March 2006

Virtualizing High Network I/O Environments

Introduction

Every enterprise can benefit from improving the utilization of IT resources. Improved utilization not only helps reduce costs but can also simplify IT infrastructure, leading to reduced operations costs. A virtual computing infrastructure can provide the features that enable these benefits. At the same time, a virtual infrastructure can heighten the availability of mission-critical applications and improve IT agility to respond to business needs. Virtuozzo software provides a high-performance, low-overhead environment for virtual computing infrastructures; the low overhead on Virtuozzo makes it ideal for I/O-intensive applications such as databases. To maintain the reliability, flexibility, and performance that a high I/O virtual infrastructure provides, physical servers hosting virtual machines need multiple, high-throughput network connections. This is an ideal environment for the use of connection-dense Intel® PRO/1000 multiport gigabit and Intel® PRO/10GbE 10-gigabit Ethernet network adapters.*

Virtualization Addresses Key Challenges Facing IT Today

IT organizations everywhere face daily challenges: Do more with less. Respond to new initiatives faster. Reduce downtime. Sound familiar? A growing number of enterprises are finding that server virtualization is a proven way to better utilize equipment, keep management, administrative, and overhead costs in check, improve availability of resources, and provide a flexible infrastructure that accommodates rapid change to business needs.

INCREASED SERVER UTILIZATION IMPROVES RETURN ON INVESTMENT

Due to a one-application-per-server deployment philosophy, most distributed servers operate at low levels of utilization. Virtualization enables the workload from multiple servers to run on the same physical server node, taking advantage of that unused processing power. Virtualization makes existing server resources available for doing more work, which allows you to get more done with what you already have. Each virtual machine provides an isolated operating environment for use by applications and users. A server virtualized in this way can be 60 to 80 percent utilized, which is expected to lower the server investment required to handle a given amount of work.

VIRTUALIZATION LOWERS TOTAL COST OF OWNERSHIP

In addition to reduced capital costs, the fewer physical servers and reduced data center complexity required by virtualization can also reduce administrative costs. Virtualized server management tools can contribute to improved administrative productivity. With a small set of server management tools that are standardized across the enterprise, training costs can be reduced. Also, fewer servers can reduce the IT infrastructure footprint and associated utility costs for power and cooling.

SWsoft

A VIRTUALIZED INFRASTRUCTURE IMPROVES APPLICATION AVAILABILITY

Application availability is essential to ongoing mission-critical operations. Server virtualization helps keep business operations running smoothly by providing the tools to migrate a virtual machine to a different physical host with a service interruption of only a few seconds—eliminating hours of downtime when planned maintenance needs to be performed.

In the case of an equipment or application failure, a virtual machine on another physical host can be provisioned and started in minutes. Backup data can be restored to the new session quickly, minimizing downtime.

The result is increased application availability, as application outages in a virtualized environment are reduced to minutes or hours instead of days.

VIRTUALIZATION ENABLES BUSINESS AGILITY

Virtualization initiatives enable a server infrastructure that can be very quickly modified to meet changing business needs or help address new opportunities. Virtualized servers can be provisioned quickly—usually in minutes instead of days—to handle peaks in demand. This ability is crucial when workloads (such as Web-based applications) are unpredictable and business operations need to respond in a rapid manner.

Introducing Virtuozzo*

Virtuozzo from SWsoft is server virtualization software that makes very efficient use of hardware and operating system (OS) resources. Virtuozzo employs a technology known as OS virtualization to define, configure, and operate multiple virtual private servers (VPSs) in a single hardware node, as shown in Figure 1.

Virtuozzo is available for Microsoft Windows Server* 2003 and multiple Linux* distributions, and runs on platforms based on Intel® Pentium® processors, 64-bit Intel® Xeon® processors, and Intel® Itanium® 2 processors. Virtuozzo supports multiprocessor configurations with up to 32 processors and 64 GB of memory.

The virtualization technique employed by Virtuozzo (see sidebar on page 3) uses a single, unmodified image of the OS to host all VPS sessions. This approach offers several additional advantages when compared to other approaches to virtualization, including:

- **Extended platform compatibility.** Virtuozzo has the same platform compatibility as the host OS and can operate on any hardware supported by the OS. No modifications are required for applications deployed on Virtuozzo. Only one image of the OS is needed for each node, regardless of the number of VPS instances running on the node.

