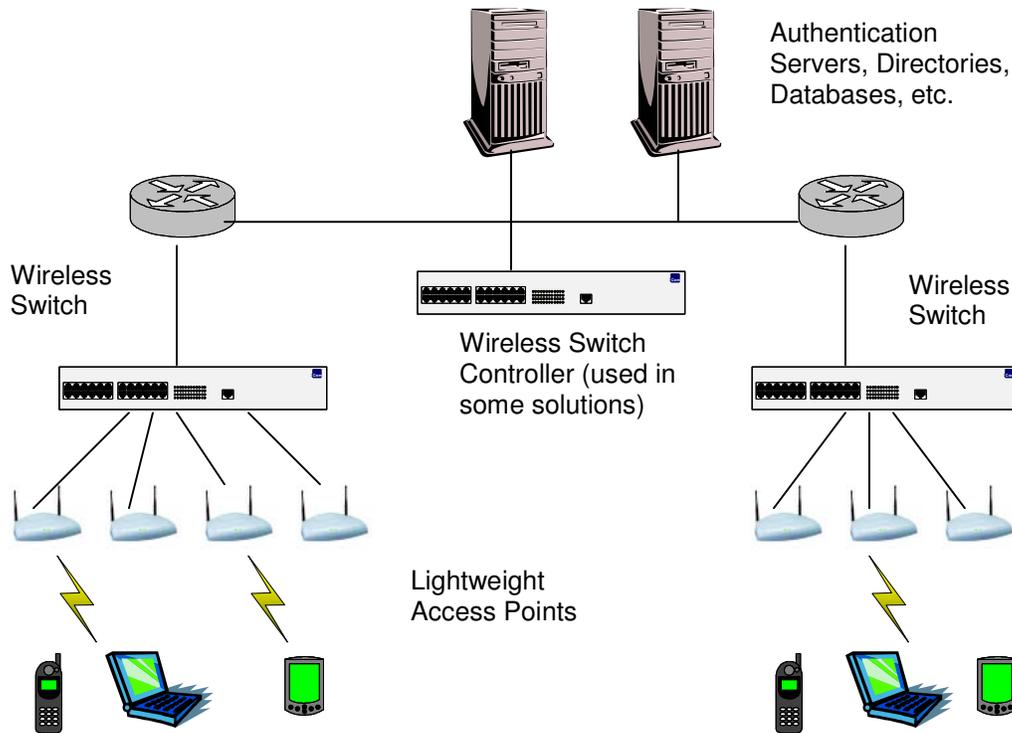


Figure 2. Gateway topology.

### 2.3 Wireless Switch Topology

In September 2002, Symbol Technologies announced a new product that had the advantages of the gateway solution such as layer 3 roaming, enhanced security features, and user/group based policy control, but also added centralized management and control of the access points. This product was the first of many new “wireless switches. An example of a common wireless switch topology is shown in Figure 3.

The access point in the wireless switch topology differs greatly from the access point in the smart access point and gateway topologies. The access points in the wireless switch topology are often called “lightweight” (or “dumb” or “thin”) access points. These APs are designed to work together with a specialized wireless switch. Lightweight access points typically only perform 802.11 wireless Ethernet to 802.3 Ethernet translation and occasionally do encryption/decryption. The wireless switch performs the other necessary functions such as logging, monitoring, filtering, authenticating, etc. as well as providing power to the AP. The access point configuration is even done at the wireless switch. When a lightweight AP is disconnected from its powered Ethernet connection, its memory is lost. This is an attractive security feature because it is useless for attackers to steal and hack into the access point, since all configuration data and encryption keys are



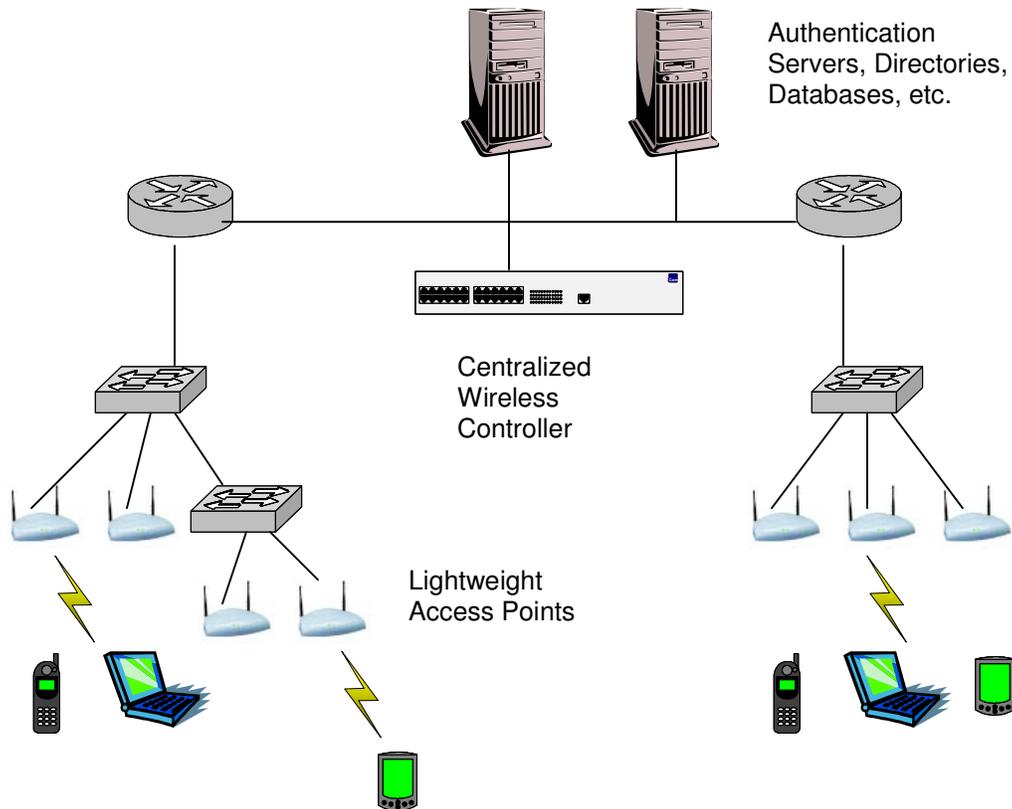
**Figure 3. Wireless switch topology.**

stored in volatile memory. Lightweight access points are also considerably less expensive than their smarter counterparts.

## **2.4 Centralized Controller Topology**

In April 2003, a wireless startup called Chantry Networks introduced a line of 802.11 wireless products based upon a routed IP architecture. This routed IP architecture differed greatly from the architecture every other 802.11 wireless company was promoting, and is the basis for the centralized controller topology. An example of the centralized controller topology is shown in Figure 4.

With a centralized controller topology, a wireless LAN controller is typically located in a data center or a distribution center somewhere near the core or distribution layer of the network. This controller manages a large number of lightweight access points at the edge of the network scattered across the campus. All packets sent to and from the wireless clients are passed through the centralized controller, as with the gateway and wireless switch solution. The centralized controller performs the same functionality that a wireless switch would perform. The difference here is that the traffic is often routed across Layer 3 subnets from the AP to the controller, and vice versa.

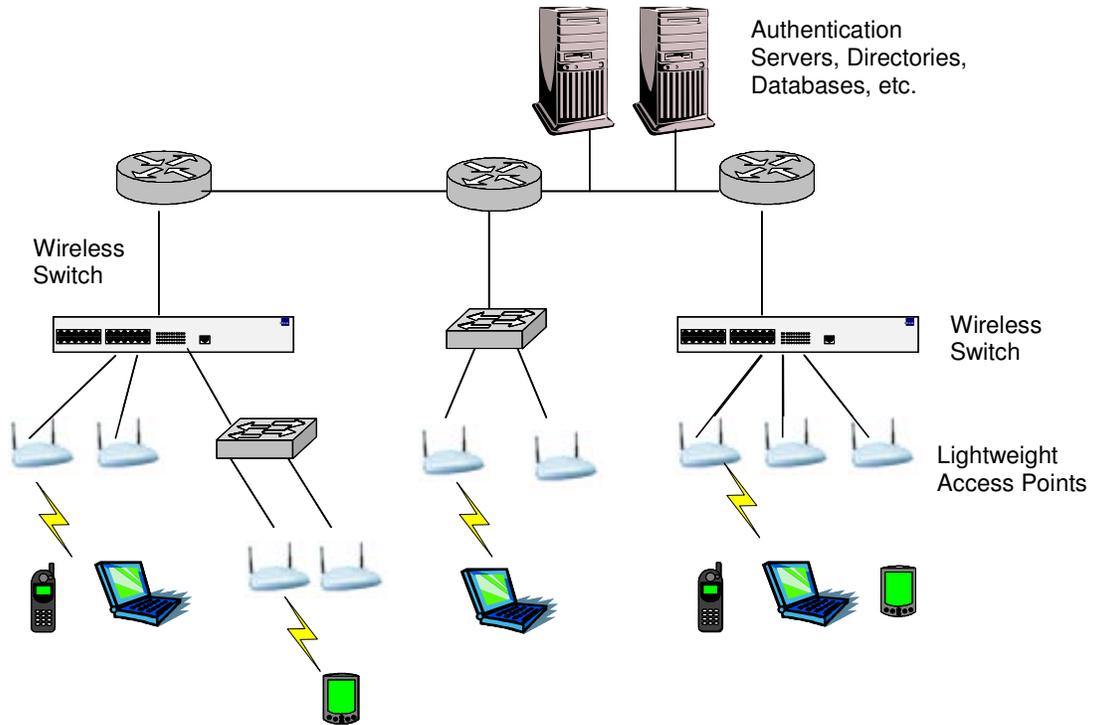


**Figure 4. Centralized controller topology.**

This topology gives network designers flexibility because access points do not need to be directly attached to a specialized wireless switch. Access points can be plugged into an Ethernet switch and will route their way back to the centralized controller for their configuration. This is especially useful for situations where only a few access points are deployed in an area

## **2.5 Hybrid Topology**

Airespace Networks, Aruba Networks, and Trapeze Networks have built Layer 3 routing capabilities into their existing wireless switch products, allowing them to function as both a wireless switch with direct connectivity and a centralized controller, thus combining the topologies discussed in sections 2.3 and 2.4. Figure 5 shows an example of this hybrid topology. In Figure 5, the access points connected to the Layer 2 Ethernet switch would be controlled by one of the two wireless switches in the picture. This hybrid topology allows wireless switches to be used in places where access point densities would be higher, and small standard Ethernet switches to be used for less dense AP deployments such as mobile office trailers and small buildings.



**Figure 5. Hybrid topology.**

This hybrid topology was deemed most suitable for use at Sandia, because of the flexibility it provides. Currently, the APs are connected to Layer 2 Ethernet switches and controlled by wireless switches in the core of the network. However, this topology lets Sandia retain the option to place a small wireless switch at the edge of the network and connect APs directly into it, if conditions dictate.

### 3 Wireless Client Evaluation

On this project, we performed throughput tests of various PCMCIA and Cardbus wireless client cards to determine suitable candidate cards for use at Sandia. We also examined installation procedures for necessary device drivers and associated user interface/configuration programs to determine ease of installation and use. Another group (TechDev, the Desktop Technology Development Department) examined wireless clients with other types of interfaces, such as PCI, USB and Compact Flash,

Although isolated screen rooms are best for controlling RF interference and ensuring repeatability [7], due to limited resources we performed our tests in conventional office-type environments. Likewise, rather than having large open spaces for testing roaming and handoff features, we made do in small, confined, areas. The client cards we tested (and many access points, for that matter) did not lend themselves to the newer testing methods [7] involving wiring the test devices together with RF attenuators, combiners, and switches.

For testing the PCMCIA and Cardbus wireless clients, we used FTP, transferring a 16 MB test file to the system containing the client device under test (DUT). The transfer was repeated at least 5 times for each DUT and the results were averaged and recorded. The more significant and/or interesting results are shown here in Tables 1-3.

We found that in general, the host computer processor type, speed, and operating system had little effect on the throughput rates. The larger factor was the make and model of client card (with associated drivers). Occasionally, but not generally, the access point influenced the throughput.

**Table 1. Throughput of 802.11b Clients from a Wired Network Server.**

<b>Make &amp; Model of wireless client card</b>	<b>Approximate Throughput (in Mbps)</b>
Cisco Aironet 350	4.3-4.6
Cisco AIR-CB21ag	5.2-5.5
Integrated (in Dell laptop)	4.5
Integrated Prism2 (in Fujitsu laptop)	3.3
Linksys WPC51AB	4.2
Netgear MA401	3.9-4.4
Netgear WAB501	4.3-4.8
Proxim Orinoco 8460-05	4.6-4.8
Proxim Orinoco 8480-WD	4.4-6.1

**Table 2. Throughput of 802.11a Clients from a Wired Network Server.**

<b>Make &amp; Model of wireless client card</b>	<b>Approximate Throughput (in Mbps)</b>
3Com 3CRPag	14.8-15.7
Cisco AIR-CB21ag	20.3-21.7
Linksys WPC51AB	18.6
Netgear WAB501	14.4
Proxim Orinoco 8460-05	19.0-21.7
Proxim Orinoco 8480-WD	15.4-18.1

**Table 3. Throughput of 802.11g Clients From a Wired Network Server.**

<b>Make &amp; Model of wireless client card</b>	<b>Approximate Throughput (in Mbps)</b>
3Com 3CRPag	9.7-15.5
Cisco AIR-CB21ag	17.6-19.9
Integrated (in HP Compaq laptop)	14.0-15.0
Integrated Intel Pro 2200BG (in Toshiba laptop)	15.0-16.2
Proxim Orinoco 8480-WD	12.9-17.5

Note that most of the devices fell into the same performance ranges. In Table 1, the performance of the Integrated Prism 2 in the Fujitsu may be artificially low. We found the Fujitsu had slower wireless performance with all wireless clients. This may be due to the Transmeta 5800 processor and the way it processes the Intel instruction set. The Netgear WAB501 a/b card, shown in Table 2, seems to have a lower than expected performance and the newer Prxim Orinoco a/b/g devices (8480-WD) produced slightly slower results than the older devices a/b (8460-05).

The software (drivers and configuration programs) for some cards installed “cleanly” but software for some of the other cards was more troublesome. Generally, the software for the Cisco and Proxim Orinoco cards installed without problems and allowed easy configuration of the DUT. Generally, the software for the Linksys and Netgear cards required tweaking and fiddling to get the card working properly, a definite disadvantage in large deployments. The 3Com card had installation problems in some cases, but once installed, operated well.

After performing testing and evaluation, we recommended the Proxim Orinoco 8460-05 (802.11a/b) card for situations demanding the highest performance and the Proxim Orinoco 8480-WD (802.11a/b/g) for overall versatility. (The Orinoco 8480-WD wireless client card is a Cardbus card, but works in older PCMCIA-slotted systems.) After further examination by TechDev, they confirmed the Proxim Orinoco 8480-WD (802.11a/b/g) card as the corporate recommendation. Since that time, the older 8460-05 (802.11a/b) card has been discontinued by Proxim in favor of the more flexible 8480-WD (802.11a/b/g) card.

The Cisco 802.11a/b/g wireless card was not available at the time of our initial client card testing. Once it became available, we tested and evaluated the AIR-CB21ag wireless client card from Cisco and concluded it could be a good alternative for systems with Cardbus (32 bit) slots. This card is not backward compatible with PCMCIA (16 bit) slots, but did exhibit consistently good performance (regarding both throughput and mobility characteristics) in modern systems. Eventually the Cisco AIR-CB21ag wireless client card was recommended by TechDev instead of the Proxim 8480-WD.

## 4 Wireless Access Point Evaluation

Some formal testing of APs was performed to learn the general characteristics and operation of access points. Some range testing was performed in the 2.4 GHz band and some load testing was performed to examine lightweight vs. fat AP behavior at 11 and 54 Mbps rates. Informal testing was performed to determine ease of configuration and management, throughput, and various vendor-dependent features.

### 4.1 Range Testing

Indoor testing of APs for range was a futile exercise because of interior wall construction, door construction, various obstacles with differing tendencies to block radio waves at various frequencies, and whether interior doors were open or closed. Only a minimal amount of outdoor, open field range testing was performed for this project, since most of our wireless infrastructure deployment at Sandia is expected to be indoors.

The AP used for this testing was a D-Link DWL-1000, 802.11b. Two clients were used. One was a Gateway laptop containing a 1.06 GHz Pentium III processor and a Proxim Orinoco 8460-05 (802.11a/b) card. The other was a Fujitsu sub-notebook with an 800 MHz Crusoe TM5800 processor and a Proxim Orinoco 8480-WD (802.11a/b/g) card.

The results of the open field testing (as measured by the Proxim client utility, in db) are listed in Table 4. Both client systems lost their connection to the 802.11b access point at a distance between 275 and 300 feet from the AP.

**Table 4. Signal Strengths at Various Distances from the AP.**

Distance from AP (ft.)	Gateway/Proxim Client (db)	Fujitsu/Proxim Client (db)
25	30	31
100	12	32
150	20	10
200	19	7
250	16	18
275	21	17
300	8	9

Both connections lost

### 4.2 Load Testing

We performed load testing of access points to observe the AP's behavior under an increasing load. We used the following procedure:

1. Benchmark each client (system/wireless interface) alone against the AP in question using IPERF.

2. Add clients, one at a time, executing the IPERF test. (Set up for the appropriate test and hit 'enter' simultaneously on all clients.) Do not change the location of the clients between the time of calibration (benchmarking) and testing of a given AP.
3. Collect data and plot both average and aggregate throughput for IPERF.
4. Repeat for the next AP to be tested until all APs have been tested.

Access Points to be tested:

- Trapeze (A mode)
- Trapeze (B mode)
- Trapeze (G mode)
- Cisco (A mode)
- Cisco (B mode)
- Proxim (A mode)
- Belkin (G mode)

This would show us the response of both heavyweight (fat) and lightweight (thin) APs to increasing load. The data is summarized in Tables 5-11.

