Speeding Up the Network: A System Problem, A Platform Solution

Pat Gelsinger
Senior Vice President
General Manager
Digital Enterprise Group

Hans G. Geyer
Vice President
General Manager
Networking and Storage Group
Intel Corporation

Justin Rattner
Intel Senior Fellow
Director
Corporate Technology Group
Intel Corporation

Table of Contents

(Click on page number to jump to sections)

SPEEDING UP THE NETWORK: A SYSTEM PROBLEM, A PLATFORM SOLUTION 3
- Overview: Updating the Outdated 3
- Today's Challenge 4
- Launching the Research 4
- First Try: A Traditional Approach 5
- Fundamental Questions 5
- In-Depth Analysis 5
- A Modern Myth 6
- Profound Implications 6
- Islands of Innovation 6
- Proving a New Solution 7
- Correcting the Vision 7
- From Research Lab to World Stage 8
- Small Changes, Big Gains 8
- A Balanced Approach 8
- Research—a Critical Advantage 9
- Common Vision and Collaboration 9
- Intel I/OAT—An Excellent Example 9
- Mushrooming Solutions 9
- In the Future 10
- Industry Support 10
- Summary 10
- More Info 11
- Author Bios 11

Speeding Up the Network: A System Problem, A Platform Solution

Pat Gelsinger
Senior Vice President
General Manager
Digital Enterprise Group

Hans G. Geyer
Vice President
General Manager
Networking and Storage Group
Intel Corporation

Justin Rattner
Intel Senior Fellow
Director
Corporate Technology Group
Intel Corporation

Overview: Updating the Outdated

There's something in your network that's very, very old. It was in existence before Sun Microsystems, Microsoft, and Compaq. Your modern, high-speed, large-bandwidth, gigabit Ethernet network is held up by this decrepit element—something that was designed before the Apple Macintosh* computer, before the Radio Shack TRS-80*, before the Commodore 64*.

Half of it actually came into being before Intel's first microprocessor, the amazing 4004, which ushered in the world of, yes, 4-bit processing (Intel supports 64-bit processing today). It even predates the world's first portable computer—the IBM Model 5100*, which weighed in at a whopping 55 pounds. This sagging element, this suite that handles your network functions, was actually born the same year the first Cray* I supercomputer went to work in 1974.

What is this shockingly aged creature that lingers on, flagging but not quite failing, as it struggles to handle communications at speeds that were not even dreamed of when it was first created, for architectures that were not yet imagined?

It is your TCP/IP protocol.

The Transmission Control Protocol/Internet Protocol (TCP/IP) is actually a suite of protocols that have been upgraded continuously since their conception between 1969 and 1973. However, the core of the protocol has remained essentially unchanged for more than a quarter century. As amazingly functional as the protocol has been, it simply cannot handle today's high-speed, high-performance networks without having to make serious changes to server and platform architectures.

Intel is now announcing Intel® I/O Acceleration Technology (Intel® I/OAT). This collection of technologies is a comprehensive solution that implements critical changes across the server platform. This includes changes to the processor, chipset, motherboard, LAN (local area network), protocol software, device-interface software, and operating system (OS).

Implemented from wire to application, Intel I/O Acceleration Technology dramatically speeds up server-network communications by optimizing the TCP/IP stack, reducing latency, redistributing functions to components that will do each type of work best, and using additional enhancements to improve server system throughput from component to application. With Intel I/OAT, servers can finally begin to take advantage of the high-speed, high-performance bandwidth that is available today.

Today's Challenge

Networks and servers are beset by a range of problems. Users consistently consume as much bandwidth as is offered them, and then demand even more. The data being processed today (such as video and audio) is increasingly both bandwidth- and CPU-intensive. The amount of data being transferred has also grown at an exponential rate, until we are now entering the "tera era," in which data is no longer being measured by gigabyte, but by exabyte (a million terabytes). Finally, the trend toward sending and receiving data in smaller, not larger, packets means that more and more server platform time is spent processing the TCP/IP header for each packet, relative to delivering each packet payload.

