

White Paper  
**Intel Information Technology**  
Computer Manufacturing  
Client Management

# Improving Manageability with OS Streaming in Training Rooms

OS streaming can deliver considerable manageability benefits to Intel IT training rooms with multi-user PCs. To evaluate performance and utilization in a production environment, we conducted proof of concept (PoC) testing in two rooms located in different buildings on an Intel campus. We found that OS streaming improved manageability and delivered fast client boot times with moderate server and network utilization, even during worst-case boot storms.

Catherine Spence, Randy Nystrom, Craig Pierce, and William Wrenn, Intel Corporation

December 2008

IT@Intel

We found that OS streaming improved manageability and delivered fast client boot times with moderate server and network utilization, even during worst-case boot storms.

## Executive Summary

OS streaming can deliver considerable manageability benefits to Intel IT training rooms with multi-user PCs. To evaluate performance and utilization in a production environment, we conducted proof of concept (PoC) testing in two rooms located in different buildings on an Intel campus. We found that OS streaming improved manageability and delivered fast client boot times with moderate server and network utilization, even during worst-case boot storms.

Streaming enables us to download software images on demand to client PCs in training rooms, rather than dispatching technicians to perform image refreshes that require the room to be taken out of service for four hours. Instead of individually managing multiple client systems, we can centrally manage one set of client images for each type of training, improving manageability and reducing cost.

For our PoC, we used a standard two-socket server based on Intel® Xeon® processor 5060 with two cores to stream the OS over the enterprise network to as many as 39 PCs in two training rooms located in different buildings. We monitored server and network performance and utilization during boot storms in which all clients booted simultaneously. We also monitored longer-term utilization during live classes.

- Clients booted very quickly, even during worst-case boot storms. Including domain authentication, boot time averaged about 95 seconds with 19 clients and 160 seconds with 39 clients.
- Good performance was provided when running the network at 1 GB between the server and the client switch, with a single 1-GB server network interface card (NIC). Bandwidth of 100 MB was adequate between the client switch and the PCs.
- Server utilization remained moderate, even during boot storms when many clients boot simultaneously from the same server.

Our successful PoC demonstrated that we can deliver good streaming performance in a production environment without heavy investment in server or network infrastructure. OS streaming improved manageability of multi-user training room PCs and could enable us to use rooms more efficiently by installing client software in about two minutes rather than several hours per room. Next steps include return on investment (ROI) analysis to determine the business case for implementing streaming across our training rooms.

# Contents

<b>Executive Summary</b> .....	2
<b>Business Challenge</b> .....	4
<b>Proof of Concept</b> .....	6
PoC Architecture.....	6
Results.....	8
<b>Key Learnings</b> .....	9
Dynamic Host Configuration Protocol Leases.....	9
Hardware Differences.....	9
Private Images.....	9
Network Performance.....	9
PoC Architecture.....	10
<b>Possible Training Room Usage Models</b> .....	10
Next: Hard Disk for Specialized Classes Only.....	10
Mainstream: Specialized or Private Images on Server.....	10
Advanced: No Hard Disks in Training Room.....	11
Future: No Devices in Training Room.....	11
<b>Conclusion</b> .....	11
<b>Authors</b> .....	12
<b>Acronyms</b> .....	12

## Business Challenge

Streaming promises the performance and versatility of the rich-client model together with the centralized management and total cost of ownership (TCO) advantages of thin clients. With streaming, client software is stored and managed on centralized servers, then downloaded on demand to the client for local execution. Because the client is a PC, local execution can preserve a user experience comparable to locally installed software.

For Intel IT, streaming also represents an evolutionary next step toward real-time, on-demand software and device-independent mobility. Our infrastructure and support model is built around our rich-client environment. This makes streaming more attractive than alternative centralized computing models. We can introduce streaming without disruption to our existing environment, while retaining the option to run locally installed software on each client. In contrast, the thin-client approach, which relies on server-based software execution, is less attractive because it would require investment in additional server infrastructure and significantly change our IT support model.

We have identified our IT training rooms as an environment that is a good fit for the manageability benefits of streaming technology. Training rooms contain multi-user desktop PCs that we need to frequently reconfigure for different classes. OS streaming could allow us to do this more quickly and at lower cost.

