



White Paper  
**Intel Information Technology**  
Computer Manufacturing  
Wireless Network Management

## Managing and Monitoring a Primary Wireless Network

Intel IT designed a highly automated approach to managing an enterprise primary wireless LAN, with a focus on service level management and monitoring. Our design includes a controller-based architecture, centralized management, template-based configuration, and a redundant, self-healing network. We closely monitor network health to detect and fix problems before they affect users. Our analysis suggests that our approach will deliver improved service to users while greatly reducing support effort and cost.

Omer Ben-Shalom, Ashok Desai, Eva Donohoe, Omri Barkay, and Sylvia Stump,  
Intel Corporation

February 2007

IT@Intel

## Executive Summary

Intel IT designed a highly automated approach to managing an enterprise primary wireless LAN (WLAN), with a focus on service level management and monitoring.

WLANs present unique management challenges due to many complex, interrelated factors, including uneven service, wireless interference, and user roaming. Because of this, we focused on service level management, rather than simply uptime, to help ensure that users consistently experience good performance and reliable service.

We designed the network to be manageable and robust with:

- A controller-based architecture and a redundant, self-healing network
- Centralized management and template-based configuration to reduce overall management effort and cost
- Network health monitoring tools that provide early warning of problems so that we can fix them before they impact users

Our analysis suggests that our approach will deliver substantial benefits in reduced management effort, faster deployment, a more consistent environment, and a better user experience over the more distributed management methods we used with our legacy WLANs.

Our analysis suggests that our approach will deliver substantial benefits in reduced management effort, faster deployment, a more consistent environment, and a better user experience.

# Contents

**Executive Summary** ..... 2

**Background** ..... 4

    The WLAN Management Challenge ..... 4

**WLAN Management and Monitoring** ..... 5

    Monitoring Network Health ..... 6

    Analysis and Reporting ..... 6

**Conclusion** ..... 11

**Authors** ..... 11

**Acronyms** ..... 11

# Background

Wireless is becoming the preferred network access method among our mobile users. Our existing WLANs are popular and widely deployed, but we maintain them as separate networks alongside the wired LANs and currently consider them a secondary means of network access.

We are developing a new architecture that integrates wired and wireless LAN infrastructures and establishes high-performance wireless as the primary access method (diagrammed at a high level in Figure 1). We are beginning to deliver data, voice, and video wirelessly to mobile users on laptops, handsets, and other devices.

We have begun a major initiative to use primary WLANs based on our new architecture at a large Intel site that consists of five buildings with about 5,000 users. This project presents many

technical challenges because it breaks new ground, both as a large-scale primary wireless network and in the converged services it delivers.

## The WLAN Management Challenge

Managing an enterprise primary WLAN presents unique challenges that require a fresh approach to traditional network management concepts. A WLAN is a complex system with many interdependent factors that affect its behavior,

