White Paper

Establishing Uncompromising Data Availability for Healthcare Organizations

Sponsored by: Dell Technologies

Eric Burgener   Phil Goodwin   Silvia Piai   Mutaz Shegewi
February 2021

IDC OPINION

One of the biggest challenges for healthcare is its data explosion. Data forms the driving force of patient care delivery and continues to proliferate as organizations accelerate their digital transformation (DX) journeys. Coherent data infrastructure helps healthcare organizations put information into context and infuse resilience into business, clinical, and operational models. The future state of healthcare will include topics like precision diagnosis, which is solely based on the consumption of data generated within the healthcare system. With the exponential growth in data from the edge, network, core, and to the cloud, healthcare organizations need the appropriate tools to store, manage, and analyze data. Secure access to vast stores of information anywhere, at any time, combined with the application of advanced analytics and artificial intelligence (AI) will enable healthcare organizations to shift from "data rich" to "data driven." Therefore, storage and data protection are vital areas for health IT leaders to consider when making enterprisewide improvements in performance, availability, scalability, agility, security, and compliance with regulatory frameworks (e.g., Health Insurance Portability and Accountability Act [HIPAA] and General Data Protection Regulation [GDPR]). Healthcare organizations need to take proactive measures and make significant changes for data infrastructure to manage current and future data growth. Given the importance of data, healthcare leaders must act now.

IN THIS WHITE PAPER

This IDC White Paper analyzes how storage and data protection play a crucial role in supporting data proliferation and DX in healthcare. Storage and data protection serve as pillars for a secure, resilient, and sustainable healthcare organization, whether a single facility or an entire health system. Consequently, health IT leaders can only benefit from becoming more familiar with and informed about the latest advancements, relevant use cases, and implications of data and storage technologies for healthcare. This White Paper includes a technology profile of Dell Technologies healthcare storage and data protection technologies as an example of a vendor with solution offerings relevant to the use cases explored.

SITUATION OVERVIEW

Data continues to grow exponentially in healthcare. According to IDC's Data Readiness Condition (DATCON) index in 2021, the healthcare data explosion will approach the 4ZB level and exceed 10ZB by 2025. Although healthcare and life science organizations currently manage on average 21PB of data (i.e., 25% less than the industry average), at the same time, they retain data almost 20% longer.

These data retention trends are due to regulatory mandates even though regulations differ around the world. Organizations tend to retain patient and financial data the longest, while they typically hold operational data for shorter periods. Regulations require some data to be retained for decades (e.g., patient health records). In contrast, other data types are retained for as little as 30 days (e.g., video surveillance data).

The growing plethora of data sources and types presents a challenge for IT organizations concerned about exhausting their resources on merely being able to keep up. Furthermore, healthcare faces a wide range of digital era stress tests alongside data proliferation — everything from aging populations to increased chronic conditions, the shift to value-based care (VBC), the rise of consumerism, and rampant cyberthreats. These challenges drive health IT leaders to look at improved storage and data protection alongside DX initiatives as a way forward to promote organizational resilience.

According to IDC, 41.4% of health IT leaders reported that they would continue to invest in technology to close gaps in their DX efforts, despite any impact of the COVID-19 pandemic. This level of digital determination in an industry frequently cited as a technology laggard reflects its need to transform. In addition, the current healthcare environment accelerated DX efforts — organizations accomplished two years' worth of transformation efforts within the first two months of the pandemic. In the pandemic, providers faced no choice but to pivot toward virtual care and do so at unprecedented speed to meet emerging frontline needs. The transformation efforts were extraordinary. Since the initial wave of the pandemic, IDC data shows that 83.9% of 1,500 respondents had a virtual care visit for the first time, 1 out of 5 downloaded a virtual care application in advance of needing it, and 72.5% used AI-based chatbot symptom checkers.

Storage and data protection underly all mission-critical workloads in healthcare. Everything from clinical and operational performance at a single facility to the resiliency of an entire organization is affected. However, data infrastructure initiatives can prove daunting. Large provider organizations can have more than 20,000 employees, with the average workforce in the range of 14,000 employees. Furthermore, clinical (e.g., electronic health record [EHR] and medical imaging), operational (e.g., claims and supply chain), and other data types (e.g., consumer and research) can vary between 7PB and 34PB per patient. Therefore, to effectively manage data and plan for future requirements, health organizations must rely on a robust, secure, cost-effective, and elastic infrastructure.

Health IT leaders should continue to direct their attention to the industry challenge of data proliferation and double down on data protection efforts. While COVID-19 redefined disruption in that survival of the fittest no longer links to organizational size or strength but resilience, and that many lessons learned will be here to stay (e.g., intelligent capacity planning, digital front door, and integrated contact tracing), all add to the data problem. Organizations must continue to find ways to contend with the expanding data feasibly and meaningfully while managing, preserving, and protecting it all.

The Cloud

Cloud computing offers healthcare organizations compelling ways to modernize IT infrastructure and undertake DX. Health information stored in cloud storage environments can help break down data silos and improve secure access across the organization or ecosystem as well as across the globe to support clinical research initiatives. Typically, organizations already have an investment in existing infrastructure, but the cloud has gained much momentum driven by its benefits alongside the ongoing transformation of skill sets, capacity, and agility of healthcare organizations. Legacy systems absorb significant resources — as high as 64% of IT budgets — and many organizations can no longer afford to do everything on premises. Furthermore, continuous advancements in functionality and services developed in the cloud add to the interest in cloud solutions for healthcare environments. Cloud technologies have also come a long way in addressing healthcare regulatory needs, while mobile device and application management tools make the cloud more secure, leading many healthcare organizations to adopt a private or hybrid cloud deployment model.