We previously performed lab tests that indicated the potential benefits of streaming for several multi-user desktop use cases within Intel, including manufacturing and call centers. However, because streaming is still a relatively new technology, there are still concerns about the impact of streaming on performance and utilization in a production environment. Downloading OS images to each client increases the load on servers and the enterprise network, particularly in “boot storms” when many clients boot simultaneously from the same server.

Any resulting delays could make streaming less acceptable to users accustomed to PCs with locally installed software.

To determine whether streaming could provide adequate performance and utilization to support live training rooms, we conducted proof of concept (PoC) testing in our production environment.

### Training Room Environment

Each year, tens of thousands of Intel employees use our IT training rooms. Training groups use the PCs in these rooms to run many different courses, covering diverse subjects such as OSs, office productivity applications, development tools, factory-specific tools, and enterprise resource planning (ERP) applications.

Each of the 41 rooms has approximately 19 client PCs. The default client OS is Microsoft Windows XP\* though some machines run Linux\*. In addition, each room contains a server that stores shared data and an image of the software on each client PC. In total, there are about 800 systems in training rooms worldwide.

Managing the training room IT environment remains labor-intensive. After many classes, we need to refresh the images on each PC so that it is ready for the next class. It takes four hours to refresh a classroom, because we need to dispatch a technician to the room, and the technician then needs to load the OS onto each PC from the image server. Technicians also need to update each client image to help ensure that PCs comply with Intel security requirements.

Besides being resource-intensive, this also limits the availability of training rooms, because we cannot schedule rooms for back-to-back training sessions. This could become even more important in the future, as we move toward a more efficient and flexible model in which rooms are shared between training and other activities. To realize the benefits of this model, we need to streamline the setup and teardown of training room systems so that we can efficiently reuse space.

### Streaming in Training Rooms

Streaming could deliver significant benefits to training rooms by improving manageability and room utilization. A server located elsewhere on the Intel enterprise network could deliver images on demand, at the start of a class, to all the clients within each training room. We could reduce TCO by managing images centrally and downloading them using streaming, instead of dispatching a technician to the room. This would also increase room utilization: We could prepare the room for a new class more quickly and easily, because refreshing client images would essentially consist of downloading a fresh image to each PC.

Streaming technology allows us to reduce the effort required to manage client OS images. We could store a total of two images on the server. We would incorporate security updates and other changes into a master update image. From this, we would create a single active image that gets delivered to all clients. The active image is different from the update image in that it contains the necessary drivers to account for hardware differences between all of the client hardware systems. Having one image simplifies security compliance: We merely need to configure each PC to point to the new image and reboot to be fully compliant.

Streaming could also simplify hardware maintenance by making it easier to quickly swap out PCs if issues occur. If we have a system failure, we can replace the PC and simply re-stream the class image to the replacement system.

### Steps in OS Streaming

The steps in OS streaming, shown in Figure 1, are as follows:

1. When powering on, clients execute a preboot execution environment (PXE) boot.
2. The client connects to the streaming server.
3. The server uses the client's media access control (MAC) address to locate the virtual disk assigned to the client, which contains the OS image that the client will load.
4. The client mounts the virtual disk and the OS image is streamed to the client.

OS streaming can use either standard or private images. For our PoC, we used a standard image. In this mode, a single active image is streamed to all clients. Each time the client reboots, it receives the same unchanged image from the server. An alternative is to use private images, in which the server stores an individual OS image for each client. In this mode, any changes made by the client are saved to the private image on the server, so they are not lost if the client reboots. The disadvantage is that the streaming server requires enough storage capacity to hold private images for all the clients it supports.

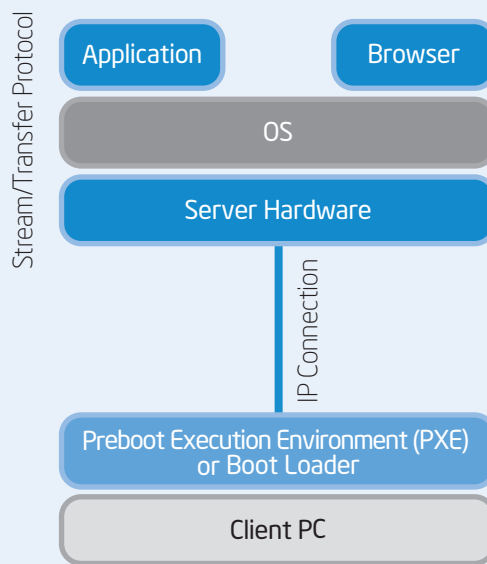


Figure 1. OS streaming.