**Table 5. Trapeze 802.11a Throughput Rates.**

Number of simultaneous clients	Client 1 (Mbps)	Client 2 (Mbps)	Client 3 (Mbps)	Client 4 (Mbps)	Client 5 (Mbps)	Client 6 (Mbps)	Aggregate (Mbps)
1	14.8						14.8
2	10.5	7.7					18.2
3	4.7	4.3	6.5				15.5
4	2.9	2.4	6.4	2.6			14.3
5	2.2	2.3	6.0	1.9	2.3		14.7
6	2.0	1.5	6.0	1.7	1.8	1.8	14.8

**Table 6. Trapeze 802.11b Throughput Rates.**

Number of simultaneous clients	Client 1 (Mbps)	Client 2 (Mbps)	Client 3 (Mbps)	Client 4 (Mbps)	Client 5 (Mbps)	Client 6 (Mbps)	Aggregate (Mbps)
1	4.1						4.1
2	2.3	2.2					4.5
3	1.5	1.4	2.0				4.9
4	1.3	0.9	1.5	1.7			5.4
5	1.0	.9	1.2	1.4	.9		5.4
6	0.8	0.6	0.9	1.1	0.8	0.8	5.0

**Table 7. Trapeze 802.11g Throughput Rates**

Number of simultaneous clients	Client 1 (Mbps)	Client 2 (Mbps)	Client 3 (Mbps)	Client 4 (Mbps)	Client 5 (Mbps)	Client 6 (Mbps)	Aggregate (Mbps)
1	6.6						6.6
2	4.1	3.4					7.5
3	2.1	1.8	4.7				8.6
4	1.9	1.8	4.1	19.2			27.0
5	1.2	1.4	3.6	19.2	2.2		27.6
6	1.4	1.0	2.5	19.1	1.5	1.5	27.0

**Table 8. Cisco 802.11a Throughput Rates.**

Number of simultaneous clients	Client 1 (Mbps)	Client 2 (Mbps)	Client 3 (Mbps)	Client 4 (Mbps)	Client 5 (Mbps)	Client 6 (Mbps)	Aggregate (Mbps)
1	12.7						12.7
2	9.1	6.5					15.6
3	4.6	2.7	8.7				16.0
4	2.8	2.1	7.0	3.3			15.2
5	2.5	2.0	6.2	2.5	2.4		15.6
6	2.3	1.6	5.2	2.0	2.3	2.1	15.5

**Table 9. Cisco 802.11b Throughput Rates.**

Number of simultaneous clients	Client 1 (Mbps)	Client 2 (Mbps)	Client 3 (Mbps)	Client 4 (Mbps)	Client 5 (Mbps)	Client 6 (Mbps)	Aggregate (Mbps)
1	4.7						4.7
2	2.5	2.6					5.1
3	1.1	1.2	3.4				5.7
4	0.9	1.0	2.9	1.0			5.8
5	0.7	0.7	2.9	0.8	0.7		5.9
6	0.7	0.6	2.6	0.7	0.7	0.7	6.0

**Table 10. Proxim 802.11a Throughput Rates.**

Number of simultaneous clients	Client 1 (Mbps)	Client 2 (Mbps)	Client 3 (Mbps)	Client 4 (Mbps)	Client 5 (Mbps)	Client 6 (Mbps)	Aggregate (Mbps)
1	12.1						12.1
2	11.3	8.1					19.4
3	4.9	4.6	7.3				16.8
4	3.6	3.7	6.9	4.3			18.5
5	2.8	2.5	7.3	2.6	2.8		18.0
6	2.5	1.7	6.0	2.6	2.2	2.7	17.7

**Table 11. Belkin 802.11g Throughput Rates.**

Number of simultaneous clients	Client 1 (Mbps)	Client 2 (Mbps)	Client 3 (Mbps)	Client 4 (Mbps)	Client 5 (Mbps)	Client 6 (Mbps)	Aggregate (Mbps)
1	12.3						12.3
2	7.0	8.1					15.1
3	2.8	2.5	10.5				15.8
4	2.3	2.2	9.0	3.2			16.7
5	1.8	1.9	7.7	2.5	2.7		16.6
6	1.7	2.1	7.3	1.9	1.9	1.9	16.8

Our observations were as follows:

- Heavyweight vs. lightweight access points did not make a significant difference.
- Generally, average throughput decreased in an exponential fashion with increasing client load. The average aggregate throughput increased with client load and levels out at about 2-3 clients, with possibly only slight increases after that point. Trapeze G mode is an exception. This is illustrated by Figure 6-Figure 11.
- Once you get above the saturation number of clients (2-3), aggregate throughput and hence, percent utilization, is about constant. Aggregate throughput is split between one “head hog at the trough” and all others. The others (other than the “head hog”) split the remaining bandwidth about equally.
- Average aggregate throughput for 2 or more clients is 25-35% of the 54 Mbps signaling rate for 802.11a; about 45-55% of the 11 Mbps signaling rate for 802.11b. One client gets the largest portion, and the remaining clients split what remains.

An approximation of the expected rate might be:

$$\text{Typical Rate} = \frac{(\text{Ave. Aggregate Throughput at Saturation}) - (\text{Bandwidth Rate of "Head Hog"})}{(\text{Number of Clients}) - 1}$$

From our tests of the Trapeze access point in 802.11a mode, we found that the maximum client throughput was asymptotically approaching a floor of 5-6 Mbps (see Table 5). This indicates that with more than just 3 or 4 clients in use, one client would probably be supported at about 6 Mbps, with the remaining clients supported equally at a slice of the remaining bandwidth. Observing from Figure 7 that the aggregate throughput of the Trapeze access point operating in 802.11a mode is about 15 Mbps, and using the approximation above to estimate how many clients could be supported (active at exactly the same time) with a floor of 1 Mbps to each client, we see that 10 clients could be supported (1 at about 6 Mbps and 9 at about 1 Mbps).

$$1\text{Mbps} = \frac{15 - 6}{n - 1}, \quad 1 = \frac{9}{n - 1}, \quad n - 1 = 9, \quad n = 10 \text{ clients}$$

If the “head hog” only received about 5 Mbps, then the other 9 clients would likely each receive about 1.1 Mbps, or potentially 10 clients could be supported at about 1 Mbps, bringing the total number of clients supported by this access point to 11. In practice a greater number of clients could likely be supported at these rates or better, as they would likely not all be active at the same moment.

To find a point where access point performance drops off drastically, or the access point completely stops functioning due to load, would take more equipment than we had available. At some point, with a heavy enough load where access point firmware limitations are being approached, it would be interesting to see if new connections are rejected or if they are accepted and possibly causing overflows in tables. Our suspicion is that before that point is reached, performance would be poor enough that additional clients would not try to associate.

In summary, we have observed performance characteristics that seem to be very similar across types and brands of access points and operating modes. Aggregate throughput of an access point levels out quickly to a constant amount and when that happens, one client gets the largest portion of that aggregate, while the remaining clients share the remaining amount approximately equally.

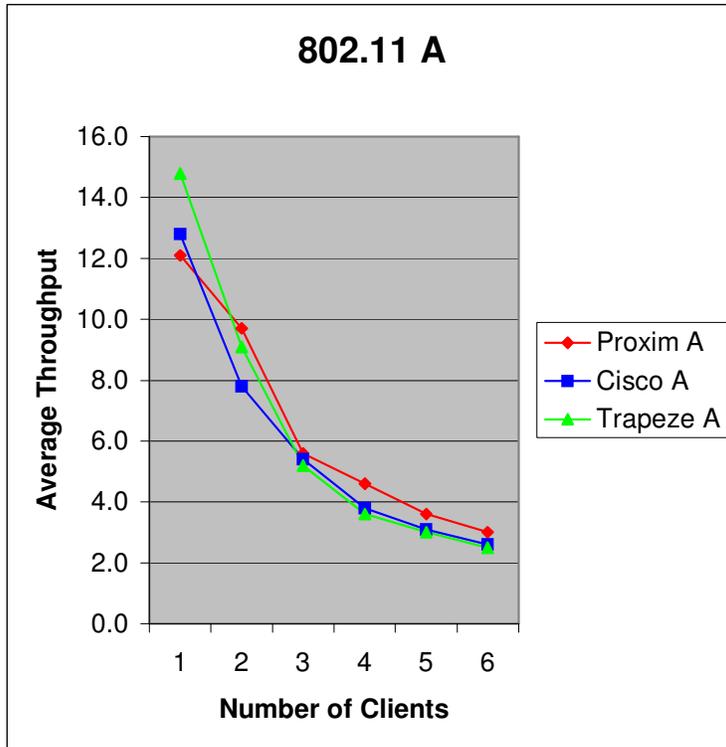


Figure 6. Average Client Throughput on 802.11a Access Points.

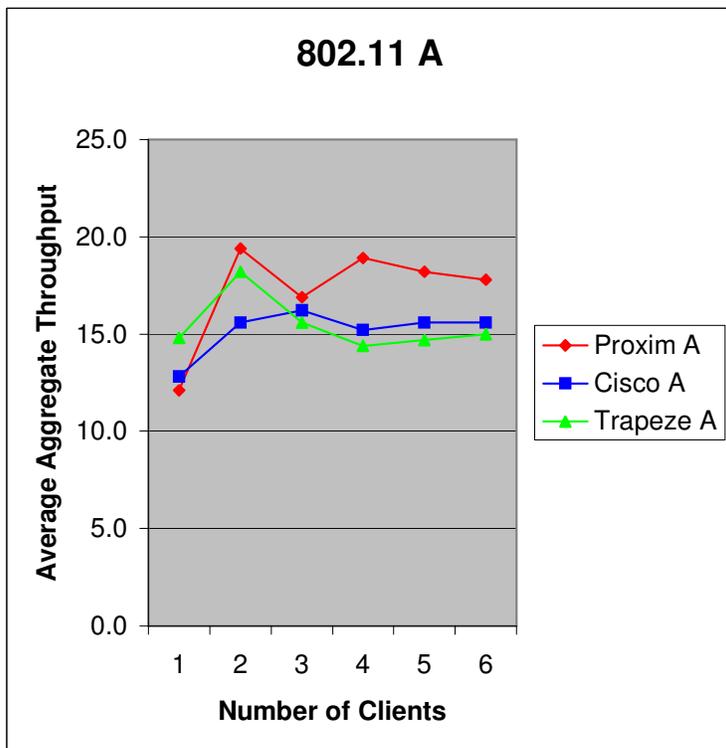


Figure 7. Aggregate Client Throughput on 802.11a Access Points.

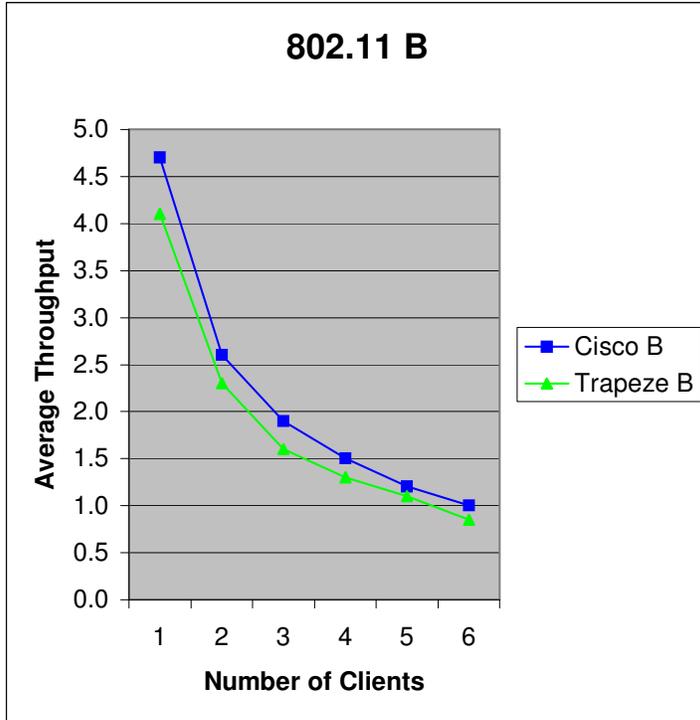


Figure 8. Average Client Throughput on 802.11b Access Points.

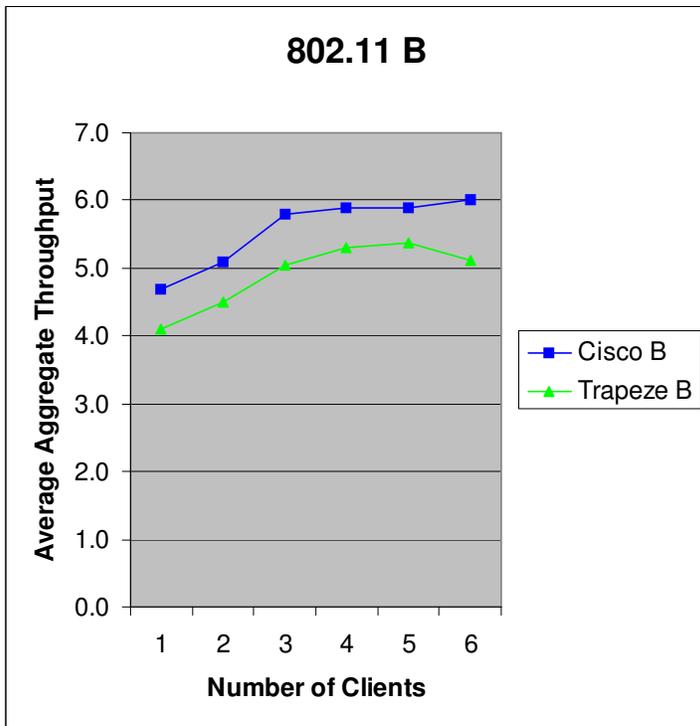


Figure 9. Client Throughput on 802.11b Access Points.

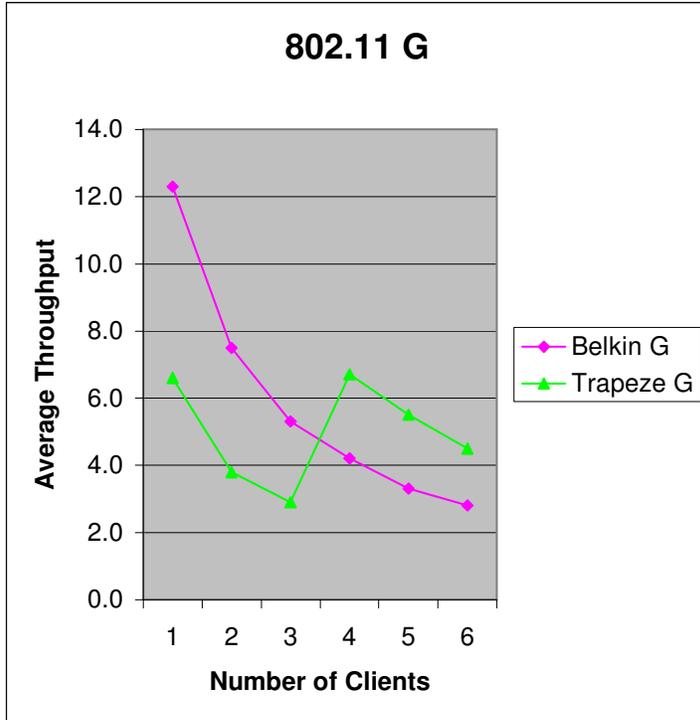


Figure 10. Average Client Throughput on 802.11g Access Points.

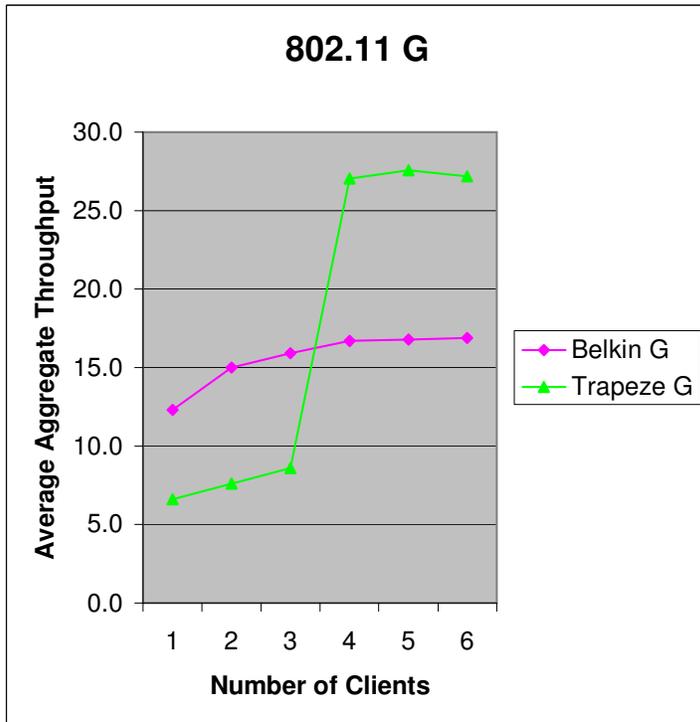


Figure 11. Aggregate Client Throughput on 802.11a Access Points.

### **4.3 Product Evaluation**

One of the first tasks for the wireless networking project was to become familiar with the 802.11 wireless infrastructure technologies in order to make choices as to what product or products would be evaluated prior to making a large-scale purchasing decision. After reading 802.11 literature and scanning the wireless market, product specifications were gathered. The team also attended network trade shows in order to meet many of the 802.11 vendors, learn of their products, and ask questions. The wireless networking team then evaluated wireless networking topologies, systems, and their components.

#### **4.3.1 New Mexico**

At Sandia, NM, several wireless networking testbeds were built. These were used to conduct much of the AP testing and client card testing. The testbeds provided the WLAN team with the requisite flexibility to test and evaluate a variety of products supporting the WLAN topologies described in section 2 of this report.

It was clear early on that Cisco was the market leader for 802.11. The team decided to purchase some Cisco 802.11 access points for testing and deploying in pilot installations. Cisco Aironet 1200 access points were purchased and tested for range and throughput, then installed in two locations for initial user testing, as described in section 8.1.1.

There was already a move in the industry away from self-contained or “smart” access points (such as the Cisco Aironet) to “dumb” access points that required a switch or controller to provide the configuration and functionality (intelligence) to the access point. The primary benefit of this technology is the lower cost of each access point due to the lower demands on the hardware. An additional and very important benefit is the ability to manage many access points from a single control point. The team recognized the benefits of this type of approach and decided that further product evaluations would focus on the “dumb” access point design. This product space was growing rapidly, with some established companies producing products and several start-ups entering the market.

Several vendors were contacted for obtaining products for testing. Access points and switches or appliances from Airespace, Inc., Aruba Networks, Chantry Networks, Symbol Technologies, and Trapeze Networks were obtained through loan agreements. The equipment was tested in a lab environment. Throughput tests were conducted to determine the maximum data rates that the access points could sustain. Rates were determined using both file transfers (FTP) and memory-to-memory transfers (iperf<sup>2</sup>). The file transfer tests (using a file of approximately 16 MB in length) represent “real life” results, or what one could reasonably expect from actual data transfers. The memory-to-memory tests are more academic in that they represent the maximum transfer rates that a particular situation can attain. This is useful for setting the ceiling of expectation (i.e.

---

<sup>2</sup> <http://dast.nlanr.net/Projects/Iperf/>. Iperf is a software tool for testing end-to-end data transfer capability. Data is written from memory to the device network interface, across the network, into the receiving interface, and to the memory of the receiving device. No transfer to hard drive takes place, which can slow the transfer process considerably.

“it’s the best it can possibly do”). A conclusion that came out of the testing was that most WLAN infrastructure products performed similarly as far as throughput on a given band was concerned. So, this would most likely not be a deciding factor in differentiating the products.

A small amount of equipment was purchased from Symbol, tested in the lab, and then deployed in a building. However, there was difficulty in getting the equipment operational, and the effort was abandoned in favor of waiting for a better solution.

A WLAN design was developed that encompassed the hybrid topology (see section 2.5) using a combination of “dumb” access points, wireless switches, and wireless controllers. It does not rely only on wireless controllers or wireless switches, but allows for either type for whatever the particular situation warrants.

From this design, an RFQ was written for soliciting bids for a large purchase of WLAN equipment to position Sandia to begin deploying wireless. The RFQ process (described in section 9), which took several months to complete, resulted in the purchase of equipment from Aruba Networks.

### **4.3.2 California**

Sandia, CA evaluated two WLAN vendors products — the Cisco 1200 series APs and the Trapeze Networks mobility system. The Cisco WLAN equipment consisted of the 1200 series APs and the associated Power over Ethernet (PoE) injectors needed. The Trapeze Networks mobility system consisted of a WLAN switch, APs that were connected to the switch (directly connected at first then indirectly connected later), and a server running the WLAN planning and management software called Ringmaster.

First, the Cisco 1200 series APs were evaluated by installing them in building MO52 after conducting a manual site survey to determine where to place them and what power and channel settings to configure. After installation, measurements were made to make sure good RF coverage was available throughout the building. WLAN packets were captured and analyzed with AirMagnet and AiroPeek to see how well the WLAN was working, and computer security was invited to watch the area for any problems. After making sure that the WLAN worked as it was supposed to, about 6 to 8 MO52 pilot users were invited to start using the WLAN via a Virtual Private Network (VPN) connection to the Sandia Open Network (SON), with Sandia Restricted Network (SRN) access coming later after computer security was satisfied that the risk was acceptable. The initial pilot consisted of user machines running Windows, Mac OS, and Linux so part of the evaluation was to see how well these different operating systems worked on the WLAN using the required VPN client software. The test users provided regular feedback as to how the WLAN was working for them.

The results were that the WLAN worked fine but the Linux VPN client had some problems. As newer versions of the Cisco VPN client for Linux became available, the client became more stable and there were fewer-to-no complaints concerning the VPN

client or the WLAN. During the time that the MO52 users were exercising the Cisco-based WLAN, work began to deploy and evaluate the Trapeze Networks equipment.