IDC data shows that 58.4% of healthcare organizations use cloud solutions to support production workloads and services, with private cloud preferred over public cloud. Furthermore, "on-premises enterprise private cloud" is the current model of choice (69.3%). More than half of the respondents (53.3%) now see "improved IT security" as the most significant benefit in a private cloud. Much of the apprehension in healthcare around cloud has largely dissipated with private cloud. With a private cloud, resources are shared within a single or an extended enterprise; the enterprise places restrictions on access and defines and controls the level of resources dedicated.

Healthcare organizations realize not all workloads can be migrated to the cloud. A few reasons include regulatory, logistical, and performance impacts and patient safety. For those workloads that can be migrated, multiple forms of IT services are required to support them. This need for greater flexibility is driving the interest of healthcare organizations toward hybrid cloud (i.e., usage of IT services including IaaS, PaaS, and SaaS apps and SaaS-SIS cloud services across one or more deployment models using a unified framework). Hybrid cloud can maintain sensitive protected health information (PHI) in a private environment while benefiting from public cloud scalability. Sometimes, the term *multicloud* gets used as an alternative to hybrid cloud. However, IDC sees multicloud as a description of an organizational strategy or architectural approach regarding the design of a complex digital service that involves the consumption of cloud services from more than one cloud service provider.

With patient lives becoming more "sensor prone" in their exposure to the edge (e.g., IoT and biometrics), the cloud offers immense opportunities to integrate that data to close care gaps, drive predictive analytics, and offer personalized care. As more health IT leaders view the cloud as a means for IT modernization, the switch to cloud will reshape organizations as far down as to the patient level. The impact will encompass current/future applications, consumption models, and where data gets accessed as well as how it is presented, in what setting, and to whom.

Finally, successful cloud adoption strongly correlates to vendor capabilities, particularly in terms of security, compliance, and technical and industry expertise. Given the evolution of healthcare toward more integrated, dynamic, and proximate service models (e.g., home, retail, urgent, and virtual care), the cloud will be vital to eliminate any guesswork from common scenarios such as end-user growth, application upgrades, data migrations, business continuity, consolidations, and the alignment of operations to outcomes. Finally, evolving interoperability standards for fast healthcare interoperability resources (FHIR), microservices, X12, HL7, and API are homogenizing data, enabling hybrid cloud and multicloud solutions to scale enterprise capabilities further to unlock and utilize data value.

Healthcare Digital Transformation

Healthcare DX can no longer be ignored; in fact, it is necessary to survive in current times. According to IDC, healthcare DX refers to "the application of 3rd Platform–related technologies (cloud, big data, social, and mobile) to fundamentally improve all aspects of healthcare." Healthcare DX is not only a matter of modernizing existing health IT systems but rather a paradigm shift toward platform thinking that combines multiplied innovation to enhance experiences and improve performance. DX relies entirely on digital determination, data, and an evidence-driven culture of knowledge workers and leaders that can amplify the productivity and future readiness of an organization. Therefore, the disruption and opportunities that DX presents in healthcare are immense.

Compared with other industries, healthcare is traditionally considered a laggard in adopting IT and the use of digital mediums to transform. However, healthcare for many years faced mounting pressures that drive the need to prioritize DX, pressures that continue to manifest. In the United States, for example, the ongoing shift in reimbursement models from volume- to value-based programs drives increased demand for new revenue streams. These programs aim to improve care for individuals, improve the health of populations, and reduce costs. Therefore, organizations require quicker and more convenient access to clinical, claims, and operational data. The aging healthcare workforce and projected shortages in specific medical disciplines (e.g., radiology and primary care), coupled with the rise of consumerism, further drive the need for new service models and care offerings (e.g., virtual care and teleradiology) powered by DX capabilities. Patients with one or more chronic conditions, aging populations, and growing healthcare expenditures further emphasize a need for healthcare organizations to undertake DX initiatives to address current and future needs.

As a result, healthcare DX is trending. In an IDC survey, 94.1% of respondents indicated that their organization has begun or is well underway in applying health IT to advance DX initiatives. In comparison, only 5.9% of respondents report that their organizations have done no DX-related work yet. Key DX areas include telehealth, big data, analytics, AI, workflow optimization (e.g., EHR and revenue cycle management [RCM]), IoT, mobility, clinical communication, patient messaging, "patient of one" initiatives (e.g., consumer 360 views and longitudinal health records), workspace collaboration, automation, AR/VR, robotics, and cybersecurity. According to the same survey, the top goals for healthcare DX included introducing new business models (28.7%), improving operational efficiency (25.9%), and meeting consumer expectations (24.9%). Finally, funding DX appeared to be a strategic imperative with different tactical approaches. IDC identified primary sources of DX funding in healthcare as dedicated DX budgets (38.4%), line-of-business (LOB) budgets (32.9%), and the IT organization (28.8%).

Cybersecurity

Volume, variety, and velocity of data have long been on the increase in healthcare as workloads expand across the edge, network, core, and to the cloud. Therefore, protecting information, such as PHI, forms a vital function for health IT. The industry witnessed increased cyberattacks recently, with rogue actors often targeting organizations in unique and diverse ways, from phishing schemes and socially engineered cyberfraud that prey on individuals to ransomware and brute-force attacks aimed at entire health systems. The Health Insurance Portability and Accountability Act and General Data Protection Regulation introduced stringent requirements and responsibilities pertaining to cybersecurity in terms of data storage, retention, protection, privacy, and ownership.

