



White Paper
Intel Information Technology
WLAN Design

Implementing Quality of Service for Voice over Wireless LANs

Intel IT developed a Quality of Service (QoS) implementation that lets us carry voice traffic over wireless LANs (WLANs) to support our transition to a new, unified network architecture in which wireless becomes the primary network access method for data, voice, and video. We are using client software to manage resources and prioritize traffic on laptops, so they can effectively run soft phones alongside other resource-intensive applications. This approach has significantly improved voice quality and provides a basis for expanded QoS capability going forward.

Omri Barkay and Omer Ben-Shalom, Intel Corporation

August 2006

IT@Intel

Executive Summary

To support voice traffic over wireless LANs (WLANs), Intel IT is developing an effective approach to Quality of Service (QoS) as part of a unified network architecture that supports data, voice, and video and establishes high-performance WLANs as the primary access method. By doing so, we expect to increase agility and productivity, eliminate excess infrastructure, and reduce costs.

We have found that the client QoS package has a significant effect, enabling high-quality voice even when a laptop is running other resource-hungry applications.

One of the biggest challenges for this new architecture is providing QoS to support high-quality voice traffic over WLANs, which have more bandwidth limitations than wired LANs. One important, though often ignored, aspect of QoS is resource management on laptops and other clients, so they can run soft phones alongside other applications that compete for client resources.

We are using client software that reserves resources and prioritizes packets on the client, ensuring adequate resources for voice traffic. We have found that the client QoS package has a significant effect, enabling high-quality voice even when a laptop is running other resource-hungry applications.

As we move forward, we are investigating a broader approach to helping ensure end-to-end QoS over the corporate network. We envisage several elements:

- Fine-grained, application-level control over QoS
- Trusted clients
- Infrastructure support for QoS
- A central policy server

We believe this approach will create a solid foundation not only for voice, but for a completely converged service infrastructure.

Contents

Executive Summary	2
Background	4
The Bandwidth Challenge	4
The QoS Challenge	5
The Solution	5
Existing Network QoS Mechanisms	5
QoS on the Client	6
Results	7
A Broader View of QoS	8
Application-level QoS Control	8
Achieving Client Trust by the Network	9
End-to-end QoS	10
Central Policy Control	10
Conclusion	11
Authors	11
Acronyms	11

Background

Intel is working toward a new network architecture aimed at empowering our increasingly mobile users (diagrammed in Figure 1). To date, we have maintained separate voice and data networks. Now, we are working to converge data, voice, and video onto a unified network infrastructure. By doing so, we expect to increase agility and productivity, eliminate excess infrastructure, and reduce costs.

We also believe that wireless will become the preferred method among our users for accessing these services. Currently, WLANs are popular and have been widely deployed at Intel, but they are slower than the wired LAN and nearly always used as a secondary access method. WLANs are maintained as separate networks alongside the wired LANs, which is expensive.

Our new architecture will integrate wired and wireless LAN infrastructure, and establish high-performance wireless as the primary access method. We have begun a groundbreaking initiative to use primary WLANs based on our new architecture within large Intel campuses housing thousands of users. Our vision is that data, voice, and ultimately video will be delivered wirelessly via laptops, handsets, and other devices to mobile users.

