WHITE PAPER

Preparing for Next-Generation Medical Imaging Data and Analytics

Bioinformatics

Carestream Health's cloud leadership and use of Intel® technologies helps health systems reduce PACS/RIS operations costs, manage data growth, and prepare for a new era of medical imaging data analytics.

Ishai Tal
Head of Platform Architecture, Carestream Health

Ketan Paranjape
General Manager, Life Sciences, Intel Corporation

Cloud computing and non-volatile memory solid-state drives are spurring a new era of medical imaging data analytics

Executive Summary

The rapid growth of medical imaging data and the rising need for collaboration across the care team are making cloud computing an increasing necessity for many healthcare enterprises. At the same time, increases in data analytics make imaging data more valuable as a tool to drive research and improve patient care.

Intel and Carestream Health are collaborating to ensure practical, high-performance solutions for picture archiving and communications systems (PACS) and radiology information systems (RIS) workflows, in both on-premises infrastructure and secure external clouds.

In a recent collaboration, Carestream Health demonstrated that new Intel® Solid-State Drive (SSD) Data Center (DC) Family for PCIe® P3700 featuring Non-Volatile Memory Express™ (NVMe™) provided a dramatic threefold increase in throughput for a critical portion of the Carestream Vue™ workflow.[1] This increased throughput can provide a more responsive experience and greater productivity for clinicians.

By taking advantage of Carestream Health's leadership in cloud computing for PACS/RIS, including its broad use of Intel® technologies and its rapid adoption of Intel innovations, health systems can reduce operations costs and manage data growth while preparing for a new era of medical imaging data analytics.

Driving the Business Need for Scale and Cost Containment

A perfect storm of business challenges, emerging opportunities, and rapid technology advances is reshaping the medical imaging data landscape. Together, these changes are making cloud computing an increasing necessity—and a critical opportunity—for hospitals, clinics, radiology practices, and other healthcare enterprises.

In the challenges column, value-based payment models, the increasing global demand for healthcare services, changing regulatory climates, and the emergence of the empowered patient are creating a business imperative to deliver higher-quality care more efficiently and at lower cost. These trends raise the need for healthcare providers to securely and quickly share patient health information, including medical imaging data. Doing so can improve collaborative care, enhance patient satisfaction, and lower costs by reducing duplicate tests, procedures, and diagnostic studies. Increasingly, it is a necessity to meet regulatory requirements.

Table of Contents

Executive Summary 1
Driving the Business Need for Scale and Cost Containment 1
Enabling Technologies: Cloud Computing, Big Data Analytics, and NVMe SSDs 2
Industry-Leading Solutions for PACS/RIS Workflows and Healthcare Analytics. 3
Assessing the Impact for PACS . . . 4
The Next Breakthrough: 3D XPoint™ Technology 4
Move Forward. 5
Reap the Rewards. 6
Learn More. 6

At the same time, the volumes and complexity of medical imaging data are skyrocketing. By one estimate, medical images consume up to one-third of global data storage.[2] Study sizes are expanding. PACS and imaging databases must handle larger images; high-definition formats; 3-D video; and varied modalities such as MRI, CT, PET, and others. All this data must be captured, stored, and analyzed by means of increasingly sophisticated algorithms. It must be securely shared in multiple formats and often in near-real time among authorized members of the healthcare team, including patients. It must be tagged intelligently to enable meaningful analysis, and it must be archived for long-term retention. It's not surprising that medical imaging storage consumes a rising percentage of the budget and floor space in the healthcare data center.