After seeing how well the Trapeze mobility system worked at their headquarters, we decided to take them up on their offer to deploy their system at Sandia, CA for evaluation. At that time, the most desirable feature of the Trapeze system was the planning tool looked to be much easier to use than a traditional site survey and seemed like it would be a real time saver. Buildings MO50/51 were selected because of their central location in the Property Protection Area (PPA), the fact that it would be a small, but non-trivial deployment (12 APs, hence quite manageable), and that there were some willing users located there. After planning the AP locations and configurations, and installing the Trapeze hardware, the system was tested for RF coverage, throughput, security, and usability. Not only was the RF coverage verified with tools like AirMagnet and AiroPeek, but also with a WLAN-enabled laptop connected to the SRN via VPN running a continuous ping to an offsite location. This test was conducted while walking around the buildings and roaming between APs. It gave an indication that there was good coverage and the VPN wouldn't drop while roaming between APs. When the laptop roamed from one AP to another, a few pings might drop but the connection would stay up and the VPN client was unaffected. Throughput was measured while connected via the VPN over various wireless network interface cards (NICs) using a tool called Chariot from NetIQ and also by timing the Windows file share transfer of a large (approximately 1 GB) directory. The measured throughput varied slightly between NICs but was normally around 4.5 Mb/s for 802.11b and about 18 Mb/s for 802.11g. After testing the new Trapeze system, users in MO50/51 were invited to begin using the WLAN and provide feedback often. About 6 users began using the WLAN in MO50/51 and the feedback was infrequent, but those who did provide feedback seemed to have no problems except for an occasional issue with the VPN client. Through all of the WLAN testing at Sandia, CA, the common factors that caused problems were usually the VPN client or the client firewall.

The Sandia, CA networking department had access to some year-end funds in FY03 and used some of it for wireless networking equipment, so the first purchase of Trapeze Networks equipment was made at that time. The equipment that Trapeze Networks loaned Sandia, CA for the MO50/51 evaluation was purchased, as well as some additional hardware. After purchasing the Trapeze Networks equipment in the summer of 2003, Sandia, CA began deploying it in favor of the Cisco 1200 series equipment and eventually replaced all the Cisco APs with Trapeze APs. During the one year period from when the Trapeze Networks evaluation equipment was installed and the funds were allocated to purchase wireless equipment for "production" networking, several more locations in the Sandia, CA Property Protection Area were outfitted with the additional wireless equipment from Trapeze Networks. In summer 2004, when funds were allocated to purchase and deploy "production" wireless networking equipment, Sandia, CA chose to remain with the Trapeze Networks, based on the features, equipment cost, and compatibility and experience with existing equipment (see section 9).

## 5 Wireless Local Area Network Architecture

This design reflects the SNL/NM wireless networking system but does not directly address remote wireless networks at home and on travel. Also, the SNL/CA wireless networking system is a separate but integrated effort and while not fully addressed here, is referenced and partially described.

### 5.1 Background

Stand-alone access points are self-contained units such as the consumer type found in homes, as well as enterprise-class units. Enterprise-class stand-alone access points are relatively expensive, and difficult to manage since each unit must be managed separately, or managed as a whole with a separate software package. Also, if a stand-alone access point is lost or stolen, an adversary could potentially extract network and security configuration information stored in the access point.

In contrast to stand-alone access points, OSI Layer 2 or Layer 3 wireless systems employ “dumb” access points that do not store significant configuration information within them. They obtain configuration information from a wireless switch or controller, and lose the configuration when power is lost (such as would happen if a dumb access point is stolen). This also reduces the cost of each access point since each does not require the memory or processing power needed to support various services such as a web server, SSH server, etc. An additional benefit of the wireless switch or controller technology is the built-in centralized control of the access points. These systems control all attached access points.

SNL has chosen not to deploy stand-alone access points for reasons of cost, security, and manageability. The most versatile solution that meets cost, security, and manageability requirements is the hybrid approach (see section 2.5 for a more detailed description), which SNL has chosen to deploy. The hybrid approach provides wireless switches where concentrations of access points are high enough to warrant a switch, and provides wireless controllers where low access point density favors routing “dumb” access points to a central location.

The basic concept of the WLAN at SNL is an isolated network that provides wireless access service to the SNL wired networks, using the SON as the transport infrastructure. The SON is chosen because wireless networking is treated as remote access, since the airwaves are beyond our physical control. Ideally, a physically separate infrastructure should be utilized for wireless access; however, this is not practical from a cost standpoint with such a large enterprise as SNL. A compromise is to use the SON infrastructure. This will require augmentation of the SON in areas where existing SON infrastructure is inadequate for a particular wireless installation, and will require additional SON infrastructure where there is none available for a particular wireless installation.

An isolated portion of the SON contains the wireless access points and switches, separated from the remainder of the SON by a firewall (or router access control list -- ACL)<sup>3</sup>. Access to the remainder of the SON and to the SRN is by VPN only. VPN access provides privacy by strong encryption (Triple DES), and authentication by two-factor SecurID or CRYPTOCARD, using SNL's existing VPN service.

Due to the limitations of Wired Equivalent Privacy (WEP) and the lack of a suitable Layer 2 protection replacement at the time of implementation, no Layer 2 protection is provided to the wireless network. Authentication and encryption are provided at Layer 3 by the VPN.

## **5.2 Wireless Network Usage Policy**

Connection of wireless access points to the SRN is not allowed. A separate external network must be used for wireless devices (e.g. SON). Wireless network users must have a computer and wireless network card registered in NWIS and have a wireless device that meets the configuration requirements consistent with the security design.

### **5.2.1 Acceptable use guidelines**

**Wireless Network Access.** Wireless devices can be used to connect to the wireless network if requirements for device registration and proper use are met. All wireless devices must be registered in NWIS, and users must be aware of proper use, including knowing where wireless use is permitted and where it is not permitted.

**Wireless Network to SRN.** The wireless network is outside the SRN and beyond physical control (air waves); therefore, Laboratory-approved remote access methods must be used (e.g., 2-factor authentication and VPN software) to gain access to the SRN.

**SRN to Wireless Network.** Pinging wireless devices from the internal network is allowed. For access to mobile devices on the wireless network, authorized users can use secure shell over their internal network connection to the target device. Dual active network interfaces (one wired on internal network and one wireless on wireless network) are not allowed.

**Internet to Wireless Network.** No direct access is allowed from the Internet to SNL wireless clients.

### **5.2.2 Ad Hoc Networks**

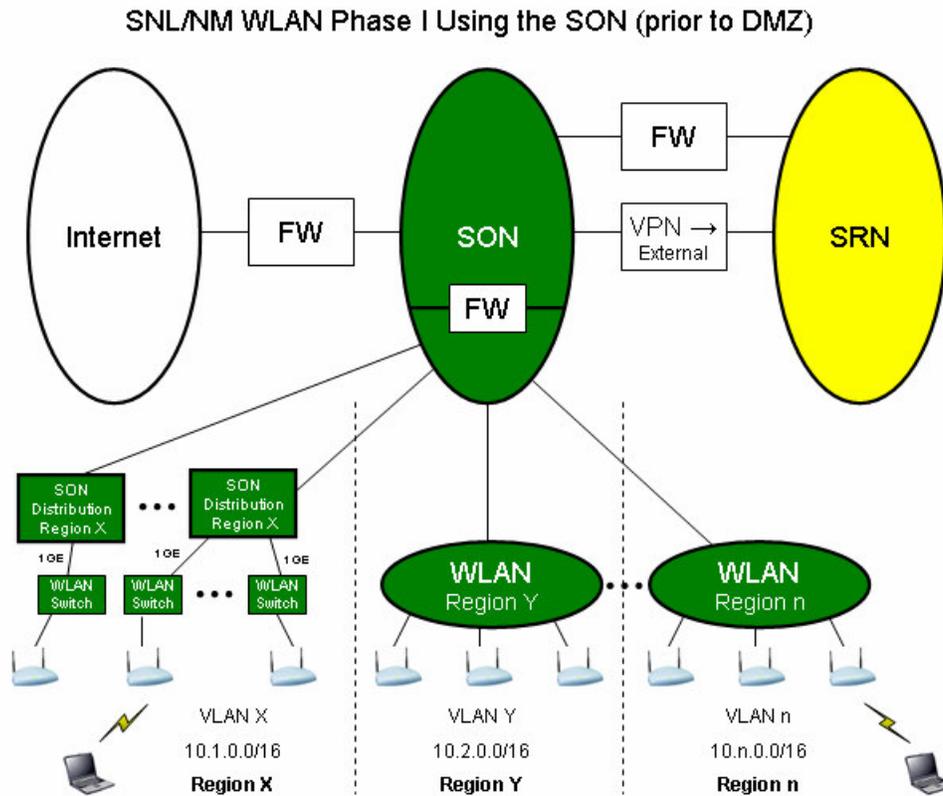
Ad hoc networks, also known as peer-to-peer networking, are networks where clients communicate directly with each other, not through an access point. These networks are not permitted due to the difficulty in enforcing necessary security practices and policies.

---

<sup>3</sup> A firewall is preferable for this application due to its stateful packet inspection and logging capability. However, for Phase I, it would be simpler to implement a router ACL, and save the firewall for Phase II and the DMZ.

### 5.3 Wireless Design

The SNL/NM WLAN design is modeled after the SNL/CA WLAN design, but extended to encompass the larger NM campus. The CA model is a single VLAN built on the SON, serving the entire CA campus. This single WLAN VLAN isolates wireless devices from all other SON traffic. While realizable at CA due to the relatively small campus size, it does not scale at NM. Instead, by utilizing the “right-sized” design approach as specified in the Network Architecture and Design document [13], the CA model can simply be duplicated over several NM “right-sized” regions as shown in Figure 12.



**Figure 12. WLAN Architecture.**

In this model, a single SON VLAN is dedicated to the WLAN within each region, functioning as a wired backbone for wireless access points. This SON VLAN separates wireless devices from the remainder of the SON, negating the need to operate a single WLAN VLAN over the entire enterprise network. Each WLAN region would correspond to the wired LAN regions.

Figure 12 shows three WLAN regions, X, Y, and n, where n is any number of duplicated regions as necessary to service all desired areas. Region X is further broken out to show how access points connect to WLAN switches, which in turn connect to SON distribution switches. A single VLAN in Region X keeps all access points and wireless switches

isolated at Layer 2 from all other SON traffic<sup>4</sup>. Isolation at Layer 3 and above is accomplished by firewall or router ACL.

A wireless client attempting to associate to an access point will initially authenticate to the access point using open authentication<sup>5</sup>. After successful association, the client will have access outside the WLAN only to DHCP for IP address assignment, ICMP ECHO (ping) for testing, and the VPN concentrator for access to the SRN. No access into the WLAN will be permissible from the Internet. This restriction is provided by a firewall or router ACL at the point where the WLAN meets the SON (at Layer 3 and above).

Access to the SRN is accomplished by running the client VPN and authenticating the user using SecurID or CRYPTOCard. This is the same VPN client as used for remote access (Cisco), so no further software development or deployment is necessary. WLAN users simply enable VPN access through WebCARS if this has not been done already for wired remote access through the VPN.

### **5.3.1 User Access**

Referring to Figure 12, a WLAN client must first associate with an access point. The WLAN client will receive an IP address such as 10.1.0.1 from DHCP. Once the IP address is assigned, the client will have access only to the VPN concentrator as a result of the firewall or router ACL that limits access. The user then starts the VPN software, establishes contact with the VPN concentrator, and authenticates using SecurID or CRYPTOCard. The client is now connected to the SRN. The path from the client to the SRN is encrypted using Triple DES.

### **5.3.2 Compatibility with CA**

The NM and CA WLAN efforts are being closely coordinated so that a final solution will allow interoperability between the sites.

## **5.4 Throughput Considerations**

**802.11a** – signal rate is 54 Mbps. Data rate is approximately ½ the signal rate, or about 22 Mbps. Typical measured throughput is approximately 15-21 Mbps.

**802.11b** – signal rate is 11 Mbps. Data rate is approximately ½ the signal rate, or about 6 Mbps. Typical measured throughput is approximately 4-6 Mbps.

---

<sup>4</sup> See <http://www.cisco.com/go/safe>, Cisco's comprehensive set of guidelines for ensuring VLAN security. Layer 2 separation is a compromise since there are known vulnerabilities. A completely separate wireless network is preferable but not realizable.

<sup>5</sup> The initial pilot implementations at NM used WEP. However, this was discontinued after discussions with SNL/CA in order maintain compatibility between the sites. WPA is becoming available across many products and could be phased in as a replacement for WEP. 802.11i will eventually replace both WEP and WPA.

**802.11g** – signal rate is 54 Mbps. Data rate is approximately ½ the signal rate, or about 22 Mbps. Typical measured throughput is approximately 13-20 Mbps. However, mixing both 802.11b and 802.11g clients results in an 802.11g signal rate of about 11 Mbps due to design constraints of 802.11b/g compatibility. This reduces overall throughput for an 802.11b/g system. It is anticipated that 802.11b clients will eventually disappear in favor of 802.11a and 802.11g, thus eliminating the constraint on 802.11g performance. However, until that day comes, 802.11g will potentially suffer from the mixed environment.

With the VPN client, WLAN throughput drops considerably due to the software encryption overhead placed on the client device, and the sharing of bandwidth with other VPN clients at the VPN concentrator. Wired connections see a decrease in throughput from about 25-33 Mbps down to about 10 Mbps. Wireless connection throughput ranges from 4 to 8 Mbps. This rate is largely dependent on the particular WLAN client machine: the faster the machine, the better the WLAN/VPN throughput. As can be seen here, the VPN will throttle the throughput rates of 802.11a and 802.11g, even though the observed rates of the various cards listed in section 3 are considerably higher.

#### **5.4.1 Wireless Clients**

Because a wireless access point is a shared-medium device, the maximum throughput must be divided amongst wireless clients. For example, for an 802.11a system, the maximum throughput of 20 Mbps must be divided by the number of simultaneous clients. If 10 wireless clients are associated to an 802.11a access point, then each client could experience an average throughput of 2 Mbps. For an 802.11b access point, 10 clients would experience an average throughput of 600 Kbps. Our load testing experiments (section 4.2) revealed some inequities (regarding client throughput) in sharing the transmission medium.

As a result of the sharing of bandwidth, it is recommended that 802.11a access points be limited to approximately 10 simultaneous wireless clients, with a maximum of approximately 20 clients. For 802.11b, only approximately 5 simultaneous wireless clients can be effectively supported, with a maximum of approximately 10 clients. For 802.11g, the numbers are the same as for 802.11a. However, the combination of 802.11b and 802.11g must be considered, which unfortunately keeps the recommended number of wireless clients at approximately 5.

#### **5.4.2 VPN Concentrator**

A potential bottleneck for the wireless system is the VPN concentrator that provides access to the SRN. The current concentrator is a Cisco VPN 3030, with 50 Mbps maximum encrypted throughput and shared between wireless access and Internet access. A single wireless switch, fully populated and running at maximum throughput, will easily overrun the VPN concentrator. This can be mitigated by clustering several existing VPN 3030 concentrators in a load-sharing configuration.

Three VPN 3030 concentrators in a load-sharing configuration can support up to 150 Mbps of encrypted traffic. Additional concentrators can be added to the cluster as needed, including more capable concentrators such as the VPN 3060 or 3080 (however, all Cisco VPN concentrators are currently limited to Fast Ethernet interfaces; no Gigabit Ethernet interfaces are available).

## 6 Wireless LAN Security Approach

### 6.1 SNL Wireless Security Design

The WLAN at SNL is an isolated network that provides wireless access service to the SNL wired networks, using the SON as the transport infrastructure. Figure 13 shows the high-level design. The wireless network is built on an isolated portion of the SON. This isolated portion is separated from the remainder of the SON by a firewall.

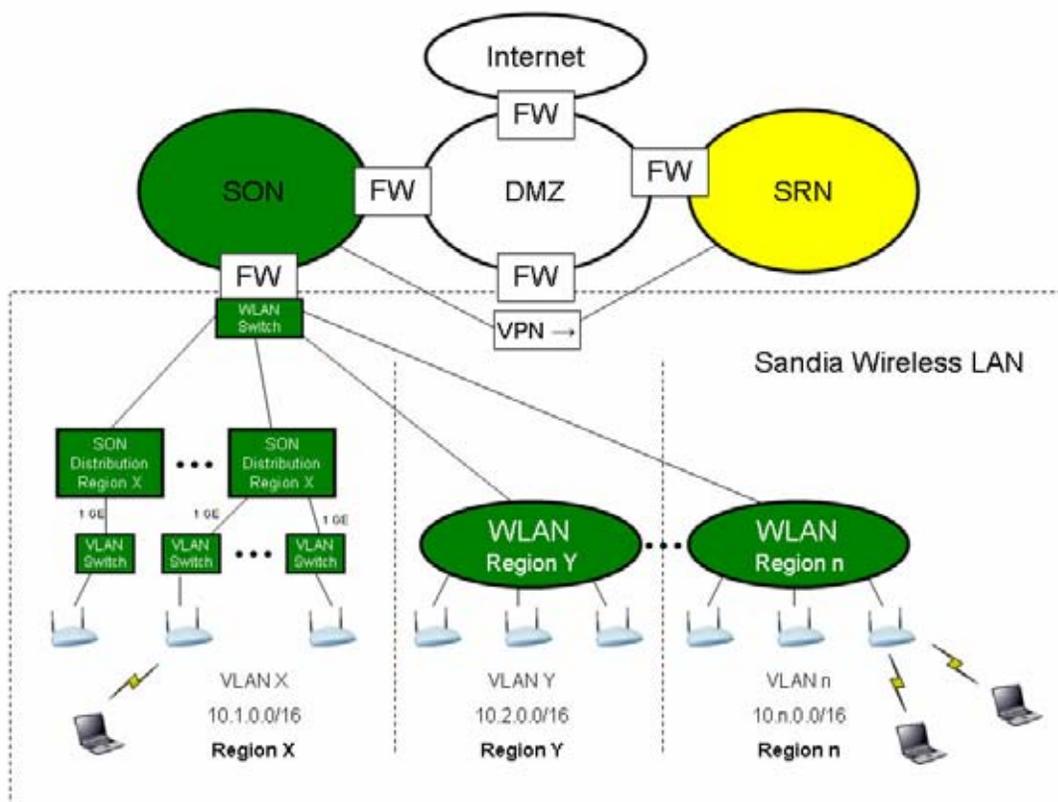


Figure 13. SNL Wireless LAN Design.

Wireless networking at Sandia is treated and protected like remote access, since the RF airwaves are beyond our physical control. Wireless access to the SRN and the non-wireless portion of the SON requires a VPN. VPN access provides privacy by strong encryption (3DES), and authentication by two-factor SecurID or CRYPTOCARD, using SNL's existing VPN service.

A single SON VLAN is dedicated to the WLAN within each region, functioning as a wired backbone for wireless access points, as shown in Figure 13. Each WLAN region corresponds approximately to the existing wired LAN regions. Figure 13 shows three WLAN regions, X, Y, and n, where n is any number of duplicated regions as necessary to

service all desired areas. Region X is further broken out to show how access points connect to LAN switches, which in turn connect to SON distribution switches, then to the wireless switch (wireless controller). A single wireless VLAN in Region X keeps all access points and wireless switches isolated at Layer 2 from all other SON traffic. Isolation at Layer 3 and above is accomplished by firewall.

## **6.2 User Access**

Referring to Figure 13, a WLAN client that associates with an access point will receive an IP address such as 10.1.0.17 (for example) from DHCP (if the client's MAC address is registered in Sandia's NWIS database). Once the IP address is assigned, the client will have access only to the VPN concentrator, as a result of the firewall that limits access. The user then starts the VPN software, establishes a connection to the VPN concentrator, and authenticates using SecurID or CRYPTOCard. Following these steps, the client is now connected to the SRN. The path from the client to the SRN is encrypted using 3DES.

Any wireless client that does not possess the proper VPN credentials, such as the group key and valid two-factor token card, will not be able to access any network resources.