For two decades, improvements in server platform performance matched the demands of network trends. For example, the shift from 10-megabit (Mb) Ethernet to 100-Mb Ethernet took long enough that Intel processors and supporting chipsets easily kept up with the rate of network improvement.

Recently however, in just a few years, the industry moved to gigabit Ethernet (GbE) speeds, and is even now shifting to 10 GbE. While the rates of improvement for processor performance have kept pace with network throughput, the performance of the memory subsystem is no longer keeping up with the rest of the server system. Today's Intel® processors perform at significantly higher speeds than memory. This means that latency (delay) increases as processors wait and wait through hundreds of idle cycles while memory works at its slower read, write, and copy tasks.

Disparity results in latency, latency causes inefficiency, and inefficiency means slow going. Ironically, as processor speeds increase, latency actually worsens.

Yet servers and networks are the foundation upon which we depend for the increasing use of Web services, the greater adoption of IP addresses for consumer goods, the introduction of sensor networks and RFID (radio-frequency ID) tags, and an increasing momentum to use IP storage. Since nearly all of the data required for those applications moves between systems in the form of network packets, network bottlenecks are approaching critical levels.

"The problem is there," said Hans Geyer, vice president and general manager of Intel's Networking and Storage Group. "Packet processing has to get accelerated, because the servers are already bogged down. What we need is to find the best approach to solve that fundamental problem."

Launching the Research

The driving force that launched Intel's current avenue of research and development was the simple fact of that disparity: network bandwidth performance has been improving at a higher rate than Moore's Law. That disparity, caused by the faster arrival of packets from the network, is beginning to impact the performance of servers that handle the very transactions coursing over today's increased bandwidth.

With that realization, small teams of Intel researchers began to look at ways to address that disparity. For example, one of the technologies that was closely examined was TCP/IP offloading. Another technology was RDMA (remote direct memory access).

The RDMA protocol allows one computer to transfer data directly into a specific destination memory on another computer. It minimizes demands on bandwidth, reduces the overhead for processing packets, and reduces latency, but it has two significant drawbacks. First, it requires modifications to the protocol stack. Second, it introduces significant management issues for data-center operations.

The other technology, TCP/IP offload engines (TOEs), are the conventional industry method currently being used to try to address network I/O bottlenecks. An offload engine works by taking some of the burden of processing data packets and "offloading" that work to another processor. This supposedly frees up the main processor (CPU) to do other tasks. If TOEs worked as promised, they should deliver a significant increase in server network I/O efficiency.

First Try: A Traditional Approach

At first, Intel took the traditional approach to the problem of sluggish server performance and network I/O bottlenecks. Intel researchers assumed, as the industry had, that conventional TOE wisdom was right: asking the CPU to run the TCP/IP protocols was an inefficient use of this valuable resource. Researchers knew up front that offloading was not likely to offer a dramatic improvement over Moore's Law, but they thought an enhanced TOE design could still relieve some of the load on the servers.

"I, myself, certainly suffered from that assumption for a number of years," explained Justin Rattner, Intel Senior Fellow and director of the Corporate Technology Group, which contains most of Intel's research labs. "So when we began this project, we looked at a variety of offload solutions. We developed some internal experimental offload processors, and we worked with a number of external vendors of offload processors. We even explored the use of Intel's IXP series of network processors as offload engines."

Intel researchers were piqued by a basic question: With all the speed in today's processors and gigabit Ethernet, why hasn't TOE or any other industry solution provided the expected increase in performance? Industry has the processors to handle the traffic. They have the network pipes—the bandwidth—to handle the traffic, and they have the operating systems to handle the traffic. So why isn't the performance of server applications reaping the benefits of today's greater bandwidth?

Researchers began to suspect that some of the approaches proposed by industry, such as offloading, were interesting but would not accomplish their objectives. They also suspected that the traditional approaches would not scale well as traffic volume and network bandwidth increased.

