Case Study

Predictive Enterprise
Intel® Xeon® processors
Intel® Server Board
Embedded technology

Accelerating Data Compression with Intel® Multi-Core Processors

Data Domain incorporates Multi-Core Intel® Xeon® processors and an Intel® Server Board into new products to help customers reduce the cost and environmental impact of data centers while facilitating the move toward a Predictive Enterprise

Challenge

- **Optimize throughput:** Integrate processor and motherboard technologies that can help customers optimize the throughput of Data Domain products by improving the performance of inline data deduplication
- **Enhance efficiency:** Use and produce the most energy-efficient and environmentally friendly products to help Data Domain customers reduce power, cooling, and real estate costs while minimizing the environmental impact of running their data centers
- **Build a Predictive Enterprise:** Adopt technologies and establish vendor relationships that can help Data Domain and its customers create energy-efficient IT environments that are strongly aligned with business goals, promote growth, and enhance the return on IT investments

Solutions

- **Accelerate performance:** Data Domain is incorporating Multi-Core Intel® Xeon® processors and an Intel® S5000PSL Server Board into its products to capitalize on multi-core processing capabilities while establishing a reliable motherboard standard that accommodates specific slot requests
- **Go "green":** Using Intel® technologies in Data Domain products helps create an energy-efficient, environmentally friendly product that enables customers to significantly reduce the costs of running their IT infrastructures
- **Sense, predict, and act:** Intel technologies and the Intel roadmap help Data Domain better sense, predict, and act on market changes. By integrating Intel technologies into its products, Data Domain enables its customers to develop a data-intense computing environment that has the enhanced flexibility to manage resources dynamically

datadomain

"The Intel® processors have helped us to produce energy-efficient appliances and storage arrays that strongly support our green initiative."

Dan McGee, Vice President of Engineering, Data Domain

"Multi-Core Intel® Xeon® processors dramatically accelerate our inline data deduplication technology. Using Multi-Core Intel Xeon processors, our new products deliver significantly higher throughput than their single-core predecessors."

Dan McGee
Vice President of Engineering
Data Domain

Data Domain is the leading provider of deduplication storage systems for a broad range of nearline storage applications, including disk backup, archiving, and network-based disaster recovery. The company's unique data deduplication and replication technologies help customers reduce the capacity and costs required to store growing volumes of data while also accelerating data recovery. Data Domain is integrating Multi-Core Intel® Xeon® processors and an Intel® Server Board into its new products to optimize throughput by accelerating the performance of its inline data deduplication technology. These energy-efficient products can help customers reduce the costs, size, and environmental impact of their data centers while enabling them to work toward more flexible, adaptive IT infrastructures.

Assessing the Situation

Data Domain is dedicated to helping organizations address the challenges of storing and managing fast-growing volumes of data. "Organizations are generating more data, and at the same time, they are storing multiple copies of that data to help ensure disaster recovery and to adhere to compliance regulations," says Dan McGee, vice president of engineering at Data Domain. "Backing up and restoring information from tape is too time consuming, but simply adding disk-based storage capacity increases storage acquisition, real estate, cooling, and management costs."

Data Domain enables organizations to reduce the physical storage capacity they require while accelerating backup and recovery times through its unique deduplication technology. "Typical data compression technologies eliminate content redundancies within each individual file. Data deduplication looks for redundancies across an entire file system and compresses data based on those redundancies," says McGee. "Consequently, data deduplication can reduce the amount of storage a company needs and increase the speed of storing and restoring it."

Because Data Domain's data deduplication operates in real time as information is moved to storage, Data Domain needs to select hardware components that can deliver optimum throughput. "Businesses often have a very narrow backup window and a large amount of content to back up," says McGee. "We need to make sure that they can move that information quickly."

Beyond selecting the fastest available processors, chipsets, and motherboards for its solutions, Data Domain needs to ensure that the company will have the flexibility to use a variety of hardware platforms in the future. "When Data Domain first set out to develop hardware products, we evaluated all the processors and motherboards available," says McGee. "We wanted to use technology that would work with the widest array of hardware platforms. We are committed to maintaining that flexibility for the future."