# Proof of Concept

To assess the viability of OS streaming for the training room environment, we conducted PoC testing in our production IT environment and IT training rooms. During our PoC, we streamed the client OS from a server over the Intel IT enterprise network to as many as 39 PCs in two rooms. We measured performance and utilization. We developed best known methods (BKMs) that we can use in future production streaming implementations.

We conducted two tests:

- We tested a worst-case scenario in which all PCs boot simultaneously. This could happen at the start of a class or when an urgent security patch is applied. This creates a boot storm that maximizes the load on the server and network. We measured boot times with and without authentication, and monitored server CPU and network utilization during this critical initial period.

- We also gathered more long-term server and client performance and utilization data by using streaming for live classes over several days.

## PoC Architecture

Our PoC architecture is shown in Figure 2. We streamed an OS image from a server in one building to two training rooms in different buildings on the same campus.

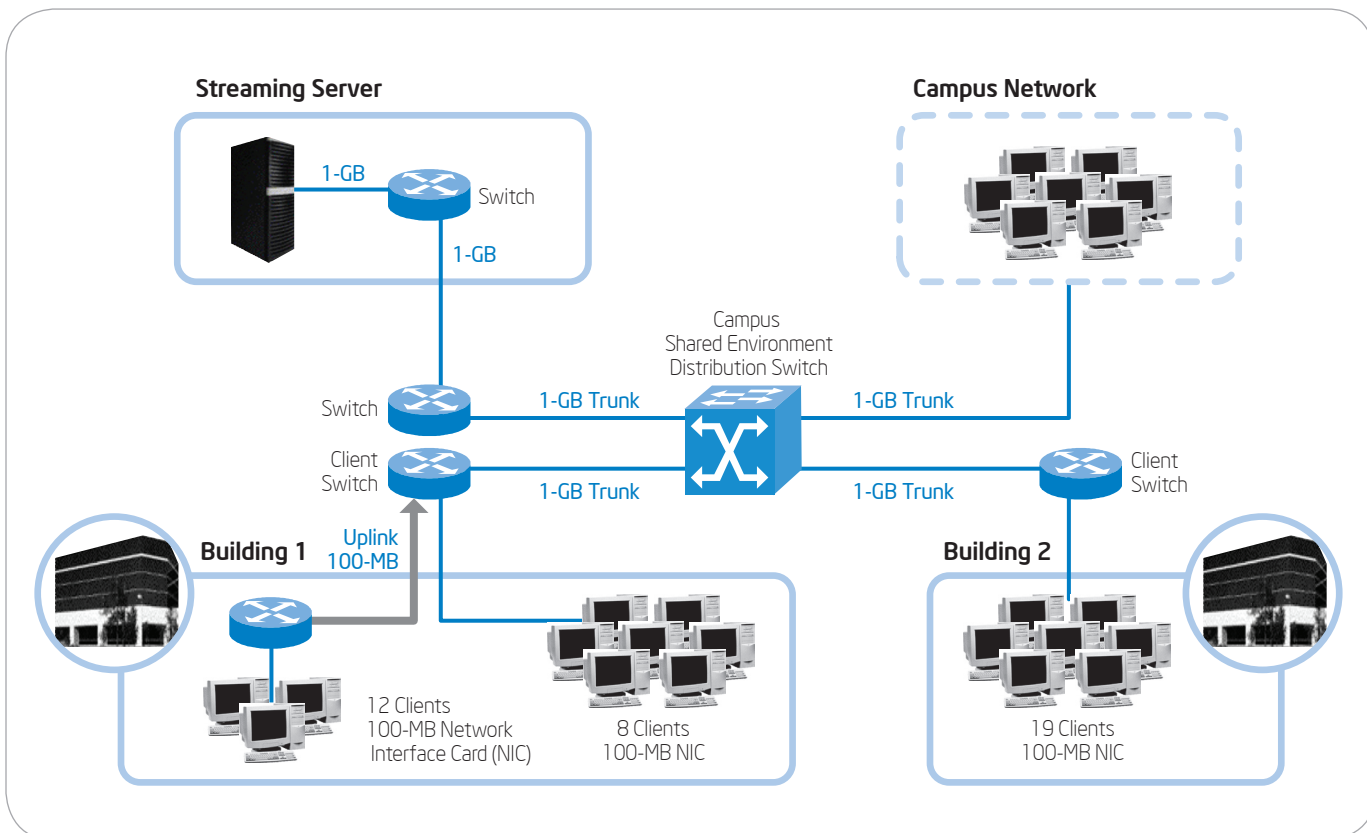


Figure 2. Proof of concept architecture.