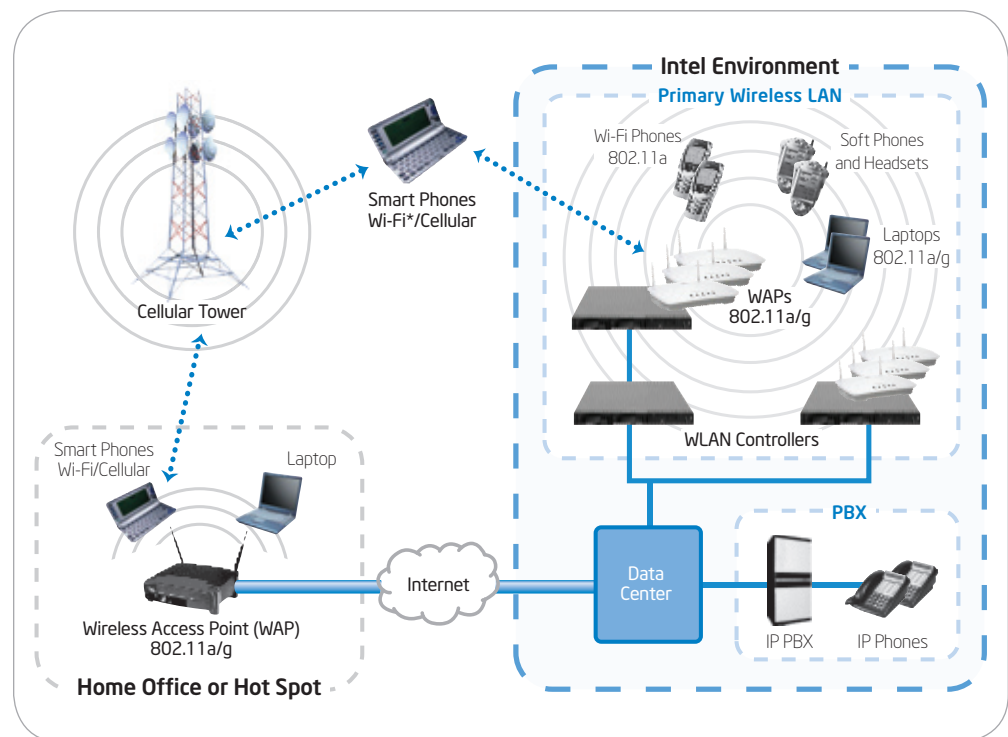


Figure 1. Our primary wireless architecture.

including traffic flows, network topologies, network protocols, hardware, and software.

- The service level may vary at different locations due to factors such as the uneven nature of the medium, a user's distance from an access point (AP), and the level of multi-path fading.
- Wireless interference has a profound impact on network performance. Because interference is highly variable and unpredictable, identifying and managing its impact remains a big challenge.
- System capacity is limited. With wired networks, it's relatively easy to provision extra capacity. With WLANs, provisioning new capacity is not straightforward because of the limited wireless spectrum and the effect of multiple wireless transmitters using the same spectrum. For example, increasing the density of APs does not necessarily increase capacity due to co-channel interference.

- Users can roam. With wired networks, users are stationary; with wireless networks, users can roam between locations. As they do so, they may experience performance degradation as they move away from an AP—or they may degrade the performance of other users when they connect to a different AP, increasing its load.

These challenges mean that managing WLANs requires a different focus than managing wired LANs. Because wired LANs provide consistent high performance, we are mostly aiming for maximum uptime with maximum mean time between failures (MTBF). With WLANs, we could achieve 100 percent uptime, yet still provide poor service to some users due to dynamically varying factors such as interference and roaming. Therefore, we need to focus on service level management: providing good coverage, keeping the network free from interference, and delivering adequate bandwidth to all users all the time. Our wireless LAN service level management model is shown in Figure 2 on page 7.

## WLAN Management and Monitoring

We designed our primary wireless network with overall manageability in mind. We reduce management effort by using an architecture based on WLAN controllers, by increasing availability through Layer 2 redundancy, and by centralizing management and monitoring.

With our architecture, we reduce the number of managed entities by managing AP controllers rather than individual APs. This approach also automates AP configuration, creates a self-healing network, and optimizes service levels. We expect that centralized maintenance with automatic upgrades, installation, and adjustments to the system will reduce overall management effort. We are evolving our network monitoring processes using tools that provide early warning of problems so that we can fix them before they impact users.

Centralized management and monitoring allow us to:

- Create user accounts with configurable permissions
- Fulfill scheduled upgrades and infrastructure configurations in a single job
- Implement consistent configurations across the enterprise by defining configuration templates and deploying them to controllers

Our analysis suggests that centralizing management and monitoring will deliver considerable savings in time and resources. We estimate that we may be able to cut the time and effort required for individual tasks by as much as 80 percent. Overall, we believe that the combination of system management, redundancy, and a self-healing architecture may cut the effort required to manage our enterprise WLAN by as much as 50 percent.

## Monitoring Network Health

Our approach places considerable emphasis on monitoring WLAN network health using a set of indicators. It includes a focus on WLAN performance, which can be affected by many factors, including interference.

We gather most of the monitoring information directly from the controllers, storing it in a data warehouse. We collect other data from clients, servers, and back-end devices such as Voice over IP (VoIP), authentication, and IP management servers. Our key parameters include interference, noise, signal-to-noise ratio, and controller and AP utilization levels. We extract and report the data using both commercial tools and tools that we developed ourselves.

We collect several categories of monitoring data in order to manage different aspects of the WLAN environment.