Therefore, the stakes are exceedingly high in cybersecurity from a reputational and legal standpoint. HIPAA violations, for example, can significantly damage an organization's reputation in addition to the enforcement of hefty penalties. As DX continues to increase digital touch points and expand IT attack surfaces, healthcare organizations must be mindful that there is always room for improvement. A recent IDC survey of over 260 healthcare and life science organizations revealed that only 60% felt they had a handle on the growth and management of their data while understanding what data required protection and how to go about protecting it. The growth of the digital footprint presents highly lucrative targets for cybercriminals and poses a critical concern for healthcare IT. Consider the potential vulnerabilities combined with the black market value of electronic health records. The value of health records can add up to $1,000 per patient due to its richness as a source of data and PHI.

Leaders should understand how storage and data protection offerings can advance cybersecurity strategies. The end state should be a realization of maximal cyber-resilience with dynamic monitoring and continual improvement. Integrated solution offerings evaluated in light of HIPAA-HITECH, SOC2, and HITRUST standards will be well positioned to support cybersecurity initiatives. All offerings must manage and protect workloads around the clock in a healthcare data environment. Healthcare leaders should also seek to include KPIs related to cybersecurity training, system updates, role-based access controls, regular risk assessments, and data recovery mechanisms.

Data-Driven Use Cases in Healthcare

Next-Generation Clinical Documentation

The electronic health record paved the way for healthcare to leap into the digital era. With the EHR, clinical documentation has come to form the most critical workload in health IT. The initial wave of EHRs enabled physicians with digital versions of patient charts to securely access real-time patient information for documentation, record keeping, and decision-making functions at the point of care, in addition to broader use in other clinical, financial, legal, and administrative functions. However, documentation requirements as part of the EHR workload can lead to widespread frustration, as many physicians have cited the technology as a contributor to burnout. A technology leap that drove innovation forward came with DX, the shift to VBC, and the rise of consumerism.

Enter "next-generation clinical documentation" marked by a shift in EHRs from traditional systems of record for routine data entry to nontraditional systems (or platforms) of engagement for value-added activities. These activities range from clinical decision support (CDS) to patient engagement, revenue optimization, and interoperability. Next-generation EHRs have distinct improvements in system "usability" (e.g., personalized care planning, population health dashboards, and ride-share integrations), "openness" (e.g., longitudinal health records, OpenAPI/HL7 FHIR integrations, developer communities, and application marketplaces), "experience" (e.g., role-based/specialty workflows, enterprisewide mobility, and patient engagement), and "intelligence" (e.g., predictive analytics, CDS, and AI-based virtual assistants).

In addition to the adjacent impact in areas like RCM and medical imaging, next-generation EHR innovation has led to the emergence of more ambient interfaces (e.g., gesture-based controls, voice-driven capabilities, and smart speaker automation) to position the EHR in the background. IDC research revealed that 53.4% of providers had increased organizationwide IT spend on EHR optimization as a top priority. EHRs also ranked first for cloud-based workloads (63.2%). For next-generation EHRs to emerge further and meet future demands, they must scale without compromise. Therefore, reliable clinical documentation equates to reliable storage and data protection.

Remote Patient Monitoring

Remote patient monitoring (RPM) technologies collect data passively from a body area network (BAN) to monitor vital signs and physiological data outside traditional care settings. Examples include implantable sensors, continuous glucometers, blood pressure monitors, pulse oximeters, trackers, wearables, smartphones, and edge gateways. The resulting influx of data that streams continuously from the edge can be life changing when aggregated to enable reliable, integrated access by patients, providers, and life science companies. As a result, medical intervention becomes timelier and patient relationships are strengthened through such things as personalized care plans and joint/emergency decision making.

RPM has mostly seen adoption in managing chronic conditions (e.g., diabetes, hypertension, and heart failure). Chronic conditions cost up to 3.5 times more to treat and account for 80% of admissions. When combined with AI, RPM can identify real-time issues at the edge, such as patient falls. In this manner, RPM helps patients maintain control, get timely intervention, avoid further complications or emergencies, and improve short- and long-term outcomes. Many opportunities abound for RPM, considering there are about 16 connected devices per hospital bed in the United States, 3 million data points in an average clinical trial, and an ever-expanding healthcare datasphere, particularly at the edge.

As more activities become biometrically monitored, the criticality of data will increase alongside an uptick in IoT data. For RPM to create value, data must be secure, accurate, accessible, and interoperable, making storage and data protection vital. The way forward for RPM will be through novel approaches that better integrate its data into IT infrastructure. The global rollout of 5G, for example, will accelerate the speed by which users remotely access a BAN and support transmission of large volumes of RPM data securely. Wrist-worn wearables show much promise in proactively detecting and alerting individuals to the existence of an illness (e.g., elevated resting heart rate or changes in infections), enabling them to take preventative measures early on to protect themselves and others. To that end, IDC predicts that 7 of the 10 leading wrist-worn wearables companies will have released algorithms capable of early detection of potential signs of infectious diseases including COVID-19 and the flu by the end of 2021.

Next-Generation Sequencing

Next-generation sequencing (NGS) determines the sequence of DNA or RNA genetic material to study genetic variation. This genetic variation could be associated with diseases or other human biological phenomena, making it critical for healthcare and life science organizations to understand, which will continue to drive precision medicine, research, and the future of drug development. NGS has already significantly increased the global amount of raw and processed genome sequencing data from research in genetics, oncology, microbiology, and reproductive health.

NGS workflows typically consist of:

- **DNA library construction.** Genetic material gets processed into short fragments and then ligated to technology-specific adaptor sequences to form a fragment library with barcodes.
- **Clonal signaling/polymerase chain reaction (PCR) amplification.** This uses PCR to amplify the genetic fragment library into clones that increase the signal detected from each target during sequencing.
- **Sequence library/sequencing.** The library gets sequenced and read by using technologies such as Illumina (Solexa), Roche 454, and Ion Torrent Proton sequencing.
- **Data analysis.** This involves analyzing massive quantities of data through analytics tools via primary (processing of raw signals), secondary (read filtering, trimming, alignment, and variant calling), and tertiary (interpretation and extraction of meaningful information) analysis.