## **6.3 Cyber-Security Threats, Risks, and Mitigations**

### **6.3.1 Insertion Attack**

An insertion attack is a security breach that takes place when unauthorized devices connect to the wired or wireless network without proper security process and review. The unauthorized device could be a client or an access point. If an access point is not correctly configured or connected using the proper network topology, an unauthorized client may associate with the access point and gain access to the network.

Access points are fairly inexpensive and can be easily purchased at many electronics stores. The ease of buying and deploying access points often makes it attractive for individuals in an organization to put up their own access points and plug them directly into the wired network. This is known as a rogue access point because it is not supported by the networking staff. Rogue access points present a security risk because they are often not properly secured and can allow unauthorized access to the network. An intruder gaining access to the facilities could also connect a rogue access point to the wired network to allow remote access to the network from outside the physical facilities.

### **6.3.2 Insertion Attack Defenses**

Sandia has established policies prohibiting the addition of access points without prior approval from Computer Security. Continuous rogue access point detection and containment will be performed by the production wireless system in many areas across Sandia's Albuquerque and Livermore campuses. If a rogue access point is attached to a

wired network at Sandia, the deployed production wireless system in the area will detect the rogue, alert the necessary personnel, and contain the rogue if requested.

Any areas not covered by automatic rogue access point detection will be manually swept to ensure that no rogue access points exist. Manual detection is the least efficient rogue access detection method and will only be done when automatic detection is not feasible.

### **6.3.3 Interception and Monitoring Attack**

Sandia's wireless LANs use RF to provide a wireless connection to wireless enabled devices. As previously mentioned, RF signals cannot be constrained to a room or building in a practical manner. An attacker could easily intercept and monitor RF signals from the organization's parking lot or even miles away with a high-gain directional antenna.

Passive monitoring of the WLAN presents a greater threat than insertion attacks. With insertion attacks, an intruder gaining access to the network can be tracked down and preventive measures can be taken. Passive monitoring allows an attacker to capture traffic without being easily detected. If an attacker is able to intercept authentication credentials, the attacker could then access the network as an authorized user, making it very difficult to detect the intruder. The attacker could also opt to continue undetected as a passive monitor.

### **6.3.4 Interception and Monitoring Attack Defenses**

Sandia requires all SNL wireless clients to use a VPN to sufficiently protect the traffic from interception. Figure 13 shows that the Sandia wireless LAN will be placed on a separate network from the SRN with a firewall and VPN concentrator separating the two networks. No user services will be available on the wireless network. This will mitigate the threat if an intruder were capable of attaching to the wireless network.

Strong two-factor authentication that incorporates "something you have" and "something you know" is used to increase the protection of the network. Security tokens creating random passwords with limited lifespan and blacklisting offer good protection. If an attacker were to intercept a valid user's authentication credentials, the limited lifespan of the password would limit the window of time for the attacker to gain access to the network using the intercepted credentials, while the blacklisting prevents reuse of any captured credentials.

The client radios will be configured to communicate with access points only, known as infrastructure mode. This helps to insure that two-factor authentication and VPN encryption are always used. Also, mobile devices such as laptops and tablet PCs will be required to have a software firewall to protect the data stored locally from attackers who may try probing the client devices.

### **6.3.5 Session Hijacking and Man-in-the-Middle Attacks**

Session hijacking allows an attacker to assume the identity of an authorized client. With session hijacking, an attacker issues a disassociate command to a valid client and then assumes the valid client's identity with the access point. When the client receives the disassociate command, it will disconnect from the network.

To implement a man-in-the-middle attack, an attacker places a rogue access point in the vicinity of a wireless LAN, with a stronger signal than a legitimate access point. A client will then attempt to associate with the rogue access point because the signal is stronger. At the same time, the attacker relays all messages to the legitimate access point, impersonating the valid client to the legitimate access point. At that point, the attacker can monitor all traffic or take over the user's session.

### **6.3.6 Session Hijacking and Man-in-the-Middle Attack Defenses**

Session hijacking and man-in-the-middle (MITM) attacks are very difficult to implement against a VPN. It is infeasible to discover a long, random VPN group key using a brute-force or other cryptographic attack, and without the group key the attacker cannot initiate a session with the VPN concentrator. Also, since man-in-the-middle attacks require rogue access points to be in the vicinity of the wireless LAN, rogue access point detection schemes will be utilized in order to identify and contain an attacker's access point. Once the attacker's AP is contained, the MITM threat is nullified.

### **6.3.7 Denial of Service (DoS) Attacks**

Denial of Service (DoS) attacks prevent authorized users from accessing legitimate services by overwhelming the network with illegitimate traffic. DoS attacks are common in wired networks and can also be applied to wireless networks. Whether intentional or unintentional, wireless networks can be jammed from other RF sources. Equipment using the 802.11b and 802.11g standards operates in the 2.4 GHz Industrial, Scientific, and Medical (ISM) band. Other RF equipment such as cordless phones, microwave ovens, and Bluetooth transmitters also operate in this band. Large numbers of these devices can cause enough interference with the WLAN to render it inoperable. Malicious jamming could also be done in order to deny service to the wireless LAN. Equipment using the 802.11a standard operates in the 5 GHz frequency band. This band is currently less cluttered than the 2.4 GHz band, so equipment using the 5 GHz band will typically receive less incidental interference from neighboring devices.

A DoS attack can be performed on an access point by flooding it with association requests. An access point or wireless client can be denied by saturating it with bogus packets. Also, an attacker or employee can configure their client to duplicate the IP or MAC address of another legitimate client, causing a disruption on the network.

### **6.3.8 Denial of Service Attack Defenses**

The physical environment of the facilities where a wireless LAN will be deployed will be examined using RF auditing tools. RF auditing tools will identify possible sources of interference that may interfere with the WLAN. All RF usage at Sandia must be approved by the frequency coordinator. The frequency coordinator identifies and investigates possible interference prior to the deployment of RF technologies. Access points can be positioned and directional antennas can be used to avoid interference.

An intrusion detection system could be installed on the wireless LAN to identify an excessive number of association requests or DHCP requests. Excessive association or DHCP requests often point to a malicious attempt to deny access to the wireless network.

It should be noted that denial of service attacks on the wireless LAN are not as destructive as the other attacks noted above. This is because a malicious denial of service attack intends only to prevent access to the wireless LAN, not obtain access to the network or intercept sensitive information. This must be kept in mind when determining what level of protection Sandia wants to deploy to defend against a denial of service attack.

### **6.4 Wireless LAN Security Summary**

The wireless LAN at Sandia is implemented as a production level protected wireless network at Sandia that satisfies current and future customer needs for mobility and portability. Wireless is treated like remote access – only a VPN with two factor authentication allows access to the SRN and SON, as shown in Figure 13. In addition, no user services are located on the wireless network.

All data transmitted across the wireless LAN is protected by a Triple-DES encrypted VPN. This protects against interception and monitoring. Rogue access point detection and containment will be performed in order to protect against MITM attacks and insertion attacks.

## 7 Rogue Detection and Containment

Unauthorized (or rogue) access points are a threat to Sandia data and/or networks. Rogue access points are unauthorized access points that are deployed by Sandians who do not know the rules regarding wireless deployment, Sandians who know the rules but decide to disobey them, or by malicious attackers wishing to connect to the Sandia networks or attract Sandia wireless clients. In the case of the knowing or unknowing Sandian, rogue access points are a threat because they could provide a non-secure path to the Sandia data networks. If a malicious attacker had physical access to a Sandia building, the attacker could connect a hidden access point to the network in order to connect to and navigate the Sandia network from an adjoining building or parking lot (or even further with specialized equipment). An attacker may also try to lure Sandia laptops to his or her access point in order to attempt to steal credentials or attempt to exploit a system vulnerability. The Sandia host-based firewall and anti-virus is a mitigating factor in this case. Another mitigating factor for all of the threats listed above is a wireless intrusion detection technique called rogue access point detection and containment.

Rogue access point detection is a process in which wireless sensors (air monitors) and access points are used to listen for 802.11 wireless signals and report the existence of access points to software that discerns between authorized and unauthorized access points. Many systems can also report whether or not any wireless clients are connected to an access point.

When a system detects a rogue access point, it generates a wireless intrusion detection alarm. If a wireless client connects to the rogue, the system may, according to configuration, contain the rogue by spoofing the BSSID of the rogue access point and sending a very large number of deauthenticate or disassociate packets to the wireless client. When this happens, all standard 802.11 wireless clients will be unable to pass wireless network traffic through the rogue.

Both of the vendors of wireless networking infrastructure equipment selected for use at Sandia, Aruba Networks and Trapeze Networks, supply rogue access point detection and mitigation capabilities [5] [6] [21] in their products. This could also be supplemented by third party, systems, such as those from AirDefense [1] or AirMagnet [3].

Extensive testing performed at Sandia has shown rogue access point detection and containment to be very effective in SNL wireless deployments. To test this, the wireless team moved an access point considered to be a rogue by the wireless infrastructure to many locations within Sandia wireless-enabled buildings. The team also used a wireless client to connect to the rogue access point to test rogue containment. The client was able to associate with the access point, but after a few “pings” were passed from client to rogue access point, the containment disallowed nearly all other wireless traffic from the client. The formal test results are stored in Web FileShare under the Wireless LAN Pilot Monthly Report documents.

## **8 Pilot Wireless LANs**

Several wireless LAN pilots were established in New Mexico and California to test WLAN design and equipment. These also served to help establish policies and procedures.

### **8.1 New Mexico Pilots**

Pilot wireless LANs were installed initially outside of the Limited Areas, and later, in close coordination with DOE/NNSA, inside Technical Area 1.

#### **8.1.1 Pilots Outside the Limited Area**

The initial pilot WLAN installations were in the medical building (831) and the trailers housing the medical applications developers (T53/T54/T55/MO249). These pilot WLANs used the Cisco Aironet 1200 access points.

##### *8.1.1.1 Medical Building Pilot*

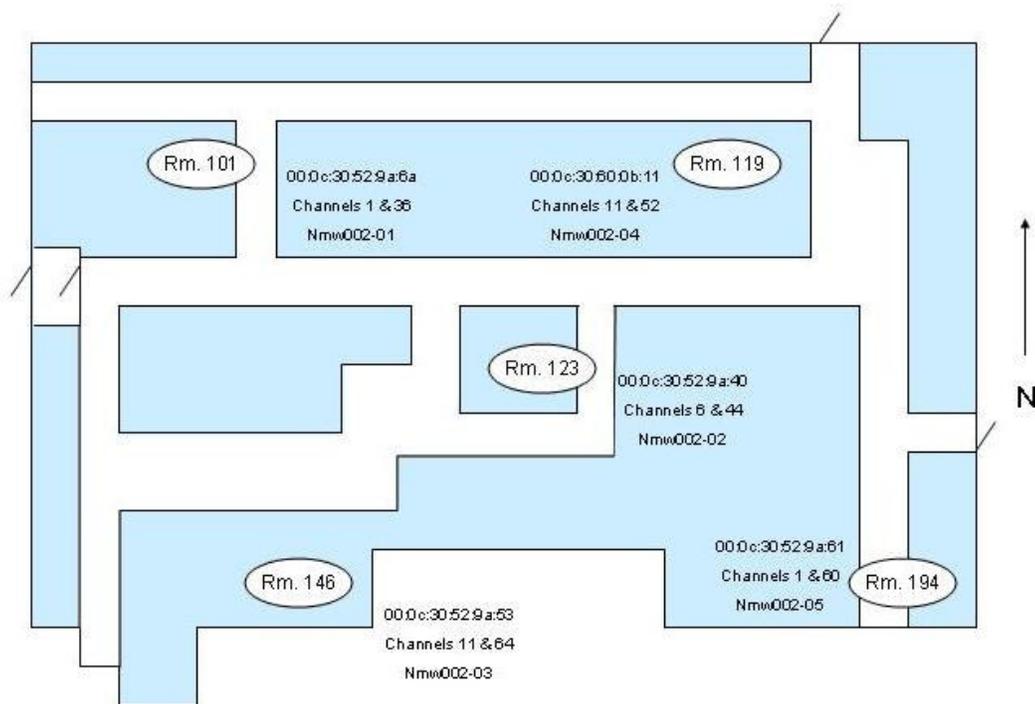
Wireless access in Building 831 is intended primarily for physician and support staff use of portable devices such as tablet PCs and PDAs. This is a growing trend in the medical profession. Doctors can have immediate access to patient information at virtually any location within the facility.

The 831 pilot WLAN was composed of both 802.11a and 802.11b radios in order to study the pros and cons of each technology. While 802.11b is very pervasive in both the commercial and consumer space, 802.11a is designed for greater throughput in a quieter frequency band (e.g. no microwave ovens), providing potentially greater performance for enterprise-class installations. Wireless enabled clients in Building 831 utilized 802.11a/b or 802.11a/b/g interface cards, having the option to associate with either radio type (a or b).

A site survey was conducted by the WLAN team to determine the optimum locations and number of Access Points (APs) for the Medical facility. The first step was to obtain a building diagram, and use this to approximate the AP locations. Then a walk-around was performed to determine feasibility of the preliminary AP locations. Two constraints dictated the locations: proximity to a drop box for wiring the AP to the network, and a suitable mounting point (preferably a ceiling location). Drop boxes were utilized in order to minimize the effort required to wire the access points to the network, as this was a pilot implementation and changes may be necessary.

The Access Point locations are shown in Figure 14. While it is desirable to locate the Access Points such that coverage is optimum, the actual locations chosen are a compromise based on available space and proximity to drop boxes. For example, the

Access Point in the far southeast corner should be located further north and east, centered in a hallway, however no drop boxes were available there<sup>6</sup>.

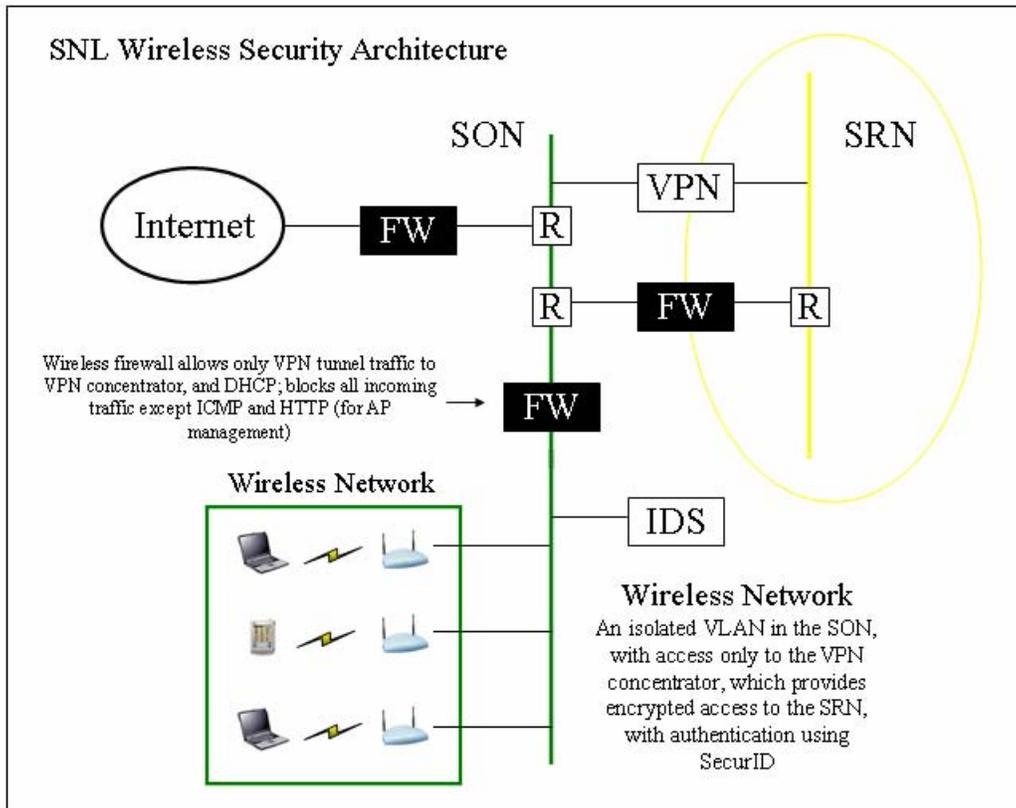


**Figure 14. Layout of Access Points in Bldg. 831.**

Figure 15 shows the security architecture for the pilot. The architecture was designed to maximize security, which also minimizes the potential for unauthorized WLAN access. Because a wireless network extends beyond the physical boundaries that protect a wired network, the wireless network will be treated as a remote-access network. As such, two-factor authentication, and strong encryption, was required, and supplied by a VPN layered over the WLAN. The VPN also overcomes the weak WLAN-native encryption (i.e. WEP). However, the two can be combined to provide additional security (the layered approach).

The architecture shown in Figure 15 was designed to isolate the WLAN from the remainder of the wired network and restrict access into and out of the WLAN. The primary service allowed out of the WLAN is the VPN, and this access was restricted to the VPN concentrator only. DHCP service is also allowed out of the WLAN, in order to dynamically provide IP addresses to WLAN clients. ICMP is allowed both into and out of the WLAN in order to provide simple but necessary communications such as “pings” that help determine the health of the network. HTTP is also allowed into the WLAN to allow us to remotely manage the APs.

<sup>6</sup> An Ethernet cable could have been run from the Access Point located in the hallway to a nearby drop box, but this would possibly require conduit, and certainly Facilities involvement, which was to be avoided for this pilot in order to maintain flexibility.



**Figure 15. Security Architecture.**

With these access restrictions in place, the only path out of the WLAN was by VPN (except DHCP and ICMP), which is an authenticated service. If an unauthorized individual gains access to the WLAN by associating with an AP from outside the building (as one example), that individual will have access only to the VPN concentrator, which is a trusted service that will only allow users that authenticate with a SecurID or CRYPTOCard token and PIN. This is the same trust in the VPN concentrator that pertains to the Internet, where the concentrator is exposed to unauthorized connection attempts from anywhere in the Internet world, and is thwarted by two-factor identification.

An additional layer of protection is the use of WEP, which while relatively easily broken, does send the signal that this AP is closed and not available for public access. Any attempts to break into this AP are deliberate and unlawful. (This is analogous to locking one's house; a break-in is still possible by breaking a window or possibly kicking in the door, but at least the message is clear; this house is not open for unauthorized access, and any attempt to do so is against the owner's will.) One drawback to the use of WEP is the need to utilize a shared secret for generating a key.

The 831 WLAN was constructed by installing an SON Ethernet switch in the 831 wiring closet for wiring to provide network connectivity to the WLAN. This switch was purchased specifically for this project. The switch was a Cisco 3550-24PWR, designed

to provide power to the Access Points using Power over Ethernet, thus avoiding having to run power to each Access Point. The 3550 switch was connected to the SON backbone via multi-mode fiber. Five 10/100 Ethernet ports were set to Virtual Local Area Network (VLAN) 91. SON subnet 132.175.91.0 was used for the WLAN pilot. An access control list (ACL) was applied to the router interface serving subnet 132.175.91.0. This ACL allowed IPsec and IKE traffic off of the subnet (VPN), ICMP in and out, DHCP in (for providing clients with an IP address), and HTTP in for managing the access points remotely. A complete description of this pilot can be found in a separate white paper [16].

Performance and usage data was collected from tests and operation of the medical pilot WLAN. Although there were some performance issues to be resolved, the 802.11a/b service was available and usable for access to the SRN. Eventually the pilot WLAN was extended to cover trailer T13, providing access for more medical personnel.

#### *8.1.1.2 Medical Application Developers Pilot*

The other initial WLAN pilot location in New Mexico was the trailer complex housing the medical applications developers (T53/T54/T55/MO249). This area was piloted so the software developers could have a development environment similar to the environment in which their applications would be operating.

The network and security architectures, and equipment (Cisco 1200 APs) were the same as those used in the medical building pilot. Again, usage and performance was monitored to collect data regarding the pilot.