While the expansion of imaging data creates management and infrastructure challenges, it also brings significant opportunities. By integrating medical imaging data with information from the electronic health record (EHR), claims data, scanned documents, genomic data, lab results, patients' personal health records, and other sources, healthcare organizations establish a rich, dynamic pool of data from which to gain a deeper understanding of health and wellness. When healthcare organizations succeed in merging Digital Imaging and Communications in Medicine (DICOM) and non-DICOM images, bringing imaging data out of its traditional silos, and analyzing it effectively, the resulting analytic insights can help:

- Deliver high-quality, personalized care and improve outcomes by presenting providers with a comprehensive view of patients' health and history
- Improve patient satisfaction and outcomes by supporting shared decision-making and enhancing patient-engagement strategies
- Increase provider and enterprise efficiency and reduce costs
- Assess treatment outcomes and identify innovative diagnostic and treatment approaches

Enabling Technologies: Cloud Computing, Big Data Analytics, and NVMe SSDs

Several major technology advances can help healthcare enterprises meet these challenges and take advantage of the opportunities, primarily cloud computing. Using cloud services, along with virtualized local infrastructure, healthcare enterprises gain a more scalable, sustainable way to store, retrieve, archive, and analyze images and imaging data while reducing operations costs for PACS/RIS systems. Well-implemented cloud solutions also foster secure collaboration among diverse organizations.

The methods and technologies of big data analytics provide breakthrough capabilities for the analysis of medical imaging data and other unstructured healthcare data. Big data methodologies evolved to handle analytics scenarios in which high volumes of data, of varied types, are being generated at high velocities. These scenarios are generally beyond the capabilities of legacy infrastructure and relational databases, so big data technologies use distributed processing frameworks such as Apache Hadoop* software, non-relational analytics databases, new analytics applications, and distributed infrastructure. Frost & Sullivan predicts that data management and analytics tools for medical imaging in the United States alone will expand from USD 39 million in 2014 to USD 82 million annually by 2019, growing at a compound annual growth rate of 16 percent.[3]

Solid-state drives, particularly high-performance drives designed to be connected through a PCI Express (PCIe)

bus and accessed using the NVMe interface, are a valuable infrastructure element for processing, delivering, and analyzing images and imaging data. NVMe is an open industry-standard interface developed by Intel and other vendors to take advantage of PCIe's low latency and parallelism.

Based on non-volatile flash memory, enterprise-class SSDs can be attached to a specific server, server cluster, workstation, or laptop to accelerate data movement to and from the processor. They can also be used in a shared disk array to alleviate network bottlenecks. In each case, the SSDs can significantly increase database and analytics throughput, performance, and reliability compared to older SSDs and traditional SATA hard disk drives, at an increasingly comparable cost.[4]

Industry-Leading Solutions for PACS/RIS Workflows and Healthcare Analytics

Carestream Health is a leading provider of medical and dental imaging and IT solutions, as well as X-ray imaging systems for non-destructive testing, and advanced materials for the precision films and electronics markets. Its comprehensive family of web-based and cloud-based solutions for PACS and RIS, along with its a "one-workstation" philosophy of user interface design and reporting, empower health organizations to optimize productivity and maximize investments.

Within the broad Carestream Health product line, Carestream Vue™ for Cloud-Based Services is a fully managed, strategic solution for PACS, archiving, collaboration, and portal services, to deliver new ways for clinicians to collaborate resulting in better patient care.

Organizations using Carestream Vue for Cloud-Based Services can:

- Avoid capital investment and reduce total cost of ownership (TCO) by as much as 30 percent with predictable, pay-as-you-go operating costs
- Enjoy the high performance, flexibility, and reliability provided by the latest technologies
- Empower radiologists, clinicians, and referring physicians with access to imaging information and tools, anytime from anywhere
- Enjoy a highly available and secure infrastructure that is monitored and supported by IT experts with clear service-level agreements (SLAs)
- Stay in control through proactive reporting on usage, activity, and status of data
- Implement a robust, cost-effective disaster recovery solution

Carestream has a strong focus on technology innovation and deploys Intel technologies throughout its IT infrastructure to provide optimal performance and reliability for its customers. It also uses tools such as Intel® VTune™ Performance Analyzer to optimize its software environment.

