

# How Healthcare IT Meets the Need for Cost-Effective, Highly Available, High Performance Storage

---



Copyright 2018 ©DataCore Software Corporation – All Rights Reserved



# TABLE OF CONTENTS

CLICK THE SECTION TO JUMP AHEAD

Introduction: Addressing Healthcare’s Critical Need for Fast and Reliable Data Availability and Accessibility .....	3
ANALYST ARTICLE	
How Best to Protect and Access Your Company’s Data .....	5
CASE STUDY	
Maimonides Medical Center .....	9
ANALYST ARTICLE	
Data Storage 101: How to Determine Business-Critical App Performance Requirements .....	16
CASE STUDY	
Hanover Hospital .....	20
SOLUTION BRIEF	
Solving Healthcare IT Challenges with DataCore Software-Defined Storage.....	25
Additional Resources .....	32
Webinar: Maintaining Mission Critical Levels of Availability in Healthcare	
Request a Live Demo	
Learn More About Healthcare Data Storage Solutions	
Appendix	
SDS Data Sheet .....	36



## **ADDRESSING HEALTHCARE'S CRITICAL NEED FOR FAST AND RELIABLE DATA AVAILABILITY AND ACCESSIBILITY**

Today's hospitals and health systems are taking advantage of an explosion of new technologies to help run their facilities more efficiently, improve patient care and keep costs contained. But with these new IT innovations comes an influx of data and a growing responsibility for hospital IT teams to manage and store that data while ensuring that it is secure and always available.

Every vendor promises real-time, always-on hospital data access, but in reality, it's not a checkbox item. For starters, there are important differences between data protection architectures and technologies. Hospitals need to consider which one is easier to implement and how much flexibility is offered with each approach. What manual intervention is required to fail-over, and what is the impact to care teams and potentially life-saving applications? How simple is the fail-back to primary systems and restoring the system to the state before an incident?

In addition to the need for high-availability, skyrocketing storage costs are eating away at already lean hospital IT budgets that are challenged by the move to value-based care. As healthcare organizations continue to expand their infrastructure to accommodate new digital technologies, the volume of data collected and stored is expected to continue to increase accordingly.

### **How DataCore Can Help Your Healthcare Organization Overcome the Most Complex Storage Challenges**

DataCore software-defined storage is ideal for helping healthcare organizations address these complex challenges by forming a transparent virtualization layer across diverse storage systems in order to maximize the availability, scalability and performance of all storage resources. The company provides hospitals with 24-7 uninterrupted data access, massive architectural scale-out and

game-changing application performance. With DataCore, hospitals and other providers can improve performance without requiring a rip and replace of existing systems or the purchase of new storage solutions. Rather, by providing a virtual storage pool that dynamically assigns workloads to optimal storage tiers based on performance requirements, healthcare organizations are empowered to scale up or scale out applications while efficiently using all of their existing resources.

With data being the lifeblood of a modern health care system, DataCore is the trusted provider for over a 1,000 hospitals and health systems today. Its patented technology eliminates storage bottlenecks with adaptive parallel I/O optimization, enables zero-downtime synchronous mirroring, and provides a true hardware-agnostic architecture – resulting in flexibility, resource efficiencies, and cost savings. Having pioneered software-defined storage and now expanding its technology leadership to hyperconverged infrastructures, the company empowers healthcare IT organizations to achieve always-available, high-performance and highly efficient data, which is critical to the core mission of providing care and saving lives.

# **How Best to Protect and Access Your Company's Data**

---

ANALYST ARTICLE

ARTICLE BY: Trevor Pott, Virtualization & Cloud Review

# HOW BEST TO PROTECT AND ACCESS YOUR COMPANY'S DATA

As much as preserving our data matters, access to that data is equally important.

**In 2018**, you have to work to find an organization that isn't utterly dependent on computers. Data has become as important as any tangible asset. As much as preserving our data matters, access to that data is equally important. Uptime matters.

One of the most important assets any organization has is its data. Unfortunately, as we are all aware, many organizations ignore data protection until it's too late. Ransomware became such a problematic epidemic specifically because of the habitual neglect that organizations have for basic backups.

In 2016 Dell EMC ran [a survey](#) on data protection. This survey reported that the average cost of data loss was \$900,000 USD, and the average cost of downtime was \$555,000 USD. As the cost of both data loss and down time have only gone up throughout my lifetime, it's reasonable to assume that if you were to run the same survey today, those costs would only be higher.

Unfortunately, backups are easy to see as a nice-to-have

instead of a must-have, and this is exactly the approach both organizations and individuals seem to take by default. For this reason I've written a seemingly endless number of articles about backups. I'm sure I'll write many more.

Yet backups aren't the only aspect of data protection that organizations ignore. Storage availability is all too frequently neglected, often with punitively expensive results.

## That Speed of Light Thing

Backups exist to protect against data loss. This loss can be because data is deleted or overwritten, or because the storage device upon which it rests has been destroyed. Because backups primarily exist to protect against statistically unlikely events, restoring from backups is rarely expedient.

In addition, there often exists a gap in backup data called the recovery point objective (RPO). RPO is a measure of the time between backups. More important, RPO is a measure of the maximum amount of data that might be lost if you have to roll back to the last backup.

Data is lost when rolling back to a backup because it takes time for data to travel between the production system and the backup location. Even if you're using continuous data protection (CDP), which is essentially real-time streaming of writes from the production location to the backup location, the speed of light says that there will always be a delay in getting data from A to B.

That delay represents the theoretical minimum RPO of any backup solution. In many cases the delay is only a few milliseconds, but it only takes a few milliseconds to store dozens or even hundreds of transactions on a Web



site, or a critical update to a human resources record.

Within a single datacenter, the distance that data would have to travel between any two systems is so short that the travel time is less than the speed at which storage devices perform transactions. In many cases, real-time replication of data is possible between two systems within the same city, with a distance between the two systems of 100km being the generally accepted maximum distance.

100km isn't very far, in disaster avoidance terms. A power failure, hurricane, earthquake or other event could easily affect both locations. For this reason, backups are typically stored at much greater distances from the production system, even though this means that rolling back to the latest backup could mean data loss: data is sacrosanct, and it must be protected against even statistically unlikely events.

When the primary storage has a little lie down—regardless of the cause—avoiding the costs of downtime requires that there be a secondary storage array to take over.

### High Availability vs. Backups

Backups are incredibly useful for recovering deleted files, but—because of that RPO gap—far less so for running workloads. That being said, the equipment on which running workloads operate can, does and will eventually fail. Nothing lasts forever.

Because the cost of downtime is so high, investment in technologies that allow for real-time replication of data between two or more storage devices with no data loss (an RPO of zero) is usually called for. This is high availability (HA). As already discussed, the distance between storage devices in an HA solution is largely dictated by the speed of light.

For many organizations, HA consists of two storage devices located physically one on top of the other on the same rack within a datacenter. For some organizations, HA will be performed between storage devices located at two different datacenters that are situated on a Metropolitan Area Network (MAN). In the finance industry in particular,

it's common to have two datacenters on a MAN providing HA, and a third datacenter in another geographic location as the disaster recovery site.

### Data Migration Affects Uptime

It's easy to understand HA in terms of device failure. Bobby Breakit trips over his untied shoelaces, takes a header into the storage rack, and makes the primary storage array go boom. Bad Bobby, do not break it.

There are, however, a number of other perfectly routine and innocuous reasons to invest in HA; updates, for example. Everything in IT needs to be updated eventually. Many HA storage solutions can be set up to seamlessly switch between physical devices, allowing administrators to apply updates to the secondary device while the primary one continues to serve workloads. Fail back over to the other device, and you can apply workloads to that one, too.