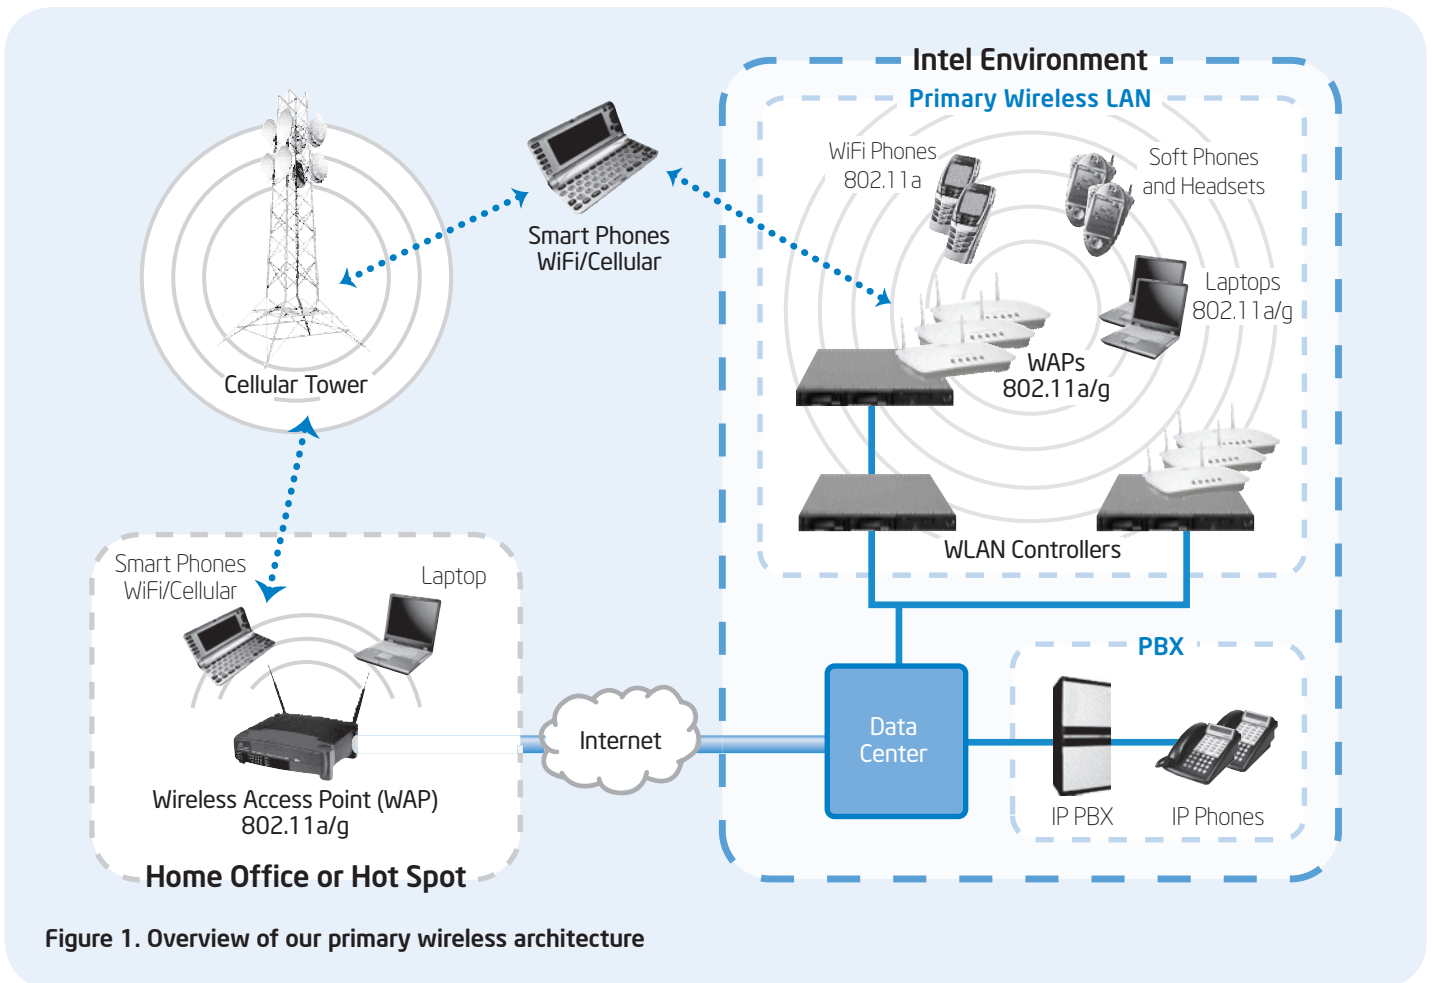


Figure 1. Overview of our primary wireless architecture

Converging voice, video, and data onto a single network is challenging, however, especially when wireless is the primary access method.

The Bandwidth Challenge

With traditional data-only networks, accessed through wired LANs, there was little requirement to prioritize traffic. If bandwidth became a problem, frequent improvements in technology made it easy to add capacity—increasing “quantity of service”—until the balance between infrastructure and bandwidth consumption was restored.

The introduction of WLANs upset this balance.

WLANs provide less bandwidth than wired LANs. Mobile devices must share the bandwidth provided by access points, instead of getting their own dedicated, switched, wired LAN connection. Moreover, the WLAN protocol suite carries a substantial overhead to compensate for the noisy medium.

The ability to add infrastructure to overcome bandwidth limitations is limited by the potential for co-channel interference between neighboring access points and mobile devices.