### **8.1.2 Pilots Inside the Limited Area**

Early in 2005, Sandia, NM received permission from NNSA/AL to deploy and operate pilot wireless LANs in three buildings inside the Limited Area for a finite period of time. The wireless pilot was authorized for several months, through June 25, 2005. The locations permitted included all of the building where the network design and operations personnel are located, and conference rooms in one of the newest buildings supporting collaborative engineering efforts.

A more complete discussion of these pilot WANs can be found in the white papers by Witzke [29] and Wiener [23]. After the authorized pilot period, these pilot networks were authorized to move to production status.

## **8.2 California Pilot**

Starting in March of 2003, Sandia, CA began piloting a WLAN in building MO52 located in the PPA near the northeast corner of the site. After getting approval for a wireless “testbed” in MO52, two people with walkie-talkies conducted a manual site survey using a battery-powered AP, a laptop running the Cisco Aironet site survey utility (included with the old Aironet client utility software that came with a Cisco Aironet 350

series wireless NIC), and a PocketPC running AirMagnet. We ended up deploying three Cisco 1200 series APs with 2 dbi omni directional antennas and the power turned most of the way up to cover the rather small mobile office building. After getting the WLAN in MO52 up and running and several users testing it, a few additional WLAN's were installed around site using the Cisco 1200 series APs. Security was provided by a private SON VLAN (dedicated to the WLAN) routed to our newly installed De-Militarized Zone (DMZ), ACLs, a controlled DHCP server, VPN encryption, and SecurID authentication.

During the spring/summer of 2003, Sandia, CA contacted Trapeze Networks to begin testing a WLAN switch with "dumb" APs to be located in buildings MO 50 and MO 51. The industry seemed to be moving away from the "smart" APs like the Cisco 1200 series that were not easily managed when large numbers of them began to be deployed, so it seemed reasonable to begin testing these new products. Trapeze Networks, a company whose headquarters was less than 10 miles from Sandia, CA, was on the leading edge of WLAN switch/dumb AP technology, and had the author of a premier 802.11 wireless networking book working as the service engineer for the California Bay Area. They were eager to work with Sandia; therefore it was easy to form a good business relationship with them quickly. The APs installed in MO50/51 were directly connected to the Trapeze wireless switch since layer 3 connectivity between the AP's and switches was not yet supported by Trapeze Networks. The deployment went very smoothly and it became immediately clear that the planning tool part of the Trapeze management software ("Ringmaster") was going to be a real time saver, over doing a traditional site survey to determine where to place APs and at what power and channel to set them. It was at about this time that Sandia, NM and Sandia, CA began working together to share lessons learned and formed the wireless networking project team. The Sandia, CA networking department had access to some year-end funds and purchased the equipment that Trapeze Networks had loaned to Sandia, CA for the MO50/51 pilot, as well as some additional hardware. After purchasing the Trapeze Networks equipment in the summer of 2003, Sandia, CA began deploying it in favor of the Cisco 1200 series equipment and eventually replaced all the Cisco APs that had been deployed, with Trapeze APs. Officially, it was during the spring of 2004 that the WLAN pilot ended but at Sandia, CA the pilot was treated more like a production network from the start since they were being supported by production networking department personnel from the time of initial installation.

## 9 Wireless LAN Equipment Selection and Acquisition

At SNL/NM facility a “best value” type procurement was used. In this type of procurement, although price is not included among the ranked criteria, it is evaluated. The best overall value is determined by comparing differences based on the offeror’s relative capabilities and price/cost in relation to all other offers received. In a best value procurement, Sandia might not award the contract to the vendor offering the lowest price; award may be made to a responsible offeror whose proposal offers a greater value.

The statement of work for this Request for Quotation (RFQ) read as follows:

*Contractor shall provide the following, on an as-needed basis, to Sandia National Laboratories: Wireless Local Area Network (WLAN) infrastructure equipment; wireless switches and appliances; tri-mode dual band 802.11 a/b/g/ access points; management, control, diagnostic, and monitoring software; site survey software, and hardware/software maintenance.*

The mandatory requirements were as stated:

*The system proposed shall include access points and appliances to support up to 2000 simultaneous users in a secure, reliable, and maintainable manner. Provide a description of how the proposed system will meet all of the mandatory requirements listed below. The system proposed shall meet or exceed the following mandatory performance requirements.*

### *Experience*

- *Minimum of one (1) year experience providing Wireless Local Area Networks*
- *Provide examples of major systems installed and operating*

### *Performance*

- *Layer 2/Layer 3 roaming (layer 2 – ability to seamlessly roam from AP to AP within a subnet; layer 3 – ability to seamlessly roam from AP to AP across subnets)*

### *System Architecture*

- *Switched architecture (AP switch) and support routed AP appliances*
- *Support simultaneous tri-mode, dual-band AP (802.11a/b/g)*

### *System Security*

- *Multi-channel rogue AP detection (the ability to detect and report an unauthorized AP) while simultaneously serving clients*
- *Rogue AP containment (the ability to prevent an unauthorized AP from accepting connections from WLAN clients)*
- *Client authentication (WEP, WPA-PSK, and 802.1x)*

- *Layer 2 encryption (WEP, WPA, and AES)*
- *Layer 3 VPN (IPSec) pass through compatibility*
- *No security-significant information (encryption keys, Radius shared secrets) stored on an AP after removal of power*

#### *System Management*

- *Performance and diagnostic monitoring (at a minimum, throughput and associations through graphs or text reports)*
- *Integrated management (all APs, switches, and appliances can be managed and monitored from a central point)*
- *SNMP v2 monitoring*
- *Policy management (the ability to assign security, Quality of Service, and VLAN attributes to a group or individual)*
- *Dynamic AP power and channel management*
- *Integrated site survey tool (including the ability to import graphical representations of building floors)*

#### *Electrical Power*

- *A system that has the AP's powered by Power Over Ethernet (PoE) according to the 802.3af specification*

#### *Standards*

- *Wi-Fi Certified (802.11a/b/g) support*
- *Plenum-rated AP*

The evaluation criteria included the following features and factors:

- Wireless Intrusion Detection (WIDS)
- FIPS 140-2 Certification of Cryptographic Modules
- MAC Address-Based Authentication
- Packet Capture of Wireless Packets
- IGMP Processing on Wireless Switches
- Discovery Protocol (Compatible with Cisco's CDP)
- SNMP v3 monitoring
- Reliability and Maintenance Concept
- Security Features and Capabilities
- Availability/Delivery Schedule – provide delivery schedule for proposed units
- Maintenance Plans Available
- Availability of Training for SNL Personnel

After the responses were evaluated, the contract was let to Network Presence, LLC to supply equipment from Aruba Networks. The initial order was received in September 2004. Aruba Network wireless LAN equipment was first installed at Sandia, NM later that month, in September 2004, replacing equipment from the Cisco pilot installation in medical (buildings 831 and T-13).

In the late summer of 2004, Sandia, CA purchased about \$92K worth of WLAN equipment from Trapeze Networks. The order consisted of a combination of APs, wireless switches, PoE injectors, external directional antennas, and associated maintenance. The order was placed through purchasing as a sole source based on the fact that the new equipment needed to be compatible with existing equipment. The equipment was received shortly after the order was placed.

## 10 WLAN Deployment

Following the prototypes, plans were made for enterprise-wide deployment of Wireless LAN technology. Deployment started in buildings outside of the Limited Areas, and once approved by NNSA, continued with buildings inside the Limited Areas.

### 10.1 NWIS Entries for Aruba Access Points

The naming convention provided by Computer Security is based on the RTIW form for each wireless installation. For NM, each RTIW is designated with a name in the form of nmwxxx, where “nm” refers to New Mexico, the “w” refers to wireless, and the “xxx” is the serial number of the RTIW. Each access point for a particular RTIW will be consecutively numbered as a “-xx” following the RTIW designator. For example, the medical installation at NM is designated “nmw002”. Each access point for “nmw002” is named “nmw002-01”, “nmw002-02, etc. This technique allows Computer Security to run reports on wireless installations for auditing purposes.

Figure 16 shows the access point basic information, such as name, owner, location, manufacturer, etc. The vendor serial number is entered as the Aruba location ID, such as 1.1.1. The model AP-52 is entered in the “or other Model” field. The OS is entered as “boot system file” because this was the closest match, and the OS version is entered as “unknown” since there is no match for the Aruba version and we don’t really care to track the version number here.

The screenshot shows a web form titled "Edit a Machine" with several fields and buttons. At the top, there are buttons for "Change", "Delete", "Reset", "Copy", "Cancel", and "Back". Below these are navigation links: "Network Interface", "IP Addresses", "Security Information", "Classified Configuration", "Network Functions", "Installed Apps", "Installed Hardware", "Roles", "Supp Notes", "Chg Log", "Initiatives", and "History". The form fields are as follows:

Machine Name:	nmw002-01	Machine Status:	Active	FA Status:	No Red
Machine Owner:	HAMILL, MICHAEL J.	Name Lookup	Login Lookup		
Reapp Dispatch #:		Machine Location:	SNL in NM		
Building:	831	ELT Bldg Lookup	Room:	101-C	ELT Bldg/Rm Lookup
Manufacturer:	aruba	Kind:	access point		
Preferred Model:		or other Model:	AP-52		
* Property #:		Vendor Serial #:	1.1.1		
OS:	boot system file	OS Version:	unknown		
Alternate OS:		Alternate OS Version:			
Megs of Memory:		Gigs of Disk:			

\* You need to enter a property number for Sandia owned equipment.  
Request Correction to Machine Information  
No SMS automated hardware information available for this machine.

Figure 16. Access Point Information.

Figure 17 shows the interface cards and IP address for this AP. Each radio has its own interface and hardware address in addition to the wired interface.

**Installed Network Interface Cards for Machine nmw002-01**

Card Name	Interface Type	Hardware Address	DHCP Enabled	Changed By	Last Changed
802.11a	wireless	00:0B:86:82:1F:F8	No	MILLER,MARC M	Oct 29 2004 10:08:30 883AM
802.11g	wireless	00:0B:86:82:1F:F0	No	MILLER,MARC M	Oct 29 2004 9:36:54 280AM
fastethernet	fast ethernet	00:0B:86:C0:21:FF	No	MILLER,MARC M	Oct 29 2004 9:37:31 127AM

[Go Top](#)

---

**IP Addresses for Machine nmw002-01**

IP Address	DNS Name	Type	Box	Port	Card Name	Changed By	Last Changed
132.175.91.200	nmw002-01.sandia.gov	canonical	OTHER	OTH	fastethernet	MILLER,MARC M	Oct 29 2004 9:53:06 097AM

[Go Top](#)

**Figure 17. Access Point Interface and IP Information.**

There will also be an alias listed in NWIS that gives a “traditional” Sandia name to the access point. The alias will incorporate the location (e.g. SA, for Sandia Albuquerque), the equipment manufacturer (e.g. AW, for Aruba Wireless Networks), the equipment type (e.g. 9, for an access point), the VLAN number, and the AP number within that VLAN. If this were the first access point on VLAN 91, the alias in NWIS would be SAAW99101.sandia.gov.

### **10.2 NWIS Entries for Trapeze Access Points**

At SNL/CA, the Trapeze Networks access points are registered in NWIS the same as at SNL/NM but with a few minor differences. First, the prefixes of SNL/CA access points begin with “caw” (California wireless) instead of “nmw”. Then, the “caw” prefix is followed by the serial number of the RTIW. Next, the building number is placed in the machine name, followed lastly by the particular access point number. The access points end up with a name like “caw016-916-1”. The NWIS access point entries also contain detailed information concerning the location, serial number, model number, manufacturer, etc.

After entering the access point name and detailed information, the unique hardware addresses of the access point are entered into the NWIS record. For each Trapeze Networks access point, there will be at least four hardware addresses registered — one for each of the two wired Ethernet interfaces and one for each of the two radios (802.11a and 802.11b/g) for each SSID the access point services. The SNL/CA access points utilize DHCP for IP address assignment, so the “DHCP Enable” field will be set to DHCP (dynamic) for the Ethernet interfaces.

### **10.3 Site Survey Methodology**

A WLAN site survey is the process of determining how many AP's will be needed for the desired WLAN site, where to place the AP's in the proposed area, how the AP's should be configured, and if there are any RF obstacles or noise sources that will need special attention. Following are the methods used to conduct a WLAN site survey at the Sandia California and New Mexico facilities.

#### **10.3.1 California**

SNL/CA WLAN site surveys should be conducted using a combination of a visual survey of the WLAN site, interviews with WLAN site customers to determine usage needs and locations, an automated RF planning tool to speed up the survey process, a battery powered access point and RF measurement tool for verification, and a comprehensive WLAN survey, security, and performance monitoring tool such as the AirMagnet product.

##### *10.3.1.1 WLAN Customer Interviews*

Interviewing the customers of a proposed WLAN site (sometimes called a "Pre Site Survey") is very important and should be the first step of any WLAN design. At a minimum, the WLAN designer needs to know how many customers will use the WLAN and at what times, what type of data throughput do the customers expect or need, if there are any special considerations (e.g. proximity to classified processing or safety considerations), the physical size and construction of the proposed WLAN area, and the potential for a future increase of users or remodel of the WLAN area. By completing an accurate pre site survey, the customer can reduce the likelihood that the WLAN will have to be redesigned in the near future and save the WLAN designer time in the initial design.

##### *10.3.1.2 Visual Survey*

A visual survey is simply the process of an experienced WLAN designer walking around the proposed WLAN site, verifying the information from the pre site survey and filling in any gaps of information that the customer might have inadvertently left off. A visual survey will also give the designer an opportunity to determine the type of construction, investigate any special areas (e.g. lead walls in a radiology lab would wreak havoc with RF propagation), and make distance measurements of the WLAN area and the proximity to the nearest buildings and/or classified processing areas. If the area is unfamiliar to the designer, then a WLAN analyzer tool (e.g. AirMagnet) will be useful to determine if there are any other WLAN's operating in the vicinity that may cause interference.

##### *10.3.1.3 RF Planning Tools*

A good WLAN RF planning tool can save the site surveyor much time and shoe tread but must be used in conjunction with an RF monitoring tool for verification. An RF planning tool is a computer program whose input is an electronic drawing of the building or area

where the WLAN will be installed, as well as information as to the construction of the building (e.g. are the walls concrete or wood). Other parameters include the number of users in the area and the expected data throughput. The output is a drawing of where the APs should be placed and how they should be configured (i.e. channel and power settings). Until such time the WLAN designer has acquired sufficient experience using a particular planning tool, the placement of APs should be verified. This is accomplished by using a battery powered AP, placing it in the various locations specified by the planning tool, and verifying that the RF coverage is sufficient for the customer needs.

#### *10.3.1.4 Traditional WLAN Site Survey*

Until recently when RF planning tools have become more prevalent, the common method for a WLAN site survey was the use of a battery powered AP and an RF measuring tool to map the coverage of a proposed WLAN area. The site surveyor would move the AP around the site and meticulously record RF measurements from a measurement tool. This tool was usually just a laptop with a wireless NIC that had a site survey utility to allow raw RF measurements. After hours or days of measurements, the site surveyors would complete their measurements and know where to place the APs and at which channels and what power level to set them. The advent of RF planning tools has saved WLAN designers much time in the site survey process but the traditional method is still the most accurate and should be used for verification of the output from the RF planning tool. It shouldn't take too long to go back and make some measurements to verify what the planning tool has indicated. It will save time and money in rewiring a building if the APs need to be moved significantly from the location indicated by the RF planning tool if it is wrong or if the designer input faulty data into the planning tool.

#### *10.3.1.5 AP/User Ratio*

If the WLAN designer uses a good RF planning tool, the tool will “ask” for the number of users and a minimum baseline throughput, then calculate the number of APs accordingly. The SNL/CA facility is guided by the thought that there should be no more than 20-30 users to one AP. An over subscription ratio might also be factored in to account for the fact that not all the users will be using the WLAN at the same time. At Sandia, users are accustomed to a high level of service on the wired network; therefore the WLAN should be designed with a lower user-to-AP ratio—probably somewhere between 10 to 20 users per AP.

### **10.3.2 New Mexico**

The SNL/NM Site Survey methodology is much like that used at SNL/CA, with one major difference. SNL/NM is employing a “dense AP” or grid technique, whereby a fairly large number of APs are deployed in an area to ensure coverage without an extensive site survey or fine tuning of AP placement.

Prior to deployment, SNL/NM employs the same first two steps as SNL/CA: WLAN Customer Interviews and Visual Survey. The interviews and visual surveys are essential

to ensure that customer expectations are met, and that the physical aspects of the installation are understood before attempting a design.

The RF planning stage is where SNL/NM diverges from the SNL/CA technique. SNL/NM uses equipment from Aruba Networks. Aruba's technique is to lay out APs at some relatively small interval such as 50 ft, forming a grid over the building. This dense deployment does not take into account any localized concerns like particular walls or other RF obstructions. Using a diagram of a building, the Aruba survey tool recommends locations for AP, and these locations can be adjusted to take into account the physical aspects of the building. Once the APs are deployed, they are automatically tuned (power and channel) to take into account the localized conditions. Where an AP is not covering well, its power might be increased. Where an AP is covering well, its power might be reduced. The same goes for channel assignments. Once the building is operational, the system is checked using wireless clients and a measurement tool such as AirMagnet.

This makes for a very simple deployment method compared to the traditional painstaking AP placement and measurement technique, or the automated technique that takes into account building materials and obstructions, or the absolutely efficient placement of APs. There is a downside to the grid technique: many APs. The design is not engineered for greatest efficiency (in terms of AP placement and coverage) or minimum equipment cost. However, this approach minimizes the likelihood of a redesign or having to add more APs to accommodate growth. Additionally, the money saved on the detailed site survey covers a significant portion of the cost of the extra access points.

Figure 18 shows a snapshot of a plan for a single-story building. The east end of the building does not require coverage; therefore there are no APs in this area.

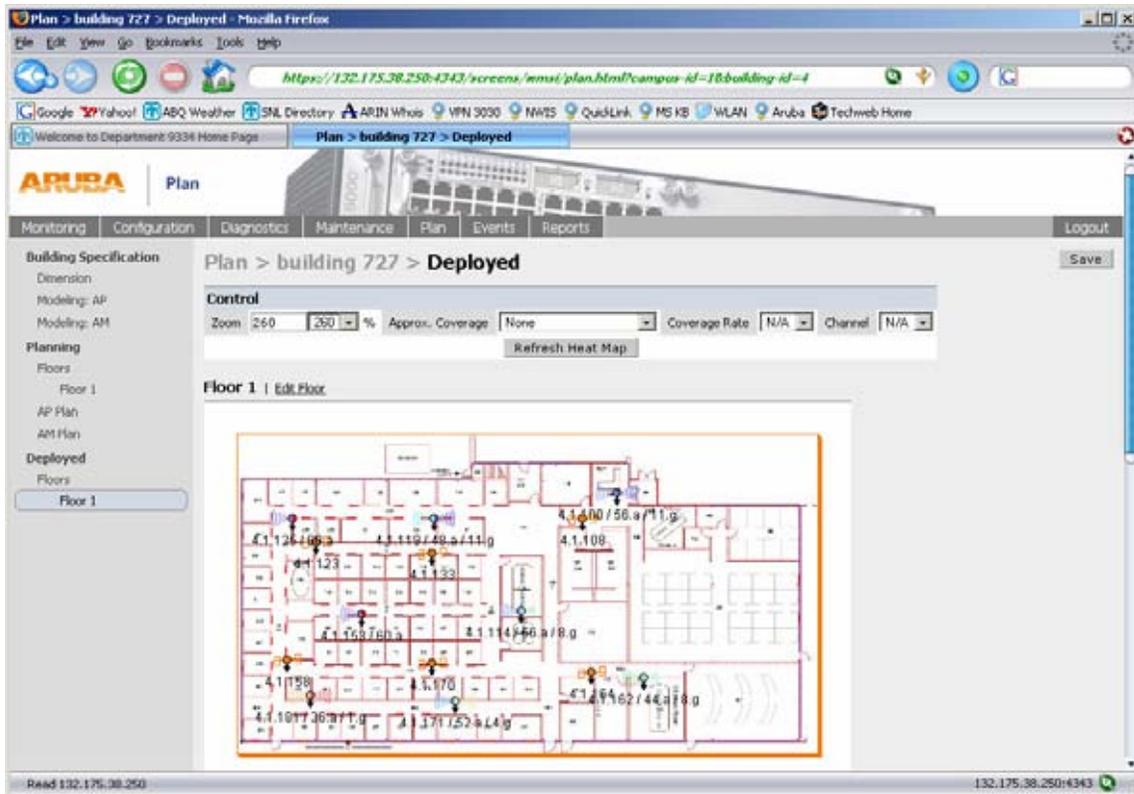


Figure 18. Aruba Site Survey Tool.

