WHITE PAPER

How Petabyte-Scale Storage Enables **Precision Medicine**

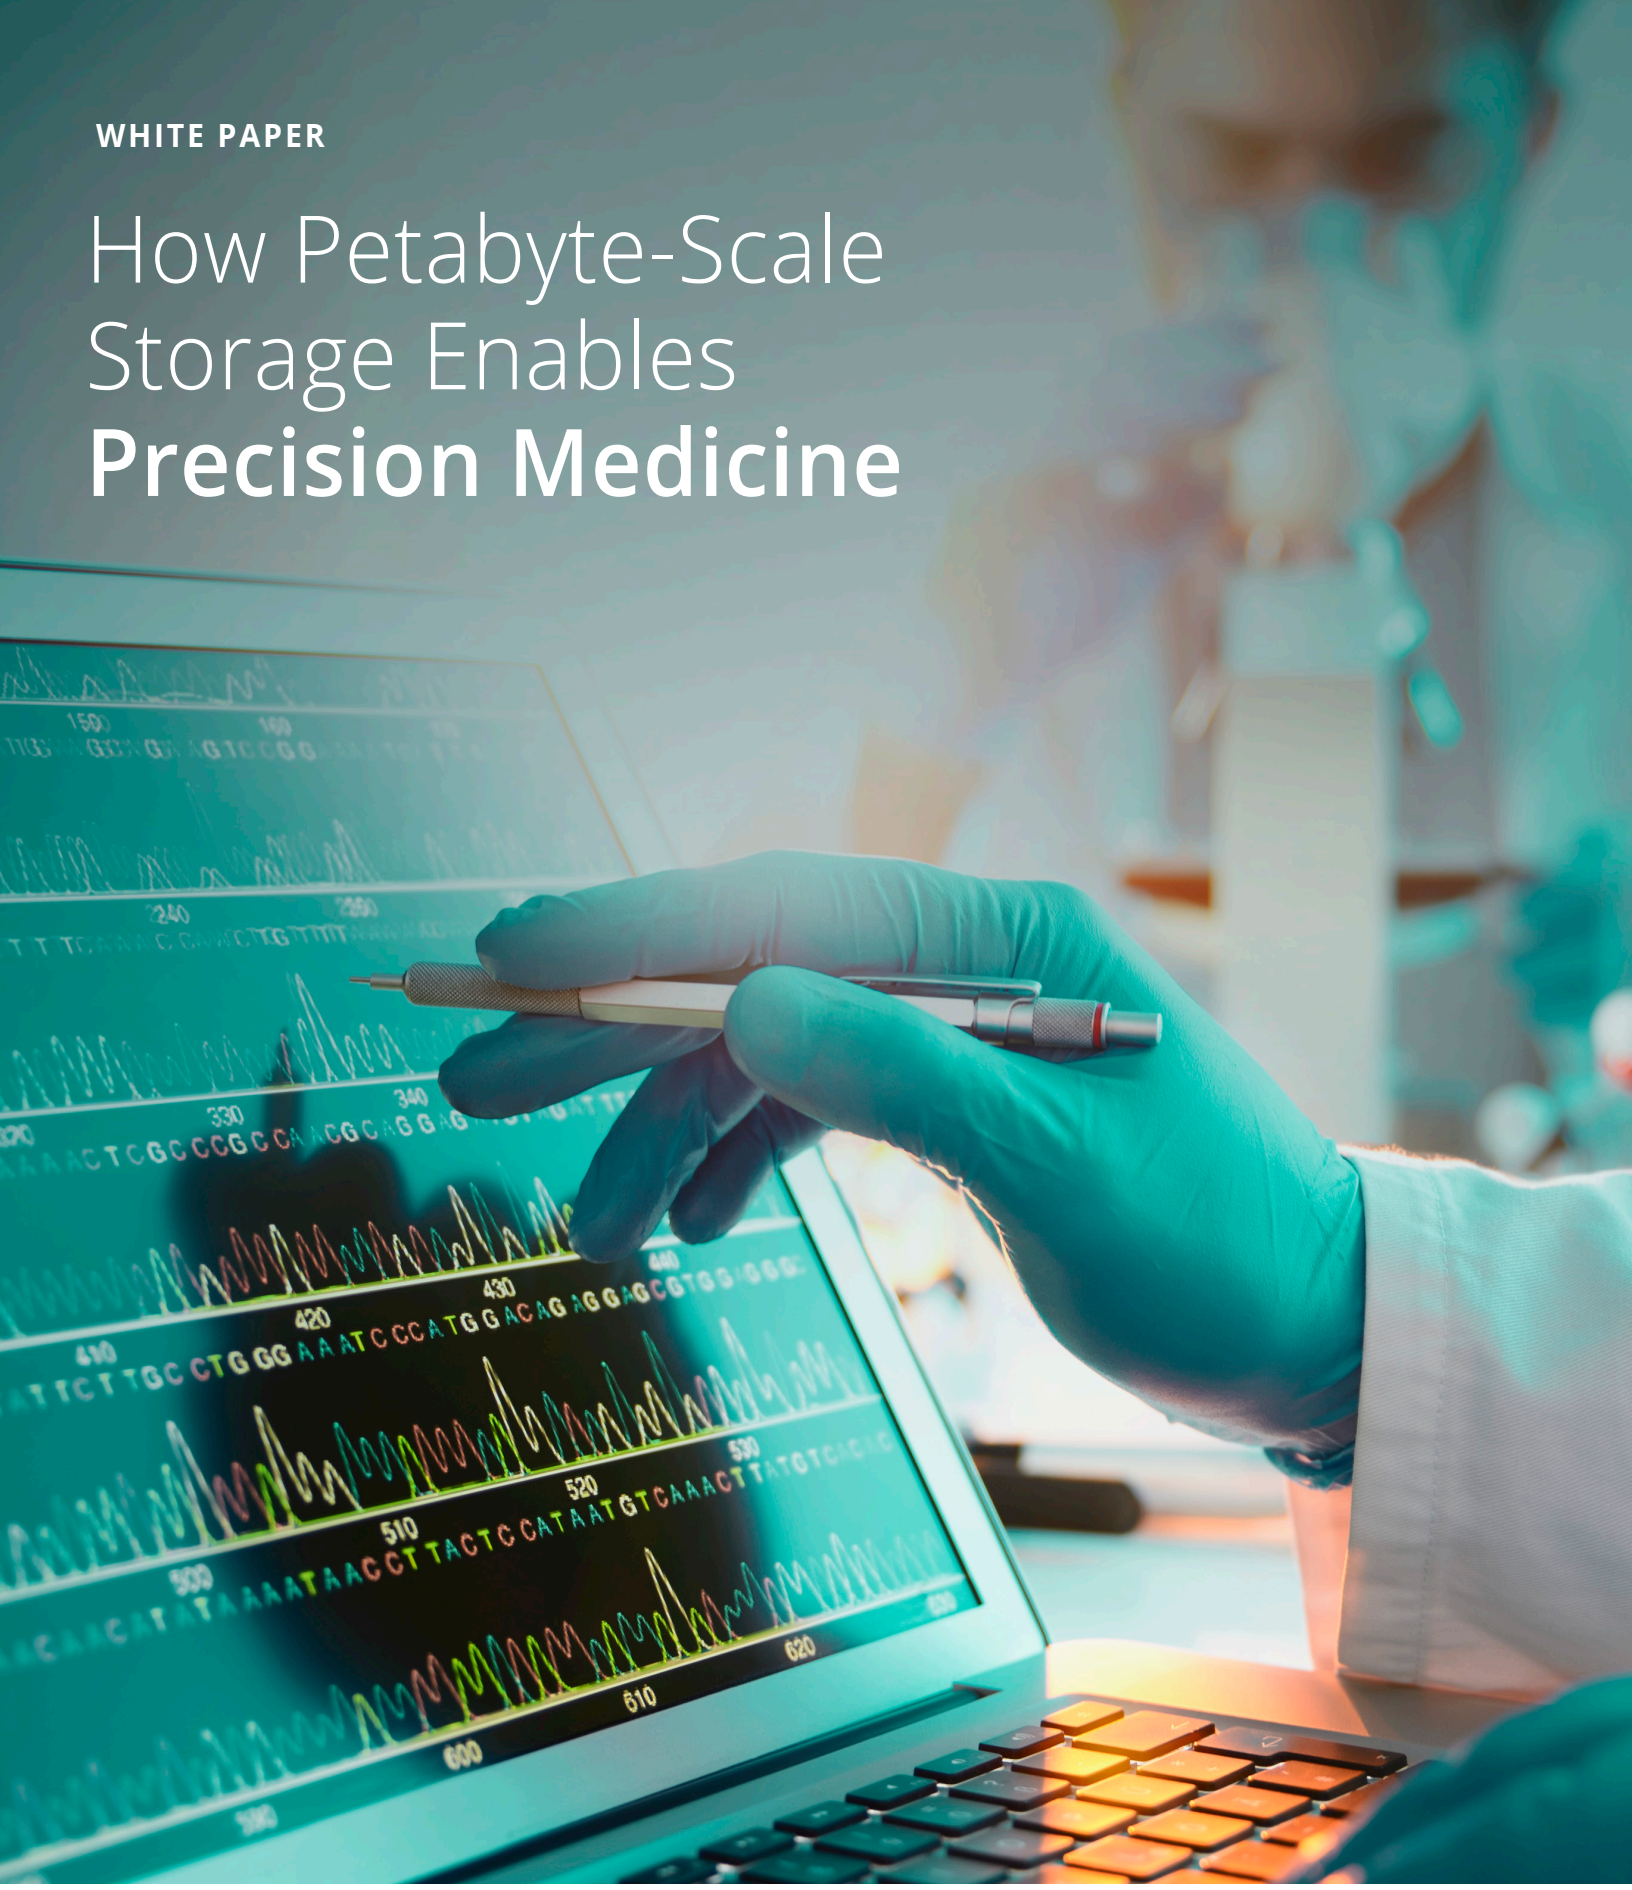

SCALE TO WIN

INFINIDAT

Table of Contents

Abstract 3
Introduction 4
Architecture Matters 5
Performance and Reliability 5
Scale and Security 6
Integration, Simplicity and Flexibility 6
Conclusion 7

Abstract

Data is the DNA of modern healthcare. As healthcare technology continues to evolve at a rapid pace, and patient data management and security evolve, emerging approaches for disease treatment and prevention—like precision medicine and healthcare content management—are becoming more necessary. Precision medicine is about moving from generic to more precise, population-focused diagnostics and treatment by factoring in data from patients' genes, environment, lifestyle factors and family history, into clinical decision-making for earlier, more accurate diagnoses, and more effective treatment and prevention. Data is at the heart of enabling doctors and scientists to execute on this mission. Additionally, rapidly changing regulations throughout