BUYER CASE STUDY

United States Postal Service Turns to HPC for Big Data and Analytics

Steve Conway
Earl C. Joseph, Ph.D.
Robert Sorensen

IDC OPINION

IDC believes that the United States Postal Service (USPS) is at the forefront of an important, accelerating IT trend. Leading organizations in the private and public sectors are increasingly turning to high-performance computing (HPC) to tackle challenging big data and analytics workloads that traditional enterprise IT technology alone cannot handle effectively. IDC uses the term *high-performance data analysis* (HPDA) to refer to workloads that are daunting enough to require HPC technology. The primary factors driving the HPDA trend are the complexity and time criticality of the most challenging big data workloads. HPC can enable organizations to aim more complex questions at their data infrastructures and obtain answers faster, even with more variables included. IDC forecasts that the global market for HPDA servers and external storage will grow robustly from $1.4 billion in 2013 to $4.3 billion in 2018. Further:

- The FedCentric Technologies/SGI HPC (supercomputing) solution employed by the USPS is unusually adept at ingesting very large, complex HPDA problems and processing them entirely in memory. Enterprise IT technology alone was unable to keep pace with the Postal Service's rapidly growing daily volumes of batch-submitted data. HPC technology has enabled the USPS not only to achieve near-real-time response rates on this expanding data but also to begin exploiting mission-critical competitive opportunities.
- Another shining example of the use of HPC technology is PayPal, which IDC estimates has saved more than $700 million because two managers with HPC backgrounds pushed for adopting this technology. We believe that organizations of all sizes that have daunting data challenges should evaluate the benefits of HPC technology, especially now that pricing for HPC systems begins at under $100,000.

IN THIS BUYER CASE STUDY

This IDC Buyer Case Study examines the use of HPC by the United States Postal Service to address mission-critical big data (HPDA) challenges that could not be adequately met with enterprise IT technology alone. The document looks at the competitive environment that created these challenges and the Postal Service's initiatives aimed at advancing its competitive standing over time as well as at related achievements to date. Finally, it provides essential guidance for the future use of HPC in the conduct of the Postal Service's business.

April 2015, IDC #255350

SITUATION OVERVIEW

Organization Overview

The United States Post Office was founded in 1775, during the American Revolution, by the Second Continental Congress and was operated by the federal government until 1971, when the Postal Reorganization Act turned it into the independent United States Postal Service. Today, the USPS exploits advanced technology, including high-performance computing, to stay competitive with other delivery services and to ensure timely, accurate delivery of 160 billion pieces of mail each year.

Challenges and Solution

The Postal Accountability and Enhancement Act of 2006

This important act significantly changed how the Postal Service operates and conducts business. The act provided new flexibility, particularly in competitive pricing for shipping services. The 2006 legislation allowed the Postal Service to respond to dynamic market conditions and changing customer needs with competitive practices that are still closely regulated but can be more agile and assertive than before.

The 2006 act arrived none too soon. By then, the USPS' shipment volumes and revenue were in steep decline. First-Class Mail volume plummeted 29% between 1998 and 2008, thanks mainly to the escalating use of email and other Internet-based communications. During the same period, competition for package delivery from FedEx, UPS, and other service providers also intensified. The Postal Accountability and Enhancement Act of 2006 allowed the USPS to begin addressing these challenges and other business challenges with greater flexibility. Improving the way the USPS handles big data would need to play a key role in this transformation.

HPDA Challenges Spark the Idea for Using Supercomputers

The Postal Service had been computerized for many years, using batch processing on business mainframes and Unix servers to handle data-intensive tasks. But around the time when the Postal Accountability and Enhancement Act of 2006 was enacted, daily data volumes were nearing the petabyte range (despite the decline in the volume of mailed items). By the end of 2006, it was taking 36 hours to process every 24 hours' worth of batch-submitted data. Clearly, the Postal Service's enterprise IT technology couldn't keep up alone and needed help.

The USPS first tried unsuccessfully to exploit barcode sorters to help solve its big data challenges. This led to the idea to buy a supercomputer, with the goal of moving the Postal Service from batch processing to real-time processing of its expanding data volumes.