When the primary storage has a little lie down—regardless of the cause—avoiding the costs of downtime requires that there be a secondary storage array to take over. In a perfect world, the handoff between the two is seamless, and running workloads never notice that their storage is now being served from a different device.

Not all storage solutions are capable of this kind of HA. Some solutions can replicate data between two devices in real time. However, failover between the two devices is not seamless, and workloads will crash if a failover event is forced. These less advanced HA solutions don't lose any of the data written to the storage devices, but require workloads to be restarted if a failover is forced, which causes at least some downtime.

This latter failover scenario is quite common when data movement is occurring between two dissimilar storage devices. Data migration activities due to datacenter upgrades, or due to space saving or archival efforts are common causes of outages. These outages are often a "hidden cost" of storage upgrades and maintenance: in addition to the cost of the equipment, you must factor in the time necessary to switch over to new devices, or repoint workloads to data that's been moved.

### Avoiding Downtime

While all of this is still true for traditional storage arrays, storage technology has come a long way in the past decade.

Data fabrics—often nebulously described under the nearly useless header of “software-defined storage”—remove the need for downtime, even when adding or removing storage devices from an organization’s datacenter.

The short version of how data fabrics work is as follows: a highly available cluster of servers acts as a data presentation layer. In turn, this data presentation layer controls any and all storage that it’s fed.

## Backups are a good and necessary thing, but because of the RPO gap, most organizations prefer to avoid rolling back to them if at all possible

The data presentation layer presents storage to running workloads. Because this presentation layer is itself an HA cluster, the presentation layer can survive the loss of a node that’s performing data presentation activities.

The data in a data fabric is spread across all storage devices made available to the fabric. This storage could be traditional storage arrays that have offered some or all of

their storage to the data fabric. And it could be in the form of whitebox servers full of disks.

The storage could also be drives attached to the nodes hosting the data presentation software itself. When a data fabric is designed in this manner, storage industry nerds call it “scale out storage.” When a data fabric not only puts storage drives in the data presentation nodes, but also allows administrators to run production workloads on those same nodes, it’s generally referred to as hyper-converged infrastructure (HCI).

Backups are a good and necessary thing, but because of that RPO gap, most organizations prefer to avoid rolling back to them if at all possible. Doing so requires HA, and data fabrics are the best solution currently available to provide HA.

The true beauty of data fabrics is that you don’t have to throw away your existing investment in storage hardware to take advantage of them. There are numerous solutions available that can marry your traditional arrays with whitebox servers, and even blend this with HCI. When your next storage refresh crops up, talk to your vendors about data fabrics: They could help avoid costly downtime, and they’re quickly becoming the mainstream storage solution of choice.



# Maimonides Medical Center

---

CASE STUDY



Maimonides Medical Center is a non-profit, non-sectarian hospital located in Borough Park, in the New York City borough of Brooklyn, in the U.S. state of New York. Maimonides is both a treatment facility and academic medical center with 705 beds and more than 70 primary care and sub-specialty programs. With a staff of nationally renowned physicians, Maimonides Medical Center strives to conduct quality research, care and education.

[maimonidesmed.org](http://maimonidesmed.org)

## Maimonides Medical Center's Quest for Operational Continuity Via Real-Time Data Accessibility

Hospitals like Maimonides Medical Center demand that critical healthcare data always be accessible, be stored compliantly, and be immediately retrievable — failure to do so can impact care, put lives at risk and subject hospitals to fines. Availability of data is paramount to a hospital's ability to deliver effective care.

Hospitals and health systems today face mounting storage challenges when it comes to securely and compliantly managing the rapidly growing amounts of patient and other business-critical data, while consistently providing real-time, always-on information accessibility. Maimonides Medical Center, based in Brooklyn, N.Y, is certainly no exception. The facility is among the largest independent teaching hospitals in the U.S. and also serves as a treatment facility with 706 beds and an academic medical center. Over 800 physicians rely on Maimonides' information systems to care for patients 24 hours a day, seven days a week.

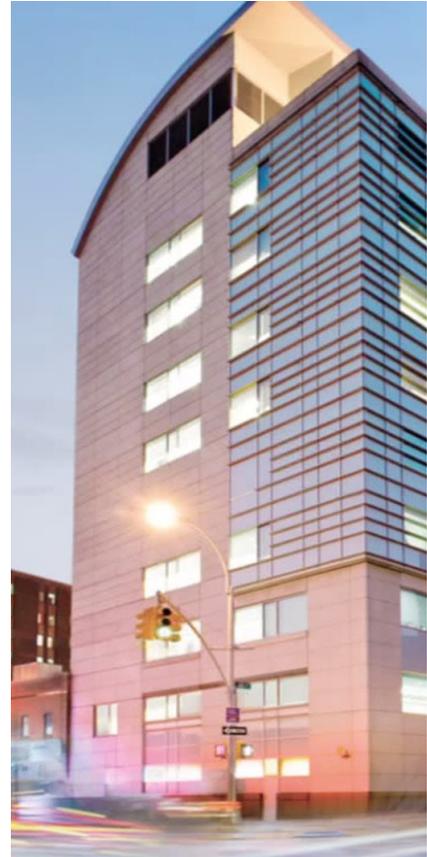
As healthcare organizations like Maimonides Medical Center continue their transition to a data-driven infrastructure, the influx of data generated continues to increase at a rapid pace with no sign of slowing down. For Rogee

Fe de Leon, head of the storage group at Maimonides Medical Center, it was becoming increasingly difficult for his team to maintain systems to keep real-time data at a doctor's fingertips while also keeping up with a growing demand for new and current Electronic Medical Records (EMR), prescription data, medical supply ordering and fulfillment, research data, clinical imaging data and voice dictation. In addition, the organization was challenged by having a number of disparate systems that were not working in conjunction with each other. This ad-hoc approach decreased efficiencies and intelligence around what departments were in need of storage resources. Historically, Maimonides Medical Center relied on storage that was directly attached to mission critical application servers, but that approach was no longer the most efficient or cost-effective solution for the hospital's mounting storage needs. As the volume of and reliance on hospital data evolved, so did the necessity for a fresh look at the organization's storage strategy. A lack of available windows for maintenance and the potentially devastating effects of downtime on the facility's patients and staff only prompted the urgency.

For Fe de Leon and his colleagues, ensuring business continuity through high data availability was a top priority for the IT initiative. This was to be accomplished by maintaining continuous, unbroken hospital operations. Maimonides Medical Center was also looking for a storage platform that was able to securely and seamlessly manage the hospital's storage resource expansion. On the data management side, the organization required an ability to consolidate storage management for mission-critical patient records, while eliminating the labor-intensive need for system administrators to micromanage the capacity requirements of life-saving applications.

### **Real-Time, Highly Available Hospital Data**

When looking for a software-defined storage partner to ensure for highly available storage in real-time, Maimonides Medical Center turned to DataCore Software, as the organizations have a long history of collaboration working together on previous storage initiatives. In this phase of the relationship, the hospital migrated to an infrastructure-wide, Storage Area Network (SAN) spread across two sites – two geographically separated SANs running as active-active data centers. The goal was to handle the data growth in the EMR application and meet existing state and federal require-



“

*High availability was the first and foremost reason for going with DataCore and for continuing with it. Now everything mission-critical to the running of the hospital is supported by DataCore's software-defined storage platform. Users not only receive faster access to data, but they benefit from more server capacity as well.*

**– Rogee Fe de Leon,  
head of the storage  
group at Maimonides  
Medical Center**

”

ments to store patient records for at least seven years or longer.