Data Domain wanted to select a vendor with the best roadmap for future technology development. "We want our software to grow along with the hardware technology," says McGee. "Early on we saw the promise of Intel multi-core processor technology for accelerating the deduplication process, and we architected our product to take advantage of multi-core processors."

Having access to a technology roadmap will help Data Domain build a Predictive Enterprise that can better anticipate and act on market changes. "To stay ahead of competitive offerings, we need to understand emerging trends and act on them quickly," says McGee. "And by building products that are ahead of the curve, we can help our customers create more flexible, responsive infrastructures."

Data Domain is also committed to using and producing "green" solutions that can reduce the costs and environmental impact of IT. "As power costs continue to rise, our customers face significant challenges in keeping operational expenditures down," says McGee. "We want to provide the most efficient products possible to help them reduce ongoing costs while also minimizing the environmental impact of data centers. Combining energy-efficient processor and motherboard technologies with our deduplication technology can help us achieve those goals."

Spotlight on Data Domain

More than 1,000 enterprises worldwide use award-winning Data Domain disk-based enterprise protection storage solutions to reduce backup costs and simplify network-based data recovery. The company's inline data deduplication and replication technologies offer data reduction rates that enable new efficiencies in enterprise data protection.

Designing the Solution

After several years of success with Intel® processors, Data Domain decided to integrate Multi-Core Intel® Xeon® processors into its product line to enhance inline deduplication performance. The company incorporated Dual-Core Intel Xeon processors into its DD510, DD530, and DD565 storage appliances, its DD580 storage array, and its DD580g gateway product. Data Domain plans to equip future-generation systems with Quad-Core Intel Xeon processors to further increase performance.

The move from single- to dual-core Intel® processors has already provided some dramatic results, including both performance improvements and power efficiencies. For example, Data Domain's DD580, which features the new Intel dual-core processor architecture, can deliver up to 800 GB per hour of aggregate deduplication throughput compared with 400 GB per hour for its previous-generation DD560.

"By using multi-core technology and the Intel® Server Board to help accelerate data deduplication, we can pack more processing power and greater storage capacity in a single system. As a result, we can help customers minimize the costs and environmental impact of their data center."

Dan McGee
Vice President of Engineering
Data Domain

McGee credits the impressive performance of the DD580 in large part to the Intel multi-core processor architecture. "Multi-Core Intel Xeon processors dramatically accelerate our inline data deduplication technology," says McGee. "Using Multi-Core Intel Xeon processors, our new products deliver significantly higher throughput than their single-core predecessors."

McGee adds, "By designing our systems to utilize the power of Intel's multi-core technology, we were also able to achieve rapid time to market and solidify our approach to deploying next-generation systems."

With intelligent power management features, the Multi-Core Intel® Xeon® processors also deliver greater performance per watt than the previous single-core processors. "The Intel processors have helped us to produce energy-efficient appliances and storage arrays that strongly support our green initiative," says McGee.

Accelerating Time to Market with an Intel® Server Board

Data Domain also incorporated an Intel® S5000PSL Server Board into new products to further enhance its deduplication performance. The S5000PSL capitalizes on the Intel® 5000P chipset's next-generation dual-processor technology to deliver high throughput with dual, point-to-point system buses, fast memory, and I/O bandwidth, plus Fully Buffered DIMM (FBD) support.

For the Data Domain product group, reliability was a key factor in choosing the Intel Server Board. "We evaluated several motherboards from another vendor, but we ran into reliability issues that would have introduced significant delays in development," says McGee. "The Intel Server Board provides the reliability we need to deliver the highest-quality products, on time."