The First Priorities: Dynamic Sorting and Revenue Protection

The initial HPDA applications targeted for supercomputing were sorting ("sortation") and revenue protection. Features of USPS operations then included:

- At the time when the RFP for its first supercomputer went out, the USPS was using scanning devices at thousands of post offices and other delivery locations in the United States and its territories to scan 4 billion pieces of mail and packages per day. This data was sent in batches to a central facility, where the Postal Service's mainframes and business servers compared the data with hundreds of billions of records to catch instances of insufficient postage, unpaid postage (e.g., xeroxing a stamp), larger fraud schemes, and other revenue-reducing

anomalies. As noted previously, the Postal Service's business computers were unable to perform this task in real time and had fallen behind the growing data volumes.

- In 2006–2007, only 1.5% of postal fraud was large scale, but these cases averaged $50,000 each to move through the federal court system. At that time, the USPS lacked the manpower to handle the small cases that together accounted for 98.5% of revenue loss to fraud.
- Dynamic sorting presented an equally daunting HPDA challenge. Picture thousands of packages arriving each day at a post office or other delivery unit. The sorting challenges go to the most experienced, highest-paid human sorters. They handle the task with error rates of 5% or less, but part of the problem is keeping up with route changes, new housing developments, and other alterations. In 2007 and the following years, in the midst of an economic recession and mounting competition, adding more high-priced human sorters simply was not an option. Pressure was in the opposite direction — to eliminate some of these existing highly paid jobs through attrition.

A New HPDA Opportunity: Package Delivery Seven Days a Week

In response to package delivery services from FedEx, UPS, and other providers, the USPS also wanted to be able to deliver packages seven days a week in major cities. This called for dynamic routing because, on Sundays, the Postal Service would need to operate with reduced staff, including fewer experienced clerks and fewer carriers to deliver the packages.

The High-Density Supercomputer Era at the USPS

Today, the USPS is already on its third generation of supercomputers, with each generation more capable than its predecessor.

Key HPC Technologies

Each HPC generation employed by the USPS has been based on an SGI UV supercomputer provided by prime contractor FedCentric Technologies (based in Chevy Chase, Maryland). The SGI UV system stands out from the cluster crowd primarily because of its large, NUMA shared memory space that is designed to scale as high as 64TB.

The Postal Service reports that the ability to process very large data problems entirely in memory can boost performance by three to six orders of magnitude. Conversely, having to process data outside of a shared memory can severely cut performance — by about 50% when moving data from one blade to a neighboring blade, by two-thirds when moving from the top blade in a rack to the middle blade, by a factor of 25 when moving data from the top blade to the bottom blade in the same rack, and by a far greater magnitude when transferring data between racks. This recognition gave rise to the USPS-coined term of "high-density supercomputing," which refers to the notion of keeping memory and processing in tight physical proximity ("affinity," in USPS terminology) and using an in-memory database to attack big data problems.

As is often the case with organizations that adopt HPC technology for advanced analytics, the Postal Service's supercomputer is used to augment, rather than replace, the capabilities of traditional business mainframes. Key technologies employed in the USPS HPC solution today include:

- An SGI UV 2000 supercomputer system with 4,096 Intel Xeon processor cores and 32TB of shared memory
- The Oracle TimesTen In-Memory Database

- The FedCentric Memory-Centric Database (MCDB) accelerator
- The FedCentric Memory-Centric Database with GPU accelerator

The PASS System

The Postal Service's supercomputer, located in the Minneapolis suburb of Eagan, Minnesota, is connected with the USPS' passive adaptive scanning system (PASS) units in 15,000 post offices and other delivery facilities in the United States and its territories around the world. The PASS units scan packages and transfer data to the supercomputer via the Internet. The supercomputer analyzes the data, comparing it with existing information in billions of records. In addition to performing revenue protection functions, the supercomputer sends sorting and routing instructions back to the delivery facilities.