Eight direct Fibre Channel switches support this infrastructure, which stores, moves, and protects electronic records for ambulatory, obstetrical, and gynecological services. A little less than one petabyte (PB) of managed storage (“virtual disks”) is used as tier one storage. This storage serves the critical applications for medical records and imaging, including the Picture Archiving and Communications System (PACS), which is the most storage-intensive application. For Maimonides, meeting Health Insurance Portability and Accountability Act’s (HIPAA) requirement for audit trails has been a fairly straightforward process, keeping more log files for a longer period of time. The disaster recovery requirements of HIPAA have been met by replicating the patient data to a second SAN.

The most pressing requirement for Fe de Leon and his team was high data availability for the hospital’s operational and care consistency as four pairs of software-defined storage nodes have been running since 2005. Each pair represents 250 TBs of mirrored, virtual storage capacity, and the physical storage capacity available behind the pairs is almost twice that, or 500 TBs per pair. Total storage capacity is approximately 2 PBs. By centralizing the management of all storage resources as a scalable, fully redundant virtualized pool, the hospital ensures 24/7 access to critical information. As a result, Maimonides Medical Center has eliminated lapses in data availability from hardware failure and storage maintenance.

“High availability was the first and foremost reason for going with DataCore and for continuing with it,” said Fe de Leon. “Now everything mission-critical to the running of the hospital is supported by DataCore’s software-defined storage platform. Users not only receive faster access to data, but they benefit from more server capacity as well.”

## The Virtualized IT Landscape

The drive to deploy virtualized servers at Maimonides Medical Center was no different from most healthcare organizations across the country. The hospital's IT department proactively recognized that sheer demand for servers and for applications from every group within the medical center made sense, both practically and economically, and as such, the organization has significantly ramped up virtualizing servers over the past several years. The medical center now maintains 150 VMs across their systems. Out of a total of 12 VMware ESX hosts, six hosts are clustered into a production environment and the others are clustered in a development environment. The medical center also has numerous physical Microsoft servers, which are clustered between the two sites for the sake of business continuity. The balance of the hospital's hosts run IBM AIX (Unix) and Red Hat Linux, all obtaining their storage from the virtualized storage platform.

DataCore's software currently runs on IBM / Lenovo x3650 servers. These standard x86 machines have been deployed with approximately 20 Fibre Channel ports, each using a combination of Emulex and QLogic host bus adapters (HBAs). All critical systems run entirely on Fibre Channel topology. In the most recent deployment of software-defined storage, the configuration virtualizes one pool of fast Fibre Channel disks on X-IO arrays, as well as two vast pools of high-density, lower-cost, SATA disks on IBM arrays.

"At Maimonides we don't have the luxury of downtime... that's why we use DataCore, we are able to maintain our uptime at 100%," states Walter Fahey, the chief information officer at Maimonides Medical Center.

## Maimonides Medical Center Today

Maimonides Medical Center partnered with DataCore to implement the company's software-defined storage platform to eliminate single points of failure and ensure for the continuous, reliable data access. This storage capability is critical to the hospital in terms of operations as well as the ability to consistently deliver the highest levels of care.

“

*At Maimonides we don't have the luxury of downtime... that's why we use DataCore, we are able to maintain our uptime at 100%*

**- Walter Fahey,  
the chief information  
officer at Maimonides  
Medical Center**

”



*In terms of DataCore serving as the storage area network backbone for the hospital, you need only know that all of the hospital's medical records, all of its clinical records and of all of its administrative records reside on the DataCore virtualized SANS which serve as the hospital's 'de facto' virtualized storage platform. All of the applications we rely on to run the hospital – including billing – are on DataCore.*

**– Rogee Fe de Leon,  
head of the storage  
group at Maimonides  
Medical Center**



Today, the ability to have rock-solid business continuity remains the overriding benefit that Maimonides Medical Center derives from software-defined storage. DataCore makes it possible for the hospital to metro-cluster applications between two different sites as if they were co-located. With this approach, if for any reason one site happens to be offline (as a result of a planned or unplanned outage), the organization's IT systems will remain up and running – ensuring for reliable, continuous business operations at all times.

According to Fe de Leon, “In terms of DataCore serving as the storage area network backbone for the hospital, you need only know that all of the hospital's medical records, all of its clinical records and of all of its administrative records reside on the DataCore virtualized SANS which serve as the hospital's 'de facto' virtualized storage platform. All of the applications we rely on to run the hospital – including billing – are on DataCore.”

### **Maimonides Medical Center's Software-Defined Storage Goals**

- A storage platform that that was able to securely and seamlessly manage the hospital's storage resource expansion.
- The ability to consolidate storage management for mission-critical patient records.
- Elimination of the labor-intensive need for system administrators to micromanage the capacity requirements of life-saving applications.
- The ability to ensure for business continuity through high availability by maintaining continuous unbroken hospital operations.

### **Maimonides Medical Center Deployment at a Glance**

- DataCore Managed Capacity: 1 Petabyte
- Number of Users: 5,000
- Total Number of Physical Servers within the IT Infrastructure: 400+
- Number of Virtual Servers: 300
- Primary Server Vendor: IBM / Lenovo
- Storage Vendor: IBM and X-IO

## About DataCore Software

DataCore is a leading provider of software-defined storage and adaptive parallel I/O software that makes use of the most efficient and cost-effective modern server platforms to solve the most urgent problem of the storage industry: I/O bottlenecks. The storage virtualization and hyper-convergent virtual SAN solutions offered by DataCore facilitate storage management and free users from being tied to a specific manufacturer by providing hardware-independent architecture.

DataCore's software-defined storage platform revolutionizes storage infrastructures while working towards a software-defined data center of the next generation, with greater utility, better performance, higher availability and easier handling. Further information can be found at [www.datacore.com](http://www.datacore.com)

For additional information, please visit [datacore.com](http://datacore.com) or email [info@datacore.com](mailto:info@datacore.com)

© 2018 DataCore Software Corporation. All Rights Reserved. DataCore, the DataCore logo and SANsymphony are trademarks or registered trademarks of DataCore Software Corporation. All other products, services and company names mentioned herein may be trademarks of their respective owners.



# **Data Storage 101: How to Determine Business-Critical App Performance Requirements**

---

ANALYST ARTICLE

ARTICLE BY: Trevor Pott, Virtualization & Cloud Review

# DATA STORAGE 101: HOW TO DETERMINE BUSINESS-CRITICAL APP PERFORMANCE REQUIREMENTS

Measuring IOPS, latency and throughput are important for determining the characteristics of physical storage that's added to a modern data fabric.

**Storage is not one-size-fits-all**, and determining what the right storage is for your needs is difficult. There are many storage vendors out there, each selling a combination of complex engineering, bitter experience and marketing pixie dust wrapped up together into a saleable product. So how do you determine which bit of storage is the best for the job, and what measurements can help you accomplish this task?

As with most things, when you attempt to determine the relative value of storage there are both quantitative and qualitative measures to consider. Quantitative measures of storage are what can be readily determined through benchmarking. Input/output operations per second (IOPS), latency and throughput are the big benchmark measurements.

Qualitative measurements are more subjective. These include things like ease of use, or the importance and utility of various storage features to the organization considering the storage.

## Quantitative Measurements

IOPS, latency and throughput are highly interdependent measurements. As a general rule, the higher the IOPS the lower the throughput, and the more stress put on a system (IOPS or throughput), the higher the latency.