NGS can now manage millions of variant calls but requires proper infrastructure. High-performance computing (HPC) can deliver the compute power for complex data analysis and interpretation, while highly scalable storage manages NGS to ensure compliance and usability of the genomic data. Researchers often keep files for future studies and do more comprehensive retrospective analyses on the data later in time. Furthermore, NGS data requires high redundancy as the data is often never deleted and is usually annotated. Therefore, storage and data protection for NGS demand long-term petabyte-scale strategies, technologies, and capabilities to protect the sensitive and personal nature of genomic data. The future for NGS will be via HPC combined with high-throughput, scalable, and cost-effective data storage and protection systems with integrated software applications that manage petabyte-sized volumes of redundant, sequencing, and genomic pipeline data sets.

Medical Imaging Transformation

Medical imaging transformation relies heavily on scalable storage and data protection to drive operational and diagnostic imaging workflow performance improvement. Imaging workflows can be complicated due to the need to quickly retrieve and reliably search vast amounts of stored and archived images with metadata across EHRs, picture archiving and communications systems (PACSs), and vendor neutral archives (VNAs) in standardized formats. Furthermore, medical images can also quickly inundate IT infrastructure due to large file sizes: radiography (10MB per image), computed tomography (250MB–1GB per exam), magnetic resonance imaging (10–300MB per exam), ultrasound (30–50MB per exam), and digital pathology (2–3GB per slide). In addition, the following new imaging techniques and longer-term studies are creating richer, more detailed, and voluminous scans as well as longitudinal image libraries that add even more data:

- **Enterprise imaging (EI)** refers to enterprisewide platforms that strategically centralize, manage, and govern Digital Imaging and Communications in Medicine (DICOM) and non-DICOM imaging data from multiple sources across different settings. EI transforms medical imaging through strategies and technologies that can unify the data silos that often result from imaging devices and solutions commissioned at various units, departments, and facilities over time. EI delivers more unified and efficient medical imaging environments, workflows, and life cycles, resulting in a higher likelihood by providers to extract value from data, especially when combined with imaging analytics and AI. In the United States, 30.9% of healthcare providers increased spend on medical imaging and archiving, with 60.4% of that spend dedicated to EI. One common challenge for IT when undertaking EI is to figure out how it can be used to treat non-DICOM data as IT ties the various "-ologies" together (i.e., radiology, cardiology, oncology, gastroenterology, neurology, nephrology, dermatology, ophthalmology, and pathology) or risk data growth outpacing platform scalability.
- **Encounter-based imaging (EBI)** occurs when image acquisition was not the intended purpose of a clinical encounter or there was no prior order for a medical image to be taken. EBI has seen increased utilization in dermatology, wound care, infectious disease management, and emergency department triage adding to imaging data growth. EBI enables more holistic approaches since there is no delineation between the encounter and imaging procedure, yet also requires robust integration, data management, and workflow optimization. Images need to get attached to the patient record so that the encounter is linked for future reference. Therefore, acquired images must be seamlessly integrated with associated metadata about the patient, encounter, and imaging procedure, which many departments cannot do because they still rely on "order-based imaging" and legacy functionality for image reports.
- **Digital imaging in pathology** has also witnessed an uptick in utilization lately, leading to additional imaging data growth. Whole slide imaging sparked a shift in pathology as a first step to unlocking the potential to integrate image features into high-dimensional analysis for personalized and precision medicine. This development leads to new opportunities, including digital collaboration, integration with EHRs, and AI-based diagnostic support solutions. Digital pathology will drive a need for extensive storage and information life-cycle management systems. Organizations will further seek options that guarantee fast access to high-quality digital imaging data to support pathologists when reviewing cases through essential functions from analysis to consultations for pathology image management.

Artificial Intelligence

AI shows vast potential to transform healthcare and countless application areas that require the right data infrastructure. AI holds much promise to impact clinical and operational use cases ranging from the data-driven use cases covered previously in this White Paper to AI-assisted/automated diagnosis, conversational chatbots, virtual triage, care automation, symptom checkers, prescription auditing, analytics (e.g., medical imaging, population health, claims, and third-party data), robot-assisted surgery, drug discovery, gene analysis, medical device studies, care/service automation, and fraud detection. Cutting-edge use cases include AI-enabled command centers for intelligent capacity planning and AI-embedded modern control towers for supply chain management. However, AI-driven predictive analytics in operational workflows represents the majority of use case adoption today.

Advances in AI have been accelerating. Yet even as AI can be a real differentiator, the technology faces barriers for adoption. IDC data reveals that only 17% of AI initiatives are in production, 15% are in development, and 17% are in a proof-of-concept stage, while over half of AI initiatives, or 51%, have failed. Therefore, there has been some skepticism, risk adversity, and regulatory challenges. Anecdotally, the closer the use case gets to clinical processes, the greater the barriers to adoption due to the increased risk level associated with the clinical sensitivity of the data, the processes, or the outcomes. The FDA recently took steps to streamline the regulation of AI algorithms to lessen the burden on the industry.

Healthcare organizations need to have an infrastructure strategy in place to develop an effective AI program. AI infrastructure extends beyond technology to also include resources, policies, and people. Furthermore, AI runs on data from internal and external IT systems. Therefore, data protection is critical as data flows in and out of algorithms, necessitating IT infrastructure excellence and good data governance. Clinical data is complex and extensive. AI needs a scalable and high-performance infrastructure that supports unpredictable and exponential data growth along with flexible data management from on-premises to hybrid and multicloud instances.