Our streaming server was a standard two-socket server based on Intel Xeon processor 5060 (3.2 GHz) with two cores per processor. The server had 8 GB of RAM and a single 1-GB NIC, as shown in Table 1. Our clients were desktop PCs with Intel® vPro™ technology, with Intel® Core™2 Duo desktop processor E6400 (2.13 GHz), 2 GB of RAM, and a 100-MB NIC, as shown in Table 2.

### Test Procedure

We built a single active virtual image on the server to support all clients. We configured each desktop to boot across the network and stream the OS from the server. We also created an image that resided on each client’s hard disk, purely as a backup.

### Boot Storm

To create a boot storm that generated the maximum load on the server and client, we ran a server script that remotely booted all clients. With network boot enabled, each client automatically looked for an active image. The image was then streamed to all clients. Another server script captured start and stop times for each client

download, allowing us to calculate load times for each client. We measured server CPU and NIC utilization during the streaming process using the Microsoft Windows Performance Monitor\* utility.

We conducted several tests:

- **Streaming to 19 clients in one room.** We ran each test twice with full domain authentication and twice without domain authentication.
- **Streaming to a total of 39 clients (19 in one building and 20 in the other).** We ran each test twice with full domain authentication and twice without domain authentication.

For comparison, we also measured how long it took client systems to boot using locally installed software.

### Production Use

We assessed performance during production use by using streaming to support six-hour live classes over a one-week period. We captured performance and utilization at one-minute intervals throughout classes conducted on three consecutive days.

**Table 1. Server Configuration**

Processor	<ul style="list-style-type: none"> <li>▪ 2 x Intel® Xeon® processor 5060 (3.2 GHz)</li> <li>▪ 2 cores per processor</li> </ul>
Chipset	Intel® 5000P chipset
RAM	8.0 GB
Network Interface Card	Single 1 GB
Disk Capacity	100 GB
OS	Microsoft Windows Server 2003*

**Table 2. Client Configuration**

Client Platform	Desktop PC with Intel® vPro™ technology
Processor	<ul style="list-style-type: none"> <li>▪ Intel® Core™2 Duo desktop processor E6400 (2.13 GHz)</li> <li>▪ 2 cores per processor</li> </ul>
Chipset	Intel® Q965 Express chipset
RAM	2.0 GB
Network Interface Card	Single 100 MB
Disk Capacity	80 GB
OS	Microsoft Windows XP*

We monitored network, memory, and CPU utilization on all 39 clients. For the server, we monitored disk I/O in addition to network, memory, and CPU utilization. For each one-minute interval, we averaged the data gathered at the same time on each of the three days.

## Results

We successfully streamed a single image to up to 39 clients during our boot storm tests and were able to support live classes using streaming. This enabled us to manage the streamed clients remotely, without dispatching a technician to the room.

### Boot Storm

Clients loaded their OSs very rapidly, even when we were simultaneously loading OS images on up to 39 clients over the network from a single server.

When streaming to 19 clients in a single training room, load times averaged approximately 95 seconds, including domain authentication. For 39 clients in two rooms, the average was approximately 160 seconds.

These times represent a dramatic improvement over our current training room multicast process, which takes 40 to 80 minutes, plus an additional 15 to 30 minutes for security and other software updates, before the clients are ready for use. Reinstalling the Intel build from a CD takes even longer—two to four hours per system. Streaming over the network was only slightly longer than the average of about one minute, without authentication, for loading locally installed software from the hard drive. Load times were about 30 percent longer, on average, in the tests using authentication than in the identical tests without authentication.

The fast load times we observed with OS streaming demonstrate that we could increase training room utilization and reduce support costs. We could stream the OS to each PC at the start of a class, rather than having to make

the training room unavailable for four hours while a technician refreshes the client OS images as is currently the case.

Server utilization remained moderate throughout the boot storm. This indicated that we can deliver good streaming performance with a standard two-socket server.