With the right benchmarking, it's possible to make some straightforward graphs of basic storage performance. On one side of the scale there's extreme IOPS with minimal throughput, on the other side there's extreme throughput with minimal IOPS. This effect is more dramatic with magnetic storage media, but the basic principle applies to solid-state storage, as well.

Storage obeys certain basic commands: read, write, modify and delete. These commands represent inputs and outputs (I/Os), and IOPS is a measure of how many of these



commands can be performed by a storage solution in a given second.

The commands issued affect the numbers obtained. Flooding a storage solution with 100 percent read requests will result in a different IOPS reading than 100 percent write requests. Magnetic storage media will respond differently to modify requests than solid-state media. The block size of the I/O requests (measured in [kibibytes](#)), the number of simultaneous requests, as well as the randomness of the requests, also impact the results obtained.

At first glance this may make IOPS seem a highly subjective measurement, but it's not. While a configuration of 70 percent write/30 percent read/64K block size/100 percent random will differ greatly from a configuration of 50 percent write/50 percent read/16K block size/50 percent

random, most storage solutions should perform identically under the same configuration.

This makes it important to pay attention to the testing profiles used by vendors or reviewers, and to compare storage solutions using the same testing profiles. A headline benchmark result of 1 million IOPS, for example, may have been achieved using a benchmark profile designed to give a maximum IOPS instead of one designed to mimic “real world” usage scenarios. In addition, each organization’s “real world” usage will differ from the next.

With all of that said, if you take a given storage profile and throw it at multiple storage solutions, the result is a good understanding of the performance characteristics of the underlying storage. Identical storage profiles for benchmarking form the basis for a rational comparison between storage solutions, however, the raw numbers can only tell so much about the overall value of that storage solution to an organization.

## Qualitative Measurements

In a perfect world, storage solutions are benchmarked both with all of their advanced storage features off and on. Features such as data efficiency, tiering, caching and so forth all affect how storage will perform under various circumstances, and they’re all highly variable based on the data being used for testing.

Consider two storage arrays: one has hardware-assisted deduplication and compression, the other performs these data efficiency tasks entirely in software. Let us assume that with data efficiency turned off, these arrays perform identically.

Running a benchmark against these systems with data efficiency turned on you might discover that both solutions perform identically up to a specific threshold. Once data volumes reach this threshold one of the arrays will reach the maximum amount of data it can perform data efficiency operations upon per second, and a demonstrable difference between the two arrays will have appeared.

Similarly, the arrays could use completely different approaches to different data efficiency. Different approaches to data efficiency have different performance costs, and produce different results in terms of data reduction. The effectiveness of data efficiency at the scale of a single storage array is also often quite different when compared

to solutions that span multiple devices and perform data efficiency tasks across the entire solution, instead of just at the level of an individual array.

Different types of storage have different levels of reducibility. RAW images, for example, compress more readily than JPEGs. Virtual desktop infrastructure (VDI) VMs generally provide a much higher level of deduplication than a collection of VMs all running different OSes.

In addition to this, storage features such as data efficiency typically compete with one another for array resources. Features like data tiering take processing power; if a storage solution is busy working on data efficiency, it might have to delay tiering tasks or vice versa. Which advanced storage features are enabled—and how they’re used—can make a noticeable difference on the performance delivered, even when comparing two arrays that would perform identically with those features off.

In a perfect world, storage solutions are benchmarked both with all of their advanced storage features off and on.

## The Advanced in Advanced Storage Features

At first glance, advanced storage features would seem to be quantitative. Features like data efficiency are affected by multiple variables, but if you can control enough variables they should hypothetically perform the same every time. There is some truth to this, but it’s also more complicated than that.

Fifteen years ago, how a storage solution performed data efficiency probably would’ve been a quantitative. The only data efficiency most storage used was compression, and an individual storage array would apply the same compression algorithm to all data, all of the time. That was then, this is now.

Today, storage tries to be smart. Some solutions will test a small piece of data to see how compressible it is before deciding whether or not to compress the whole data stream, or which algorithms to apply. Deduplication comes in flavors, and may only be applied on certain tiers of data, when the array is idle or in response to other parameters.

Tiering of data between different storage media—or even between different arrays, sites or clouds—can occur based

on any number of criteria, and the criteria themselves can change as the array learns storage patterns and adapts which features it implements under which circumstances. The smarter storage becomes, the harder it is to predict.

## Profiles, Policies and SLAs

Storage solutions are no longer discrete, self-contained items. An individual SAN or NAS is often just one part of a larger whole. When multiple arrays are joined together with server-local storage, cloud storage and who knows what else, an organization's storage becomes a data fabric.

A single data fabric today could easily join multiple on-premises sites' worth of storage, services provider storage and public cloud storage.

Data fabrics store data on multiple devices. These devices can use multiple physical storage media, be located across multiple sites, and even across multiple infrastructures, where the different infrastructures are owned and operated by different organizations. A single data fabric today could easily join multiple on-premises sites' worth of storage, services provider storage and public cloud storage.

Data fabrics typically have the ability to add or remove storage in a non-disruptive fashion, meaning that the overall physical composition of the data fabric is itself a variable. This changes how you must measure data fabrics, both qualitatively and quantitatively.

Hyper-converged infrastructure (HCI) and scale out storage are examples of simplistic data fabrics. HCI and scale out storage both take storage located inside individual servers, lash all that storage together and present it using a single interface. In the case of HCI, workloads are run on the same nodes that supply storage to the fabric, while scale out storage dedicates nodes to storage alone. Data fabrics can get much more complicated, however, and consist of any or all storage that an organization uses.

Because data fabrics are a logical construct instead of a fixed physical asset, you would rarely attempt to measure the

performance of the whole fabric. Instead, profiles, policies and service-level agreements (SLAs) are set in the data fabric, and tests are performed to see if the fabric can deliver.

What is the maximum ingestion rate of data into the fabric? Can the fabric deliver a LUN with x IOPS, y throughput and z latency? How many of these can it deliver simultaneously? Does the fabric warn when it has been asked to deliver performance beyond its capabilities, and how does it do this warning?

The response of fabrics to change is a critical measure. If you manage to ask more of a fabric than it's currently capable of delivering, how smoothly will it absorb additional storage into the fabric, and how quickly will that help meet the storage demands, especially if the fabric is overburdened? Does the fabric perform parallel I/O, or is everything funneled through a single orchestration node?

## Fabric Softener

The existence of data fabrics doesn't make classic qualitative measurements of storage irrelevant. Data fabrics are an orchestration layer that smooshes together multiple storage devices, layers a universal set of advanced storage features on top, and then presents a single storage interface for storage consumers. This makes storage easier to use, but it does not magically solve performance problems.

In order for data fabrics to do their thing, they must be supplied with adequate amounts of task-appropriate storage. If the fabric is supplying storage with unacceptably high latencies, you might consider adding NVMe solid-state drives to the mix. If capacity is a problem, but performance is fine, then perhaps a big box of magnetic hard drives is the ticket.

Measuring IOPS, latency and throughput are important for determining the characteristics of physical storage that's added to a modern data fabric. That said, extended proof-of-concept testing with real world workloads and copies of real data are equally important for teasing out the qualitative performance characteristics of a storage solution, especially when the data fabrics used start doing spooky things like tiering solid data to the public cloud, or dynamically changing data efficiency approaches based on load.

# Hanover Hospital

---

CASE STUDY



# Hanover Hospital

## A Storage Efficiency Case Study

Drastically Reducing Time Spent on Routine Storage Tasks and Reducing Storage Costs.