the world are affecting the management of all healthcare data. Infinidat removes data management barriers from this level of data interaction by removing isolated islands of storage and allowing much more data to reside on a single, high-performance, highly available platform.

Introduction

In healthcare, speed is everything. Every half-second spent waiting for a medical record to be retrieved from a rolling nurse's kiosk makes a difference in the quality and timeliness of service that a patient receives. Like every industry today, healthcare is struggling with data that is growing at an exponential rate. Until recently, healthcare data was largely "human"-created. Today, as Internet of Things (IoT) becomes pervasive, more and more data is "machine"-created, coming from billions of sensors and other automated data collection points. This is especially true in healthcare, which is trending towards patient wearables and mobile medicine. Add to this the technological advancements in image processing (e.g., 2D to 3D) and healthcare data growth predictions eclipse those of many other industries.

Storing and protecting this mountain of data is pushing information infrastructure and budgets well past their limits, yet represents only part of the challenges associated with data at scale. In the fast-paced connected world, simply storing and protecting data is not enough. This data must also be analyzed, manipulated and made available to an increasing number of people and systems as quickly, efficiently and cost-effectively as possible.

These challenges are magnified within healthcare. The advent of electronic health records (EHR) and the increase in collaborative healthcare data have literally become a matter of life and death. Healthcare providers must integrate data from patient systems with other providers, payers, pharmaceutical suppliers, researchers and partners, while guaranteeing the data adheres to strict compliance and privacy mandates. The stewards of all this data must pull off this balancing act at a scale never imagined, under tremendous cost-containment pressure and with no margin for error.