Even with the relative slowness of the Internet, round-trip duration for the data averages only 50-100ms within the continental United States and just 225ms for a location as distant from Eagan, Minnesota, as Guam. The Postal Service says with that level of performance, it can provide near-real-time responses for the PASS system's 15,000 scanners at peak levels as high as 10 million packages per hour. The new SGI UV 2000 supercomputer should be able to boost that capability even further.

GPGPUs and Geospatial Data

The USPS plans to take dynamic routing much further in the coming years in order to provide real-time routing solutions based on information from handheld devices and GPS. The goal is to increase delivery efficiencies by addressing the major uncertainties that affect the first and last mile — especially changing traffic and weather conditions.

This geospatial data is a good fit for array analysis, and general-purpose graphics processing units (GPGPUs) are tailor-made for this. GPUs were originally designed to handle array analysis for visual processing, that is, pixelation. Their GPGPU cousins perform similar analysis on arrays of numbers, analogous to the vector processors that dominated supercomputer designs from the 1960s until about 1995. Standard x86 CPUs excel at instruction-level parallelism (ILP), while GPGPUs are far better at data-level parallelism (DLP).

This is one major reason why the USPS' SGI UV 2000 supercomputer contains both Intel x86-based processors and GPGPU accelerators. Another reason is the USPS' goal of high-density supercomputing. Each of the USPS' Intel CPUs contains 15 cores, while a GPGPU can include as many as 3,000 cores (albeit with much narrower threads). A supercomputer with, say, 1 million x86 cores would need far more space than one with 1 million GPGPU cores.

Results

Major Achievements to Date

Outstanding USPS accomplishments with supercomputer technology include the following:

- The USPS moved from batch processing to stream and complex event processing, delivering near-real-time results and capacity for up to 15,000 scanning devices at post offices and processing facilities across the United States and its territories.
- The USPS recorded five-nines (99.999%) of availability for the supercomputer-based system.
- The USPS augmented the sorting work of senior clerks with specialized knowledge, enabling some of these highly paid positions to be eliminated through attrition.

- The USPS used geospatial technology and inferencing to accurately predict and report real-time events.

Seven-Day Package Delivery Gets Real

The USPS also used near-real-time analytics and reporting to perform dynamic routing and create new USPS revenue opportunities through Sunday delivery and same-day package delivery. In November 2014, the Postal Service announced that it would be delivering packages seven days a week in major cities and high-volume areas from November 17 through Christmas Day in response to expected double-digit package volume growth. Actual numbers weren't yet available at the time this document was written, but the USPS had expected to process 450 million to 470 million packages in the 2014 holiday season – about 12% more than in the previous holiday season.

ESSENTIAL GUIDANCE

IDC offers the following guidance to USPS management:

- **Remain patient.** HPC is a proven game-changing technology, but it won't solve every problem overnight. It took years for the USPS to get into a situation where data volumes became too large to handle even as new competitive challenges were arising from other delivery services. It will take time for the USPS to master all these challenges. The good news is that the Postal Service has already made some impressive achievements through the use of HPC and is at the forefront of this important trend.
- **Look for other areas of the business where HPC could help.** The HPC experience that the USPS has amassed is a major competitive asset in relation to other delivery services. The Postal Service should actively evaluate its whole operation to search for new areas where HPC could provide a competitive advantage.
- **Stay on top of the fast-evolving HPC technology scene.** We are in an HPC era marked by unusually high technology innovation. It will be important for the USPS to track these developments closely, including through participation in HPC industry meetings where trends and experiences with emerging technologies are actively discussed.