Hanover Hospital is an independent, not-for-profit community hospital and part of Hanover HealthCare PLUS network of services. Hanover hospital is also a member of the UPMC Pinnacle. The hospital is located in Hanover, Pennsylvania and manages 6,000 patient visits, 190,000 outpatient visits, with 600,000 lab tests and 90,000 imaging scans.

Within healthcare, IT is under enormous pressure to increase storage capacity, improve resiliency and accelerate performance – all while managing costs. Hanover Hospital is one of more than 1,000 healthcare customers that have trusted DataCore Software to virtualize its storage infrastructure.

Supervisor of Technical Support Douglas Null is charged with providing technical leadership for the IT infrastructure, including servers, storage, networking, security, as well as technology and application conversion.

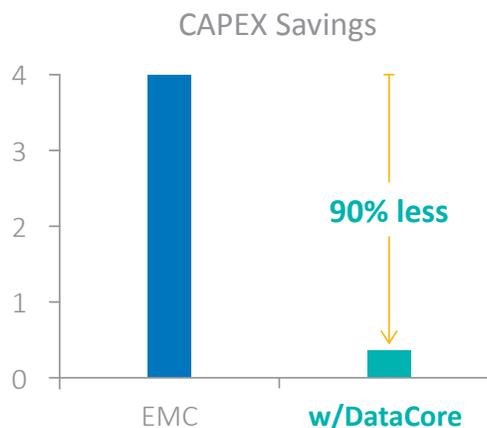
*“If we had gone down the existing path we were on with our traditional SAN, it would have been a very lengthy process to acquire the storage and to deploy the storage to ESX servers in the environment,” stated Null. “We wanted something we could quickly scale, and is simple to manage – all while not impacting performance. Additionally, we did not necessarily want flash storage if we did not need it.”*

### ABOUT THE CUSTOMER

Hanover Hospital is a not-for-profit community hospital which is part of a larger system called Hanover HealthCare PLUS Network, dedicated to the promotion of wellness, preservation of health, and the provision of diagnostic and therapeutic services to the people of the Greater Hanover, PA area. Hanover’s mission is to be the provider of choice for hospital-supported acute care, diagnostic, therapeutic, rehabilitative and wellness services required to support its network of care.

[www.hanoverhospital.org](http://www.hanoverhospital.org)

### COST REDUCTIONS WITH DATACORE



Null and the rest of the IT team at Hanover faced numerous challenges with the incumbent storage vendor. For one, new departmental applications needed to be deployed quickly. In addition, these applications required a significant amount of storage capacity, which was beyond what the current storage array could support. A storage upgrade using the incumbent vendor was considered. However, the high cost for-performance consideration quickly made this an unviable option. The performance concern was compounded by the fact that any new storage needed to have the ability to support even more applications in the future.

The IT team at Hanover needed a solution that helped them meet their application requirements. They needed the ability to quickly scale, keeping the applications available far more efficiently than they had been able to do so in the past. Hanover Hospital did not want to be tactical and go down the same old path of deploying new applications on a replicated, traditional two-array SAN configuration. Instead, the IT team began looking for a smarter, more strategic option.

### BECOMING SOFTWARE-DEFINED WITH DATACORE: REDUCING TIME SPENT ON ROUTINE STORAGE TASKS AND REDUCING STORAGE COSTS

Hanover Hospital trusted DataCore to virtualize its storage infrastructure because the hospital did not have to upgrade its existing, underlying storage – thereby making a CapEX outlay much lower. Hanover started with a single DataCore installation and over a handful of years the hospital added synchronous mirroring that stretched storage availability between its two on-campus data centers.

Today, DataCore's adaptive, self-learning and healing technology is taking the pain out of manual processes and helps deliver on the promise of the software-defined data center through its hardware agnostic architecture. With DataCore, Hanover Hospital has drastically reduced the time spent on routine storage tasks and has reduced storage costs – all while increasing capacity utilization and the performance of its applications. What's more, high-availability storage at Hanover has significantly reduced the time it takes to provision storage and systems.

## TOP THREE REASONS FOR SELECTING DATACORE

- **DataCore was 90% less than the total price of EMC**
- **Quickly scale performance for VMs and applications**
- **Active / active, zero touch failover and failback**

## HANOVER'S JOURNEY WITH DATACORE:

Null reports that from a performance and applications perspective, he has seen a lot of improvement over the years as far as base features go with DataCore solutions. The results attained with DataCore are impressive, including:

### COSTS

- 90% less in CAPEX
- 80-90% less in OPEX

### OPERATIONAL SIMPLICITY

- Provisioning went from four hours to five minutes

### CAPACITY

- 60% overprovisioned

### PERFORMANCE

- 1K IOPS from SATA drives with DataCore

### AVAILABILITY

- No downtime for last four years despite storage failures and outages

## THE SOFTWARE-DEFINED APPROACH DELIVERS BETTER OPERATIONAL SIMPLICITY AND BETTER CAPACITY MANAGEMENT

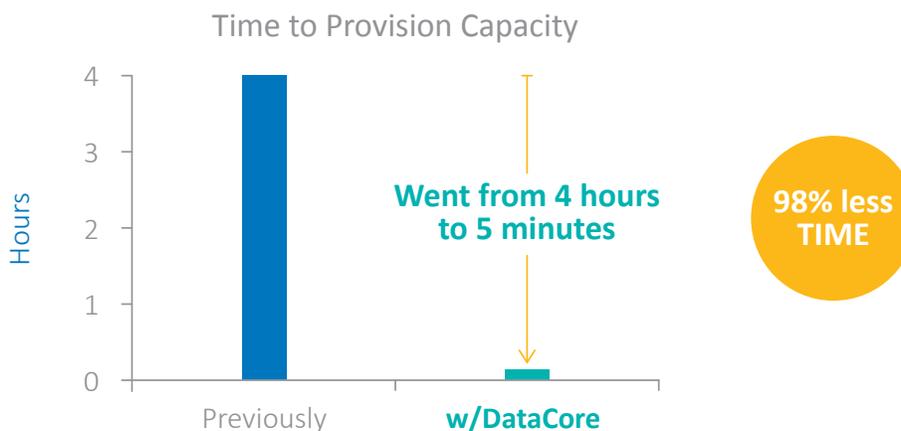
One of the other benefits of software-defined storage with DataCore that has significantly helped Hanover Hospital has been the overall operational simplicity of storage management. Provisioning a DataCore environment is truly straightforward and, according to **Null**, is “essentially a two-step process.” This software-defined approach took provisioning from four hours down to five minutes. This is because DataCore SANsymphony™ excels in automating storage-related tasks – whereas the steps needed to provision storage in alternate ways are often daunting. With a traditional SAN set in an active-active configuration, Null and others at Hanover would be tasked with configuring an equipment storage solution on one side, then configuring the equipment storage solution on the other side. Then, they would have to configure the replication between the two sites. They would also have

to configure journals and logs. Finally, they would need to configure Site Recovery Manager (SRM) from VMware that is used to failover the VMs from site to site.

“As you can see, we save a lot of time in even the most simple configuration by provisioning storage in both of our data centers with DataCore,” **Null** emphasized.

What’s more, Hanover now has the ability to better manage capacity through thin-provisioning. According to Null, the IT environment at Hanover is about 60% overprovisioned. This benefits the organization by allowing it to not have to add new storage when new systems come online. Existing storage can be allocated to new systems immediately and overall storage can be expanded down the road. When new systems are introduced, the IT team simply thin-provisions virtual volumes of storage and presents them to the ESX hosts as well as to the VMs. By doing so, Hanover saves both time and money.