Fundamental Questions

Unable to verify the improvements they expected, Intel researchers stepped back and began to examine the problem differently. They set aside the component approach and began to look at the problem from a broader, more-fundamental point of view to determine exactly what was happening in the system.

While some researchers continued to work on the offloading problem in this context, others wanted to see if they could improve the way network packets were handled. They began a detailed investigation into each step required in a system to move a TCP/IP packet from wire to the point at which the packet payload was available to the application.

Intel began by asking basic questions: How exactly were packets moved from the network to the application, from one platform subsystem to another, from one component to another, and from one memory location to another? Where was the time being spent? The investigation began with a broad study of IP traffic in terms of storage traffic, Web server traffic, connection management traffic, database traffic, and so on.

In-Depth Analysis

The first step in the investigation was to perform in-depth trace analysis to understand exactly where time was being spent to process each packet. This included complete systems modeling and simulation to find out where the bytes were going, exactly how many CPU cycles were needed to pack and unpack the packets, how many bus cycles were being used to transfer words in and out, and so forth.

Researchers then looked at the system architecture to see exactly where those periods of greatest delay—greatest latency—were occurring. They soon learned that the key to any real solution would be the ability to deeply and fundamentally understand the full complexity of how IP traffic operated inside a platform.

There was internal competition, innovation, and more than one team working to discover or prove the better technical approach to the answer. What they learned was more than surprising. It not only changed the researchers' viewpoints of packet processing, but it completely changed the direction of Intel's product development with regards to network acceleration.

A Modern Myth

When Intel began this project, the research teams were not trying to change Intel's direction or deliver a new technology, but simply to get a very deep understanding of input-output (I/O) workloads and IP traffic.

At the time, industry thought that the most difficult issue to resolve was header management for IP packets. What researchers had expected—and what conventional thought and the offloading approach had perpetuated—was the assumption that most of the inefficiencies in server network I/O (input/output) were caused during the processing of that base TCP/IP protocol. This meant that an offload approach should improve system efficiency significantly.

Instead, what researchers found was an eye-opener. The main problem with throughput of packet processing cannot be solved by offloading, because the main problem is not protocol processing inside the CPU. The actual, underlying problem of slow throughput is latency caused by data movement. Specifically, the worst latencies are caused by the movement of data within memory, and by the movement of data between system memory and the network interface card (NIC) buffers.

"Once the data came in, we were all—even those who were passionate that offloading was a better approach—surprised at how poor that approach really was," admitted Rattner. "None of us had realized there was such a large gap between what offloading promised and what was realized. It was clear that offloading was a poor solution for customers."

Profound Implications

The more researchers looked at their results, the more they realized that no one component could be identified as an area to enhance by itself to improve server performance. Instead, improvements would be required across the platform to address the fundamental issues of slow processing of TCP/IP packets.

This meant changes, not just to the CPU, but to the chipset, LAN, protocol software, device software interfaces, and the operating system. It was believed that such an approach would yield dramatic improvements in processing all sizes of TCP/IP packets, which are the bulk of today's Internet traffic. It was also believed that this approach would benefit a host of other network protocols.

The change in direction was clear: a platform approach was the right way to deliver any real kind of solution for network I/O acceleration for servers, not only for Intel, but for customers now and in the long run.

"I'm not easily swayed to the latest fad," said Pat Gelsinger, senior vice president and general manager of Intel's new Digital Enterprise Group (DEG). "It's easy to view this shift toward a platform solution as just the latest fad, but it's actually not a fad at all. There are solid technical, architectural underpinnings to this strategic direction. This is a great example of how in-depth research, new capabilities, and well-thought-out answers to serious problems can drive a total platform solution."

Islands of Innovation

The team performed more in-depth research to confirm the surprising results that pointed them toward a platform solution, in particular because of a nagging question. In terms of performance, offload engines run at much lower clock speeds than Intel processors do today. What was it about the offload approach, they asked, that allowed those relatively slow processors to deliver so much performance against an Intel processor with its much higher clock speeds and much higher instruction performance—thus justifying the latency myth? They wanted to know if they could harness that capability into a core Intel platform.