With our initial WLANs, the impact of these restrictions was limited. Because these WLANs were used as a secondary access method, it wasn't necessary to support all services. Users could always use the wired LAN for services, such as backup, that consume a lot of bandwidth.

However, as the WLAN network evolves to become a primary access method, largely replacing the wired LAN, we need to support all existing data services.

The QoS Challenge

Carrying voice and video over WLANs imposes additional requirements. Voice and video are real-time in nature, and highly latency-sensitive. Network delays that might be imperceptible when surfing the Web can significantly affect voice quality.

To provide high-quality voice, we need to implement QoS over the network, prioritizing voice traffic so that it doesn't suffer delays and other problems.

However, we found that we need to address three issues in order to assure QoS.

- **Robust and reliable infrastructure.** This is provided by our new architecture, described in the companion paper, “Architecture and Design of Primary Wireless Network.”
- **Prioritization of traffic on the network.** For this, we need QoS standards to prioritize packets as they pass over the wired and wireless portions of the network.
- **Ensuring adequate resources on client devices.** The need to ensure client resources has not been considered as widely as other QoS issues within the industry—but we have found that it is critical. In many cases, users will access data, voice, and video services from a single client machine such as a laptop or handheld device. For instance, we are providing users with laptops with soft phone applications installed. A laptop typically runs many different simultaneous processes, and so voice has to compete with these other processes for the client's resources. If insufficient resources are available, even briefly, voice quality will degrade.

The Solution

In our wireless campus initiative, we are using a client software package to ensure resource allocation and also to prioritize packets transmitted from the client. This approach, used together with a robust infrastructure and prioritization over the network, has resulted in a significant improvement in voice quality.

Existing Network QoS Mechanisms

We have used QoS mechanisms on parts of the wired network infrastructure within Intel for some time. Historically, these mechanisms have been used in situations where bandwidth was scarce, mainly in the WAN. As we move forward, we expect to apply QoS mechanisms to prioritize voice and other traffic as it passes over the WLAN portion of our unified network. Our intention is that the QoS marking we apply to data originating on the WLAN will be retained to prioritize the data, when necessary, as it moves over our wired network.

For prioritization within the wired LAN, IEEE defined the 802.1p standard for QoS at the Ethernet media access control (MAC) layer, Open Systems Interconnection (OSI) Layer 2. The 802.1p standard defines eight priority queues, identified as the Class of Service (CoS), using a three-bit priority field, which was added to the 802.1Q virtual LAN tagging standard.

Layer 2 marking is lost as packets leave the LAN, however. Therefore, to retain QoS as packets travel over campus networks and the WAN, we use differentiated services (DiffServ), a common network layer QoS mechanism. It relies on differentiated services code point (DSCP) marking in the Type of Service (ToS) byte of the IP header. This allows per-hop decisions and policing of the marking.

As we replace wired LANs with wireless, we are adopting emerging wireless QoS standards. The

WLAN equivalent of 802.1p is 802.11e, which has recently been certified but is not yet extensively supported. In the near term, we expect to use the interim IEEE solution, Wi-Fi* Multimedia (WMM), which is becoming widely available.

With WLANs, unlike switched, wired LANs, clients have to share the bandwidth provided by each access point. This requires that the MAC layer handle medium access prioritization. The priority is set by altering the expected amount of time a station waits for medium access, depending on the traffic service type. This is called a contention window.

WMM defines four access categories, or service levels, where the top category—voice—has the shortest contention window and, therefore, the best statistical chance of transmitting frames first.

QoS on the Client

Our client package prioritizes transmitted packets and helps ensure that adequate client resources are reserved for voice.

Packet Prioritization

We use a service that exploits the Microsoft Windows* traffic control application programming interface (API) to prioritize traffic for transmission. It defines traffic flows and filters traffic based on the transport layer ports used by different applications, assigning matching packets to specific flows. We define two flows: promoted and demoted. The promoted flow is used for latency-sensitive real-time traffic such as

voice. The demoted flow is for backup and other lower-priority applications. In each case, we override the packets' default Layer 2 and Layer 3 priority marking and set bandwidth floors and ceilings to provide the bandwidth needed for each flow. All traffic that does not match these filters is treated by default as normal best-effort traffic and, therefore, becomes a medium-priority flow, as shown in Figure 2.

For handheld 802.11 Wi-Fi phones, the device marks the packets using WMM and the infrastructure grants all of its packets voice-type service.