### Radio Frequency Management

Management of the radio frequency (RF) environment focuses on finding a balance between WLAN capacity and fidelity. We aim to reduce co-channel interference to provide optimum capacity and high performance. Data in this category includes:

- Interference from 802.11 WLAN sources
- Ambient noise
- Rogue APs

### Infrastructure and Radio Management

We aim to make sure equipment is optimized and configured to our design specifications. We monitor the use or load on part of the system. Data in this category includes:

- Channel utilization
- Client counts per radio
- 802.11 statistics, such as packet errors, retransmits, and drops
- Quality of Service (QoS) statistics for high-priority data streams such as voice traffic, including jitters and queue delays.

### Client Management

Tracking client data such as performance allows us to monitor the network from a user's perspective. Data in this category includes:

- Coverage, signal strength, and signal to noise ratio (SNR)
- Association speeds
- Client count per service set identifier (SSID)
- Client throughput

## Analysis and Reporting

We generate analysis and reports to assess network health based largely on thresholds that we set for our selected parameters. We continually refine these thresholds based on our experience. An example might be that an AP should have 15 or fewer connected users. We track the percentage of time that the system meets our expected thresholds. We correlate this information to create useful views of the state of our environment and we generate health reports based on these views. These focused health reports enable administrators to quickly assess the state of the network. Key reports are listed below.