That type of question is part of a classic cycle of discrete innovation, core integration of those new capabilities into the platform, discrete innovation, core integration of capabilities, and so on.

"Like so many other architecture innovations over the years, specialized processors get built to address a particular performance problem. They're like islands of enhancements," explained Rattner. "Over time, as we come to fully understand how they are able to deliver that performance improvement, we can go back to the base platform and incorporate those ideas and innovations at a much lower cost, typically with performance that is as good as or better than the discrete component, and without any need for specialized hardware or software."

Continuing their examination of different innovation approaches, Intel researchers identified a number of technologies that could improve packet processing if they were implemented into the server platform. All of these technologies would allow the core platform to meet the demands of a very high-performance network without having to call on an external, specialized coprocessor to handle the workload. In a neat twist, offloading went from being an interesting approach, to driving the in-depth research toward a much better solution.

“It was as a direct result of building on that foundation, that in-depth knowledge of offloading, onloading, and IP traffic, and being able to look at the problem from a total platform point of view, that we were able to identify these enhancements and technologies that we're calling Intel I/O Acceleration Technology,” explained Rattner.

Proving a New Solution

Once the direction of the results was realized, researchers were able to start focusing on potential solutions.

Several important enhancements grew directly out of the lab’s investigations. For example, researchers had discovered that offload engines performed better than Intel processors, not because TOEs addressed a root problem, but because Intel processors were using a TCP/IP protocol stack based on algorithms and code that were over 20 years old. The inefficiencies inherent in that outdated stack consumed CPU cycles that could be better used elsewhere.

Once the stack was upgraded, streamlined, and optimized for modern computer architecture, the traditional approach of using an offload engine offered little real benefit for I/O acceleration. That new, optimized protocol stack is now one of the primary features of Intel I/O Acceleration Technology.

“Our lab people provided convincing evidence that you can run the protocols efficiently on the main processor,” said Rattner. “They proved that you can do it at competitive levels, and that you can meet the demands of today’s fastest networks. Any time you go against conventional wisdom in computer design and win—that’s very exciting.”

Offloading engines would still be useful in the back-end server and database server market segments, which are approximately two to five percent of the server market as a whole. But the place where offloading really belonged was in storage acceleration, not in network I/O acceleration. With this realization, the platform solution was looking better and better.

"Our researchers hadn't started down the platform path intentionally,” explained Gelsinger. “They weren't looking or even wanting to go down that path. It was their research that led them to that answer."

Correcting the Vision

With the platform approach confirmed by more and more research, the team now had to convince the internal Intel business units that there was a better way to improve server performance than simply by improving clock cycles or processor speed. They also had to convince business units that they could deliver this technology in a timely manner.

As the offload engines already in the marketplace proved unable to deliver on their promised performance, the business groups began taking heed. Fully convinced by Intel’s exhaustive research, they realized that their own proposed TOE offerings were simply not going to be as competitive as the alternative acceleration technology now being proposed by the research teams. They could see that introducing yet another, improved offload engine would still not satisfy the market segment requirement nor resolve the underlying problem.

"The more we got into designing our own offload engine, the more we realized that it was the wrong approach,” explained Geyer. “It cost too much money, it wouldn't scale well with the number of connections that needed to be handled, and it didn't take advantage of the power already available in our processors. A platform solution scales better, it's cheaper, and it uses the ever-increasing performance of our microprocessor set for network I/O.”

From Research Lab to World Stage

As the research became more widely known inside the company, the team—led by Don Newell and Alan Crouch—that looked into Intel I/O acceleration technology became experts on network acceleration. They suddenly found themselves pulled out of the research labs, put on the road to customer sites, and convincing the industry of a new, provocative and comprehensive approach to a critical-mass server performance problem. What had started with a simple question had completely shifted Intel's vision—and the industry's—of packet processing for server network I/O.