When the server was streaming to 19 clients, server CPU utilization ranged from 15 to 35 percent, and utilization of the server's 1-GB network connection reached approximately 50 percent for approximately 100 seconds. When streaming to 39 clients, server CPU utilization ranged from 20 to 40 percent, and network utilization reached approximately 60 percent for approximately 100 seconds. After the boot storm, server network utilization stabilized at less than 15 percent in all the tests.

### Production Use

Server and client utilization remained low during production use, confirming that our server and client hardware was adequate for supporting streaming.

For the server, CPU utilization generally remained at less than 15 percent, with network utilization below 15 percent.

For clients, utilization of the 100-MB network connection remained well below 1 percent, CPU utilization remained below 30 percent, and memory utilization ranged from 15 to 18 percent.

### Instructor Feedback

Feedback from one classroom instructor indicated that system performance with streaming was comparable to using a locally installed OS. The instructor indicated that responsiveness, speed of operation, functionality, and ease of access were all similar to using an OS installed on the hard drive of each PC. Overall, the instructor indicated a slight preference for streaming, suggesting there is value in centralization when performance is roughly equivalent.

# Key Learnings

We gained considerable experience during our PoC, which will help us in future streaming initiatives. Key learnings included:

## Dynamic Host Configuration Protocol Leases

Because we retained an image on each PC's hard disk as a backup, each PC required two directory entries and IP addresses: one for the streamed image and one for the image on the local hard disk. We needed to expand the number of Dynamic Host Configuration Protocol (DHCP) leases allocated to the rooms to help ensure that there were enough IP addresses.

## Hardware Differences

Probably the most important training room hardware consideration is uniformity. Since a common image is streamed to multiple machines, it is important that we have identical client machines to help ensure the proper drivers are contained in the image. This eliminates the need for the OS to detect new hardware once the final image has been created. Having identical hardware in each room reduces the effort required to create and deploy an active streaming image to the training rooms.

We did, however, develop a technique for accommodating hardware differences among the PCs used in the PoC, by creating an image that includes all the drivers used by any of the PCs. Though this includes several manual steps, it greatly improves manageability because we only need to manage one image. This also saves storage space because we do not have to store a separate image for each client.

1. Start with the master update image.
2. Boot the first machine using private mode.
3. Detect any new hardware and install the necessary drivers.

4. Save the interim image.
5. Repeat steps 2 to 4 for each PC in turn, using the same interim image and updating it each time. At the end of this process, the image can support all hardware used by any machine in the room.
6. Save the final image as the active image.

We also suppressed unnecessary plug-and-play error messages that instructed the user to reboot in order to make drivers take effect, even though the drivers were already installed. We accomplished this by editing the Microsoft Windows\* registry using information obtained from a Microsoft support Web site.

## Private Images

Some training requires students to reboot during the class. Each time this happens, it resets the OS on the PC. To retain user configuration settings and local data after a reboot or during classes that last several days, we need to maintain a private image on the server for each client PC.

## Network Performance

We tried different server network configurations and compared the performance. We found that setting the server NIC to 1 GB resulted in a solid user experience and efficiently utilized the NIC. A slower 100-MB server NIC setting resulted in longer load times of approximately 20 minutes, with caching keeping the network bandwidth high for several minutes after booting all systems within the room. Dual, load-balanced 1-GB NICs did not deliver any significant performance advantage. For clients, we found that a 100-MB connection seemed adequate.

This is important, because many training rooms have 100-MB LANs. Our results suggest that the local LAN does not need to be upgraded to 1 GB in order to run OS streaming.

We therefore determined that a good network configuration is a single 1-GB NIC for the server, with a 100-MB LAN within the training room and 100-MB NIC for each client. We used this configuration during the PoC.

## PoC Architecture

Our model, using a central server for streaming to several training rooms in multiple buildings across a campus, worked well for the scenarios we tested. Even with a modest server configuration, performance was adequate in a worst-case boot storm with 39 clients.

# Possible Training Room Usage Models

Our PoC used a conservative approach with a parallel hard disk image installation for redundancy. This takes advantage of OS streaming products' ability to allow a choice of loading an OS from the server or from the hard disk. The hard disk image can serve as a backup in case the streamed OS gets corrupted or other problems ensue. This can ease concerns for IT support personnel, especially during the initial transition to streaming. The flexibility to load from the hard disk is also useful for classes that save their work from day to day or manipulate the OS running on the device.