## 10.4 Wireless Network Design

Wireless system design activities begin with an assessment of a customer request for services. Preliminary research, such as a visual inspection of the area, and conceptual system data are considered. A list of questions to determine customer needs and requirements is compiled. A project meeting or customer interview is scheduled to obtain and document responses to the above mentioned questions. Data gathered at this meeting will include but not be limited to the following:

- ✓ Operational results required and expected;
- ✓ System performance expectations;
- ✓ Vendor equipment preferences (if any);
- ✓ Project Timeline requirements;
- ✓ Budgetary issues and concerns;
- ✓ Documentation and process expectations.

Upon completion of the customer interview, response data is evaluated and the foundational engineering document (Customer Requirements) is prepared and submitted to the customer for review and solicitation of any change suggestions.

After completion of the Customer Requirements document, conceptual design work begins. Technology, vendor platform, and performance testing and network architecture data and considerations are studied and analyzed. Additional APs are included, to be placed in monitoring mode and serve as detection points for rogue access points and other

wireless intruders. A conceptual design is then prepared and submitted for peer review. During the peer review, an assessment of design concept viability is made and peer suggestions and input are documented.

Next, production design work begins. The production network design is prepared addressing all system deployment, network integration and architectural considerations and concerns adhering to ISO processes and “engineering best practices.” A final production design review is conducted and peer endorsement is obtained. At this point the design documents are placed and posted in appropriate document resource data bases and delivered to tier2 or network infrastructure contractors for execution of the implementation activities.

### ***10.5 Installation Issues***

Several issues related to installation of the access points needed to be resolved during this project. One dealt with whether the APs should be visible (below the ceiling) or hidden (above the ceiling). Another dealt with power to the APs. To maintain flexibility in these decisions the team requested that the APs be Plenum-rated and that they support the IEEE Power over Ethernet specification.

Hiding the APs may make installation easier and reduces the possibility of theft of the APs. Making the APs visible allows a user to easily determine that wireless coverage might be available in a given area. Seeing the APs also provide a visual cue to people regarding the usage of wireless transmissions and the security precautions they may need to take. The project personnel decided to mount the APs below the ceiling, where they are visible.

Power could be delivered to the APs over the network cable (PoE, or Power over Ethernet, the IEEE 802.3af specification) or through an externally connected power supply. If the APs were installed above the ceiling, it would be easy to connect an external power supply to them, if they were close enough to a power outlet. If not additional power cabling would need to be installed. If the APs were installed below the ceiling, powering them with external power cables may prove very awkward, and PoE would likely be much more convenient. The project personnel decided to use PoE. If for any reason the communication switch to which the AP was connected could not supply the required power, an external power supply could be plugged into an electrical outlet and power could be put onto the network cable via a power injector.

## **11 Conclusions**

Sandia National Laboratories can in some ways be considered an early adopter of wireless LAN technology. Although Sandia was not among the very first large enterprises in industry to use wireless LANs, they were the first Nuclear Weapons Complex (NWC) lab or plant to deploy production wireless infrastructure in exclusion areas. Sandia was also an early adopter of the thin AP type solutions (sections 2.3-2.5).

Sandia developed a staff of well-trained wireless networking personnel over the course of this project, whose skills and knowledge were assembled from a combination of formal training and project experience. Lessons were learned along the way and a path forward into the future was charted.

### **11.1 Lessons Learned**

Specific lessons learned on this project include:

- Keep communication lines open with all interested parties, but especially those in the security fields (e.g. cyber security, local NNSA office, etc.). Our frequent communications with various security personnel smoothed the path towards wireless LAN deployment in the exclusion areas.
- Keep the users informed at all times regarding security rules during the piloting and deployment of new technologies. Sometimes rules change or vary somewhat between the pilot and production phases, base on information, experience, or knowledge gained from the pilot. To make the project a success, we had to help users steer clear of security infractions.
- Continue to work towards a sustainable, long term process for the deployment of new technologies. On this project the development team focused heavily on the technology, while not putting enough emphasis on production operation procedures.
- Work closely with other laboratories and plants. Sharing information and experience can maximize efforts.

### **11.2 Future Directions**

Future directions for this project, along with continuing to expand the production wireless LAN, should include monitoring technology advances and evaluating their potential for improving wireless networking at Sandia. Two specific technologies, 802.11i and Wi-Max, are covered in more detail later in this section. 802.11i should be deployed at Sandia in accordance with DoD and NIST recommendations. Wi-Max technology should be monitored as it matures, for possible use in certain applications at Sandia, as noted in section 11.2.2.1.

Other technologies such as wireless USB and ultra wideband technology should be monitored for their possible application to efforts at Sandia. Higher speed standards,

such as 802.11n, should be tracked as possible paths forward to higher throughput in Sandia's wireless LAN. Wireless Wide Area Network (WWAN) technology, such as the cellular telephone data technologies (e.g. EV-DO, EDGE), should be monitored for application to, and integration with Sandia's wireless LAN infrastructure. Wireless voice over IP (VoIP) should be investigated for potential use at Sandia, particularly in exclusion areas where there is WLAN coverage.

### **11.2.1 IEEE 802.11i<sup>7</sup>**

IEEE 802.11i is a standard approved by the Institute of Electrical and Electronics Engineers in the 802.11 series, concerning wireless LANs. Amendment 6, 802.11i, deals with medium access control (MAC) security enhancements. The IEEE 802.11i standard specifies a new generation of security measures for wireless LANs, to remedy the security shortcomings of the earlier 802.11 standards.

#### *11.2.1.1 Current Conditions at SNL*

Sandia National Laboratories instituted wireless networking at the New Mexico and California sites over the past two years. Due to the limitations of wireless security standards at the time of deployment, the system was designed to use non-authenticated wireless access, with authentication control and data privacy applied through use of a layered VPN, which supplied two-factor authentication, data encryption, and data integrity.

While this was an acceptable security solution, it was far from ideal. Because there is no authentication for the wireless system itself, it is possible for non-authorized persons to connect to the system. For this reason, the system was built on an isolated, open infrastructure having no direct access to the internal network. Access to the internal network is possible only through the authenticated VPN.

Ideally, the wireless system itself should require mutual authentication between the user and the wireless system, and provide strong encryption, rather than having to rely on a layered VPN that provides only one-way authentication, impacts performance, and complicates the system. With the advent of IEEE 802.11i, this is possible. With 802.11i, a wireless system can utilize a number of strong authentication protocols based around the EAP protocol. Data privacy is provided either by the strong encryption standard, AES, or a backward compatible, improved version of WEP called TKIP. These elements provide a system that is strong yet usable.

Additionally, 802.11i is compatible with smart cards, which are being adopted by Sandia for user authentication in place of reusable passwords. With the user already authenticated at time of machine login using a smart card, attachment to the wireless system can be accomplished using the same credential (digital certificate) without additional user interaction.

---

<sup>7</sup> Major portions of this section were excerpted from the whitepaper, "A Recommendation for the Use of 802.11i with Sandia National Laboratories' Enterprise Wireless Network" [15]

### 11.2.1.2 802.11i Evaluation

802.11i defines two new types of wireless networks called Robust Security Network (RSN) and Transitional Security Network (TSN). A true RSN allows only RSN-capable devices to connect and requires devices to support a number of new functionalities over the old Wired Equivalent Privacy (WEP) standard. However, because many people will want to upgrade hardware over a period of time and use pre-RSN equipment during the upgrade, the 802.11i standard also defines a TSN in which both RSN and WEP systems can operate in parallel.

When the security flaws with WEP began to surface, the major Wi-Fi manufacturers decided that security was so important to users that it had to move as fast as possible to create a replacement for WEP. Furthermore, they realized customers would not replace all their existing Wi-Fi equipment in order to switch to RSN. The Wi-Fi Alliance adopted a new security standard called Wi-Fi Protected Access (WPA). This standard is a subset of the RSN draft, which only includes parts of RSN that legacy WEP hardware could support with a simple software upgrade. Once the 802.11i standard was accepted, the Wi-Fi Alliance implemented a WPA2 certification process to ensure all WPA2 certified Wi-Fi equipment is 802.11i compliant<sup>8</sup>.

Confidentiality, integrity, mutual authentication, and availability are important issues for wireless LAN security. The integrity of all messages must be ensured, but we can split the problems of confidentiality and mutual authentication into two separate sections. We will deal with authentication in a wireless network first.

### 11.2.1.3 Authentication Using 802.1X and EAP

Two authentication mechanisms are possible in the 802.11i specification. In one mechanism, the possession of a Pre-Shared Key (PSK) authenticates the peers. A 128-bit encryption key and another distinct 64-bit Message Integrity Code (MIC) can be derived from the PSK. This mechanism, using a static key distribution, is intended for smaller organizations, thereby incurring a high management overhead cost for larger organizations.

Authentication and key management can also be handled by a combination of 802.1X and EAP (Extensible Authentication Protocol). IEEE 802.1X is a port-based network access control framework. Initial 802.1X communications begin with an unauthenticated supplicant (i.e., client device) attempting to connect with an authenticator (i.e., 802.11 access point). The access point responds by enabling a port for passing only EAP packets from the client to an authentication server located on the wired side of the access point. The access point blocks all other traffic, such as HTTP, DHCP, and POP3 packets, until the access point can verify the client's identity using an authentication server (e.g., RADIUS). Once authenticated, the access point opens the client's port for other types of

---

<sup>8</sup> Note that in this document, WPA2 is used interchangeably with 802.11i. There are in fact some minor differences between WPA2 and 802.11i, but those differences do not come into play in the context of this report.

traffic. Mutual authentication of both the access point and the client is important because man in the middle (MITM) attacks are possible when only one party in the communication is authenticated.

The Wi-Fi Alliance's WPA standard defined two authentication mechanisms. The first, WPA-Personal, includes WPA-PSK (pre-shared key). This is a simpler version that does not support 802.1X or a separate RADIUS server for mutual authentication. The second WPA authentication mechanism, WPA-Enterprise, includes all of the features of WPA-PSK plus support for 802.1X RADIUS authentication and is appropriate in those cases where a RADIUS server is deployed.

#### 11.2.1.3.1 EAP

EAP defines a set of messages that are used to make the introductions between the supplicant and authenticator and to allow mutual authentication. These messages are used with a number of upper-layer authentication methods. The RFC defining EAP, which is only nine pages excluding references and acknowledgments, defines only four different messages: request, response, failure, and success. It is a lightweight protocol with the meat being the specific authentication methods. EAP also allows two parties to exchange information that is specific to the authentication method they want to use. The content of these authentication-specific methods is not defined in EAP. In fact, they can be completely proprietary authentication methods or newly invented ones. EAP's ability to handle part of the communication in a standardized way and part in a specific way is the key to its extensibility.

A number of RFCs have been written detailing how to use EAP with specific authentication methods. For example, there is an RFC specifying how to use Transport Layer Security (TLS) over EAP (EAP-TLS), another draft stating how to use Tunneled TLS (TTLS) over EAP (EAP-TTLS), one describing Protected EAP (EAP-PEAP), and many other RFCs detailing authentication methods.

Three TLS-based protocols have been developed for use with EAP and are suitable for deployments in wireless LANs. The TLS protocol, which is described in IETF RFC-2246, is based on the SSL 3.0 Protocol Specification as published by Netscape. Full TLS includes authentication and confidentiality; however RSN and Wi-Fi Protected Access (WPA) use only the authentication portion of TLS relying on a separate protocol for confidentiality. The TLS authentication method uses asymmetric cryptography, in the form of digital certificates, to authenticate both parties and fits well into the EAP/IEEE 802.1X framework.

#### 11.2.1.3.2 EAP-TLS

EAP-TLS uses a TLS handshake as the basis for authentication. It is well documented and has been analyzed quite extensively. Study of the protocol has not yet revealed significant weaknesses in the protocol itself. (Several implementations have suffered from bugs, however.)

TLS authenticates peers using digital certificates. In EAP-TLS certificates are used to authenticate in both directions. The server presents encrypted information along with a certificate to the client, and after proving the server's possession of the private key associated with the validated certificate, the client presents a similar packet to the server. Naturally, the certificate may be protected on the client by a pass phrase or PIN, or stored on a smartcard, depending on the implementation. One flaw in the EAP-TLS protocol, noted by many observers, is that the identity exchange proceeds in the clear before the exchange is encrypted, so a passive attack could easily observe user names.

Digital certificates are the Achilles heel of EAP-TLS. The use of certificate authentication of clients mandates a concurrent or previous PKI rollout. If you do not already have a PKI in place, the additional work involved in issuing and managing certificates is quite large. In comparison with other PKI-enabled protocols, EAP-TLS may impose a greater certificate management overhead because of the need to revoke certificates as users have wireless LAN access revoked.

The bottom line is that EAP-TLS is secure, but the requirement for client certificates is too large of a hurdle for most institutions. Fortunately, Sandia already has a PKI in place.

#### 11.2.1.3.3 EAP-TTLS and PEAP

Both EAP-TTLS and PEAP were developed in response to the PKI barrier in EAP-TLS. Client certificates are not ideal for user authentication for a variety of reasons. Other methods of user authentication are as secure as certificate-based authentication but without the high management overhead. Both EAP-TTLS and PEAP were designed to use established authentication mechanisms while retaining the strong cryptographic foundation of TLS.

The structure of EAP-TTLS and PEAP are quite similar. Both are two-stage protocols that establish security in stage one and then exchange authentication in stage two. Stage one of both protocols establishes a TLS tunnel and authenticates the authentication server to the client with a certificate. (EAP-TTLS and PEAP still use certificates to authenticate the wireless network to the user, but only a few certificates will be required, so it is much more manageable.) Once that secure channel has been established, client authentication credentials are exchanged in the second stage.

EAP-TTLS uses the TLS channel to exchange "attribute-value pairs" (AVPs), much like RADIUS. (In fact, the AVP encoding format is very similar to RADIUS.) The general encoding of information allows an EAP-TTLS server to validate AVPs against any type of authentication mechanism. EAP-TTLS implementations today support all methods defined by EAP, as well as several older methods (e.g. CHAP, PAP, MS-CHAP and MS-CHAPv2). EAP-TTLS can easily be extended to work with new protocols by defining new attributes to support new protocols.

PEAP uses the TLS channel to protect a second EAP exchange. Authentication must be performed using a protocol that is defined for use with EAP. In practice, the restriction to EAP methods is not a severe drawback because any "important" authentication protocol would be defined for use with EAP in short order so that PEAP could use it. So far Microsoft has been a big proponent of the PEAP standard, leaving EAP-TTLS support in Windows to third party vendors.

EAP-TTLS and PEAP are extremely similar, however there are some differences between the two that we outline below:

- Both EAP-TTLS and PEAP set up a TLS session for the first round of authentication.
- EAP-TTLS exchanges "attribute-value pairs" (typically with a RADIUS server) inside the TLS tunnel.
- PEAP performs a second EAP session inside the TLS tunnel.

#### 11.2.1.3.4 Conclusions on Authentication

Selection of an authentication method is the key decision in securing a wireless LAN deployment. The authentication method drives the choice of authentication server, which in turn drives the choice of client software. Fortunately, selecting an authentication method compatible with 802.11i is a reasonably straightforward endeavor. The static key distribution solution of PSK is not feasible unless the organization is extremely small. Unless you have a well-oiled PKI already deployed, bypass EAP-TLS to avoid the client certificate headaches. Though there is not a large technical difference between the EAP-TTLS and PEAP protocols, EAP-TTLS has a number of minor advantages. In addition to a moderate degree of flexibility at the protocol level, both open source and commercial EAP-TTLS products are now available supporting a much wider variety of client operating systems.

An important point to keep in mind is that man-in-the-middle attacks are possible, as pointed out in the EAP working group draft "The Compound Authentication Binding Problem", when one-way authenticated tunnels are used to protect communications of one or a sequence of authentication methods. For this reason mutual authentication must be done at every phase of the authentication process.

#### 11.2.1.4 Confidentiality

The 802.11i standard includes two link-layer encryption protocols. The first, Temporal Key Integrity Protocol (TKIP) was designed to work on pre-802.11i hardware. The

second, Counter Mode/CBC-MAC Protocol (CCMP) using Advanced Encryption Standard (AES), was designed from the ground up to offer the highest level of security possible.

#### 11.2.1.4.1 TKIP

Temporal Key Integrity Protocol (TKIP) exists for one reason: to allow legacy WEP systems to upgrade their security. This is the reason TKIP was created, and this requirement guided the design throughout the standardization process [10]. TKIP is based on the RC4 stream cipher -- the same as WEP. TKIP however provides important data encryption enhancements including a per-packet key mixing function, a message integrity check (MIC) named Michael, an extended initialization vector (IV) with sequencing rules, and a re-keying mechanism. Through these enhancements, TKIP addresses all of WEP's known vulnerabilities.

However, weakness is predestined due to the limitation of re-using legacy hardware. Although the TKIP key mixing function has stronger security than the WEP key scheduling algorithm, it is possible to find the MIC key given one of the per-packet keys and its corresponding packet. Furthermore, security is broken for the duration of a Temporal Key (TK), if two different per-packet keys are discovered having the same IV. Finally, the Michael algorithm provides only 20 bits of security, which means that an attacker can forge a packet after  $2^{19}$  attempts, on average.

#### 11.2.1.4.2 CCMP

To provide confidentiality CCMP uses AES in counter mode. For integrity, CCMP uses Cipher Block Chaining Message Authentication Code (CBC-MAC). This cipher provides the best security currently available in commercial products. In addition, CCMP protects some fields that were not previously encrypted. The additional parts of the IEEE 802.11 frame that get protected are known as additional authentication data (AAD). AAD includes the packet's source and destination addresses. This protects against attackers replaying packets to different destinations.

#### 11.2.1.4.3 Conclusions on Confidentiality

The vulnerabilities pointed out above do not mean that TKIP is insecure, only individual parts are weak. However, these weaknesses show that TKIP was built upon a weak foundation, and could lead to the discovery of vulnerabilities in TKIP. It is recommended that any organization that does not have legacy WEP hardware deployed or that requires a high level of security, should only allow AES-CCMP devices to connect to their RSNs. Transitional networks should deploy both TKIP and CCMP encryption with a plan to migrate all users to CCMP in the future.

### 11.2.1.5 802.11i Conclusions

Although Sandia National Laboratories' current wireless system is deemed safe, it does not protect against unauthorized associations and relies entirely on a layered VPN solution for protection, which impacts performance and adds complexity. With the advent of 802.11i, the wireless network can be protected to the highest degree current commercial technology allows without impacting performance, and sets the stage for potential removal of the layered VPN. Initial implementation of 802.11i will most certainly be hampered by client configuration issues and software problems, but the advantages far outweigh the challenges that will be encountered. With the removal of the overlay VPN (assuming eventual approval), there will be a single, integrated system rather than a wireless system and a layered VPN system.

### 11.2.2 Wi-Max

WiMax is a broadband wireless Metropolitan Area Network (MAN) technology that fits between wireless LANs (such as 802.11) and wireless WANs (such as those based on cellular networks). WiMax can serve as a backhaul for cellular systems and Wi-Fi hot spots, and as a "last 100 meter access" solution for residential and business broadband services. WiMax will face stiff competition in the United States from DSL and cable providers, but regions such as South America and parts of Asia will likely welcome wireless broadband [30].