"This was a cross-group effort that began with creative dissonance, and ended with phenomenal collaboration between the labs and the business groups," said Gelsinger.

Small Changes, Big Gains

Surprisingly, researchers and developers found that a great deal of new logic or software was not required to implement the acceleration technologies. A series of relatively minor, but important changes could be made in existing platforms to enable the I/O acceleration solution. This included the rearchitecture of the most critical workload element that servers operate on today: the IP stack.

"We just needed to change little pieces in key places inside the platform, in the network controller, in the chipset, and in the CPU," explained Gelsinger. "It was still a bit tumultuous," he admitted. "Critical CPU, chipset, and network component projects were already in-flight and running down the course. Those projects were on schedule and working hard. They were doing just fine. They weren't asking for more features from the lab."

However, the result has been everything hoped for. The combined benefits of these small, critical changes offer a dramatic improvement in the overall performance of the platform. Expectations are that, with Intel's protocol acceleration technologies, the delivery of packet data to and from server applications can be improved by 30 to 50 percent, and possibly by as much as a factor of 10.

"This notion of protocol acceleration—which is the basic objective here—will produce as many, if not more benefits to storage," pointed out Rattner. "Storage, and in particular, network-attached storage, is also increasingly based on the TCP/IP protocol. As we improve protocol performance and enable IP-based storage, we'll also be opening the door for greater use of that IP storage."

A Balanced Approach

Business units across Intel began to distribute the acceleration innovations to the development areas and product lines where the enhancements would be best implemented.

For example, when it comes to packing and unpacking data, the traditional approach is offloading again, which requires a separate microprocessor or controller, along with the associated separate memory. It's a solution that adds $20 or $30 in unnecessary components to each system. In contrast, Intel has incorporated that capability—which still requires fast processing—into a software routine that runs on the standard Intel CPU.

"Since Intel's CPUs continue to get faster and more powerful, either because of MHz or additional cores, that additional performance can be applied to some of the packet processing costs," explained Geyer. "The CPU is in the system anyway, so that work is done for free."

For example, tasks that have nothing to do with the CPU—such as moving blocks of memory or moving blocks between memory and the network-controller buffers—those tasks are best done in a DMA (direct memory access) controller that sits in the motherboard chipset. Other simple processing tasks, such as network-oriented processing, have been shifted to the Ethernet controller. The chipset (via DMA) now handles the data transfers between memory and memory or between memory and the network controller.

"It is a much more balanced approach in the system," said Geyer. "Tasks are assigned where they can be done most efficiently, with the least amount of effort—meaning the least amount of cost, both in terms of processing time and component cost."

Research—a Critical Advantage

Intel is a large company, but size does not mean slow to move, change direction, or innovate. Intel's size actually gives the company two important advantages over competitors when it comes to delivering real solutions.

First, Intel is able to assign significant resources to the research labs. Rather than support research simply to improve CPU or chipset design, Intel has both the people and the resources to conduct in-depth investigations into real problems. Whereas some competitors use research to justify a technology, Intel can use research to find out why a technology does—or doesn't—work, and whether it is justified at all for today or in the future. Intriguing research—external, not just internal—can be supported or encouraged, innovation from other areas and industries can be examined, detailed experiments conducted, assumptions corrected and myths debunked, and alternative approaches studied.

"We've spent a lot of time building up our analysis modeling, our simulation, and our prototyping capabilities," said Rattner. "Not just for CPU clock speed or chipset performance, but across the platform, and at in-depth levels, such as investigating the complete paths of the I/O space. As a result, we're able to perform these very deep analyses, identify the fundamental problems, and then create their solutions. That's pretty impressive to me."

Common Vision and Collaboration

Intel's in-depth research makes it clear that no one technology can solve the packet processing bottleneck. What is required is collaboration among manufacturers of all the various platform components.