The Intel Server Board was also able to accommodate particular hardware requirements. The S5000PSL features four PCI Express (PCIe) slots, eight Fully Buffered DDR2 533/667 MHz DIMMs, up to six SATA 3 GBps ports, and dual gigabit Ethernet ports with Intel® I/O Acceleration Technology (Intel® I/OAT). "The motherboard is not just the home for the chips, it's also where the PCI bus lives and where we plug in all of our I/O cards," says McGee. "We have some very specific needs in terms of slot requirements, and the Intel Server Board was able to provide us with what we needed."

Adopting the Dual-Socket Intel Xeon 5000 Sequence Server Board as a standard has helped Data Domain significantly accelerate time to market for new products. "Now that we have a reliable, standard server board architecture, we've been able to cut development time for new products substantially," says McGee. "In the future, as we move from dual-core to quad-core processors, we will be able to do so very quickly because we can continue to use the same board."

"We're also integrating a different Intel Server Board into an upcoming Data Domain product, and because the Intel Server Boards share the same architecture, we don't have to do much more porting work," says McGee. "Compared with the development time required with the previous motherboard we used, we've cut time to market from nine months to just three. As a result, our customers have faster access to advanced Intel technology that integrates easily into their existing IT environment."

> "Early on we saw the promise of Intel® multi-core processor technology for accelerating the deduplication process, and we architected our product to take advantage of multi-core processors."
>
> Dan McGee
> Vice President of Engineering
> Data Domain

Building a Foundation for a Predictive Enterprise

For customers, the move to multi-core processors and an Intel® Server Board in Data Domain products means faster backup and ultimately more time for other projects. "By accelerating the deduplication process, our new products will let our customers back up those mounds of data in far less time," says McGee. "That will enable their IT administrators to focus more on innovative projects and less on storing copies of data."

Working together, Intel and Data Domain technologies are also helping customers reduce total storage costs. "With data deduplication, companies typically require 10 to 30 times less storage capacity they would otherwise need," says McGee. "As a result, customers will need fewer media servers to push data to storage and fewer storage devices to retain that data."

In addition, Intel® technology helps Data Domain deliver a "green" product that can dramatically reduce customers' power, cooling, and real estate costs. "By using multi-core technology and the Intel Server Board to help accelerate data deduplication, we can pack more processing power and greater storage capacity in a single system. As a result, we can help customers minimize the costs and environmental impact of their data center," says McGee. "Intel technology also helps us compare very favorably with other solutions. A DD580, for example, typically uses less than two percent of the power and heat per TB of a comparable VTL. It can also reduce the footprint by as much as 90 percent."

Data Domain has gained much more than just high-performance technology from its ongoing relationship with Intel. The company has also gained a reliable partner. "The success of every product launch depends on our technology partners," says McGee. "They need to deliver the technology they promised, on time, without fail. Intel always delivers, reliably and predictably—and in the world of hardware, those qualities are essential."

Key Technologies

- Multi-Core Intel® Xeon® processors
- Intel® S5000PSL Server Board
  - Intel® I/O Acceleration Technology
  - Fully Buffered DDR2 DIMM memory

Meanwhile, having access to the Intel roadmap has put Data Domain on the path toward creating a Predictive Enterprise, in which technology and businesses goals are well aligned. "We are able to work with our core software and introduce products with the multi-core platform in a very short amount of time," says McGee. "We do not have to do a major re-architecting of our software. What made that possible was the knowledge of the multi-core platform that Intel provided well in advance. We decided to develop our products so we can scale their performance as Intel scales its multi-core technology. It has already proven to be a very good decision for both us and our customers."

Data Domain intends to continue its close relationship with Intel in the future. "We're going to continue to evolve as Intel evolves," says McGee. "As processor speeds and multi-core technology scales, we will be able to consider possibilities that weren't available in the past. We know that Intel technology and our own unique intellectual property will help us produce the best solutions for our customers."

Find a business solution that is right for your company. Contact your Intel representative or visit the Intel® Business/ Enterprise Web site at intel.com/business.

1007/YMB/TDA/XX/PDF 318434-002US