Figure 1. The Virtuozzo environment*

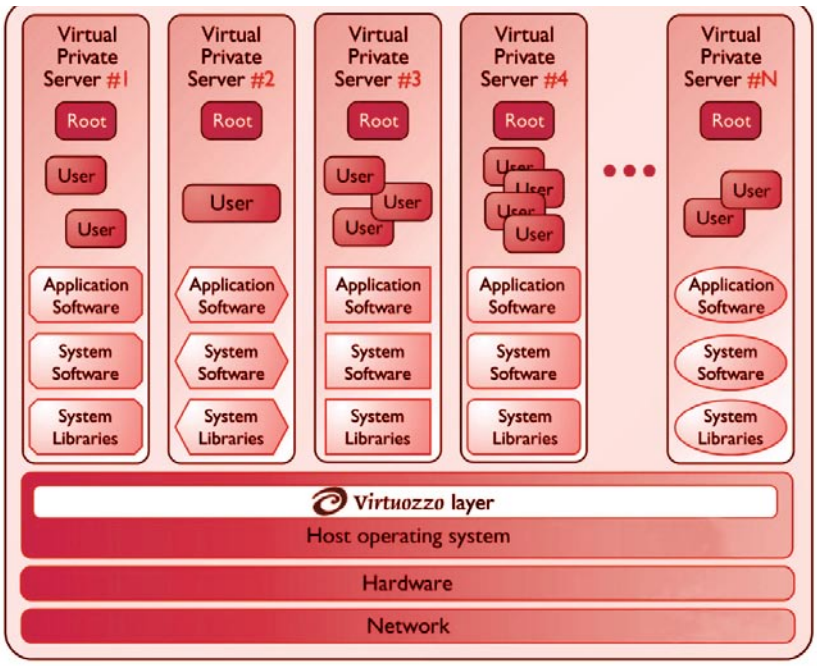

- **Rapid provisioning.** Virtuozzo instantiates a VPS session in several seconds because the underlying OS is already running. This enables rapid provisioning of underutilized server infrastructure.

Virtuozzo Adds to the Benefits of Virtualization

When used in a data center virtualization initiative, Virtuozzo provides additional benefits in the areas of processor efficiency, operational cost reductions, application availability, and dynamic administration features.

VIRTUOZZO IS HIGHLY EFFICIENT

The lightweight virtualization architecture that Virtuozzo is built upon incurs little overhead. An application running within a Virtuozzo VPS typically exhibits between 1 to 3 percent performance overhead when compared to a native, nonvirtualized application. The low performance overhead enables I/O-intensive applications such as databases to perform well in virtualized instances. This very low overhead (in essence, processor cycles spent to provide a virtualized environment) means that compute resources are devoted to application workloads, users, or additional VPS instances—not the process of virtualization.

VIRTUOZZO CAN REDUCE VIRTUALIZED SERVER OWNERSHIP COSTS

Virtuozzo can reduce costs in several ways. Since OS virtualization uses a single operating system image, software licensing costs are minimized. Virtuozzo also offers a complete set of management capabilities that enable single-console control of all virtual servers on all nodes by an administrator, eliminating the need to visit every virtual server to perform administrative tasks. This ability can reduce the amount of time spent (and thus, costs incurred) by IT staff on routine system management activities.

ZERO-DOWNTIME MIGRATION INCREASES AVAILABILITY

Situations such as routine hardware or software maintenance often necessitate moving a virtualized application from one physical server to another. Typically these types of maintenance activities or migrations are scheduled for non-work hours, but Virtuozzo provides the capability to move an available VPS to another physical server without any interruption in service or availability. The Virtuozzo Zero-Downtime Migration is unique among virtualization products because it does not require any storage area network (SAN) investment—this capability can be used on any Virtuozzo server, associated VPSs, and applications at any time.

COMPACT DATA MODEL REDUCES RESTORE TIME, PROVIDING

THREE APPROACHES TO VIRTUALIZATION

The concept of a virtual machine (VM) is not new. VMs have been available in mainframe computing environments since the late 1960s. Over time, a variety of techniques have been developed to create VM environments. The following approaches are commonly found in an Intel® processor–based environment:

Hardware virtualization provides a software interface that is identical to the underlying hardware. This requires complete hardware emulation via a layer of software (the virtual machine monitor or VMM) that runs between the operating systems and the hardware. The VMM usually features compatibility with a discrete set of hardware devices and the ability to run multiple, different operating environments simultaneously.