Here is where the second benefit to Intel's extensive resources becomes apparent. Intel is not a single-component company. Intel already researches, tests, designs, and produces all components of a platform: CPU, chipset, motherboard, and LAN silicon. Intel writes drivers, provides open-source code to developer communities, and has strong alliances in industry with other component, operating system, and application vendors.

Other, more limited manufacturers (such as a server manufacturer, a processor manufacturer, or an I/O component manufacturer) are able to pursue only one piece of a technology solution. These manufacturers do not have the resources or cross-technology expertise to explore and develop a holistic, platform solution. Instead, they must enter into partnerships to attempt to provide an end-to-end solution. That's a process hampered by political and technological bumps and bruises, as each company tries to work with outside parties that have different business interests. With each member in a partnership vying for optimum benefit, it's difficult to create a balanced solution.

In contrast, with the resources to work with other organizations, such as Microsoft and the Linux* community, and the ability to direct internal groups to work together to resolve complex problems, Intel can offer customers a common vision and truly comprehensive solutions.

Intel I/OAT—An Excellent Example

Intel I/O Acceleration Technology is a perfect example of this cross-technology approach. Three major Intel corporate entities were involved in this project: the Networking and Storage Group (NSG), the Corporate Technology Group (CTG), and the Digital Enterprise Group (DEG).

"This type of platform approach is something only a company like Intel can do," said DEG general manager and senior vice president Gelsinger. "Our presence at multiple places in the platform allows us that vision, and allows for this type of total solution, from front to back, from application to wire."

Mushrooming Solutions

Intel I/O Acceleration Technology is not the only major new technology being offered by Intel. Intel is also exploring or offering technologies that will enable or enhance remote management of systems, as well as security and virtualization:

- Intel® Active Management Technology
- Security and encryption technologies
- Virtualization technologies

"It's as if we breached a threshold when we realized the potentials of a platform solution," said Gelsinger. "All of a sudden, major new technologies are popping up like mushrooms. Each time we start to explore one of those areas, we find golden opportunities to develop or add even more new capabilities to our platform to solve all sorts of other industry problems."

Rattner agreed, saying, “Once you have this framework, this concept of an end-to-end platform solution, there are a whole variety of I/O enhancements and optimizations that suggest themselves for networking, and not just for today, but for the future.”

Intel’s next step is developing the Intel I/O Acceleration Technology to its full breadth. That work includes exploring and developing capabilities such as network stack affinity/partitioning, direct cache access (DCA), lightweight (software) threading, and queue-based interface.

"One of the greatest architectural challenges we now face,” admitted Gelsinger, “is making sure that all these new technologies work together in a very strong, consistent, complementary, and architecturally cohesive way."

With Intel now focused on a platform approach, the collaboration needed between business units and researchers to provide that cohesion will be even easier.

In the Future

Much of Intel’s research into future Intel® Architecture (IA) technologies will also enhance network I/O acceleration capabilities. For example, some of the technologies to come, such as multi-core and multithread architectures, allow multiple applications to run on multiple processor cores. This generates the need for better and faster I/O. Fortunately, the very availability of those multi-core and multithread architectures will allow Intel to build an even more efficient protocol service.

For example, one of the top items on Intel's laundry list of 10 or 15 technologies to evaluate is iSCSI, a networking standard that may offer significant improvements for data transfers to and from storage facilities.

Industry Support

Intel’s research has already been well received in the industry. For example, after seeing Intel's extensive research, Microsoft quickly recognized that the platform approach would allow their customers to see significantly better application performance. It would also allow Microsoft customers to scale servers with a solution that can take advantage of multi-core processors. Microsoft is now working closely with Intel to develop Intel I/OAT support for the Windows* operating system.

“Microsoft needs to be able to offer customers operating-system support coupled with processors that, together, deliver data faster to applications. Windows* Server and Intel I/O Acceleration Technology are a great solution to this challenge,” said Jawad Khaki, Corporate Vice President of Microsoft.

Intel has also engaged the Linux community on the thinking behind the implementation of this platform solution.

