

# Enhancing Postal Efficiency

*Analyzing big data with Intel® Xeon® processors helps streamline operations, reduce costs, and drive new revenue opportunities*



Scot Atkins,  
Program Manager, Mail Information  
Technology, High Density  
Supercomputing (HDSC), Geospatial,  
and Real-Time In-Memory Database  
(IMDB) Systems,  
United States Postal Service

## At a Glance:

### Project

- Facilitate near-real-time decision making to streamline mail tracking, sorting, routing, fraud detection, geospatial mapping and analytics, and more.

### Accomplishments

- Moving from batch processing to stream and complex-event processing, delivering near-real-time results and capacity for up to 15,000 devices at post offices and processing facilities across the United States and its territories.
- Anticipating five-nines (99.999%) availability upon project completion.
- Expecting to significantly reduce costs by using automation to augment the sortation work of senior clerks with specialized knowledge.
- Using near-real-time analytics and reporting to conduct dynamic routing and create new USPS revenue opportunities through Sunday delivery and same-day delivery.
- Using geospatial technology to accurately report real-time events and predict events through ge-fencing and automated data

### Key Technologies

- SGI® UV™ supercomputer with Intel® Xeon® processors deployed with professional services from FedCentric Technologies
- FedCentric Memory-Centric Database (MCDB\*) Accelerator
- FedCentric Memory-Centric Database with GPU Accelerator technology (MCDBg)\* Oracle TimesTen In-Memory Database\*

### Lessons Learned

- Organizations could improve performance from three to six orders of magnitude by implementing a non-uniform memory access (NUMA) environment with highly affinitized data management and using Intel® processor-based servers, FedCentric MCDBg GPU Accelerators and MCDB technologies, Oracle relational and in-memory databases, a Linux\* operating system, solid-state drives, embedded application programming interfaces (APIs), and other technologies that bring analytics closer to the processors.

Today's U.S. Postal Service (USPS) relies on advanced technology to stay competitive with other delivery services and to ensure timely and accurate delivery of 160 billion pieces of mail every year. "For years we have used midrange and high-performance computing (HPC) resources as well as mainframes to help move and track U.S. mail across the globe," says Scot Atkins, program manager for High Density Supercomputing (HDSC™) and Real-Time In-Memory Database (IMDB) Systems at the USPS. "Resources such as these can also be helpful in identifying delivery 'exceptions,' including mail that might have insufficient postage, fraudulent duplication of postage, or security-related issues."

The volume, variety, and velocity of data that the USPS needs to process are growing rapidly. "The move to electronic postage, growing

demand for analysis and tracking, geospatial analysis, and new reporting requirements mean that we process a tremendous amount of data every day," says Atkins. "Unfortunately, the midrange servers we were using several years ago couldn't keep up. They required 36 hours for processing each 24-hour batch workload. We decided to take a new approach that would avoid batch processing and move to more real-time streaming and complex event processing."

Atkins's team began exploring alternatives to the existing infrastructure. "To accelerate performance, we wanted to implement in-memory and GPU data processing. We needed a hardware platform that could provide large-scale, globally addressable memory in a fairly massive, parallelized infrastructure," says Atkins. "Ultimately, we also

decided to move to an industry-standard environment, using a Linux operating system running on x86 processors."

### **Building a High-Density Supercomputer Based on Intel Xeon Processors**

To select and deploy new computing resources, the USPS IT group worked with FedCentric Technologies, a solutions provider that specializes in solving big data challenges with high-density computing and memory-centric approaches. "We first began discussions with FedCentric in 2007," says Atkins. "The FedCentric team understood our requirements and offered deep expertise, not only in supercomputing but also with Oracle databases, our current standard." Together with FedCentric, the USPS IT group implemented a new supercomputer in 2010 using an SGI UV 1000\* system equipped with Intel

*"Early performance data suggests that we could see a fourfold increase in performance per blade by moving to these new servers with the Intel® Xeon® processor E5 family compared with previous-generation processors in prior servers. By delivering faster real-time decisions to more postal units, we can improve the efficiency of tracking, sorting, routing, reporting, analytics, and business intelligence while also providing capabilities to proactively detect fraud and identify delivery exceptions."*

*Scot Atkins,  
Program Manager, Mail Information Technology,  
High Density Supercomputing (HDSC), Geospatial,  
and Real-Time In-Memory Database (IMDB) Systems,  
United States Postal Service*

Xeon processors. This system, which can provide the in-memory processing the USPS requires, includes 1,024 cores and 16 TB of shared memory. The USPS runs an Oracle TimesTen In-Memory Database on a Linux operating system. The infrastructure capitalizes on FedCentric Memory-Centric Database (MCDB) Accelerator technology to enhance response time and throughput compared with traditional disk-based systems.