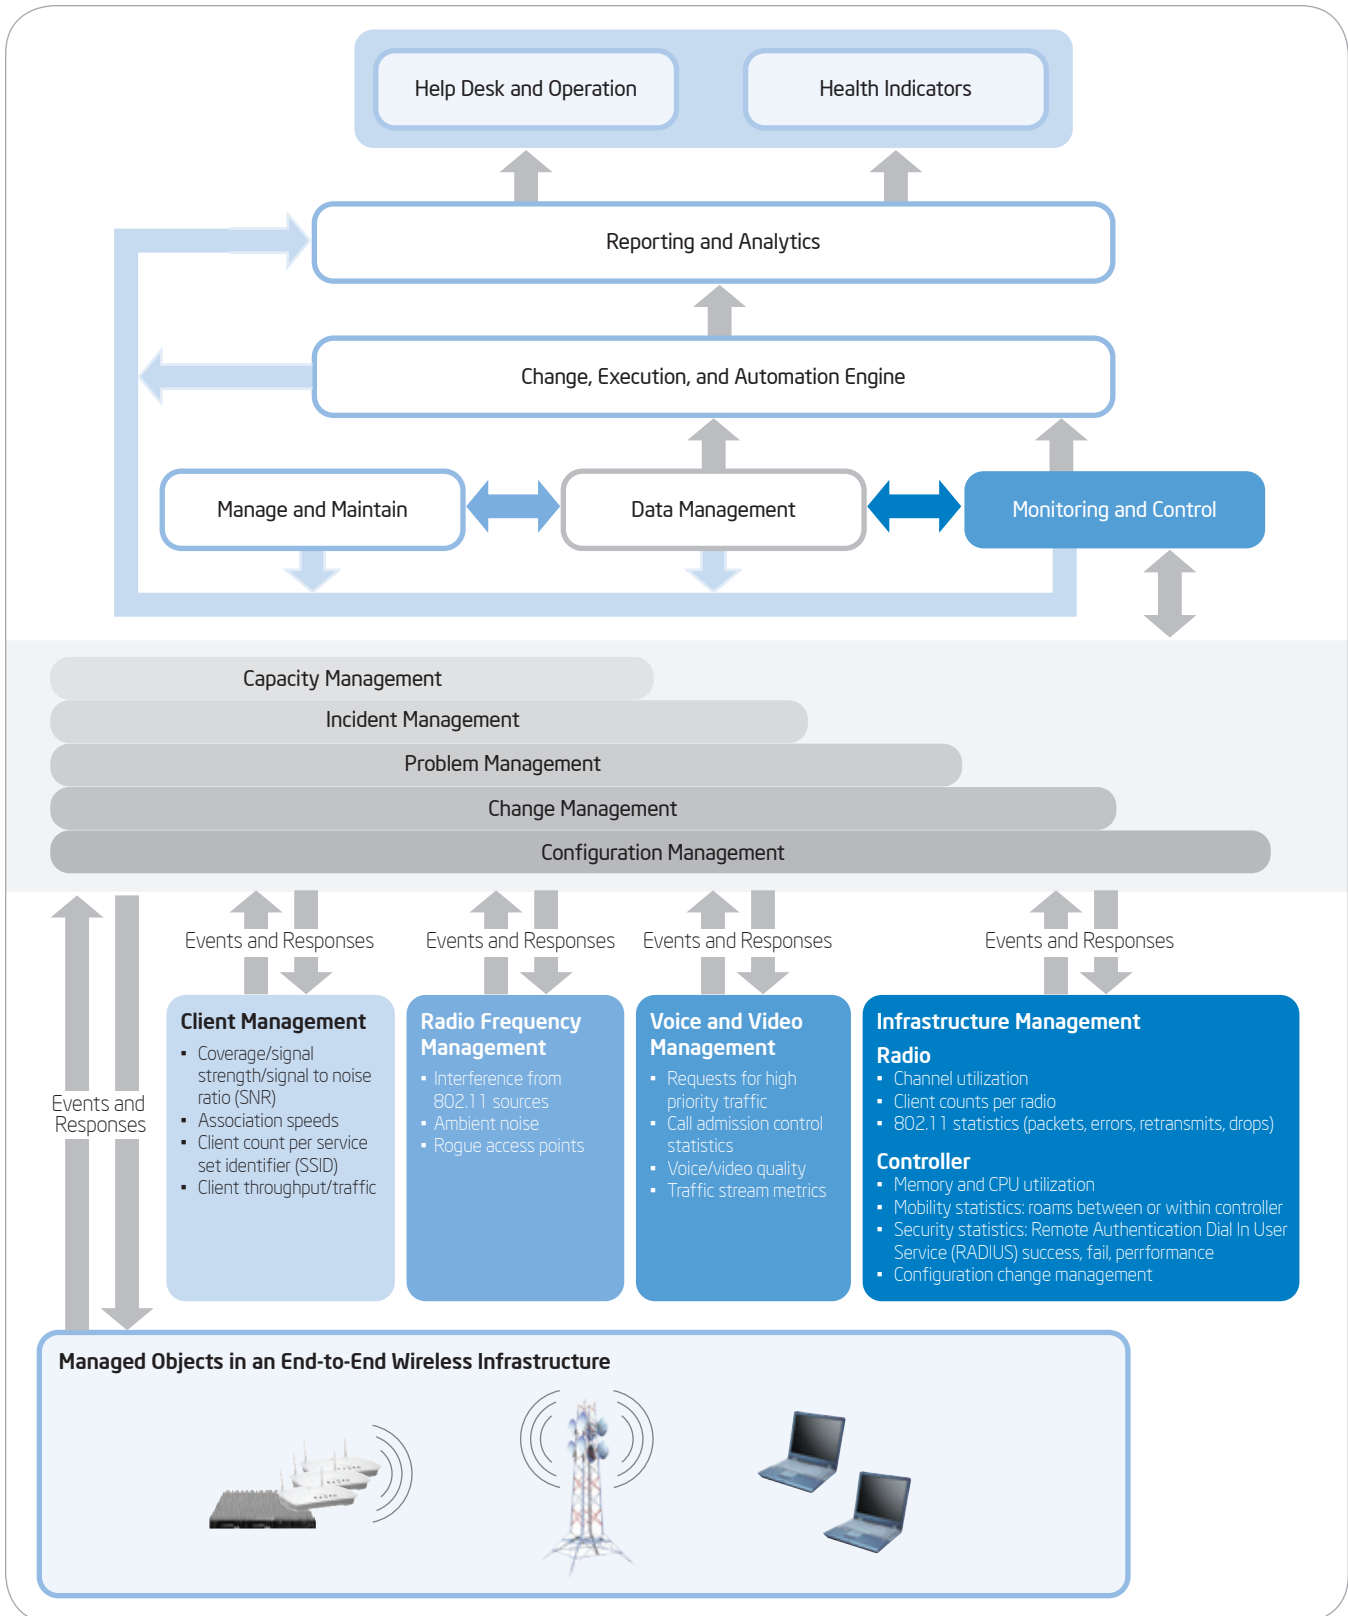


Figure 2. Our wireless LAN service level management model.

### Daily Client Counts by Service

This report shows use of specific WLAN services each day, providing a measure of how much the user population is taking advantage of wireless. This also helps track the users' transition from LAN to WLAN over time. We can produce

reports showing overall use of the 802.11a primary wireless band and the 802.11g legacy band used to support legacy devices, as shown in Figure 3. This provides a daily estimate of wireless use by campus building.