WiMax carries IP (Internet Protocol) and provides fixed or portable (but not mobile) wireless solutions. (The distinction between portability and mobility has been discussed elsewhere [28].) It can be used at rates between 40 and 72 Mbps at ranges up to 5-15 miles. The point-to-point mode of WiMax may even be usable up to 50 miles. It can integrate into both 3<sup>rd</sup> generation (3G) mobile networks and wireline networks.

WiMax is defined by the IEEE 802.16d (also known as 802.16-2004) standard. This is based upon and improves on the 802.16 and 802.16a standards. An additional standard (802.16e) is in progress to support mobility, but the first WiMax deployments will be fixed point to fixed multi-point.

There are many frequency bands where WiMax could be implemented, but the three areas where initial WiMax implementations are focused are the 3.5 GHz licensed band (300 MHz of spectrum bandwidth from 3.3 to 3.6 GHz), the MMDS bands (licensed spectrum at 2.500-2.690 GHz and 2.700-2.900 GHz), and the unlicensed 5.8 GHz band (Upper Unlicensed National Information Infrastructure – U-NII – band, 5.725-5.850 GHz). Some cordless telephone systems use the 5.8 GHz Upper U-NII band so local interference problems may occur, as with 802.11b/g and the cordless phones that operate in the 2.4 GHz ISM (Industrial, Scientific, and Medical) band.

Security is built into WiMax with PKI certificate-based encryption. Ramana Mylavarapu points out [17] that there are many potential vulnerabilities in WiMax. Time will tell how many of these vulnerabilities are practical.

There are very few WiMax products on the market at this point. Most are still under development. It was expected that long-distance carriers like AT&T and MCI would use WiMax to bypass local service provider networks [14], but now that both of these companies (which are the two largest long-distance providers in the United States) are being acquired by local phone companies, their enthusiasm for WiMax may decrease.

#### *11.2.2.1 WiMax Recommendation*

At this point Sandia should not be in a hurry to adopt WiMax. Currently Sandia does not have a compelling need for the point-to-multipoint “last 100 meter” capabilities that WiMax provides. Sandia currently has wireless bridging technology in place (although WiMax might be able to provide greater bandwidth) where point-to-point backhaul solutions are needed. WiMax, being a new technology that has not yet had a “trial by fire”, will need some time for the user community to expose security holes and for the vendors to repair the flaws.

Sandia should monitor the development of WiMax. As the technology and standards mature, Sandia should re-evaluate possible uses of WiMax, especially for wireless bridging applications. When the standard for mobility has been completed and implemented, Sandia can revisit WiMax in light of their current technology deployments and needs, to determine if there would be any benefit to implementing mobile WiMax, possibly for outdoor use between buildings.

### **11.3 Summary**

This project team has deployed Wireless Local Area Networking at Sandia National Laboratories! The personnel on this project team, with the help of many colleagues and vendors, have learned a great deal about wireless networking. They experimented with, and evaluated many aspects of wireless LAN hardware and systems. They selected and procured the equipment and other materials for Sandia’s wireless LAN infrastructure. This team deployed pilots, both outside and inside of the exclusion areas. They developed security procedures, and started production deployment of wireless local area networking at Sandia. These people also developed a path forward for wireless networking at Sandia, into the future.

## Appendix A      Diffused Infrared Equipment

Early in the project we examined Diffused Infrared (DFIR) as a possible technology for use inside the limited areas, where RF communication technologies might not be permitted. Diffuse radiation is scattered out of the direct beam by refraction and reflection, as opposed to direct beam, which comes in a direct line from the source. DFIR fills an enclosed area, using the room's surfaces (e.g. walls, ceiling) and surfaces of objects in the room (e.g. tables, cabinets) to bounce IR signals between the transmitter and the receiver. (A further discussion of DFIR and a comparison to direct-beam IR can be found in [28].)

DFIR systems are not able to support the higher data rates that RF systems are able to support, but IR is much easier to contain than RF. Because IR signals do not penetrate walls, floors, etc., and are heavily attenuated by glass, DFIR could provide wireless data networking in confined or bounded indoor areas. It would be possible to inadvertently couple information (as covered in the books by Cooper [8] [9]) to an infrared access point, but the effects would be limited by the inability of the IR waves to penetrate the walls, floor, and ceiling, thus being confined to the room.

We found a company called Spectrix Corp. who claimed to have a system [19] that operates at 4 Mbps. It included PC cards for laptop computers (that fit into a Type II PCMCIA slot), IR "antennas" that serve as access points, and Wireless Hub Routers that serve as controllers and protocol gateways for the system. Each Wireless Hub Router can control eight or 16 antenna units and contained a 10/100Base-T uplink to connect into a wired Ethernet network. Spectrix uses a proprietary MAC and PHY layer. They employ a proprietary protocol between the PCMCIA card and the Wireless Hub Router, rather than the IR PHY layer of the 802.11 specification. Beyond the Wireless Hub Router, standard Ethernet is used. This system had a specified error rate of  $10^{-6}$  over a 1000 square foot area, but could vary with the environment.

To determine if this system would serve our needs adequately, we used OPNET and ACE to perform modeling and simulation of likely-use scenarios. Having previously determined there was strong interest in having the ability to process incoming and outgoing email while in meetings [28], we constructed scenarios involving reading email with no attachments, reading email with small (109 KB) attachments, reading email with large (2.1 MB) attachments, and sending email with no attachments from a conference room. We modeled these scenarios for 1, 2, 3, 5, 8, and 12 simultaneous users. Since the modeling software did not have a 4 Mbps IR module, we felt using the 802.11 2 Mbps module would be a conservative choice. The models showed that due to various buffer constraints and bottlenecks, the entire 2 Mbps bandwidth could not be used. The results are summarized in Table 12.

**Table 12. Data from OPNET Model Simulating 2 Mbps LAN.**

Number of users	Read email (no attachment) 0.091 Mbits		Read email (small attachment) 0.874 Mbits		Read email (large attachment) 18.288 Mbits		Send email (no attachment) 0.749 Mbits	
	Bandwidth used (Mbps)	Elapsed time (sec.)	Bandwidth used (Mbps)	Elapsed time (sec.)	Bandwidth used (Mbps)	Elapsed time (sec.)	Bandwidth used (Mbps)	Elapsed time (sec.)
1	0.01	16	0.10	18	0.25	71	0.08	17
2	0.02	17	0.19	18	0.70	103	0.08	26
3	0.03	17	0.28	18	1.5	54	0.12	26
5	0.05	18	0.48	17	1.6	143	0.24	26
8	0.08	18	0.60	26	1.6	180	0.58	26
12	0.12	19	0.80	55	not run	not run	0.90	27

Note that this presents the worst case, that is when all 5, 8, 12, or however many users each click on “read mail” at the same instant. Not everyone in a conference room will likely try to retrieve a message or send a message at exactly the same time. Total elapsed times under 30 seconds could be provided for 12 simultaneous users reading or sending emails with no attachments or 8 simultaneous users reading retrieving emails with small (103 KB) attachments. It would not be advisable to retrieve emails with large attachments (2.1 MB) in an area employing this system. Based on these results and usage constraints, we felt this would be an acceptable system for conference rooms in areas where RF transmitters are not permitted for security reasons.

We specified equipment to build a pilot DFIR wireless LAN covering several conference rooms, several office areas, and several laboratory areas. When we contacted Spectrix (the only known vendor of DFIR LAN equipment) for pricing and ordering information, we found that they had apparently gone out of business. We were not able to find other DFIR LAN equipment in production. Because of this setback, we increased efforts to permit RF wireless LANs in areas that are within 100 feet of classified processing.

## **Appendix B      Video Compression over Wireless Links**

Any review of applications for wireless should include high bandwidth applications. One such application of interest to Sandia is high quality video-over-IP (IP here refers to Internet Protocol). This class of applications attempts to transport a high resolution computer generated video (such as computer screen) made up of 1280x1024 pixels and up to 60 frames per second over an Ethernet network and reconstruct the images at the original frame rate at the remote site. The required bandwidth for this data stream is approximately 2.5 Gbps (Gigabit per second, or  $10^9$  bits per second).

Sandia has been actively developing hardware to deliver such high quality video for many years. Most recently, Sandia has developed hardware that has been successfully transferred to the commercial company, Logical Solutions, which is marketing the device under the name "Global Link." The Global Link system consists of two devices, a transmitter that captures the video from a computer's graphic card and sends UDP video packets over IP on a Gigabit Ethernet network. The receiver device takes these video packets and regenerates a facsimile of the original video, at the original frame rate. The Global Link was designed to operate optimally over either a direct connect (point-to-point) or switched Gigabit Ethernet network.

The Global Link was designed with considerations given to the IP network characteristics. For instance, video data is transported over Ethernet using RTP (Real Time Protocol) and UDP (User Datagram Protocol) and IP (Internet Protocol). RTP version 2 is defined in RFC 1889. Specific RTP formats are generally described in RFCs, however, no RFC has been submitted for consideration for the Global Link format as of this writing. RTP was chosen for this application because it is the standard Internet Protocol for transport of real-time data (including in this case, video data). RTP is carried in UDP packets.

UDP is a connectionless transport protocol, i.e. fire and forget. There is no guaranteed delivery of the data from the transmitter to the receiver. This results in specific network behavior that must be considered when constructing a protocol for any application using this transport mechanism, such as video-over-IP. Under conditions of congestion, network switches are programmed to drop UDP packets first, preserving as much of the TCP (Transmission Control Protocol) traffic as possible. The result is, for a protocol to work well with the network, each packet must be self-contained and the loss of any particular packet's data, due to it being dropped by the network, must not affect any other packet. Therefore each packet contains an address describing the location where the packet's data should be placed on the screen at the remote receiver. To understand the effect of the network on the delivery of video and subsequently the quality of the reconstructed screen, one must understand the format of the video being transported.

## Transport of Video Packets

The key to a successful video-over-IP protocol is the way a screen is subdivided and the data is enclosed in the Ethernet packet. The screen, of course, is made up of pixels, each pixel described by three color values (Red, Green and Blue, i.e. RGB). Each color is described by 8 bits resulting in 24 bits per pixel. In this context, the Global Link system divides the screen into blocks of data called *segments*. A segment is defined as a linear group of pixels on a single horizontal line. In the 1280x1024 format, for instance, a segment is one fourth of a horizontal line or 320 pixels. Each segment consists of a total of 960 bytes. There are 4096 pixels on the screen, numbered 0 to 4095. Segments are ordered starting in the upper left hand corner of the screen (segment 0) and proceeding left to right and top to bottom to the lower right corner (segment 4095). The size of the segment is chosen to maximize the data in the Ethernet packet, thus improving the network efficiency, while keeping the data in each packet the same size, thus making it easier to pull apart the incoming packet and store the data in the receiver's video frame. The size of the segment is dependent on the screen size and aspect ratio. Below is a table (Table 13) of proposed segment sizes for several standard screen formats.

**Table 13. Frame Formats.**

Format	HOR x VERT	Segment Size (*)	Segments per Frame	Segments per Horiz Line	Bytes per Segment
VGA-60	640x480	160	960	2	960
SVGA-60	800x600	200	1200	2	1200
XGA-60	1024x768	128	3072	4	768
SXGA-60	1280x1024	160	4096	4	960
UXGA-60	1600x1200	160	6000	5	960
QXGA-60	2048x1536	128	12288	8	768
QSXGA-60	2560x2048	160	16384	8	960
1920-60	1920x1024	160	6144	6	960
1920A-60	1920x1200	160	7200	6	960
2000-60	2000x2000	200	10000	5	1200

\*The segment size is specified in double pixels (dp) which are represented by 48 bits (6 bytes). A pixel is represented by 24 bits (3 bytes), in the normal red-green-blue format.

## Data Packet Format

The data packet is enclosed in a specialized version of the standard RTP format. The RTP fields are defined as follows:

Version (VER) = 2

Marker bit (M) = marker bit is set to one for the last segment of a video frame

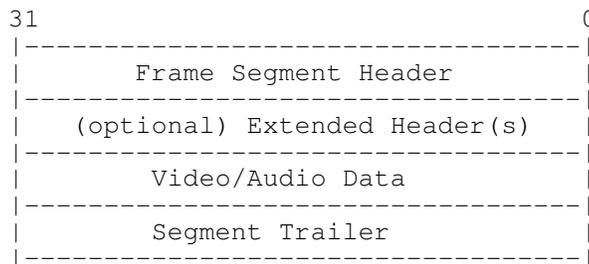
Payload Type (PT) = Defined according to RFC 1890. Suggested PT = 24 with a clock rate of 90000 Hz.

Sequence Number (SEQ) = A 16 bit count that is incremented with each RTP packet.

Timestamp (TS) = 32 bit 90 kHz timestamp representing the time of the first segment of each frame. All RTP packets of the same frame have the same timestamp.

Synchronous Source ID (SID) = 16 bit value is the source ID of the video stream.

The RTP payload (Figure 19) consists of a Frame Segment Header (FSH), Extended Header (EH), Video Data (VD) and a Segment Trailer (ST). The FSH contains information such as the header type, screen format, frame count and segment count. When a packet is received, the segment number is used to insert the segment data in the video buffer.



**Figure 19. RTP Payload Description.**

## **Updating the Screen**

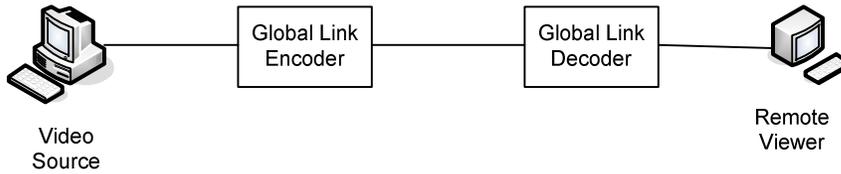
Two triggers are used to determine when to send a segment of data to the receiver. The first method sends a segment of data if a change is detected in any of the pixels in the segment. The result is that even a single bit change will result in an entire segment being transmitted. While this method is bandwidth wasteful, it was chosen because it results in more rapid insertion of the update data in the receiver. Several methods of compressing the data in the segment (only sending the modified data) have been proposed, but are outside the scope of this paper. A second, independent trigger is used called the seed process. Seeding is used to update the screen during startup and to renew an entire screen after extended network outages. The seed process uses pseudo random number generator to send a segment of data regardless of whether a change was detected. The most commonly used parameters results in a complete screen update approximately every 2 seconds. In this way, every pixel on the screen is updated at least once every two seconds, constantly correcting any pixels lost in transmission.

## **Configuration**

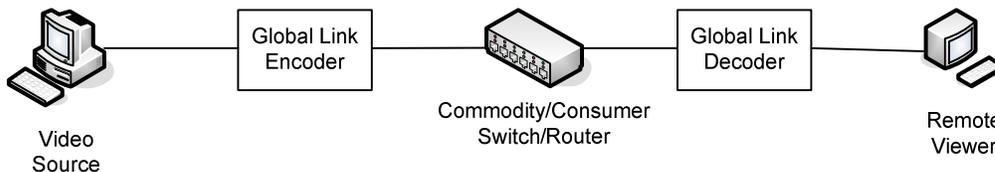
Three network configurations were used to conduct the tests. Also specific Global Link hardware and software configurations were used in the testing. These configurations are described in the following sections.

## Network Configurations

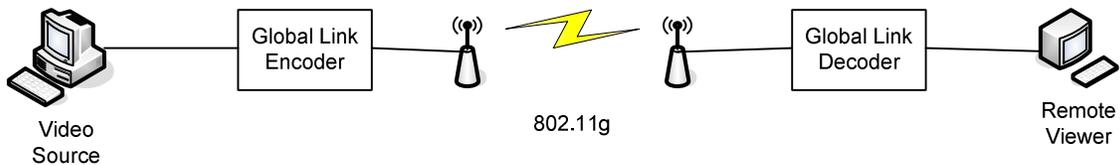
The three network configurations used for the video testing were direct connection, connection through a commodity/consumer-grade router/switch, and connection over a wireless bridge. These are shown in Figures 20, 21, and 22.



**Figure 20. Direct Connection Network Configuration.**



**Figure 21. Consumer-Grade Router/Switch Network Configuration.**



**Figure 22. Wireless Bridge Network Configuration.**

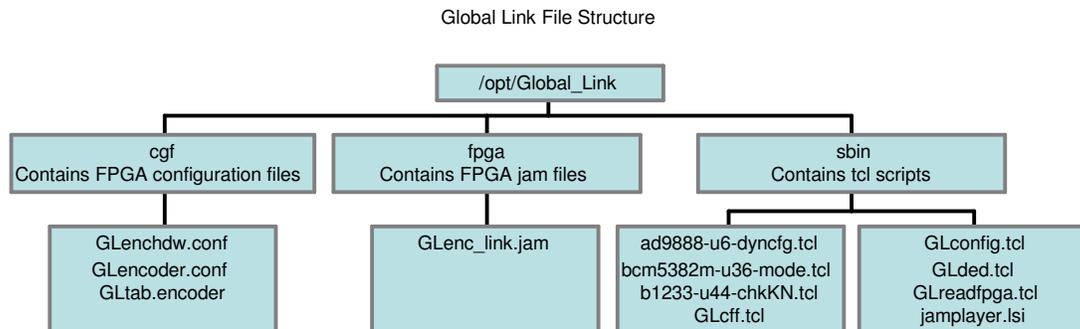
## Hardware Configuration of the Global Link

The Global Link Encoder and Decoder consist of two separate boards, one for video processing and one with a PC running an embedded Red Hat Linux. The video board contains the input and output video sections, an Altera 1S30 Stratix FPGA chip for actual video processing and SDRAM for video frame storage. The FPGA also connects to Ethernet through a separate MAC chip. The Linux connects to a 10/100 port on the embedded Ethernet switch chip and thus has access to the single 1 GigaBit Ethernet (GBE) network connection at the rear of the box. Linux can also access a JTAG port on the video board, which is used to push a JAM file and configure the FPGA and an I2C port, which is used to load the configuration parameters into the FPGA.

## Software Configuration of the Global Link

The embedded Linux processor serves three functions. The first is to download the FPGA configuration JAM file into the FPGA on power-up over the JTAG port. The second is to download the configuration parameters into the FPGA registers over the I2C port. The third is to make a connection between the decoder and the encoder over which it sends keyboard and mouse commands and data.

The configuration files are found in the following tree structure, shown in Figure 23.



**Figure 23. File Structure for Global Link Configuration Files.**

The files shown here are only the ones related to programming the FPGA and configuring the registers in the FPGA. These files are over ridden by a similarly named file located in the /etc/opt/Global\_Link tree. If a file is found in the /etc/opt tree, then this file is used. If not, the file in the /opt tree is the default.

## Configuration Registers

For this set of experiments, the `/etc/opt/Global_Link/cfg/GLenchdw.conf` file is used to change the packet rate of the encoder. An example of this file is shown here.

```
#####  
###  
# Logical Solutions Inc  
# Project:           Global Link  
# Name:             GLenchdw.conf  
# Description:      Encoder Hardware Config File  
#                  Hardware Configuration for the Video  
section  
# Location:         Encoder  
# Date:            2004-09-20-1114  
# History:  
#####  
###  
  
# device for the KM host emulation (physical KM)  
KMhstSP=/dev/ttyS4  
  
# device for the KM device emulation (cpu KM)  
KMdevSP=/dev/ttyS5  
  
# device for the serial to i2c converter  
BL233SP=/dev/ttyS6  
  
# port for the parallel port  
JamIO=0x20c0  
  
# local unit's mac address  
SouMAC=$MACPROM$  
  
# interpacket gap / control  
IpgIP=255.1.1.1  
  
# FPGA configuration file  
FPGAfile=GLenc_link.jam
```

A number of encoder parameters are controlled by the four bytes in the Inter Packet Gap, IpgIP. As delivered from the factory, the IpgIP = 255.1.1.1. This parameter represents a 32 bit value in four octet IP address format. Understanding the operation of these bits takes some explanation.

First, the IpgIP consists of the following fields.