Client Resource Management

We use a tool to guarantee that the client reserves sufficient processor and memory resources to supply the needs of an application such as the soft phone. This is done dynamically; when the soft phone is in need of the resources, the tool interacts with the operating system to guarantee the resources are available. When the soft phone doesn't require the resources, they are available for other applications to use.

Results

In initial testing, the client package significantly improved voice quality.

We measured the effect of the client package using the Mean Opinion Score (MOS) test. This is an industry-standard measure of voice quality, based on listeners' assessments of the quality of standard voice phrases transmitted over a network.

We tested the quality of voice transmitted over the WLAN by laptops with soft phones installed. To assess the effect of the client package, we stressed each laptop's resources by running resource-hungry applications in two separate tests: a virus scan configured so that it attempted to use 100 percent of available CPU resources and an FTP file transfer that generated heavy network traffic. We assessed the effect of these applications on voice quality, with and without the client package running.

When the laptops were not running other resource-intensive applications, voice transmission was of high quality whether the client package was running or not, as shown in Figure 3. Voice quality was rated at 4.0 or better on the MOS scale, a score that indicates all listeners were satisfied to very satisfied with the

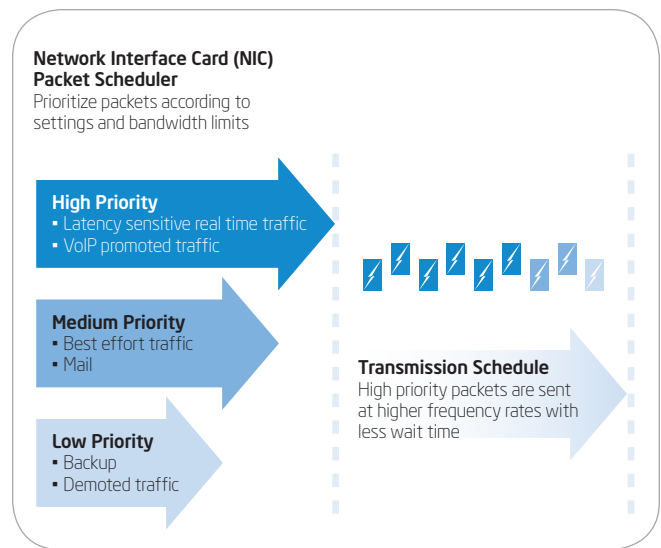


Figure2. Packet prioritization on the client.

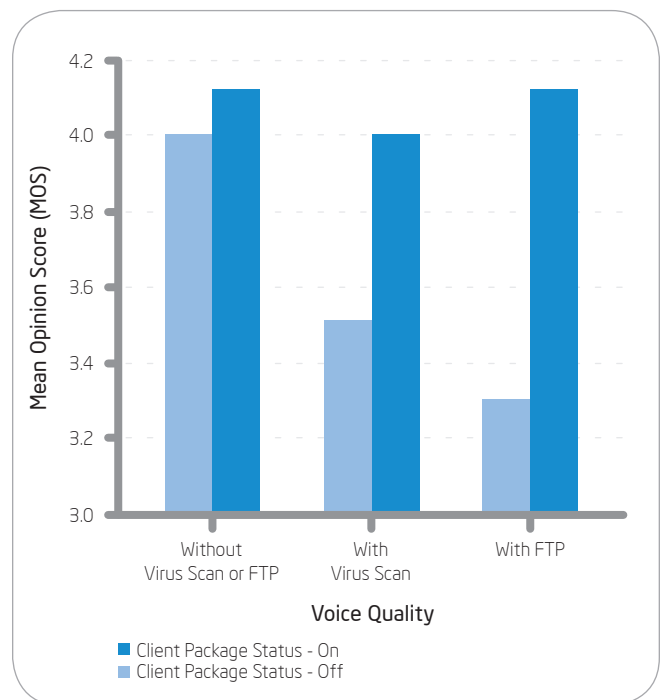


Figure 3. Effect of client package on voice quality.

quality, and close to the maximum that is typically achieved with standard voice-encoding hardware.

The client package made a big difference, however, when we ran other resource-intensive applications on the client. With the client package running, voice quality remained equally high, even with the virus-scanning package running or a file transfer taking place. Without the client package, voice quality dropped sharply.

The client package effectively reserved resources for the soft phone and prioritized voice traffic, so that voice quality remained high even when other applications tried to dominate the laptop's resources. This suggests the package will play an important role in providing users with reliable, high-quality voice conversations through their laptops.