Paravirtualization offers a software interface to virtual machines that is similar, but not identical, to that of the underlying hardware. This requires operating systems to be explicitly ported to run on top of the VMM, but may enable the VMM itself to be simpler, and the virtual machines that run on it to achieve higher performance.

OS virtualization manages user and application access to the operating system by creating discrete, isolated VPSs. In the case of Virtuozzo, all VPSs run atop a single OS kernel. This virtualization technique does not require any hardware emulation and corresponding processor overhead, making more processor cycles available to applications and users. Device compatibility matches that of the native OS, and the same operating environment is available in each VPS.

A COST-EFFECTIVE DISASTER RECOVERY SOLUTION

Should an unanticipated event occur and render the host node unavailable, an online backup can quickly provide the source data for restoration to a functioning VPS on a different node. A Virtuozzo VPS backup contains the applications running on the VPS, application data, any files that are required by the applications, and user data. There is no need to restore the OS itself, since a functioning OS is present on the target VPS. This reduces the amount of data that needs to be transferred, and thus the time needed to perform the restoration of a VPS. A shorter recovery period means that application downtime is minimized. The use of shared storage architecture (such as a storage area network or network-attached storage devices) can be used to replicate servers and provide an always-available disaster recovery solution, although Virtuozzo does not require the use of shared storage.

MANAGEMENT TOOLS PROVIDE COMPREHENSIVE DYNAMIC CONTROL

Virtuozzo includes a full set of management tools that provide real-time creation and control of any VPS on any node running any operating system supported by Virtuozzo, all from a single console. Thus, an administrator can manage both Microsoft Windows* and Linux environments using a single tool. Provisioned VPSs can be started very quickly in a Virtuozzo environment—usually in seconds because the OS is already running.

Once a VPS is operating, the Virtuozzo Management Console (see Figure 2) allows administrators to dynamically allocate individual VPS resources (amount of VPS server memory, disk space, processor time, and so forth) without service interruption. This capability gives system managers granular, real-time resource control, resulting in increased infrastructure agility without requiring a shutdown of the VPS or the entire hardware node.

A VPS can also be copied, or cloned, in real time via the management console. Cloning a VPS can be useful in a development or a production testing environment to observe the effects of a software upgrade or update on regular server operation.

These management capabilities enable administrators to rapidly respond to business needs or opportunities, shifting resources dynamically, on a flexible infrastructure made possible by Virtuozzo.

Figure 2. Virtuozzo components*

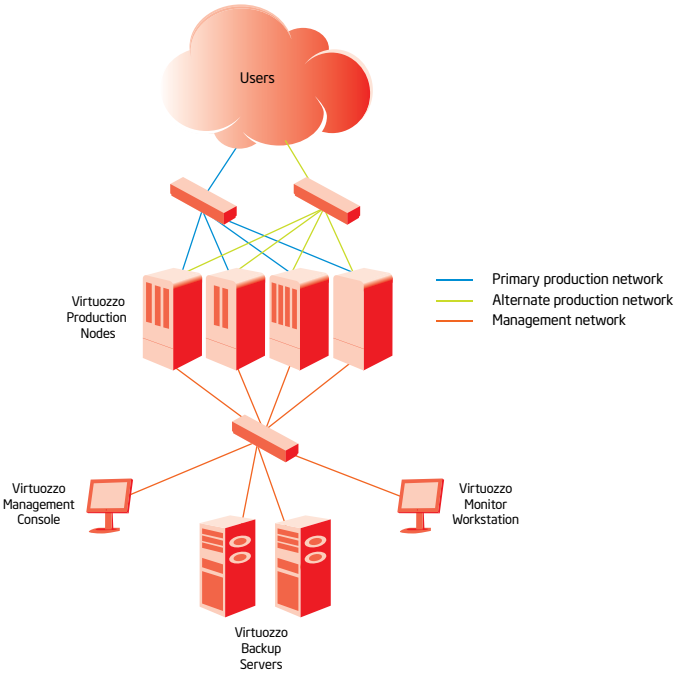

In a Virtuozzo environment, users connect to one or more VPS nodes via a primary production network. To ensure high availability, a spare Virtuozzo server is available for VPS migration, and an alternate production network path exists to each node. The Virtuozzo Management Console and the Virtuozzo Monitor Workstation are attached to the VPS nodes via the management network, which also connects to the Virtuozzo backup servers.