INFRASTRUCTURE CONSIDERATIONS IN THE DIGITIZED HEALTHCARE ERA

NVMe

Nonvolatile memory express (NVMe) is a storage protocol that is rapidly replacing the use of SCSI for performance-sensitive workloads. NVMe is optimized for use with solid state media, utilizing its performance and capacity resources much more efficiently than SCSI. NVMe enables significantly more parallelism in storage processing than SCSI as well (a feature that is particularly important in this era of multicore CPUs) and leads directly to a much more efficient IT infrastructure overall.

In addition, NVMe opens up the use of emerging persistent memory technologies like Intel's 3D XPoint (pronounced "crosspoint")-based Optane solid state disks (SSDs) and NVMe over Fabrics (NVMe-oF) host connections. Although not all healthcare organizations will need NVMe, IDC recommends that IT managers understand what it has to offer as they plan out storage infrastructure refreshes since it is highly likely they will need the level of performance provided by NVMe for at least some workloads in the near future. Since NVMe delivers much lower latencies and higher performance density than SCSI, its use can often result in much more streamlined and efficient IT infrastructures that require far fewer application servers and cost less than those based on SCSI. In many large-scale clinical settings, there will be not just a performance but also an economic argument for the use of NVMe technology.

Data Integrity, Protection, and Resiliency

To ensure that healthcare data is always accurate and accessible, healthcare IT departments will first want storage systems to support data integrity features like the use of error correcting code (ECC), checksums, and the T10 Data Integrity Field (DIF) standard. Storage systems will also need to be both highly available and resilient (i.e., able to transparently recover from various types of failures without putting data integrity at risk or cutting off data access). Highly available storage systems have both local and remote recovery capabilities. From a local point of view, ensuring that systems can transparently recover from various failures requires hardware redundancy (i.e., dual controllers, redundant fans, and power supplies) as well as software features like host multipathing and RAID or erasure coding that is configurable to different levels of resiliency. Hot pluggable, field replaceable units allow failed components to be replaced without impacting service availability. Features like these ensure that systems will continue to operate despite disk, controller, or other component failures.

To enable data recovery when needed, healthcare IT managers must institute snapshot-based data protection regimens that enable rapid local recovery as well as create remote copies of data to make available in the event of catastrophic failures like the loss of an entire storage system or site. Synchronous and asynchronous replication provide options that can maintain remote copies while giving IT managers an ability to trade local application performance off against potential data loss on recovery, as necessary. "Air gapped" remote copies use a combination of snapshots, purpose-built backup appliances, and replication to maintain recovery copies that provide protection against data corruption or malware attacks. IT managers can also rely on more traditional options such as using backup software to do backups to virtual tape libraries or physical tape.

Cloud can also play a critical role in data protection, whereby purpose-built backup appliances "cloud tier" data to the cloud — that is, automatically replicate data from on-premises infrastructure to the cloud. These cloud repositories ensure data survival for any on-premises disruption and, when properly air gapped, can enable effective recovery from ransomware or malware attacks. These cloud copies can also be potentially leveraged for data analytics, testing, training, and other productive use cases.

Storage Efficiency

Given the very large storage capacities most healthcare providers will be maintaining, storage efficiency technologies like thin provisioning, compression, and data deduplication can be particularly compelling because of their ability to maximize capacity utilization. All these technologies minimize physical storage capacity consumption while maximizing the actual amount of data that can be stored. The savings associated with storage efficiency technologies will vary with different data types, and storage administrators should be aware that medical images will be much less reducible than more text-heavy content. Yet the use of storage efficiency technologies will be important, given the amount of data that has to be managed. Even low storage efficiency ratios can result in significant infrastructure savings.

Storage Tiering

In terms of scalability, healthcare IT managers should consider both a storage platform's native storage capacities and the platform's ability to automatically tier storage to archive tiers such as object-based storage or the cloud. With regulations like HIPAA requiring healthcare providers to retain medical records for the life of the patient, an effective archive strategy is a must. Archive tiers will be where most patient data is stored, so they must be highly scalable and offer very aggressive dollar-per-gigabyte costs, yet keep data easily accessible to meet different use cases including clinical documentation review, analytics, and research.

The cloud can offer additional options for tiering and archiving data that can be cost effective, particularly for data that must be retained over long periods and is not likely to be accessed frequently. Object-based storage tiers are generally the preferred platform for low-cost, long-term data retention, and the S3 API is the industry-standard way to communicate and exchange data with object tiers. Most healthcare providers will prefer to at least have the option to implement cloud environments to give themselves access to more types of cloud-based services and to ensure they can benefit from price competition between different cloud providers.

Security

Maintaining the privacy of patient data is of the utmost importance to patients and healthcare providers alike. The same is true for many life science and pharmaceutical companies that supply products and services for the healthcare space and for universities and other organizations doing research. The rise of regulations like GDPR is largely driven by privacy concerns. To ensure data remains secure and protected from hackers and other "bad actors," healthcare IT managers need the flexibility to encrypt their data, often both at rest and in flight. The gold standard for this is AES 256 encryption. Regulation in many private sector industries requires at least 256-bit encryption, and the federal government stipulates it when FIPS 140-2 compliance is required. While HIPAA requires at least 128-bit encryption, IDC recommends the stronger 256-bit approach.

Healthcare IT managers that want to use both encryption and data reduction (i.e., compression and/or deduplication) together will need to ensure that their storage systems can support that approach. Depending on their implementation, some storage systems allow the use of one or the other but not both together. Also, the ability to apply encryption without impacting storage latencies is a critical concern for many performance-sensitive workloads. Hardware-assist approaches can be used by storage vendors to speed encryption and/or compression times, whereas many NVMe-based systems are fast enough to use software-based encryption and/or compression without impacting application performance. This level of data security must extend to data protection systems, whether encrypted backups in flight or at rest, encrypted purpose-built backup appliances, or encrypted tapes.