Healthcare data is highly diverse, and is comprised of a wide variety of data types including emails, clinical notes, diagnostic test results, medical images, family history and massive volumes of ancillary research data, which further complicate the management of this data. In addition, from a data management perspective, healthcare content management (HCM) is the emerging replacement for disparate and more recently API, standardizing EHR and PACS/ Imaging applications. HCM promises to combine this content into a new application that allows users to view images from the EHR, and also bring them together in one platform and add workflows. It's clear from this description that shared storage is at the heart of healthcare content management. Until recently, and despite advances in storage technology, there was no unified solution that could reduce complexity, lower cost, eliminate risk and seamlessly integrate into healthcare enterprises to manage this data at scale.

> ***As health IT systems evolve and supercomputer power becomes increasingly available to research institutions, the cost and time involved in genomics will decrease enough to allow more organizations to benefit from this research—assuming that healthcare organizations implement the data storage and warehousing capabilities to cope with the massive volumes of big data these activities produce.[1]***

In order to reach these goals, new storage innovations are required to achieve optimum results at lower costs to the organization. These innovations include greater data access and availability, higher reliability, consistently high performance, greater areal density, simpler manageability, and more robust application ecosystem integration. Shared storage is at the heart of healthcare content management.

There is a smarter way to think about storing and managing healthcare data at scale—overcoming these challenges not only for today, but well into the future.

Today, systems tasked with capturing, storing and analyzing data at scale present a unique set of challenges for healthcare organizations. They are often too expensive, fundamentally unreliable, perform poorly or inconsistently,

[1] http://healthitanalytics.com/news/ibm-expands-footprint-into-healthcare-big-data-analytics

and consume too much floor space and energy. As a result, modern storage infrastructures have become a patchwork of disparate, underutilized point solutions, trade-offs and compromises.

Infinidat overcomes all of these challenges by delivering the next generation of storage innovation and value. InfiniBox® delivers sustained performance that often exceeds all flash systems by 100 times by servicing as much as 85-90% of all read I/Os from DRAM. InfiniBox's innovative system design and software architecture provides seven nines (99.99999%) availability, which is 100-to-1,000 times more reliable than typical incumbent arrays. InfiniBox also offers as much as 50% greater storage density than arrays utilizing traditional RAID data protection schemes, and provides the ability to store multiple PBs in a solution that occupies a single floor tile, and consumes less than 2W/TB of power at full operational load. Finally, InfiniBox does all of this at a disruptive price point with flexible consumption and pricing models that enable healthcare businesses to operate more efficiently and competitively.

Architecture Matters

Infinidat has developed a storage software architecture which boasts over 120 patented and patent-pending innovations. It is not enough for storage systems today to deliver the best performance, reliability, scale or manageability. They must also maximize each of these characteristics and enable organizations to achieve all their goals, without compromise or increasing costs. Infinidat is able to deliver all of these capabilities in our enterprise-proven storage solution family, InfiniBox.

> ***In genomics and precision medicine, the amount of information about rare diseases, inherited traits, and responses to interventions available in a minuscule strand of DNA is simply staggering.[2]***

Infinidat has developed a revolutionary storage software architecture. InfiniBox is built with the power of three in mind, featuring industry leading N+2 active-active-active redundancy for all components, and an adaptive grid architecture that enables superior performance and reliability. InfiniBox consists of industry-standard hardware components that are coupled with our enterprise-proven software. The core storage software architecture features an innovative declustered RAID-like algorithm, called InfiniRaid®, that is highly scalable and offers a dual-parity RAID schema that solves both the performance and rebuild time challenges of high-capacity drives. Additionally, the architecture delivers hyper-efficient in-memory data management, highly optimized SSD utilization and hotspot-free automatic data dispersion. These capabilities enable InfiniBox to deliver on performance, reliability, flexibility and scale at a cost otherwise not attainable by any competitive offering.

Performance and Reliability

In order to help healthcare providers deliver faster and better service, they choose Infinidat. InfiniBox's DRAM and SSD architecture is key to providing the fastest possible response times, enabling doctors to see more patients and spend less time waiting for data to be retrieved from legacy, latency-bound storage arrays. InfiniBox's software engine is powered by a series of innovative, patented caching and machine learning algorithms. First, it starts with the concept that the system has 8-to-10 times more cache than traditional arrays. Next, instead of brute forcing cache dumps to keep up with application workloads, InfiniBox's software engine intelligently coalesces incoming writes and flushes to persistent disk every five minutes.