A Broader View of QoS

The QoS approach we are taking today meets our requirements for wireless voice over the campus network. However, as we move forward to deploy this architecture more widely, we expect that we will need to adopt a broader approach, with additional QoS tools and mechanisms to impose more fine-grained, application-level control over quality, achieve end-to-end QoS, and manage and configure QoS policies.

Application-level QoS Control

While the client QoS package we use today meets our initial requirements, it has limitations. It is based only on network and transport layer information, and the filtering rules used for assigning traffic to different flows limit flexibility.

Ideally, client-side QoS should be based on a more detailed understanding of network traffic. The most important indicator of this is the process that generated the traffic and the way traffic is used within the application that created the process.

As an example, a commercial Internet-based Voice over Internet Protocol (VoIP) service behind a corporate firewall encapsulates all VoIP packets

in HTTP headers. Without knowledge of the originating process, the network will not treat these packets differently than other Web-based traffic, so the traffic won't be prioritized to help ensure QoS.

In a similar way, many other services are being transformed into Web services. Typically, these encapsulate all data in XML and then in HTTP or HTTPS. If we rely only on transport layer information, we will not be able to correctly identify the originating application and prioritize traffic accordingly.

In addition, many applications, including corporate e-mail systems, use dynamic ports, making it impossible to classify traffic from these applications using ports alone.

Therefore, we plan to further develop the current QoS package so that it can track any connection opened by an application process and use the socket information unique to that process to generate a specific rule for handling by the packet scheduler. This will enable us to prioritize traffic based on the application that created it. In Figure 4, we show a hypothetical user interface for this enhanced package

Achieving Client Trust by the Network

One issue with implementing client-based QoS is the question of trust. It is relatively easy to create new rules for the packet scheduler, so users with appropriate privileges could mark all traffic as high priority. Potentially, worms or viruses could also abuse the mechanism.

For these reasons, networks typically either disregard all client packet marking, or re-mark all client traffic. To effectively use the QoS client package, the network needs to be able to trust the marking on packets transmitted from the client.

Currently, we are assessing two main approaches for implementing client trust. The end result of both these methods is to help ensure that only

clients that meet the corporate policy are allowed to mark packets, otherwise marking will be either ignored or overridden.

Authentication and Authorization on Initial Connection

With this model, a client is checked when it initially connects to the network. This includes a security assessment, which verifies that the clients are running mandated agents and safeguards, such as anti-virus software. Only if the client passes this assessment is it allowed onto the network. While this method does not prevent client tampering after the initial connection, it at least guarantees that the client meets corporate standards for security defenses and QoS policy.

Ongoing Enforcement

This approach relies on an external agency to monitor the ongoing behavior of clients connected to the network. This agency can be a network intrusion detection or prevention system (NIDS/NIPS), or a technology that is embedded on the client machine but outside the control of the user or the operating system. This embedded technology could be based on virtualization techniques and/or Intel® Active Management Technology (Intel® AMT).

Application executable	802.1P Tag	DiffServ Tag	Bandwidth Token	Bandwidth Limit
SoftwarePhone.exe <input type="button" value="Browse..."/>	6 <input type="button" value="v"/>	Expedited Forwarding <input type="button" value="v"/>	6400 <input type="button" value="v"/>	128000 <input type="button" value="v"/>
Backup.exe <input type="button" value="Browse..."/>	1 <input type="button" value="v"/>	Assured Forwarding1 <input type="button" value="v"/>	1000 <input type="button" value="v"/>	2000 <input type="button" value="v"/>

Figure 4. Hypothetical enhanced prioritization package on client

End-to-end QoS

Client packet marking is only useful if the marking has an effect within the network. Our first step, requiring no network infrastructure support for QoS, is to use client marking in the initial WLAN uplink to determine network access priority. The next step is to complement this for traffic going in the opposite direction, by implementing downstream prioritization of packets on the final WLAN link.

However, to make full use of the packet classification and prioritization performed on the client, we need to leverage these client markings to achieve end-to-end prioritization of traffic.

For example, on the Intel WAN, a number of different priority queues guarantee proper service, even when the network becomes congested. Today, traffic is classified and directed to these queues by WAN routers, which don't trust any marking coming from the client or LAN. We intend

to change this as we move forward; making clients trusted will help us exploit client packet marking more widely within the network.

Central Policy Control

To move from a primary WLAN at an initial campus to a broader corporate deployment, we will need a toolset that allows easy deployment and central administration of client agent software and QoS policies.