Agile Approaches

As IT has become more central to driving quality patient care, it also must respond to an increasingly dynamic healthcare space. When it comes to IT, agility to data access speed and an ability to respond to requests for new IT infrastructure or new storage capacity very quickly are paramount. Systems that are simple and easy to deploy, enable rapid storage provisioning, include GUIs that implement policy-based storage management, and offer a means of scaling performance (i.e., storage controllers) and capacity independently as needed to achieve a well-balanced growth path meet the need for agility. Storage infrastructure should also include APIs that allow storage system features to be easily integrated into automated operations and, if systems will be used in VMware virtualization environments, offer VMware integration through APIs.

For healthcare organizations that are developing their own applications, the need for agility also applies to supporting newer, more agile methods of application development and deployment that maintain software quality while improving responsiveness. In this case, there are additional features in storage infrastructure that can help better support DevOps. Chief among these is support for containers (through the container storage interface [CSI]) and the Kubernetes container orchestration platform. In DevOps environments, the integration APIs mentioned previously can be used to support self-service portals that IT may create to further streamline DevOps operations.

STORAGE SOLUTIONS FROM DELL TECHNOLOGIES FOR HEALTHCARE

Dell Technologies is a $92 billion multinational technology company whose product offerings cover both personal and business computing. As a leading provider of enterprise IT infrastructure solutions, Dell Technologies' portfolio spans servers, storage, networking, storage infrastructure software, security, and cloud-based solutions (both public and private). The vendor's enterprise storage offerings include:

- Entry-level block storage: PowerVault
- Midrange unified storage: PowerStore and Unity XT
- High-end block storage: PowerMax and PowerFlex
- Scale-out file storage: PowerScale
- Scale-out object storage: ECS
- Data protection: PowerProtect
- Cloud storage: Dell Technologies Cloud Storage and Multi-Cloud Services

To help organizations more effectively consolidate workloads as they refresh storage infrastructure, many of these systems support multiple access methods, while all deliver enterprise-class performance, availability, scalability, and security. Through its VMware subsidiary, Dell Technologies provides software-defined datacenter solutions based around VMware vSphere technology, a platform that enables hybrid, multicloud environments with support for both on-premises and cloud-based virtualization infrastructure.

Dell Technologies is a major provider of storage infrastructure for healthcare environments. The vendor has established relationships with key application providers in healthcare offering integrated solutions in the EHR, medical imaging, edge, and cybersecurity arenas for healthcare providers.

Dell EMC PowerMax

The Dell EMC PowerMax is Dell Technologies' flagship enterprise storage platform. Featuring NVMe and all-flash technology and capable of delivering "six-nines plus" availability, this platform is designed for dense consolidation of an organization's most mission-critical, latency-sensitive workloads. The PowerMax includes host multipathing, thin provisioning, dual parity RAID, inline data reduction (both compression and deduplication) that does not impact application performance, quality of service (QoS) to ensure that workloads continue to meet their performance service-level agreements (SLAs) even when densely consolidated, and FIPS 140-2-certified AES 256 encryption. The system supports space-efficient snapshots, replication (both synchronous and asynchronous), remote site disaster recovery (DR), and active/active stretch clusters for the highest levels of availability. Management is through a GUI, which includes a variety of wizard-driven workflows to streamline administration and includes a CLI and REST APIs (e.g., CSI) to enable easy integration with Kubernetes and other automation and orchestration tools.

The PowerMax boasts the features needed in both virtualized and hybrid cloud environments. The array is block based, although it does support a NAS gateway, minimizing the need to maintain separate file servers in virtual environments. Supporting extensive VMware integration, the system also provides support for cloud-native, container-based environments through CSI. External tiering of data to the vendor's ECS object storage platform as well as public cloud environments (Amazon, Microsoft, and Google) is supported through S3. With its extensive feature set, the PowerMax provides the performance, availability, scalability, security, and agility needed for mission-critical, latency-sensitive healthcare workloads like EHR and other relational database-driven applications.

Dell EMC PowerStore and Unity XT

Dell Technologies offers two midrange "unified storage" arrays, the PowerStore and the Unity XT. Both these arrays can simultaneously run block- and file-based applications natively, enabling broader workload consolidation and delivering high performance and efficiency across both workload types. The NVMe-based PowerStore offers five different all-flash models to provide customers with a variety of midrange price points and includes enterprise-class storage management features like host multipathing, dual parity RAID, thin provisioning, inline data reduction (compression and deduplication), QoS, and AES 256 encryption. PowerStore delivers "five-nines" availability and supports space-efficient snapshots, replication (both synchronous and asynchronous), and remote site DR, as well as the same VMware and automation integration APIs as the PowerMax.

While PowerStore can be deployed as disaggregated storage, the system also includes the unique AppsON feature. With PowerStore's "hypervisor" mode, applications can be run directly on the storage controllers, providing very low-latency access to storage internal to the appliance. By combining compute and storage services in the same appliance, AppsON can be used to streamline edge infrastructure, simplifying these deployments down to a single piece of hardware that both runs applications and provides storage services. These small, simplified but powerful edge configurations are a good fit for medical informatics workloads. For customers requiring a midrange storage solution for EHR and relational database-driven, mission-critical, and latency-sensitive workloads and/or looking to consolidate block- and file-based workloads onto streamlined infrastructure, the Dell EMC PowerStore provides the performance, availability, scalability, security, and agility needed.