A Virtualized Environment Requires a Reliable, High-Capacity Network

To successfully consolidate server workloads onto a virtualized environment, it is essential that all server subsystems (processor, memory, disk, and I/O) can accommodate the additional workload. While Virtuozzo requires a single network connection to operate, careful attention to and planning of the networking infrastructure of a Virtuozzo environment can ensure both optimal performance and high availability.

MULTIPLE VPSS INCREASE NETWORK TRAFFIC

With multiple workloads, the network capacity needs to scale to match the requirements of the combined workloads expected on the node. In general, as long as

the node's processor is not fully utilized, the consolidated network traffic will be the sum of the traffic generated by each application. For example, if five physical servers using an average of 30 Mbps of bandwidth and running at 15 percent utilization are consolidated, then the Virtuozzo host machine running the five VPSs requires roughly 150 Mbps of bandwidth. Therefore, a network connection running at 100 Mbps will not adequately service such a configuration, and a higher-capacity network interface is necessary.

VERY HIGH BANDWIDTH IS DESIRABLE DURING MIGRATIONS AND RESTORATIONS

While a VPS can be created and started in minutes, any migration or restoration operation requires time to transfer all of the data from the source to the target. Network throughput is directly related to the time it takes to copy the data to a new VPS on a different hardware node. VPSs can have gigabytes of data to transfer, and the faster this data can be copied over the network, the more rapidly the operation can be completed—especially in the case of a migration or a restoration.

HOST NODE RELIABILITY IS CRITICAL

With multiple virtual machines running on a single node, the reliability of that node becomes much more important. Likewise, redundancy and fault tolerance are critically important for network infrastructure. The failure of a connection supporting a node could affect every user of every VPS running on the node—a costly interruption that is time-consuming to fix. Providing an alternate network path through a meshed switch fabric and using multiport adapters that support failover can mitigate that risk and improve overall reliability and availability.

All of these considerations point to designing a virtualized infrastructure's networking subsystems with the capacity and availability features delivered by multiple high-capacity network connections. This can be challenging if servers are peripheral component interconnect (PCI) slot-constrained. Small 1U servers acting as host nodes have limited expansion capability. Incorporating multiple network interfaces helps make such servers more reliable and available, as well as better able to handle the increased traffic.

Intel® Network Adapters Offer Performance, Flexibility, and Reliability for VPS Deployments

The Intel® PRO/1000 family of network adapters and the Intel® PRO/10GbE 10-gigabit adapters meet the performance, flexibility, and reliability needs of host nodes in connection-dense Virtuozzo environments. In particular, the dual- or quad-port Intel PRO/1000 multiport adapters provide the needed port density for use in slot-constrained server platforms.

PERFORMANCE

The Intel® PRO/10GbE SR and Intel® PRO/10GbE LR Server Adapters offer exceptional throughput and some of the fastest Ethernet connectivity available. Both adapters feature jumbo frame support of up to 16 KB, which reduces packet overhead for increased performance. To reduce processor involvement with the processing of network traffic, each adapter has mechanisms to delay and reduce transmit and receive frame interrupts, as well as TCP/IP segmentation and TCP checksum off-load. Intel PRO/10GbE network adapters offer the throughput necessary for VPS restorations or migrations—when time is of critical importance and very high throughput is a requirement.

To meet the needs for multigigabit traffic, the Ethernet ports on the Intel PRO/1000 family of multiport adapters can be aggregated to deliver greater total bandwidth. Server performance can also be enhanced by teaming connections on adapters with connections on other Intel® PRO Server Adapters or with LAN on Motherboard components to achieve the desired scalability. Throughput counts during a VPS restore operation; the faster data can be transferred to the new VPS, the sooner the VPS can be placed in operation.

At very high data rates, network traffic can take a toll on server processor cycles. All Intel PRO/1000 Server Adapters feature interrupt moderation, which significantly reduces processor utilization and makes more processor resources available to applications.

To increase performance in a multiprocessor environment, the Intel® PRO/1000 PT Server Adapter can balance network loads across processor cores when used with Receive-Side Scaling from Microsoft or Scalable I/O on Linux, environments both supported by Virtuozzo.