### OPERATIONAL SIMPLICITY WITH DATACORE



“ [With DataCore] we get very impressive performance and bandwidth throughput for the amount of VM servers and applications we are hosting on our environment. Plus, we have improved storage utilization since we are able to over-provision storage by about 60%, meaning we are more efficient in our ability to meet the growth and cost demands for more capacity.

– Douglas Null, Supervisor of Technical Support.

”

## A VIRTUALIZED, AGILE INFRASTRUCTURE THAT SUPERCHARGES PERFORMANCE, SCALES TO MANAGE DATA GROWTH AND ELIMINATES DOWNTIME

Hanover has seen very good performance with DataCore SANsymphony. In fact in one instance where a RAID 5 virtual volume was mirrored between the two sites, the IT team was particularly impressed. This was when the team used a tool from Veeam to monitor the VM disk metrics of a VMware ESX 5 VM; they realized that the particular application was taking well over 1,000 IOPS/second on the DataCore powered storage, however, the latency on the data stored was “nil” and the reads were never even hitting the backend physical disk because of the DRAM caching that is part and parcel of the DataCore SANsymphony.

Significantly, when the IT team at Hanover breaks down the technology that is running on the EMC storage versus the technology running the VM’s that are powered by DataCore virtualized storage platform back-ended by HP MSA, of the 266 total VMs, 140 are now running on the DataCore software-defined storage platform.

The two data centers are connected by dark fiber. Roughly 90% of the infrastructure is on VMware vSphere. The majority of the “compute” is done on 34 HP DL380 servers. On these

servers, the hospital has approximately 266 virtual machines (VMs) in operation. Hanover is using EMC and HP MSA arrays as its current storage platforms, with the HP MSA’s used specifically to support the DataCore deployment.

DataCore SANsymphony now serves as a unified storage services platform across the entire multi-site infrastructure. In particular, it is relied upon extensively for various mission-critical, clinical and enterprise applications. Examples of these include the components of the hospital’s MEDITECH EMR, clinical middleware, medical dictation and transcription services, components of an Ambulatory HIS, a Medisolv BI reporting platform, Citrix XenApp, Infrastructure Management and Monitoring Applications, multiple Microsoft product offerings such as Microsoft SQL and MySQL, as well as other applications.

Hanover Hospital reports that with DataCore SANsymphony deployed in a synchronous mirror configuration, it has realized continuous uptime through high-availability storage and has significantly reduced the time it takes to provision storage and systems.

Today, Hanover Hospital is, according to **Null** – “A happy DataCore customer with well over half of our environment running on DataCore-powered storage.”

---

## ABOUT DATACORE

DataCore is a leader in software-defined storage. The company’s storage virtualization and Virtual SAN solutions empower organizations to seamlessly manage and scale their data storage architectures, delivering massive performance gains at a fraction of the cost of solutions offered by legacy storage hardware vendors. Backed by 10,000 customer sites around the world, DataCore’s adaptive and self-learning and healing technology takes the pain out of manual processes and helps deliver on the promise of the new software defined data center through its hardware agnostic architecture.

[www.datacore.com](http://www.datacore.com)

For additional information, please visit [datacore.com](http://datacore.com) or email [info@datacore.com](mailto:info@datacore.com)

© 2018 DataCore Software Corporation. All Rights Reserved. DataCore, the DataCore logo and SANsymphony are trademarks or registered trademarks of DataCore Software Corporation. All other products, services and company names mentioned herein may be trademarks of their respective owners.



# **Solving Healthcare IT Challenges with DataCore Software-Defined Storage**

---

SOLUTION BRIEF



## Solving Healthcare IT Challenges with DataCore Software-Defined Storage

### Empowering Real-Time, Always-On Data

Today's hospitals are benefitting from an explosion of information technologies that is ushering in a new era of healthcare. With these advanced technologies, significantly more data is being generated, from a much wider variety of sources, and at a more frequent pace. All of this data must be stored, shared, and protected.

With data the lifeblood of healthcare, IT departments are challenged to adopt new storage and management strategies to handle the deluge of data. The skyrocketing costs to achieve continuous data availability, cope with exponential data growth, and provide timely data access rank among the most pressing challenges facing healthcare IT organizations today.

That is why a growing number of healthcare institutions are deploying DataCore's software-defined storage platform. Only DataCore enables today's hospitals and health systems to address mission-critical healthcare IT challenges while maximizing the availability, performance, and utilization of IT resources – allowing them to enhance patient outcomes while keeping costs low.

## DataCore Software-Defined Storage: The Ideal Platform to Address Today's Healthcare IT Needs

DataCore software is the cornerstone of the next-generation, software-defined data center. A complete, centrally managed solution that costs a fraction of expensive hardware alternatives, DataCore software abstracts and automates the data services and management of underlying storage capacity to free customers from vendor lock-in and deliver unrivaled performance and data protection. It can be applied to a critical part of the IT environment initially, then scaled out to cover the entire storage infrastructure in affordable and manageable steps. The software is independent of the storage system and resolves the incompatibilities of different generations of storage equipment, diverse models, and multiple manufacturers.

“

*High availability was the first and foremost reason for going with DataCore – and for continuing with it. Now, everything that is mission-critical to the running of the hospital is supported by the DataCore software-defined storage platform. Users not only receive faster access to data, but they benefit from more server capacity as well.*

– Roger Fe de Leon, head of the storage group at Maimonides Medical Center

”

## On-Demand Scalability

DataCore's software-defined storage platform enables seamless, on-demand scale up of storage infrastructure to handle the growing volume of healthcare data due to applications such as EHR and PACS and data retention regulations (HIPAA).

## Simple Management

DataCore simplifies management of all storage resources, freeing IT to deliver valuable new services to address emerging healthcare trends more quickly and easily.

## Cost-Effective

DataCore lowers costs by virtualizing storage resources, which removes the constraints of such elements as data locations, data types, applications, or hardware types, and combines them all in one big pool of data.

## Always-On Reliability

DataCore ensures continuous availability of data by enabling immediate, automated recovery of critical healthcare data, with no risk of downtime.



CHALLENGE	HOW DATACORE SOLVES IT	KEY FEATURES
<p>Consolidate and manage healthcare IT data from disparate systems.</p>	<p>DataCore abstracts and automates the data services and management of underlying storage capacity.</p>	<p><b>Centralized Management</b> Single-pane management across storage devices from all vendors and models, enabling Infrastructure-wide control and monitoring of pooled storage resources from one console.</p> <p><b>Storage Pooling</b> Consolidates like or unlike disks – enabling cost-efficient tiering of resources by price/performance and capacity.</p> <p><b>Thin Provisioning</b> Increases storage utilization and enables cost savings by allocating just enough space just-in-time.</p> <p><b>Data Migration</b> Transparently moves and migrates data; allows non-stop business operations during data center relocations, VM moves, system refreshes.</p> <p><b>Deduplication/Compression</b> Reduces space needed to store multiple copies of the same data.</p>
<p>Safeguard healthcare data and applications from cyberattacks, system outages, data loss, natural disaster, and human error.</p>	<p>DataCore ensures your infrastructure is always available, no matter what the issue, by mirroring your data transparently, automatically, in real time.</p>	<p><b>Sync Mirroring</b> Delivers zero-downtime, zero-touch high availability by continuously mirroring active/inactive copies between physically separate locations accessible to local/metro clusters as shared disks.</p> <p><b>Asynchronous Remote Replication</b> Delivers fast, simple, remote-site disaster recovery and migration; keeps distant copies up-to-date without impacting local performance.</p>

## CHALLENGE

## HOW DATACORE SOLVES IT

## KEY FEATURES

Ensure ultra-fast application response times and real-time data availability for life-critical healthcare applications.