LEARN MORE

Related Research

- *High-Performance Computing in Practice: HPC User Forum, September 15-17, 2014, Seattle, Washington* (IDC #251975, October 2014)
- *New Cray Urika-XA System Targets HPC-Big Data Convergence* (IDC #IcUS25203814, October 2014)
- *Emergent Use Cases in High-Performance Data Analysis: HPC User Forum, September 15-17, 2014, Seattle, Washington* (IDC #251976, October 2014)
- *Experiences with Accelerators and Coprocessors in High-Performance Computing: HPC User Forum, September 15-17, 2014, Seattle, Washington* (IDC #251973, October 2014)
- *Major Global High-Performance Computing Initiatives: HPC User Forum, September 15-17, 2014, Seattle, Washington* (IDC #251971, October 2014)
- *Lenovo Completes Acquisition of IBM's x86 Server Business* (IDC #IcUS25176214, September 2014)

- *USPS Touts Benefits of HPC for Big Data* (IDC #lcUS25123014, September 2014)
- *Worldwide Broader HPC 2014–2018 Forecast: Servers, Storage, Software, Middleware, and Services* (IDC #248835, June 2014)
- *When Massive Data Never Becomes Big Data* (IDC #lcUS24922014, June 2014)
- *Worldwide Technical Computing Server 2014–2018 Forecast* (IDC #248779, May 2014)
- *Perspectives on High-Performance Data Analysis: The Life Sciences* (IDC #248348, May 2014)
- *Global HPC Market Dynamics in 2013* (IDC #248137, April 2014)
- *Industrial Partnership Programs and High-Performance Computing: HPC User Forum, April 7–9, 2014, Santa Fe, New Mexico* (IDC #248113, April 2014)
- *International Perspectives on Industrial High-Performance Computing Partnerships: HPC User Forum, April 7–9, 2014, Santa Fe, New Mexico* (IDC #248122, April 2014)
- *Worldwide HPC Public Cloud Computing 2014–2017 Forecast* (IDC #247846, April 2014)
- *IDC's Worldwide High-Performance Computing Predictions 2014* (IDC #WC20140211, February 2014)
- *Market Analysis Perspective: Worldwide HPC, 2013 — Directions, Trends, and Customer Requirements* (IDC #244742, December 2013)
- *HPDA Pulse Results: 2013 Hardware and Storage Market Analysis* (IDC #244493, November 2013)
- *Catalyst Supercomputer Heralds Shift to More Balanced Architectures* (IDC #lcUS24437513, November 2013)
- *China Eyes 10,000-Fold Data Reduction for Internet of Things* (IDC #lcUS24392513, October 2013)
- *HPC User Forum, October 2013, Seoul, Korea* (IDC #243786, October 2013)
- *Tools and Techniques for Technical Computing in Life Sciences: HPC User Forum, September 2013, Boston, Massachusetts* (IDC #243778, October 2013)
- *Perspectives on Quantum Computing: HPC User Forum, September 2013, Boston, Massachusetts* (IDC #243777, October 2013)
- *High-Performance Data Analysis in the Life Sciences: HPC User Forum, September 2013, Boston, Massachusetts* (IDC #243774, October 2013)
- *Chinese Research in Processor Designs for High-Performance Computing and Other Uses* (IDC #243502, October 2013)
- *The Broader HPC Market 2012–2017 Forecast: Servers, Storage, Software, Middleware, and Services* (IDC #242742, August 2013)
- *IDC's Worldwide Technical Server Taxonomy, 2013* (IDC #242725, August 2013)
- *10 Things CIOs Should Know About High-Performance Computing* (IDC #241565, June 2013)
- *Worldwide High-Performance Data Analysis 2013–2017 Forecast* (IDC #241315, June 2013)
- *Top Issues for HPC Sites: HPC User Forum, April 29–May 1, 2013, Tucson, Arizona* (IDC #241463, June 2013)
- *Advanced Visualization: HPC User Forum, April 29–May 1, 2013, Tucson, Arizona* (IDC #241446, June 2013)
- *Livermore Lab Expands Industry Partnerships: Economic Security Is Vital for National Security* (IDC #240232, March 2013)

- *Advanced Research at the South Ural State University Supercomputing Center* (IDC #238225, December 2012)
- *China Confirms Plans for 100 PLFOPS Supercomputer by 2015* (IDC #lcUS23797112, November 2012)
- *HPC End-User Site Update: RIKEN Advanced Institute for Computational Science* (IDC #233690, March 2012)

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-insights-community.com
www.idc.com