The 12Gb SAS-based Unity XT provides performance, availability, and value in a midrange appliance package. Available in both all-flash and hybrid flash (i.e., including both SSDs and HDDs) configurations, Unity XT supports the same enterprise-class storage management features as the PowerStore — including an ability to meet "five-nines" availability requirements — but adds FIPS 140-2-certified encryption. Unity XT does not support federated clustering or AppsON. For environments that do not require NVMe performance, Unity XT provides a high-availability solution for EHR, medical imaging, and other clinical application workloads that delivers excellent value for the money.

Dell EMC PowerFlex

For healthcare environments looking for the flexibility of a software-defined, scale-out, and block-based solution, PowerFlex is a suitable option. PowerFlex provides enterprise-class storage management, including host multipathing, thin provisioning, RAID or erasure-coded (EC) data protection, inline data reduction, space-efficient snapshots, QoS, AES 256 encryption, and replication. Entry-level three-node configurations are highly available; nodes can be configured using either NVMe or 12Gb SAS, and clusters can be scaled up to hundreds of nodes to provide high-end storage capacity. PowerFlex enables nondisruptive scalability as well as nondisruptive multigenerational technology refresh, a feature that can extend the life cycle of the platform considerably.

PowerFlex's entry-level configurations provide opportunities to deploy PowerFlex in edge or distributed environments, but its massive scalability also makes it an excellent datacenter platform when customers want the flexibility of software-defined approaches. The clustered storage system can handle highly available, latency-sensitive workloads like EHR and other relational database-driven applications.

Dell EMC PowerScale

The Dell EMC PowerScale is a software-defined, highly available, and scale-out file system environment that can scale from small, two-node edge configurations out to hundreds of nodes that can support tens of petabytes of raw storage capacity. Running the (Isilon) OneFS distributed file system, it can accommodate all-flash, hybrid flash, and HDD-only nodes, while the SmartPools feature gives customers the opportunity to create storage pools within a cluster that use NVMe (to support extremely low latency and high throughput) or SAS (for more cost-sensitive archive requirements). CloudPools allow the seamless tiering of data to public or private cloud providers. PowerScale supports multiple access methods (NFS, SMB, S3, HTTP, FTP, HDFS, and internet protocols IPv4 and IPv6) to maximize the opportunities for workload consolidation and data sharing. All data can be simultaneously read and written through any protocol, allowing the data to be easily shared across multiple applications with different access methods.

PowerScale supports the performance, availability, scalability, security, and agility healthcare IT managers need for general-purpose file sharing, big data, and healthcare vertical market workloads. The OneFS operating environment supports inline compression and deduplication, snapshots, AES 256 encryption, and SmartLock (for file immutability and retention requirements). Replication is supported both for local data protection and for data distribution, and working in conjunction with features like SyncIQ, IT managers can create PACS caches, for example, in remote locations that are maintained in centralized core infrastructure locations. In the healthcare market, PowerScale is deployed for PACS, VNA, EHR, digital pathology, encounter-based imaging, telehealth, precision medicine, enterprise file sharing, and big data and analytics (both real time and batch oriented).

Dell EMC ECS

Dell EMC ECS is a software-defined scale-out object storage system that can cost effectively handle exabyte-class capacity requirements within a single global namespace. Sold under an appliance model, ECS provides a storage-as-a-service customer experience, with multitenancy, detailed metering, a self-service portal, metadata search and management, and billing integration for chargeback in flexible configurations that can support between 144TB and 3.8PB in a single datacenter rack. Delivering cloud economics for on-premises infrastructure, ECS supports S3, NFS, and HDFS access methods in a highly available, disaggregated storage platform that provides geo-distributed data protection and no single point of failure.

As a scale-out, object-based streaming data platform with enterprise-class capabilities, ECS is well suited to support archive requirements in healthcare as well as analytics and cloud-native applications. Entry-level configurations as small as three nodes enable cost-effective edge or remote deployments, while the system's ability to scale to hundreds of nodes makes it equally appropriate as a central archive or big data repository in core locations.

Dell EMC PowerProtect

Dell Technologies arguably has one of the broadest portfolios of data protection products in the industry. Its PowerProtect solutions complement the native snapshot and mirroring capabilities of Dell's storage systems. Whereas snapshots and mirrors can be used to quickly recover from local system outages and user error, Dell's data protection software and purpose-built backup appliances provide recovery capabilities extending to the most extensive DR and cyber-recovery requirements. These systems can meet the data protection and recovery requirements that healthcare organizations need to create a solid foundation for data availability.

The Dell EMC data protection portfolio consists of data protection appliances and software that allows organizations to protect and recover their critical data assets across edge, core, and cloud-based infrastructure. With support for traditional (physical, virtual) and modern (containers, Kubernetes, and cloud-native apps) workloads, combined with integrated cyber-resiliency capabilities, these solutions provide healthcare organizations with the end-to-end data protection needed to ensure data is protected, secure, and recoverable.

Dell EMC PowerProtect Cyber Recovery Solution delivers an air-gapped approach that strengthens and builds upon the base immutability offering through retention lock capabilities. CyberSense analytics assists with assured data recovery; CyberSense reads through the backup set, so there is no need to restore data just to determine if it is clean, thereby eliminating the risk of opening harmful data in the vault. CyberSense can also evaluate the full contents of the file, not just its metadata, to deliver superior analytics. In addition, running analytics on the data in the vault is an important component to enable a rapid recovery after an attack. Analytics helps determine whether a data set is valid and usable for recovery (free from malware) or has been improperly altered or corrupted, making it "suspicious" and potentially unusable. Healthcare organizations are common ransomware attack targets, and Dell EMC PowerProtect Cyber Recovery Solution can provide assured data recovery with a minimum of data loss.