“The Linux community has been looking for a better solution than TOE for some time for networking,” said Geyer. “They are very interested in supporting our new acceleration concepts.”

As a result, Intel will be providing open-source drivers for Linux that can be incorporated into the various Linux implementations.

Summary

Intel I/O Acceleration Technology takes advantage of Intel's expertise in LAN silicon, chipsets, processors, network controllers, software, and even operating systems. With expertise in all aspects of server network technologies, Intel has been able to create a solution that offers dramatic benefits as a whole, more so than any of the individual pieces could offer by themselves.

“We have looked at the best ways to accelerate the network I/O with the least hardware cost,” said Geyer. “We have looked at every aspect of the system, from I/O paths to CPU-intensive tasks, to data transfers, memory and the Ethernet controller. What I find most exciting is that we're taking a system problem and solving it on a system level, instead of with dedicated, separate components. This is the right implementation."

It’s also a comprehensive implementation that encompasses all aspects of industry. Because both operating system and applications manufacturers are already onboard with this approach, acceleration technology will be able to be delivered, not as a partial implementation, one component at a time, but as a holistic solution. Intel’s I/O Acceleration Technology is a perfect example of how a platform vision can solve a pervasive industry problem.

"When you get people who are passionate, who are smart, and then who deeply understand the right problem, great things happen," said Gelsinger.

With new acceleration technologies, Intel is improving the ability of servers to process packets and take advantage of the large bandwidths that are available today. This is a comprehensive approach that speeds up server I/O throughput, scales for future technologies, is easier to implement, and costs less than other, piecemeal solutions. With Intel I/O Acceleration Technology, data centers will now have a powerful, platform-wide solution for a system-wide problem, using all-Intel technology, and offering dramatic improvements to the data centers of today and tomorrow.

More Info

You can learn much more about Intel I/O Acceleration Technology on the Intel Web site, where in addition to general information you’ll find:

- An animated demo
- A technology brief
- A white paper

Author Bios

Pat Gelsinger is senior vice president and general manager of Intel's Digital Enterprise Group. Gelsinger joined Intel in 1979, and has more than 20 years of experience in general management and product development positions. He holds six patents and six applications in the areas of VLSI design, computer architecture and communications. He has more than 20 publications in these technical fields, including "Programming the 80386," published in 1987 by Sybex Inc. Gelsinger has received numerous Intel and industry recognition awards, and his promotion to group vice president at age 32 made him the youngest vice president in Intel’s history. Gelsinger a bachelor's degree from Santa Clara University, magna cum laude, and a master's degree from Stanford University, both in electrical engineering. You can read Gelsinger’s full bio on the Intel Web site.

Hans G. Geyer is vice president and general manager of Intel Corporation’s Networking and Storage Group. He joined Intel in 1980, and has since held various positions, including general manager of European Operations, FAE-specialist for computer architecture and technical marketing manager for the German region, and European marketing manager for telecom products. Prior to joining Intel, Geyer was involved in hardware and software development for intelligent and point-of-sales terminals at Siemens AG, Germany. Geyer studied computer science and mathematics at the Technical University of Munich and holds a masters degree (Diplom-Informatiker) in computer science. He is an INSEAD alumni. You can read Geyer’s full bio on the Intel Web site.

Justin Rattner is an Intel Senior Fellow and director of the Corporate Technology Group, Intel’s main research and development group. Through his career at Intel he has worked on advanced circuit, microarchitecture, architecture, and programming systems research, as well as the prototyping and deep analysis of future computer applications and workloads. Rattner joined Intel in 1973. He was named its first Principal Engineer in 1979 and its fourth Intel Fellow in 1988. Prior to joining Intel, Rattner held positions with Hewlett-Packard Company and Xerox Corporation. He received bachelor's and master's degrees from Cornell University in electrical engineering and computer science. You can read Rattner’s full bio on the Intel Web site.

—End of Technology@Intel Magazine Article—