Servers and storage systems powered by the Intel® Xeon® processor family are designed for data-intensive workloads, enterprise-grade clouds, big data analytics, and data centers of every size. The newest Intel Xeon processor E7 v3 family provides energy-efficient performance, expanded memory capacity, and enhanced reliability features. It includes the hardware-aided virtualization innovations of Intel® Virtualization Technology (Intel® VT), offering further support for software-defined infrastructure. Intel also drives a broad range of

Intel® SSD Data Center Family for PCIe* P3700 provided a dramatic threefold increase in throughput for a critical portion of the Carestream Vue™ workflow. The result: a more responsive experience and greater productivity for clinicians.

innovations to enhance data privacy and security. Intel has also worked with Cloudera*, provider of the most popular version of Apache Hadoop software, to enhance its security, performance, and management capabilities.

Converged networks with Intel® Ethernet Controllers offer industry-leading I/O virtualization innovations and performance. These controllers extend Intel VT to the network, to further enhance the software-defined data center and allow rapid provisioning of networks.

The Intel Solid-State Drive Data Center Family for PCIe brings extreme data throughput directly to Intel Xeon processors with up to six times faster data transfer speed than 6 Gbps SAS/SATA SSDs.[5] The performance of a single drive from the Intel SSD Data Center Family for PCIe, specifically the Intel SSD Data Center P3700 Series (460K IOPS), can replace the performance of seven SATA SSDs aggregated through a host bus adapter (approximately 500K IOPS). Intel SSD Data Center Family products also consume less energy than traditional drives, further reducing the data center's environmental impact.

Assessing the Impact for PACS

Technologists from Carestream Health and Intel established a proof-of-concept system to explore the benefits of the new Intel SSDs. The team incorporated the Intel Data Center P3700 Series featuring NVMe into the workflow for database management and PACS image caching.

The results were dramatic. For a data set with 4,000–5,000 CT exams, Carestream Health found that the configuration shown in Figure 1 provided a threefold increase in performance, loading 700 images per second compared to an average of 150 to 200 images per second on the SAS and SATA hard drives with 1 Gbps network connectivity.

These increases in performance can make the difference between a radiology study that loads in minutes and one that loads in seconds.

This performance can enhance the productivity of health professionals, reducing the time needed to read a study. It can also help keep pace even as studies grow more data-rich, with larger, higher-resolution images and more images per study. It is an essential enabler to meeting current and future PACS/RIS workloads.

The Next Breakthrough: 3D XPoint™ Technology

Further innovation is on the horizon. Carestream Health continues to advance its Vue family, and Intel is pushing its platform technologies forward in performance, throughput, capacity, security, and other areas. In addition, Intel and Micron Technologies have announced a revolutionary storage/memory innovation that will be highly relevant for medical imaging data and analytics, and other data-intensive areas. Intel® 3D XPoint™ Technology (pronounced *3D cross-point*) is a disruptive innovation in non-volatile memory technology that