Dell EMC purpose-built backup appliances, available in either integrated appliances (PowerProtect DP series), hardware only (PowerProtect DD series), or software only (Virtual Edition), have the integrated features and technologies healthcare organizations need to satisfy the top drivers for data protection refresh. Easy scalability up to the petabyte range, assisted with inline compression and deduplication that can deliver up to 65:1 data reduction ratios, cloud tiering capabilities, integrated solid state storage options that speed both backup ingest and data recovery, and the flexibility to accommodate various deployment models (bare metal, virtual machines, and containers), enables these systems to deliver on the promise of modernized infrastructure to meet evolving data management and protection requirements.

Dell EMC data protection solutions are fully integrated with vSphere to enable VM admins to manage data protection directly from the native vSphere UI. With Dell EMC advanced VMware integration, VMware admins are empowered to more efficiently control their own data protection strategies, resulting in faster backups and restores for virtualized mission-critical applications. In addition, PowerProtect Data Manager delivers support for vSphere 7 and Tanzu, paramount for protecting modern container-based applications and Kubernetes workloads.

Project APEX: Delivering a Flexible Range of Consumption Models

As healthcare IT managers look to make the most of limited budgets, they need more than just outright purchase consumption models for IT infrastructure. While most IT organizations were introduced to subscription-based pricing models through the public cloud, the ability to buy using a "pay-as-you-go model," offload infrastructure management to external providers, and/or move assets off the balance sheet can be very attractive for certain types of on-premises infrastructure as well.

Project APEX, which becomes available in the first half of 2021, will give customers a variety of consumption model options (both opex and capex) to meet the following requirements:

- Dell Technologies Cloud is a self-managed hybrid cloud offering that combines the power of VMware technology and Dell infrastructure to provide a consistent operating model and simplified management across public cloud, private cloud, and edge locations through a unified control plane called the Dell Technologies Cloud Console.

- Dell Technologies Cloud Storage for Multi-Cloud provides for native replication between on-premises Dell storage and public cloud-based infrastructure at Amazon Web Services (AWS), Microsoft Azure, or Google Cloud Platform. This offering supports multicloud agility with no public cloud provider data lock-in. A fully automated disaster recovery-as-a-service offering leverages VMware Cloud running on AWS public cloud infrastructure, providing an ability to rapidly recover and run workloads in the public cloud while a primary site failure is addressed.
- VMware Cloud on Dell Technologies is a datacenter-as-a-service offering that is deployed in customer premises but maintained and supported by Dell. It gives IT managers the option to offload infrastructure management responsibilities, allowing IT to more fully focus on driving business value.
- Dell Technologies on Demand provides a usage-based consumption model for any self-managed on-premises Dell Technologies infrastructure.

These programs are in addition to outright purchase, which is still an option for on-premises IT infrastructure, and apply to all Dell Technologies IT infrastructure offerings, not just storage.

CHALLENGES/OPPORTUNITIES

Healthcare IT managers have many decisions to make as they refresh storage infrastructure as part of a move to digitized medicine. It is IDC's view that, for most healthcare organizations, a hybrid multicloud strategy that offers traditional on-premises infrastructure, private cloud infrastructure, and public cloud infrastructure as deployment models for workload placement will be the right strategy. As new technology developments impact the pros and cons of each deployment model, Dell Technologies will need to keep abreast of these changes and continue to provide recommendations to its prospects and customers that help them craft the best IT infrastructure solutions for their needs.

Given Dell Technologies' broad storage infrastructure portfolio, the vendor's products have the features and capabilities healthcare customers are looking for as they modernize their storage infrastructures. The vendor offers a flexible range of deployment and consumption models as well, which means it can meet a broader range of healthcare data management use cases than most of its competitors. To date, Dell Technologies has demonstrated an ability to clearly connect storage infrastructure features and capabilities with customer requirements in healthcare, delivering the performance, availability, scalability, security, and agility that customers are looking for as part of modernization efforts. To the extent that Dell Technologies can continue along this path, the vendor has a good opportunity to stay a market leader in healthcare.

CONCLUSION

While the healthcare industry was already moving toward more digitized models for care delivery, the global pandemic has significantly accelerated that migration. The explosion in patient data, use of AI/ML technologies to assist in diagnosis as well as therapeutics, and increasing use of telemedicine for care delivery all depend heavily upon an IT infrastructure that transcends the capabilities of legacy systems. IT services are now a critical component of care delivery that must always be available and highly performant. Regulations like GDPR and HIPAA are raising the bar for security and demanding new approaches that simultaneously deal with data corruption and ransomware attacks while guaranteeing patient privacy. And the increasing pace of operations in healthcare demands a more responsive IT organization and infrastructure.

Storage platforms today must deliver performance at scale, high availability, multi-petabyte scalability, proven security, and flexibility. They must support cloud integration as healthcare providers look to the cloud for new services, lower-cost DR and business continuity solutions, and access to cost-effective capacity for long-term retention. The mismatch between legacy storage infrastructure and the requirements for digitized healthcare is driving those organizations that are currently undergoing DX to refresh their in-house storage infrastructure. Healthcare IT managers are looking to leverage storage technologies like NVMe; software-defined, AI/ML-driven automated operations; and cloud-based data protection to help them meet evolving requirements.

Dell Technologies is a leading IT infrastructure provider with a rich storage and data protection portfolio of healthcare solutions. Healthcare IT organizations will find the technologies they are looking for as part of DX, regardless of whether their data is structured (block) or unstructured (file/object), in the storage solutions offerings from Dell Technologies, available in a variety of different deployment and consumption models that give them the choice to craft the solutions that best fit their individual requirements.

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-community.com
www.idc.com