Figure 3. Typical daily client connections by 802.11 wireless bands.



## Health Indicators

Figure 4 shows four vectors of wireless performance for the 802.11a and 802.11g bands. The infrastructure grades itself in real time on a pass-fail basis, based on whether infrastructure elements such as APs meet our specified thresholds. Our report allows us to verify that we meet our performance expectations more than 99 percent of the time for all four vectors:

- Coverage, with a threshold of no coverage holes detected
- Client load, with a threshold of fewer than 15 clients per AP
- Ambient noise below our specified threshold
- Interference below threshold

## RADIUS Delta Statistics


These reports continually compare views of user authentication, based on the Remote Authentication Dial In User Service (RADIUS), at multiple layers of the environment, as shown in Figure 5. They provide views into the volume of authentications as well as authentication failure rates. They allow us to track data such as client requests, accepts, rejects, retransmissions, bad authentications, dropped authentication packets, and client timeouts. Because network access requires authentication, the health of the authentication infrastructure is a key indicator of network health.

## WLAN Controller Utilization

This tracks utilization of controller CPU, memory, and interface, as shown in Figure 6, so we can help ensure that we are not running out of resources. The data is used for both short-term and long-term capacity management and reporting.

Our health reports provide us with early warning of potential issues, such as those created by insufficient coverage or signal quality. The reports also enable us to pinpoint problems such as rogue APs and sources of interference.

Many reports are built as drilldown reports. An administrator can delve into the report to find the source of the potential problem.



**Figure 4. Typical health indicators report.**

### Interference in 802.11 Wireless Bands

The 802.11g band is susceptible to signal interference caused by other 802.11 sources. The 802.11a band is less susceptible to interference due to the higher number of available channels, so it consistently meets performance design levels for primary wireless LANs.




Figure 5. Remote Authentication Dial In User Service (RADIUS) delta statistics report.




Figure 6. Wireless LAN controller utilization report.

# Conclusion

By using a centralized WLAN architecture combined with a strong emphasis on monitoring health and performance, we have been able to mitigate the bandwidth and service level challenges of WLAN management and provide a dependable service to our users.

The controller-based architecture and centralized management reduces the number of managed entities and provides a view of the entire wireless infrastructure. Without this centralized visibility, we believe that it would not be possible to achieve RF management, which is key to optimizing performance and avoiding interference.

In addition, our experience shows that the right instrumentation and analysis can greatly help achieve the goal of providing acceptable service levels to users. In the future, we expect to integrate more data from clients and from the VoIP infrastructure. By providing us with a more complete view of the network, this will help

ensure that we continue to provide services that meet user expectations.

Preliminary analysis suggests that our approach may deliver substantial benefits in reduced management effort, faster deployment, a more consistent environment, and a better experience for users.

Because we are WLAN pioneers, we needed to put considerable effort into creating our system management suite, since comprehensive management tools were not yet available. We expect that this will change and that in the future, many of these capabilities will be available off the shelf.

## Authors

**Omer Ben-Shalom** is a wireless LAN engineer with Intel Information Technology.

**Ashok Desai** is a senior network specialist with Intel Information Technology.


**Eva Donohoe** is a network engineer with Intel Information Technology.

**Omri Barkay** is a network engineer with Intel Information Technology.

**Sylvia Stump** is a project manager with Intel Information Technology.

## Acronyms

<b>AP</b>	access point	<b>SNR</b>	signal to noise ratio
<b>MTBF</b>	mean time between failures	<b>SSID</b>	service set identifier
<b>QoS</b>	Quality of Service	<b>VoIP</b>	Voice over IP
<b>RADIUS</b>	Remote Authentication Dial In User Service	<b>WAP</b>	wireless access point
<b>RF</b>	radio frequency	<b>WLAN</b>	wireless LAN



**[www.intel.com/IT](http://www.intel.com/IT)**

This paper is for informational purposes only. THIS DOCUMENT IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NON-INFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE. Intel disclaims all liability, including liability for infringement of any proprietary rights, relating to use of information in this specification. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

Intel, the Intel logo, Intel. Leap ahead. and Intel. Leap ahead. logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in other countries.

\* Other names and brands may be claimed as the property of others.

Copyright © 2007, Intel Corporation. All rights reserved.

Printed in USA  
0207/ARM/RDA/PDF



Please Recycle  
Order Number: 315513-001US