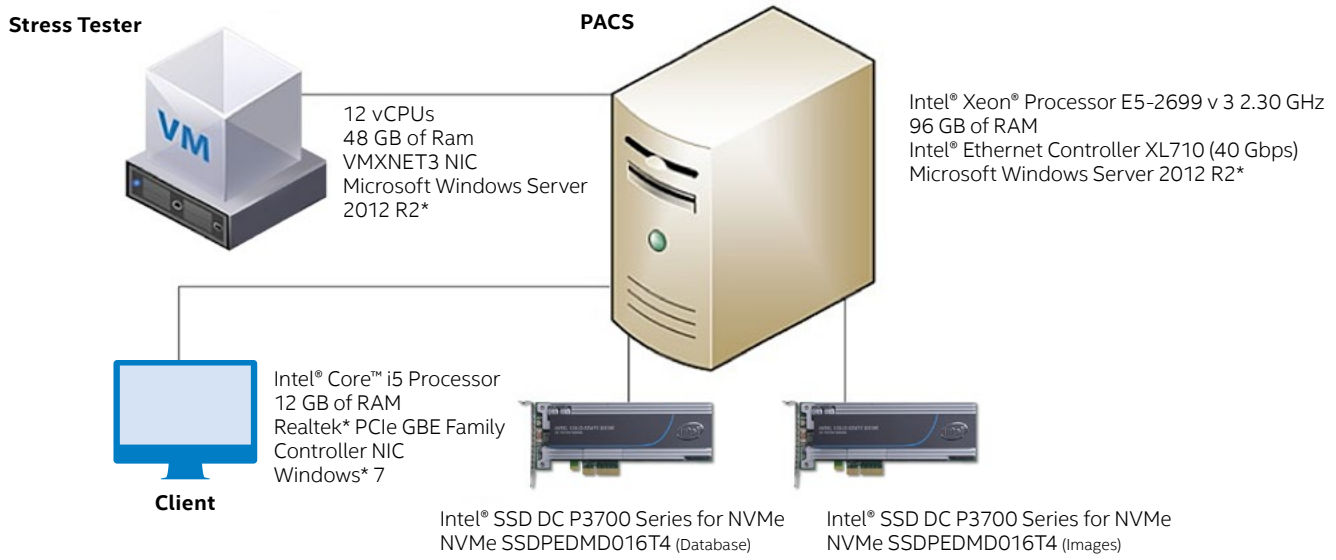

Figure 1. PACS Infrastructure with Intel® SSD Data Center Family for PCIe P3700

is 8–10 times denser than conventional DRAM memory and up to 1,000 times faster than NAND, the most popular non-volatile memory in the marketplace today.[6] Intel, Carestream Health, and other industry leaders are also continuing work to deliver practical advances in areas such as imaging informatics algorithms, collaboration capabilities, and user interface design.

These innovations will enable healthcare organizations to support more users with higher performance and throughput while controlling costs, managing data growth, and exploiting analytics opportunities.

Move Forward

With the maturing of cloud computing, rapid innovations in big data analytics methods, a new generation of SSDs based on NVMe, and the promise of 3D XPoint Technology, it's time for healthcare organizations to prepare for the new era of cloud-enabled medical imaging data and big data analytics, if they're not doing so already. We recommend a three-part approach.

Develop a Cloud Roadmap

Scalability, performance, throughput, and security will be top priorities for medical imaging data in the cloud, whether on-premises or external. To build success, select a cloud-based solution such as Carestream Vue that provides comprehensive support for PACS/RIS workflows, and is compatible with or can interface smoothly with your web-based imaging applications.

Establish policies that identify which workloads and data sets will migrate into a cloud environment. If you decide to use external cloud services, choose a partner such as Carestream Health that has specialized expertise in designing infrastructure and processes for securing imaging/healthcare data in the cloud. A vendor such as Carestream that is aggressive in adopting the latest Intel technologies demonstrates its commitment to maintaining industry-leading performance and helping you keep pace with the changing requirements for medical imaging data. An experienced healthcare cloud services partner can also provide guidance as you determine which applications and which types of data need to remain on-premises.

Modernize Onsite Infrastructure

Optimized infrastructure based on Intel technologies is the foundation for a balanced, software-defined data center that provides scalable, secure, and reliable performance for on-premises clouds and applications. Having a modern local infrastructure also allows for a much more robust interface with external cloud services.

Modernizing older infrastructure with Intel technology-based solutions can improve energy efficiency and density in the data center while reducing complexity and enhancing IT efficiency. Intel technology-based servers, storage systems, network adapters, and data center-class SSDs are designed for large data volumes, demanding data analytics, and mission-critical enterprise computing environments. In particular, incorporating Intel Data Center SSDs with NVMe can significantly increase throughput and responsiveness for PACS/RIS applications, performance-constrained databases, and other data-intensive applications. Along with high-performance servers and storage systems, your medical imaging infrastructure will benefit from 40 Gb Intel Ethernet Adapters to connect the server backbone, and 10 Gb Intel Ethernet Adapters for client workstations.

Create and Implement an Analytics Strategy

An analytics strategy should address medical imaging data and other types of healthcare data. It should result from collaboration among clinical,

By integrating medical imaging data with EHR information, claims data, scanned documents, genomic data, patients' personal health records, and other sources, healthcare organizations establish a rich, dynamic pool of data from which to gain a deeper understanding of health and wellness.

operational, business, and IT leaders to identify opportunities where actionable, data-fueled insights can improve operational excellence, clinical care, and/or the bottom line.