Infinidat's InfiniRaid delivers the most density efficient, high-performance de-clustered RAID in the industry, and includes virtualization, thin-provisioning, and compression driving better storage utilization, hotspot free, highly scalable, and highly available storage.

[2] http://healthitanalytics.com/news/precision-medicine-genomics-require-strong-data-governance

Infinidat also minimizes failure and data access vulnerability by introducing industry-leading rebuild times of less than 15 minutes for (up to) dual 6TB drive failure. InfiniRaid assures the highest reliability, lowest risk of failure and lowest risk of performance degradation available on the market.

Scale and Security

Infinidat helps healthcare companies better prepare for the coming tsunami of data from healthcare IoT. The healthcare IoT drives massive data volumes via mobile medical devices, patient wearable sensors, social media and global unstructured research data. Healthcare content management is the emerging replacement for disparate and more recently API standardizing Picture Archiving and Communication System (PACS) Imaging applications. HCM promises to combine this content into new applications and use cases that enable users to not just view images from the EHR (a functionality recently added by Epic), but to bring them together in one platform and augment them with additional workflows and supporting data.

Today's healthcare systems are not only challenged with expanding capacity requirements and the need for increased performance, but also with the need for efficient, reliable encryption of all data sets to meet regulatory requirements without huge capital expenditures. Infinidat provides the performance needed to service patients quickly, as well as unique zero latency data compression which increases system density while introducing no impact to system performance. InfiniBox can store more data per floor tile more cost-effectively than any competing solution, while reducing space, power consumption and cooling. Additionally, InfiniBox helps meet all regulatory requirements with self-encrypting drives, which again have no impact on system performance.

Integration, Simplicity and Flexibility

The InfiniBox solution provides easy integration with a customer's existing application infrastructure. InfiniBox has a robust, fully documented RESTful API that makes it easy to integrate with third-party applications and orchestration/management tools.

InfiniBox is designed from the bare metal to deliver extreme, best-in-industry ease-of-use. Things like self-service provisioning, an intuitive HTML5 GUI, and powerful CLI help to make InfiniBox simple to deploy and use. Data and experiential feedback from our clients demonstrate that one person can typically manage as much as 10PB of capacity, helping to drive very low OpEx and maximum staff productivity.

Infinidat's optional COD (Capacity On Demand) consumption model enables healthcare research and provider organizations to purchase just the capacity they require today, and provides them with an extremely fast and predictable expansion model that does not require any hardware upgrades, allowing them to meet even unpredictable data growth challenges quickly and cost-effectively. This allows them to achieve long-term stability and expand into the future with no forklift upgrades or excessive upfront CAPEX outlays.

Conclusion

In a recent HIStalk article, Amie Teske observed, "HCM encompasses much more than simply unifying ECM and EI technologies together into a single architecture to enable shared storage and a single viewing experience for all unstructured content, DICOM and non-DICOM. Just as important is workflow and how all document and image content is orchestrated and handled prior to storage and access. This is essentially the secret sauce and the most difficult aspect of an HCM initiative." [3]

> ***Infinidat's core strength is an enterprise-proven storage software architecture, designed for today's data-intensive healthcare applications that removes the barriers to implementing evolving HCM and precision medicine workloads.***

Infinidat's core strength is an enterprise-proven storage software architecture, designed for today's data-intensive healthcare applications that removes the barriers to implementing evolving HCM and precision medicine workloads. The architecture supports an extremely high level of data security and availability, delivers the highest areal density in the industry, and provides extremely high sustained performance for the widest range of workloads.

The architecture also allows for virtually unlimited scalability and carefully balances a rich enterprise-class feature set with a low and predictable TCO. In addition, Infinidat's design supports rapid, flexible, cost-effective deployments in the most demanding customer environments, fine-tuned to fulfill all healthcare requirements at petabyte scale. Infinidat delivers storage solutions that are fast, reliable, and scalable at a disruptive price that don't require organizations to compromise on their overall mission of delivering the best possible clinician experiences and patient care.

[3] http://histalk2.com/2016/10/17/readers-write-ecm-for-healthcare-advances-to-hcm-healthcare-content- management/