DataCore enables you to form a transparent virtualization layer across diverse storage systems to maximize the availability, performance, and efficiency of healthcare application workloads (EHR, PACS).

#### Continuous Data Protection and Recovery

Businesses can roll back to a previous point-in-time prior to a disaster, virus attack, or other disruptive event without doing an explicit backup.

#### Snapshots/Backups

Captures point-in-time images quickly across storage devices; Efficiently clone system images, speed up recovery, and enable off-line analysis.

#### Parallel I/O

Maximizes utilization of multi-core processors to dramatically cut latency and removes I/O as the bottleneck in application performance.

#### High Speed Caching

Accelerates disk I/O and business-critical application performance by empowering existing storage assets.

#### Auto-Tiering

Makes intelligent and automatic usage optimizations based on cost and performance across different types of storage.

#### Random Write Accelerator

Significantly increases performance for workloads characterized by many random writes, such as frequently updated databases, ERP & OLTP systems.

#### Quality-of-Service Controls

Ensures high-priority workloads meet SLAs with predictable I/O performance by regulating resources consumed by lower priority requests.

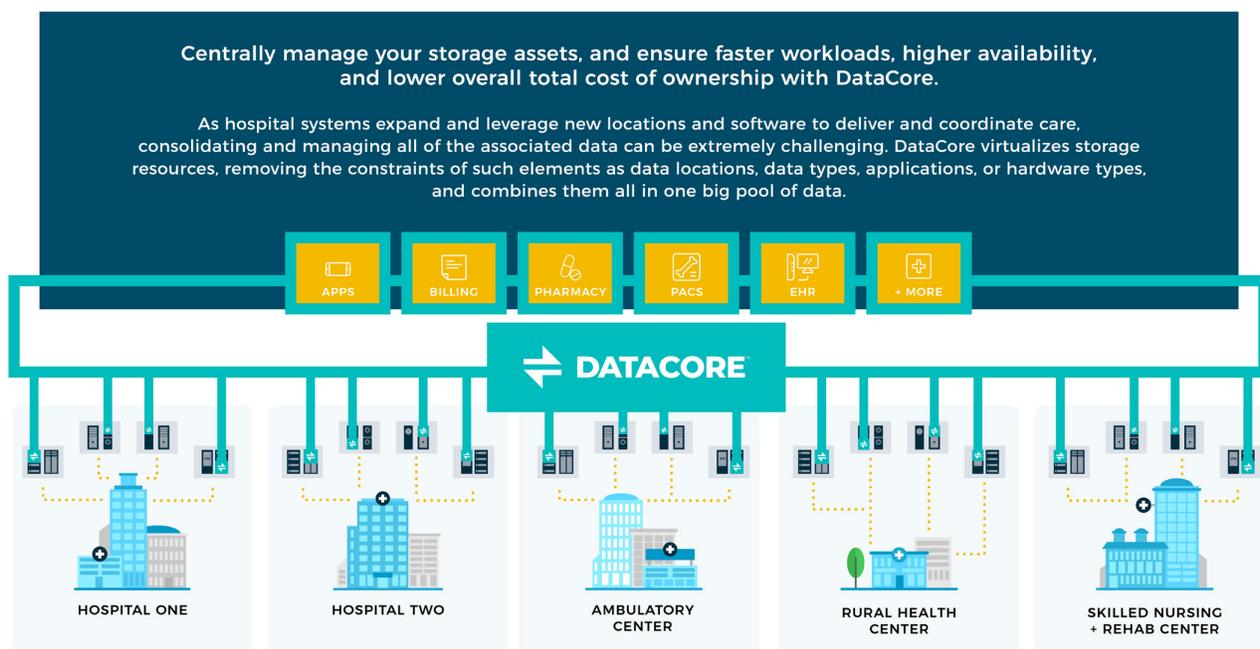
CHALLENGE	HOW DATACORE SOLVES IT	KEY FEATURES
Scale storage as needed – easily, instantaneously, inexpensively, and non-disruptively.	DataCore software enables you to improve performance without requiring you to rip and replace your system. By providing a virtual storage pool with existing and new storage systems that dynamically assigns workloads to optimal storage tiers based on performance requirements, DataCore lets you scale up or scale out your applications while efficiently using all your resources.	<p><b>Parallel I/O</b> Maximizes utilization of multi-core processors to dramatically cut latency and removes I/O as the bottleneck in application performance.</p> <p><b>High-Speed Caching</b> Accelerates disk I/O and business-critical application performance by empowering existing storage assets.</p> <p><b>Auto-Tiering</b> Makes intelligent and automatic usage optimizations based on cost and performance across different types of storage.</p> <p><b>Storage Pooling</b> Consolidates like or unlike disks – enabling cost efficient tiering of resources by price/performance and capacity.</p> <p><b>Thin Provisioning</b> Increases storage utilization and enables cost savings by allocating just enough space just-in-time.</p> <p><b>Data Migration</b> Transparently moves and migrates data; allows non-stop business operations during system refreshes.</p>

“

*The fact is that we are not tied to a particular vendor. We can always just get the best storage for what we are trying to accomplish without having to learn a new management system to for every model.*

– Dustin Fennel, chief information officer at Mission Community Hospital

”



## About DataCore Software

DataCore is the authority on real-time data. The company pioneered [software-defined storage](#) and has now expanded its technology leadership to [hyperconverged](#) infrastructures. DataCore empowers IT organizations to achieve always-available, high-performance and highly efficient data. Its patented technology eliminates storage bottlenecks with adaptive parallel I/O optimization, enables zero-downtime synchronous mirroring, and provides a true hardware-agnostic architecture – resulting in flexibility, resource efficiencies, and cost savings.

DataCore Software is the cornerstone of the next-generation, software-defined data center. DataCore's value has been proven in more than 10,000 customer deployments across traditional, hyper-converged, cloud, and hybrid environments. Visit [www.datacore.com](http://www.datacore.com) or call (877) 780-5111 for more information.

For additional information, please visit [datacore.com](http://datacore.com) or email [info@datacore.com](mailto:info@datacore.com)

© 2018 DataCore Software Corporation. All Rights Reserved. DataCore, the DataCore logo and SANsymphony are trademarks or registered trademarks of DataCore Software Corporation. All other products, services and company names mentioned herein may be trademarks of their respective owners.



# Resources

---



## WEBINAR



### **Maintaining Mission Critical Levels of Availability in Healthcare**

Bob Strachan & Giovanni Greci, Solutions Architects - DataCore Software; Gabriel Sandu, CTO - Maimonides Medical Center

Running mission-critical applications and providing continuous access to data are the lifeblood of any healthcare organization. That's why a business continuity plan is uniquely important to the healthcare sector.

Hear directly from Maimonides Medical Center's CTO Gabriel Sandu as he shares how their organization uses software-defined storage as a strategic solution to achieve their business continuity plans.

DataCore has been assisting many hospitals and healthcare organizations throughout the world in building upon their continuity planning efforts.

During this webinar, our technical business continuity experts will walk you through real-life healthcare disaster incidents. They will share how major hospitals have survived hurricanes, floods, power outages and building failures by using software-defined storage to provide uninterrupted operations to applications and critical medical data systems during crisis events.

Join this session and learn how DataCore's software empowers business continuity for healthcare organizations such as New York-Presbyterian, Mount Sinai, Maimonides Medical Center, and Englewood Hospital and Medical Center.