Ideally, this would involve a single QoS policy agent on the client, handling resource management and packet prioritization according to corporate policy and operating outside the user's control. This agent should be controlled and monitored from a central policy server, and integrate with other corporate security and policy mechanisms.

In Figure 5, we show one potential arrangement, with end-to-end prioritization of packets according to policies set at an enterprise policy server.

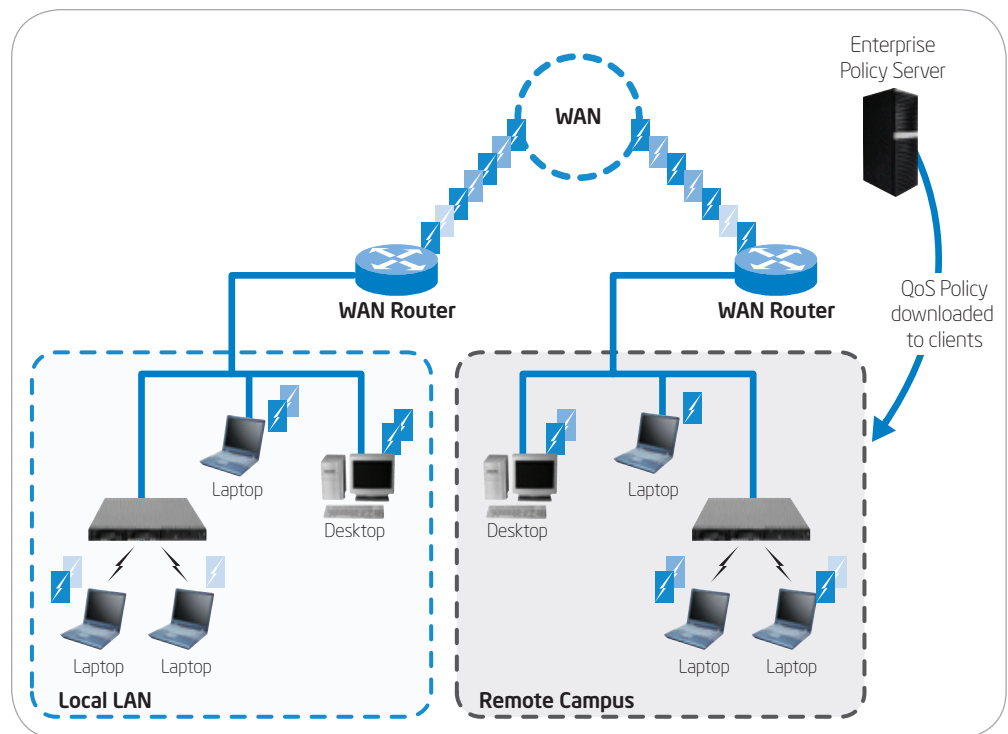


Figure 5. End-to-end QoS. In this potential arrangement, packets are prioritized according to policies set at an enterprise policy server.

Conclusion

Our approach to client-based QoS is a promising step toward achieving high-quality voice over primary wireless networks. Based on our tests so far, using a client package results in a significant improvement in voice quality.

Going forward, we are looking to develop a broader approach to QoS, implementing a central controlled policy system and mechanisms for end-to-end QoS across the network. We believe

this will create a solid foundation not only for voice but for a completely converged service infrastructure.


Authors

Omri Barkay is a wireless LAN engineer with Intel Information Technology.

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

Acronyms

API	application programming interface	NIPS	network intrusion prevention system
CoS	Class of Service	OSI	Open Systems Interconnection
DiffServ	differentiated services	QoS	Quality of Service
DSCP	differentiated services code point	ToS	Type of Service
Intel® AMT	Intel® Active Management Technology	VoIP	Voice over Internet Protocol
MAC	media access control	WLAN	wireless LAN
MOS	Mean Opinion Score	WMM	Wi-Fi Multimedia
NIDS	network intrusion detection system		



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 Corporation may have patents or pending patent applications, trademarks, copyrights, or other intellectual property rights that relate to the presented subject matter. The furnishing of documents and

other materials and information does not provide any license, express or implied, by estoppel or otherwise, to any such patents, trademarks, copyrights, or other intellectual property rights.

Intel, the Intel logo, Intel. Leap ahead, and the 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 © 2006, Intel Corporation. All rights reserved.

Printed in USA
0806/ARM/RDA/PDF

Please Recycle
Order Number: 314561-001US