FLEXIBILITY

To support unpredictable network demands, such as heavy traffic on Web servers and intranets, IT departments are increasingly segmenting enterprise networks. Deploying Intel® multiport network adapters can provide the platform flexibility to respond to evolving needs for segmentation and enable the connection to multiple switches to segment traffic on a network.

Segmentation can be employed effectively to optimize operation in a Virtuozzo-based network. A single segment can be used for production traffic. To remove single points of failure and thus increase network reliability, a second network segment with its own second, meshed network switch can provide an alternate path for normal production traffic. A third segment for VPS management traffic, as well as backup/restore operations, keeps large traffic volumes from impacting production workloads.

Migration from legacy 10 Mbps and 100 Mbps to high-speed Gigabit Ethernet networks can also be achieved easily and cost-effectively using Intel multiport adapters. Intel PRO/1000 multiport adapters all support 10/100/1000 Mbps speeds and are autosensing. This capability allows for easy, nondisruptive integration into current networks while providing for network capacity growth in the future.

For additional platform flexibility, Intel PRO/1000 multiport adapters are available for both copper and fiber optic networks. Both use a common driver technology—Intel® SingleDriver™ technology—which helps reduce IT complexity, training, and support costs.

With a flexible design that fits almost any type of PCI bus, dual- and quad-port Intel adapters include standards-based management features and wide network OS support to help ensure extensive compatibility with the latest server and networking environments. Intel adapters are compatible with full-height and low-profile PCI slots: the standard, full-height bracket can easily be replaced by the shorter, low-profile bracket for installation in high-density servers that feature low-profile PCI slots.

The Intel PRO/10GbE adapter family is the answer to high-throughput network demands because it is standards-based Ethernet. This makes it simple to integrate into existing networks and easy to manage with existing tools. The Intel PRO/10GbE LR adapter provides 10 Gigabit Ethernet connectivity to a distance of 10 kilometers over single-mode fiber, making it ideal for use in a campus setting, where the data center might be located in a separate facility. For distances of up to 300 meters, the Intel PRO/10GbE SR adapter offers 10 Gigabit Ethernet running over multimode fiber infrastructure.

RELIABILITY

In a virtualized server environment, component redundancy is a key method for providing high availability. Intel® Advanced Network Services (ANS) and Intel® PROSet Utility software allow administrators to create redundant teamed connections on all Intel PRO network adapters. Intel multiport adapters can be configured to automatically switch to a secondary link when a server's primary link fails. Switch fault tolerance and test switch configuration features help test compatibility and further increase uptime. Advanced cable diagnostics dynamically test and report network problems such as interrupts, error rates, and excessive cable lengths and automatically compensate for cable issues such as a cross-over cable or wrong pin-out/polarity configurations.

Intel multiport network adapters ensure server availability by helping avoid the cost of taking a server offline to install additional network connections when moving to a virtualized environment. Intel multiport adapters also can enable business continuity efforts by allowing IT administrators to dedicate network ports for two key components of Virtuozzo—remote storage and management.

While Virtuozzo provides comprehensive remote management capabilities for fine-grained VPS control, dual- and quad-port Intel adapters support remote hardware management standards to further reduce administrative costs. These include Wired for Management (WfM), Microsoft Remote Installation Service (RIS), Simple Network Management Protocol (SNMP), and the Desktop Management Interface (DMI).

All Intel PRO adapters come with a limited lifetime warranty.