## **RESOURCES**

**Click Here to  
REQUEST A LIVE DEMO**

**Click Here to Learn More About  
HEALTHCARE DATA STORAGE SOLUTIONS**

# Appendix

---



# Software-defined Storage Platform

The availability, performance, agility and TCO needed to meet the demands of next generation storage

When it comes to IT infrastructure, storage has fallen off the pace. Server performance has dramatically increased, but storage just hasn't kept up, causing applications to slow down. Poor storage and I/O performance are being disguised by techniques that end up spreading workloads unnecessarily across many machines under the excuse of scaling out. In addition, clustered, highly available servers have a Recovery Time Objective / Recovery Point Objective (RTO/RPO) measured in seconds, whereas recovery from storage failures may take several minutes to days of manual processes to restore data from a backup copy or, worse, a Disaster Recovery (DR) site. Lastly, a diverse mix of standardized servers can be managed from a single console, while different storage products and brands each have their own management and can't work together, even if they are from the same vendor.

The DataCore™ SANSymphony™ enterprise-class Software-defined Storage (SDS) platform provides a high-performance, highly available and agile storage infrastructure with the lowest Total Cost of Ownership (TCO).

Compared to enterprise storage arrays and other SDS products, DataCore offers the following advantages:

- **Go faster:** Faster I/O for databases, email, VDI and line of business applications means more transactions processed, more data analyzed quicker and increased customer satisfaction, leading to more revenue.
- **Always on:** Highly-available infrastructure means applications are always up, reducing business disruptions and lost sales.

The net result is faster performance, higher availability and greater cost savings across your company through consolidation.



- **Save money:** Run more workloads with better performance and availability by pooling diverse storage assets, deferring hardware refreshes and simplifying administration. This results in remarkable cost savings in both acquisition (CAPEX) and ongoing operations (OPEX, power, cooling and space).
- **Less risk:** High-performance, highly available, agile infrastructure means efficient operations, reduction in disruptions to operations and productive users (employees, customers and partners).

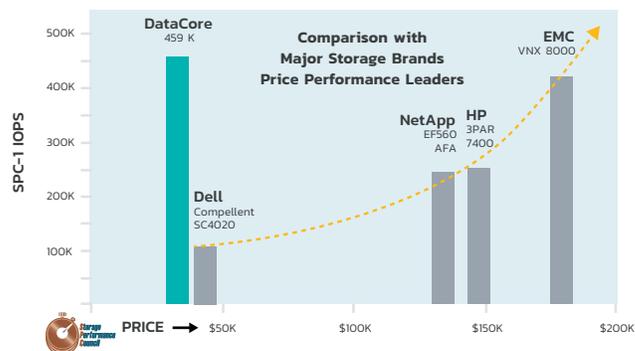
## DATACORE™ SANSYMPHONY™ SOFTWARE-DEFINED STORAGE PLATFORM

DataCore SANSymphony infrastructure software takes isolated storage devices, sometimes spread between different locations, and places them under one common set of enterprise-wide services. It pools their collective resources, managing them centrally and uniformly despite the differences and incompatibilities among manufacturers, models and generations of equipment in use. Key features include:

### Highest Data Performance

- **Fastest response time in the industry:** Utilizing DataCore™ Parallel I/O Technology, proven by Storage Performance Council SPC-1 benchmarks<sup>1</sup>, SANSymphony is 3-10x faster than any other storage infrastructure product. As a result, applications' data can be accessed, updated and stored faster, for more timely decisions and actions.
- **World leader in price performance:** At least 66% better than the next closest competitor, as shown by the SPC-1 benchmark.

### DataCore Price Performance is Off the Curve



<sup>1</sup> [http://www.storageperformance.org/results/benchmark\\_results\\_spc1\\_top-ten](http://www.storageperformance.org/results/benchmark_results_spc1_top-ten)

## Highest availability with multi-layered protection

- Zero Touch, Zero Downtime means that data is always available, including across a stretch cluster. Storage failures are handled automatically, without any downtime, as is failback.
- For DR purposes, an efficient asynchronous replication scheme assures another copy of the data is available hundreds of miles away for recovery from large-scale regional disasters.
- To protect from ransomware, virus outbreaks and malware, Continuous Data Protection (CDP) keeps all changes up to the previous 2 weeks so applications can be rolled back before the infection occurred.

## Lowest TCO

- Single pane management provides a common provisioning and monitoring framework across all storage devices, decreasing time spent on storage administration.
- Hardware-independent storage services ensure heterogeneous storage devices benefit from a consistent set of advanced services, including support for OpenStack and VVols.
- Freedom to substitute different hardware from competing suppliers when appropriate results in a cost-effective solution, unlike competitors that restrict choice.

VIRTUALIZED HOSTS		PHYSICAL SERVERS	
<b>DATAcore SOFTWARE-DEFINED STORAGE PLATFORM</b>			
AVAILABILITY	PERFORMANCE		EFFICIENCY
<ul style="list-style-type: none"> <li>Synchronous Mirroring</li> <li>Asynchronous Replication</li> <li>CDP</li> <li>Snapshots / Backups</li> </ul>	<ul style="list-style-type: none"> <li>Caching</li> <li>Auto-tiering</li> <li>Random Write Accelerator</li> <li>Quality of Service (QoS)</li> </ul>	<ul style="list-style-type: none"> <li>Storage Pooling</li> <li>Thin Provisioning</li> <li>Data Migration</li> <li>Deduplication/Compression</li> </ul>	
MANAGEMENT			
Centralized Management	Analysis & Reporting	Wvols	NAS/SAN (Unified Storage) Cloud Integration

## Risk Mitigation

- Data can be migrated to different storage devices without disrupting applications' access to data, minimizing risk.
- New technology can be integrated seamlessly, without requiring downtime. This minimizes the uncertainty associated with deploying new technologies and makes storage devices completely interchangeable.

## DATAcore CUSTOMERS REPORT UP TO



75% reduction in costs  
Lower Cost of Ownership



10x performance increase  
Faster Applications



100% reduction in storage-related downtime  
Higher Availability



90% decrease in time spent on routine storage tasks  
Greater Productivity

Source: **TechValidate**  
www.techvalidate.com



### CASE STUDY: REDUCED DOWNTIME AND PERFORMANCE GAINS IN HEALTHCARE

Englewood Hospital and Medical Center is focused on delivering safe, high-quality and efficient care to their patients. There is zero tolerance for error, from the operating room to the server room, with greater expectations for IT. As Englewood virtualized their servers, they were seeing performance bottlenecks in their applications. In addition, since they ran 24x7, storage downtime, both planned and unplanned, was disruptive to their operations. Additionally, they were expecting a 10x growth in data over 4 years.

To solve their problems, Englewood deployed DataCore SANSymphony SDS platform. The most immediate benefit was the increased performance of their applications. Their Exchange environment responded significantly faster, while their Payroll system ran in half the time and their backups went from 8 hours down to 3 hours. They've also seen the benefit of running in an active-active configuration, as their storage downtime has been eliminated, allowing them to deliver patient care seamlessly and without interruption. Lastly, Englewood has been able to better utilize their existing capacity, extending the life of their current storage assets and driving down the costs of storage.

For additional information, please visit [datacore.com](http://datacore.com) or email [info@datacore.com](mailto:info@datacore.com)



© 2018 DataCore Software Corporation. All Rights Reserved. DataCore, the DataCore logo and SANSymphony are trademarks or registered trademarks of DataCore Software Corporation. All other products, services and company names mentioned herein may be trademarks of their respective owners.