With a strategic direction in place, you can start planning the necessary steps to improve data interoperability, acquire or develop new skills, and remove cultural and technical barriers to data sharing/integration. Identifying practical use cases and undertaking small-scale proof-of-concept projects let you start gaining value as you build expertise.

Distributed infrastructure for big data analytics requires the same high levels of energy-efficient performance, throughput, and scalability as other infrastructure for data-intensive healthcare computing. On the software side, choose tools and solutions that maintain open source compatibility throughout the Apache Hadoop software stack while providing a secure, manageable analytics environment.

Reap the Rewards

Delivering high-quality, value-based care requires cost-effective operations and data-driven, collaborative decision-making. It requires managing medical imaging data efficiently, sharing it effectively, and increasing the productivity of health professionals who create, use, and share it. It requires analyzing imaging and other healthcare data, and deriving maximum value from that analysis. These are not easy tasks, but they are vital ones.

Intel and Carestream Health are driving innovations that help healthcare organizations make the most of their medical imaging data and achieve their objectives. By taking advantage of these innovations—including Intel's next-generation platform and storage technologies, and Carestream Health's rapid adoption of Intel's innovations—health systems can better manage PACS/RIS data growth while optimizing their PACS/RIS and analytics performance. They can meet today's workflow needs, move forward with medical imaging data analytics, and free resources to focus on improving the quality and cost of care.

Learn More

Learn more about how Intel and Carestream can help you make the most of medical imaging data. Talk to your Intel or Carestream Health representatives, or visit us on the Web:

- www.intel.com/healthcare/bigdata
- www.intel.com/healthcare/optimizecode
- http://www.carestream.com/medical.html

Stay abreast of continued innovations. Follow us on Twitter:

- @portlandketan, @IntelHealth
- @carestream

Join the Intel Health & Life Sciences Community for ongoing updates, discussions, and more.

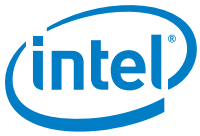

Carestream

[1] This performance was derived from a system using Intel SSD DC P3700 Series compared with a system using SAS 15k drives configured in RAID 10 for the database plus either SAS 10k drives or SATA 7.2k drives configured in RAID 5/6 for image caches.

[2] Digital Imaging in Medicine: Trends and Challenges, International Hospital Equipment and Solutions, http://www.ihe-online.com/feature-articles/digital-imaging-in-medicine-trends-and-challenges/index.html

[3] Frost & Sullivan, US Medical Imaging Industry Leaps Firmly into Big Data Realm, June 29, 2015. http://ww2.frost.com/news/press-releases/us-medical-imaging-industry-leaps-firmly-big-data-realm/

[4,5] Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Configurations: Performance claims obtained from data sheet, sequential read/write at 128k block size for NVMe and SATA, 64k for SAS. Intel SSD DC P3700 Series 2 TB, SAS Ultrastar® SSD1600 MM, Intel SSD DC S3700 Series SATA 6 Gbps. Intel Core i7-3770K CPU @ 3.50 GHz, 8 GB of system memory, Windows Server* 2012, IOMeter. Random performance is collected with 4 workers each with 32 QD.

[6] Technology claims are based on comparisons of latency, density, and write-cycling metrics among memory technologies recorded on published specifications of in-market memory products against internal Intel specifications. Learn more about 3D XPoint technology at http://www.intel.com/content/www/us/en/architecture-and-technology/3d-xpoint-technology-animation.html and http://newsroom.intel.com/community/intel_newsroom/blog/2015/07/28/intel-and-micron-produce-breakthrough-memory-technology

Intel® Virtualization Technology requires a computer system with an enabled Intel® processor, BIOS, and virtual machine monitor (VMM). Functionality, performance, or other benefits will vary depending on hardware and software configurations. Software applications may not be compatible with all operating systems. Consult your PC manufacturer. For more information, visit http://www.intel.com/go/virtualization

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. Check with your system manufacturer or retailer or learn more at intel.com