Bits	31..24	23..16	15..8	7..0
IpgIP	Control Register	Reserved	Interpacket Gap Register	Reserved

The Packetizer State Machine is controlled by the contents of the high byte of the IpgIP register, called the *Control Register*. The bit values are defined below. This register can be used to reset either the transmit or receive state machines in the IP\_PACKETIZER block of the FPGA HDL, halt video traffic and adjust the seed rate or even shut off seeding. Seeding is the sending of all screen segments regardless of whether a change has been detected. When confronted with a network that is highly lossy, seeding is essential to repair any damage to the screen image due to packet loss. The bits of the Control Register is shown below.

CTL is the upper 8 bits of the IPG. So, the ENCODER has the following definitions:

```
CTL[7] = undefined
CTL[6] = rx_reset
CTL[5] = tx_reset
CTL[4] = state_machine_reset_n
CTL[3] = undefined
CTL[2] = videodataon
CTL[1] = seed_off_n
CTL[0] = seed_slow_n
```

CTL[0] controls seed\_slow\_n. When low seed rate is slower. Changes the modulus in the counter from 31 to 127.

CTL[1] controls seed\_off\_n. When low seeding is off.

CTL[2] controls the video output. When low, the video is not transmitted. This is useful to shut off the video back stream from the decoder.

So as an example, to turn off video data, change the IPG value from 255.1.1.1 to 251.1.1.1. The Video data control bit is bit 2 of the IPG control register. To activate the new IpgIP value type in GLconfig -c C and hit return.

The packet rate transmitted by the Encoder must meet the Gigabit Ethernet specifications [20]. For 802.3z (found in IEEE 802.3-2002 specification) the time between Ethernet Frames at Gigabit Ethernet is specified to be greater than 96 nanoseconds. The third byte of the IpgIP register is the Interpacket Gap Register. This 8 bit value controls how many clock cycles of the 100 MHz clock are used as a delay between the sending of IP packets to the MAC chip. There is a minimum value which is represented by the '1' and meets GBE specifications. Numbers higher lead to longer delays and ultimately lower bit rates. A '10' is approximately 100 Mbps while a '100' is about 10 Mbps. Values larger than '100' have little more effect on the bit rate.

The delay counter is a 32-bit counter running at the 33 MHz clock rate (a period of 30.3 ns). The counter starts when the destination Ethernet address is pushed into the MAC fifo and continues as the rest of the packet is pushed into the MAC chip. That is, the counter runs all during the packet transfer and continues after it has been completed. After the entire Ethernet packet is completed, the transmit state machine waits in a delay state until the counter reaches the value defined by:

$0x0000 \& \text{IpgIP}[15..8] \& 0xFF,$

where ‘&’ is used to indicate bit concatenation.

At this point, the state machine continues and can send another packet. The minimum time that it can take to send a packet is  $0x000001FF = 511$  counts, which takes  $15.5 \mu\text{s}$  ( $511 \times 30.3 \text{ ns}$ ). During that period of time, the state machine will have transmitted  $1025 \times 8$  bits. The packet rate is then calculated to be

$$\text{DataRate} = (1025 \text{ bytes} \times 8 \text{ bits / byte}) / (15.5 \text{ E} - 6 \text{ sec}) = 530 \text{ Mbps} .$$

In the same manner, an IpgIP[15..8] setting of 0x0A yields a data rate of 96.1 Mbps and a setting of 0x64 results in 10.5 Mbps. The maximum setting, 0xFF, results in 65535 counts at 30.3 ns per count, setting the slowest data rate at 4.1 Mbps.

On both the Encoder and Decoder you can read back the IPG and differencing by using the GLconfig -c R.

Once the file has been changed, updates are pushed into the FPGA registers using the “GLconfig -c C” command. This command updates everything except the network and the JAM file. Other options are listed here.

```
- GLconfig
  -v :displays the version.
  -c C :configures all but network and jam
  -c S :configures network only
  -c J :load jam file only
  -c A :configures the ad9888 only
  -c F :configures the FPGA parmeters only
  -c R :shows the configuration including the added
```

The Encoder and Decoder were configured with the following IP Addresses.

Device Type	Linux	Video
Encoder	10.9.8.7	10.9.8.5
Decoder	10.9.8.8	10.9.8.6

## ***Network Behavior***

The quality of the video at the receiver varies depending on a number of factors including; the quality of the network, network congestion, network speed and network architecture.

The design of the Global Link is intended to insure frame integrity at the receiver. Each changed segment of a complete frame is sent to the receiver before the next frame is started. If the network is slow enough, and if the number of segments that need to be transmitted is large enough, several frames might be skipped during the time it takes to transmit the changes. In this way, video compression takes place by dropping frames. This insures that each frame appears at the receiver exactly as it appeared in the original video. This is very important for a large variety of applications.

It is obvious is that the higher the bandwidth available the higher the quality of the reconstructed video. But there are a host of network conditions, which result in artifacts on the screen.

When packets are dropped, some updates do not make it, obviously, to the receiver video frame buffer. Therefore, there are segments of the screen that are not updated properly and will display stale pixels. Since the segments are only one pixel tall, often the data is just a really small change and not noticed. However, if the network congestion is bursty and large sections of the screen have been changed from one frame to the next, there might be linear groups of segments that will be in error. Assuming that the congestion is only temporary, the random seeding process will update these segments over the next two seconds.

## **Description of the Test Applications**

The purpose of these studies is to determine the behavior of the Global Link video system operating over several different real-world network architectures including high-speed wireless infrastructure.

For each architecture under evaluation, especially stressful full motion video was used to determine how well the link handled the video IP traffic. The sample of full motion video is a 15 fps sample, 320x240 pixels in size, captured from a live NTSC broadcast. This sample was converted into a Quicktime 7 format and played, full screen, on the computer. The video content was chosen because it represents a high frame rate and a large percentage of the pixels on the screen are changing every frame thus providing the greatest stress on the network.

## **Video Quality with Wired Ethernet**

Under normal operation, the encoder and decoder are connected, back-to-back, using a single twisted pair Ethernet connection. Whether connected via a single cable or through

a gigabit Ethernet switch, the encoder transmits data at around 450 Mbps. The encoder can be configured to send packets at slower rates such as 100 Mbps and 10 Mbps used in tests 1, 2 and 3. In the back-to-back configuration at GBE rates (Table 14, entry 1), the video is nearly as good as that viewed on the monitor directly connected to the source computer. The explosions in this segment of the video are quite good.

When the encoder transmission rate was reduced to 100 Mbps (Table 14, entry 2), the video update rate begins to suffer. The full motion sample suffered the most as the screen began to show signs of slow scanning updates. The static screen example was still able to keep up with moving windows and would be quite acceptable for text editing applications or even web surfing. In the picture, you can see the transition between the new frame (top) and the old frame (bottom). In between, there are a large number of original video frames that are missing and were never sent.

When the encoder transmission rate was reduced to 10 Mbps (Table 14, entry 3), the video resembled slow-scan video, with associated frame rate reduction down to a few frames per second. This rate was unacceptably slow for viewing full motion video. However, each frame does appear on the screen without tears and in full high resolution. The static sample was usable, however it would be unacceptably slow when larger windows were being updated. In the picture, we again see the frame transition as a horizontal line.

The encoder and decoder were also interconnected through a 100 Mbps Ethernet switch (Table 14, entries 4, 5 and 6). The behavior through the switch depended on the selected transmission rate of the encoder. At all transmission rates it was evident that the Ethernet switch was having problems with the high flow rate of packets. For this application, the packets are metered to appear at a constant rate. It would appear that the Ethernet switch was designed to handle more bursty data and had a very hard time with the constant flow of packets. At the highest rates, the encoder swamped the Ethernet switch; the result being a large number of lost packets. The screen then suffered from many segments that were not updated. When the data rate was reduced, so that the switch could handle the incoming packets without loss, the screen again resembled slow scan video. In this case the frame update rate is reduced significantly. Even at the lower data rate there were a considerable number of missing packets. In (Table 14, entry 5) the screen size is reduced, however, the image is still not very sharp.

Once the data rate is reduced to 10 Mbps (Table 14, entry 6) the full screen can be updated, a frame at a time, with no loss of data.

### **Video Quality with Wireless Ethernet**

The encoder and decoder were interconnected via a wireless bridge and operated at 54 Mbps (Table 14, entries 7 – 11). The bridge devices were Netgear Model WGE101 54 Mbps Wireless Ethernet Bridge.

When we blast away at the wireless bridge (Table 14, entry 7), the wireless bridge that is connected to encoder begins to drop lots of packets. Keep in mind that the actual data rate achieved across the wireless link is significantly less than 54 Mbps signaling rate. We begin to see the four sections of the screen used for sectioning. The effect is viewing portions of two scenes and many multiple frames of each all at once. It would be impossible to use this link for anything productive.

When the encoder data rate is reduced to 100 Mbps (Table 14, entry 8) the situation improved slightly with less packet loss (since we are sending less packets in the first place). However the screen still looks messy.

Reducing the data rate to 10 Mbps (Table 14, entry 9) help a lot. Larger sections of the screen are from the same scene. However, since we are now down to 3 seconds to transmit a single screen, the update rate is not acceptable given that the screen is still broken. Note that the four segments are not as noticeable in the screen image.

When the video window is reduced to the native 230x240 size (Table 14, entry 10) and the data rate is kept at 10 Mbps, the background clears up (with transmission redundancy) and the video window frame rate improves significantly. Still the image is much cleaner.

Making the video window about one quarter smaller still (Table 14, entry 11) makes the image look quite nice. However we are approaching the point where the seed packet rate is approaching the update packet rate and further improvements are incremental.

Stopping the seeding (Table 14, entry 11) does not help because any lost packets in areas of the screen that are not directly affected by the video window might not get updated for quite some time. This would include other windows and the mouse cursor.

The results are summarized in the table below. The descriptions relate to the photos in the right most column.

**Table 14. Video Quality Comparison.**

	<b>Configuration</b>	<b>Description of the Quality of the Video</b>	<b>Photos</b>
		<i>Twisted Pair (TWP) Ethernet direct connect</i>	
1	~450 Mbps	High quality video with little evidence of frame drops. Not discernable from a directly connected monitor.	
2	100 Mbps	Video becomes noticeable slower, approximately 3 fps. Each updated frame is high quality, no lost packets.	
3	10 Mbps	Best described as slow scan video, 3 seconds per frame. Each updated frame is high quality, no lost packets.	
		<i>Connected via wired side of MR814V2 Wireless Router</i>	
4	~450 Mbps	This is only a 100 Mbps switch and many packets are being lost.	
5	100 Mbps	Much the same results, even when the video is only a small part of the screen, still have many lost packets.	

6	10 Mbps	No evidence of lost packets, screen looks great. However, video is slow scan, 3 seconds per frame like when using TWP interconnect.	
		<b><i>Connected via 802.11g Wireless Bridges</i></b>	
7	~450 Mbps	The wireless bridge connected to Encoder is dropping lots of packets at all speeds. Not suitable for viewing.	
8	100 Mbps	Still over-running the links ability to handle the data rate.	
9	10 Mbps	Believe it or not, this looks better, slower scan rate is leading to fewer dropped packets.	
10	10 Mbps, smaller changing video window	This looks a lot better, getting pretty good frame rate, still not totally clear with dropped frames.	
11	10 Mbps, even smaller video window	Now we are likely matching send rate and update rate, video looks good, but still some dropped packets.	

### ***Summary of Compressed Video over Wireless Link Testing***

The Global Link product was designed for use at Gigabit Ethernet rates and works best at these speeds. High frame rate video data is transmitted at nearly a constant rate due to the design of the hardware implemented IP stack. Ethernet switches are designed to handle packets in short burst. Generally, this is the case because some processor has to buffer up the data prior to transmission. The video encoder performs this task in hardware and so can transmit data at a flat rate. Subsequently, it is easy for the Global Link hardware to over-run the network hardware buffers (which as a rule, are not very large) and the network begins to drop packets. As the data rates are reduced, the screen updates begin to resemble slow scan video with frames being painted from top to bottom. Video packets are evenly spaced in time. If the packet rate is reduced to the point that the fifos do not fill, then the screen would be updated, segment-by-segment, from top to bottom. The frame rate will be greatly reduced, however the video frame will be transmitted in tact.

When connected through an 802.11g wireless bridge, the effective network data rate is reduced significantly. Packets are dropped even with the video bandwidth reduced below 10 Mbps. The video frame rate can be improved by making the video window smaller.

## Appendix C      Bibliography

1. Advanced Encryption Standard (AES), Federal Information Processing Standards Publication 197, National Institute of Standards and Technology, Gaithersburg, MD, November 26, 2001.
2. AirDefense, Inc., <http://www.airdefense.net>, 2006.
3. AirMagnet, Inc., <http://www.airmagnet.com>, 2006.
4. Aruba Networks, <http://www.arubanetworks.com/>, 2006.
5. Aruba 5000 Product Sheet, <http://www.arubanetworks.com/products/mobility-controllers/aruba-5000>, 2006.
6. Aruba 6000 Product Sheet, <http://www.arubanetworks.com/products/mobility-controllers/aruba-6000>, 2006.
7. Graham Celine and Charles Wright, Effective WLAN Testing Begins to Emerge, in *Wireless Systems Design*, pp. 23-26, January 2005.
8. James Arlin Cooper, *Computer and Communications Security*, McGraw-Hill, New York, 1989.
9. James Arlin Cooper, *Computer-Security Technology*, Lexington Books (D. C. Heath and Co.), Lexington, MA, 1984.
10. Jon Edney and William A. Arbaugh, *Real 802.11 Security*, Addison-Wesley, Boston, 2004.
11. Matthew S. Gast, *802.11 Wireless Networks*, O'Reilly & Associates, Sebastopol, CA, 2002.
12. Jim Geier, *Wireless LANs*, Sams Publishing, Indianapolis, IN, 2002.
13. Steve Gossage, "Network Architecture, Volume 4: Enterprise Campus Network Building Block," ver. 1.1, whitepaper, Sandia National Laboratories, September 30, 2003.
14. Elena Malykhina, Wireless Goes Faster, Farther, in *Information Week*, pp. 77-82, April 18, 2005.

15. Marc Miller, Jamie Van Randwyk, Damon McCoy, Mark Lodato, "A Recommendation for the Use of 802.11i with Sandia National Laboratories' Enterprise Wireless Network," white paper, Sandia National Laboratories, June 2005.
16. Marc Miller, Dallas Wiener, Edward Witzke, and John Long, "SNL/NM Wireless Pilot – Medical Facility, Building 831," white paper, Sandia National Laboratories, June 30, 2003.
17. Ramana Mylavarapu, Security Considerations for WiMax-Based Converged Network, in *RF Design*, pp. 20-26, August 2005.
18. Theodore S. Rappaport, *Wireless Communications*, 2<sup>nd</sup> ed., Prentice Hall, Upper Saddle River, NJ, 2002.
19. Spectrix Corporation, <http://www.spectrixcorp.com/products.html>, 2000.
20. TechFest Ethernet Technical Summary, <http://www.techfest.com/networking/lan/ethernet2.htm>, 1999.
21. Trapeze Mobility System Family, <http://www.trapezenetworks.com/en/products/index.asp>, 2006.
22. Trapeze Networks, <http://www.trapezenetworks.com/en/homepage.asp>, 2006.
23. Dallas Wiener, "Deployment and Use of Wireless LAN Technologies Within Limited Areas at Sandia National Laboratories," white paper, Sandia National Laboratories, March 24, 2005.
24. Dallas J. Wiener, *Mobility/Portability Technology Roadmap*, SAND2003-1471, Sandia National Laboratories, Albuquerque, NM, May 2003.
25. Dallas J. Wiener, *Technology Concept and Evaluation Report on the Wireless LAN Prototyping Efforts at Sandia*, SAND2004-0399, Sandia National Laboratories, Albuquerque, NM, February 2004.
26. "The WiMax Option for Delivering Converged Services," white paper, Lucent Technologies, 2004.
27. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Medium Access Control (MAC) Security Enhancements, IEEE 802.11i-2004, Institute of Electrical and Electronics Engineers, Piscataway, NJ, 2004.
28. Edward L. Witzke, *An Approach to Wireless Communications at Sandia National Laboratories*, SAND2002-3312, Sandia National Laboratories, Albuquerque, NM, October 2002.

29. Edward L. Witzke, "Unclassified Wireless Local Area Network (LAN) Limited Area Pilot at Sandia National Laboratories," white paper, Sandia National Laboratories, May 3, 2005.
30. Maury Wright, WiMax Wireless Broadband: Fixed-Flavor Questions Abound, Mobile Lurks, in *EDN*, pp 44-53, March 31, 2005.
31. George Wu, Make Way for WiMax Certified Products, in *Wireless Systems Design*, pp. 9-10, March 2005.

DISTRIBUTION:

1	MS 0453	M.R. Sjulín, 2120	1	MS 0806	B.R. Kellogg, 4336
1	MS 0630	K.E. Washington, 4600	1	MS 0806	L.G. Martínez, 4336
1	MS 0630	B.V. Hess, 4610	5	MS 0806	M.M. Miller, 4336
1	MS 0662	T. Klitsner, 4341	1	MS 0806	J.H. Naegle, 4336
1	MS 0662	J.K. Perich, 4343	1	MS 0806	U. Onunkwo, 4336
1	MS 0662	C.A. Quintana, 4537	1	MS 0806	J.A. Schutt, 4336
1	MS 0672	R.L. Hutchinson, 5616	1	MS 0806	L. Stans, 4336
2	MS 0672	L.G. Pierson, 5616	1	MS 0806	J.S. Wertz, 4336
1	MS 0788	M.J. Benson, 4334	1	MS 0806	D.J. Wiener, 4336
1	MS 0788	M.J. Hamill, 4334	5	MS 0806	E.L. Witzke, 4336
1	MS 0788	J.H. Maestas, 4334	1	MS 0806	T.J. Pratt, 4338
1	MS 0788	S.D. Olsen, 4334	1	MS 0806	T.C. Hu, 4338
1	MS 0788	M.A. Rios, 4334	1	MS 0806	T.D. Tarman, 5622
1	MS 0788	V.K. Williams, 4334	1	MS 0807	J.F. Mareda, 4537
1	MS 0788	D.B. Bateman, 4338	1	MS 0813	R.M. Cahoon, 4311
1	MS 0788	L.S. Chance, 4338	1	MS 0813	J.P. Abbott, 4312
1	MS 0788	P.L. Manke, 4338	1	MS 0813	J.P. Long, 4312
1	MS 0795	P.C.R. Jones, 4317	1	MS 0813	G.K. Rogers, 4312
1	MS 0795	D. Kilman, 4317	1	MS 0832	J.H. Dexter, 4335
1	MS 0795	A.A. Quintana, 4317	2	MS 0874	P.J. Robertson, 1711
1	MS 0799	J.A. Chavez, 4333	1	MS 1202	J.D. Tang, 5622
1	MS 0799	G.E. Connor, 4333	1	MS 1206	M.H. Johnson, 5625
1	MS 0799	M.J. Ernest, 4333	1	MS 1393	J.A. Larson, 12120
1	MS 0801	R.W. Leland, 4300	1	MS 9011	E.D. Thomas, 8965
1	MS 0801	D.S. Rarick, 4310	1	MS 9011	T.J. Toole, 8965
1	MS 0806	L.F. Tolendino, 4334	1	MS 9012	R.D. Gay, 8949
1	MS 0806	J.P. Brenkosh, 4336	1	MS 9012	M.L. Kahn, 8949
1	MS 0806	N. Dautenhahn, 4336	1	MS 9012	B.A. Maxwell, 8949
1	MS 0806	J.M. Eldridge, 4336	2	MS 9012	M.G. Mitchell, 8949
1	MS 0806	A. Ganti, 4336	1	MS 9151	L.M. Napolitano, 8900
1	MS 0806	S.A. Gossage, 4336	1	MS 9151	C.T. Oien, 8940
1	MS 0806	C.M. Keliiaa, 4336	1	MS 9152	S.J. Marburger, 2998
2	MS 9018	Central Technical Files, 8944			
2	MS 0899	Technical Library, 4536			