This approach is not an ideal long-term solution, because the need to maintain the image on the clients in the training rooms means we lose the potential manageability benefits provided by the streaming approach. In addition, we need two IP addresses for each client.

## Next: Hard Disk for Specialized Classes Only

As a possible next step, we could use the hard disk only for specialized classes. Most standard classes could be supported entirely by streaming, which would eliminate the need for an additional IP address for the local hard disk during these classes.

However, some classes may simply not be a good fit for streaming. Examples could be OS device driver or debugging classes in which students continuously reboot systems so their changes can take effect.

At any time before the specialized class, we could provision the local hard disk with the image required for the class. This would not affect our ability to stream standard classes to the same client. Immediately before the specialized class, we could simply change the settings for all the clients in the class. From the server console, we would set the class PCs to boot from the hard disk instead of streaming the virtual image from the server, and then reboot the PCs. After the specialized class was completed, we would return the server to its original setting for streaming.

## Mainstream: Specialized or Private Images on Server

We could realize additional benefits if we do not have to install images on the local hard disks. Specialized class images can be created, stored as a virtual image on the server, and streamed to

training room PCs. The hard disk is only used for local storage or as local caching for the OS.

However, if the class is manipulating the OS and saving these modifications over a series of days, then we need to make a copy of the active image for each client machine in the training room, assign each client its own image, and ensure that the image is set to private mode. This means the server must have enough storage to hold all the images. This is not an issue for training that is stored in a central repository, such as Web-based training.

We hope that streaming vendors will provide a simple model based on using a single main image and allowing each client machine to save only changes to its own private caching location, then deleting all the private cached images at the end of the class. With this model, when a client makes a change to its OS and reboots, the client's changes are still applied.

## Advanced: No Hard Disks in Training Room

Once we have established processes and maintenance to support streaming for most classes, there will be much less need for hard

disks within the room. Disk failures are a common cause of PC problems, so eliminating them could reduce training room maintenance costs and the likelihood that problems will occur. Instead, students could use USB or other external storage devices to store any local data or for specialized classes.

We can also add flexibility by using notebooks instead of desktops. This means we would not have to dedicate specific rooms for training use only. The rooms could be used as additional conference rooms, for example. When the room is needed for training, we could roll in a cart with training laptops, plug them into the LAN connections, and then stream software to the systems to initialize them for the class.

## Future: No Devices in Training Room


At some point, there might not be any need for desktops in most training rooms. Students could bring their own mobile devices and have the training image streamed to the devices in virtual containers. We could loan devices to students who do not have their own dedicated systems.

# Conclusion

Our successful PoC demonstrated that we can effectively provide streaming, with good performance, to live training rooms in our production IT environment. Clients booted quickly, even in worst-case boot storms. We were able to deliver streaming with a standard IT server, which supported up to 39 clients with low to moderate server utilization. This indicates that streaming, in contrast to the thin-client approach, would not require substantial investment in new server hardware or our IT support model.

Streaming enabled us to improve training room management by reducing the need for desk-side support. We anticipate that it will improve room utilization and facilitate more flexible use of the space. Now that we have established technical

feasibility, we plan to perform ROI analysis using the Lean Six Sigma\* methodology to verify that implementing streaming technology can indeed lower costs in our production environment.



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

## Authors

**Catherine Spence** is an enterprise architect with Intel IT.

**Randy Nystrom** is a service line manager with Intel IT.

**Craig Pierce** is a systems engineer with Intel Digital Enterprise Group.

**William Wrenn** is a principal engineer with Intel IT.

Special thanks to **Sabraish Kumar** and **Jerry Wheeler** for their contributions to the project.

## Acronyms

<b>BKM</b>	best-known method	<b>PoC</b>	proof of concept
<b>DHCP</b>	Dynamic Host Configuration Protocol	<b>PXE</b>	preboot execution environment
<b>MAC</b>	media access control	<b>ROI</b>	return on investment
<b>NIC</b>	network interface card	<b>TCO</b>	total cost of ownership

This paper is for informational purposes only. THIS DOCUMENT IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NONINFRINGEMENT, 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 Core, Intel vPro, and Xeon are trademarks of Intel Corporation in the U.S. and other countries.

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

Copyright © 2008 Intel Corporation. All rights reserved.

Printed in USA  
1208/JLG/KC/PDF

Please Recycle  
320562-001 US