In 2013, the USPS IT group began to expand the environment to accommodate more data processing and provide high availability. For the expansion, the organization is deploying an SGI UV 2000\* system, which uses the Intel Xeon processor E5-4640 and Intel Xeon processor E5-4603. The new high-availability environment uses nodes with a combined total of 4,096 cores and 32 TB of memory. It also incorporates emerging MCDBg GPU Accelerative technology.

“Using the SGI UV with Intel Xeon processors is the right fit for the Postal Service,” says Gerry Kolosvary, founder of FedCentric Technologies.

“The Intel Xeon processors deliver the raw compute power, core density, and large-scale memory capacity needed to accelerate database performance. The SGI UV system enables organizations to capitalize on Intel Xeon processors while scaling up with a single system image.”

### **Integrating with Local Mail Scanning**

The USPS IT group integrated the centrally located supercomputer with passive adaptive scanning system (PASS) units located in thousands of postal stations across the United States and around the world. The PASS units scan packages and send data to the supercomputer through the Internet. The supercomputer analyzes that data, comparing it with existing information in billions of records, and then sends back sorting and routing results as well as notifications for delivery exceptions.

“The PASS units were implemented to enhance the efficiency of sorting, routing, tracking, and reporting; improve identification of fraud and other delivery exceptions; and eliminate the need for using highly skilled senior clerks for those tasks,”

explains Atkins. “To maximize the benefit of those units, our goal was to achieve a 300-millisecond round trip or less for data to and from the supercomputer. Because most of that travel time is consumed by sending data over the Internet, we needed an incredibly fast response time from the supercomputer.”

### **Delivering Near-Real-Time Decision Making**

By using the high-density, in-memory approach with SGI UV shared-memory servers and Intel Xeon processors, the USPS IT group has been able to surpass its challenging performance goals. “We are averaging just 50 to 100 milliseconds round-trip within the continental United States, depending on the postal location. The supercomputer can even exchange data with Guam locations in under 225 milliseconds,” says Atkins. “With that performance, we can provide near-real-time responses for 15,000 PASS devices and manage 10 million packages per hour at our peak.” Atkins expects performance to improve with the integration of the new SGI UV 2000 servers equipped with the Intel Xeon processor E5 family.

“Early performance data suggests that we could see a fourfold increase in performance per blade by moving to these new servers with the Intel Xeon processor E5 family compared with previous-generation processors in prior servers,” says Atkins. “By delivering faster real-time decisions to more postal units, we can improve the efficiency of tracking, sorting, routing, reporting, analytics, and business intelligence while also providing capabilities to proactively detect fraud and identify delivery exceptions. We can provide better delivery service for customers and ultimately help stabilize postal rates.”

### Cutting Costs

By streamlining package sorting, routing, and analysis, the Postal Service’s CIO organization is also helping to cut costs. “Employing a supercomputer with the PASS units, we continue to reduce our

dependency on senior, highly skilled clerks for sorting functions,” says Atkins. “As those clerks retire, we can achieve significant yearly savings in personnel costs.”

### Creating New Revenue Opportunities

Using the new supercomputer with PASS units is also helping to open up new revenue opportunities for the USPS. “We now have the capability for much more dynamic routing of packages,” says Atkins. “The supercomputer can determine the correct carrier route if a package is to be delivered Monday through Saturday. And it can also provide dynamic routing for a reduced complement of carriers on Sunday.” Online retailers are already capitalizing on this ability to deliver packages on Sunday. Late in 2013, Amazon began offering Sunday delivery in select areas through the USPS. “Working with

Amazon and other online retailers to deliver packages is a huge opportunity for the USPS,” says Atkins. “We’ve already seen how services such as these can provide a competitive edge. While other shipping companies struggled to deliver packages on time during the 2013 holiday season, the USPS experienced a 10 percent increase in mail volume on average across the country, and a 30 percent increase in some metropolitan areas, in large part due to Sunday delivery capabilities. Our supercomputing resources are helping to increase revenue while providing exceptional service to people across the United States.”

This document and the information given are for the convenience of Intel’s customer base and are provided “AS IS” WITH NO WARRANTIES WHATSOEVER, EXPRESS OR IMPLIED, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NONINFRINGEMENT OF INTELLECTUAL PROPERTY RIGHTS. Receipt or possession of this document does not grant any license to any of the intellectual property described, displayed, or contained herein. Intel® products are not intended for use in medical, lifesaving, life-sustaining, critical control, or safety systems, or in nuclear facility applications.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations, and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to <http://www.intel.com/performance>

Intel does not control or audit the design or implementation of third-party benchmark data or Web sites referenced in this document. Intel encourages all of its customers to visit the referenced Web sites or others where similar performance benchmark data are reported and confirm whether the referenced benchmark data are accurate and reflect performance of systems available for purchase.

© 2014, Intel Corporation. All rights reserved. Intel, the Intel logo, 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.

0814/LJ/TDA/XX/PDF

330170-001US