Intel® PRO/1000 Multiport and Intel® PRO/10GbE Adapter Family Features and Benefits					
Feature	**Benefit**	**Intel® PRO/1000 PT (Dual) §**	**Intel® PRO/1000 MT (Quad)**	**Intel® PRO/10GbE SR (Single)**	**Intel® PRO/10GbE LR (Single)**
PCI Express* architecture	Enhances performance and reliability while retaining software compatibility with existing PCI infrastructure	✓			
PCI-X architecture	Established parallel bus architecture		✓	✓	✓
Intel® 82571EB Gigabit controller	Enables two gigabit connections in a single adapter, delivering increased bandwidth for slot-constrained servers and providing high performance, reliability, and low power use in a single, integrated, dual-port PCI Express Gigabit Ethernet controller chip	✓			
Two Intel® 82546EB Gigabit controllers	Enables four gigabit connections in a single adapter, delivering increased bandwidth for slot-constrained servers and providing high performance, reliability, and low-power use in a single, integrated, dual-port Gigabit Ethernet controller chip		✓		
Intel® 82597EX 10GbE PCI-X 133 MHz/ 64-bit MAC controller	High performance on widely adopted second-generation PCI host-bus interface	✓	✓		
Intel® SingleDriver™ technology	Common driver suite for Gigabit Ethernet reduces installation complexity	✓	✓	✓	✓
Load balancing on multiple processors	Increases performance on multiprocessor systems by efficiently balancing network loads across processor cores when used with Receive-Side Scaling from Microsoft or Scalable I/O on Linux*	✓			
Interrupt moderation	Delivers increased performance while significantly reducing processor utilization	✓			
Compatible with x4, x8, and x16 full-height and low-profile PCI Express slots	Allows dual-port operation in almost any PCI Express server slot, except x1 slots, and allows each port to operate without interfering with the other	✓			
Support for most network operating systems	Enables widespread deployment	✓	✓	✓	✓
Remote management support (WfM, RIS, SNMP/DMI)	Reduces support costs with remote management based on industry-wide standards	✓	✓	✓	✓
Category-5 unshielded twisted pair 4-pair cabling	Uses existing 4-pair cabling and saves rewiring costs	✓	✓		
Intel® PROSet utility for Microsoft Device Manager*	Provides point-and-click power over individual adapters, advanced adapter features, connection teaming, and virtual local area network configuration	✓	✓	✓	✓
RoHS compliant, lead free	Compliant with the new European Union directive (effective July 2006) to reduce the use of hazardous materials	✓			
Advanced cable diagnostics	Dynamically tests and reports network problems (error rate, cable length) and automatically compensates for cable issues (cross-over cable, wrong pin-out/polarity)	✓	✓		
Intel backing	Backed by an Intel limited lifetime warranty, 90-day money-back guarantee (U.S. and Canada) and worldwide support	✓	✓	✓	✓
Backward compatibility with prior generations of Intel® PRO/1000 adapters	Reduces support cost and simplifies installation and maintenance		✓		
Full-height and low-profile screw-on brackets included with each adapter	Compatible with full-height and low-profile PCI standard. Full-height bracket can be easily swapped out and replaced by the shorter, low-profile bracket, for installation in high-density 1U and 2U servers that feature low-profile PCI		✓		
Large send off-load/TCP segmentation off-load	Compatible with large send off-load feature of Microsoft Windows Server* 2003, 2000, and Windows* XP, which enables wire-speed performance with low processor utilization	✓	✓	✓	✓
Expanded advanced server features	Intel® ANS software increases uptime with redundant, teamed, connections and scales bandwidth by balancing network traffic across multiple server connections. Now offers switch fault tolerance and test switch configuration to test compatibility and further increase uptime	✓			
Compatible with Fast Ethernet	Reduces training costs and time to rollout	✓	✓		
Autosensing, self-configuring 10/100/1000 Mbps performance	Enables migration to faster networks—adapter automatically senses a faster speed and self-configures	✓	✓		
802.3 compatibility	Supports Gigabit Ethernet industry-wide networking standard	✓	✓	✓	✓
Distance		100 m	100 m	300 m	10 Km
§Also available for PCI-X slot architecture as the Intel® PRO/1000 MT Dual-Port Adapter					

Make the Connection to a Virtualized Environment

Virtuozzo software running on Intel processor–based platforms connected via Intel network adapters provides a compelling virtualization platform. Together, these products offer a way to help IT managers help their businesses do more with less and respond to initiatives faster with a flexible, cost-effective, and efficient IT infrastructure.

Learn More about These Virtualization Products

For more information about the products described in this white paper, visit www.virtuozzo.com or www.intel.com/network/connectivity.

ABOUT SWSOFT

SWsoft is the recognized leader in OS server automation and virtualization software. The SWsoft suite of products delivers proven performance, reliability, manageability, and value. Headquartered in Herndon, Va., with offices around the world, SWsoft is a privately held, high-growth company. For more information, visit www.virtuozzo.com or call +1 (703) 815-5670.

ABOUT INTEL

Intel, the world leader in silicon innovation, develops technologies, products and initiatives to continually advance how people work and live. Additional information about Intel is available at www.intel.com.

311840-001 US

0306/BY/TDA/PDF