



# Reducing Data Center Infrastructure Costs with Software-Defined Storage

WHITE PAPER




## TABLE OF CONTENTS

1. Introduction: Today's Storage Cost Crisis.....	3
2. The Rise of Software-Defined Storage.....	3
3. Determining the Cost of Today's Complex Infrastructure.....	4
4. DataCore SDS – Delivering the Industry's Best overall TCO (and fueling new investments)...	5
5. DataCore SDS TCO Summary.....	5
6. Assessing the TCO and ROI benefits of DataCore SDS in your environment.....	6
7. Conclusions.....	6

### 1. INTRODUCTION: TODAY'S STORAGE COST CRISIS

The pace of change and innovation in enterprise computing during the last decade has created enormous pressure on the underlying data storage infrastructure, which was designed for an era that predated the recent explosion of innovation. To attempt to keep up with the pace of change, IT teams have rapidly expanded storage capacity, added expensive new storage arrays to their environment, and deployed a range of disparate point solutions.



Yet, despite pouring significant investments into the infrastructure, the storage layer has remained particularly problematic and, despite representing a significant percentage of IT budgets, it continues to be at the root of many of the top challenges in IT: inability to keep up with rapid data growth rates, vendor lock-in, lack of interoperability, and most significantly, increasing hardware costs.

Most storage hardware vendors suggest that all hardware should be replaced every three years with new arrays that promise better performance and

more capacity, despite the cost and the enormous burden of replacing hardware and performing complex data migrations. Moreover, hardware-vendor pricing models aim at locking decisions for six or more years. Given that IT teams cannot continue to simply outspend the problem, it has become clear that a more fundamental solution is required to address the cost and complexity of the storage infrastructure.

## 2. THE RISE OF SOFTWARE-DEFINED STORAGE

As IT architects and decision-makers look for long term solutions to this challenge, and prepare their organizations for a future of unpredictable change, software-defined storage (SDS) is increasingly being recognized as a viable solution for the short and long term. SDS is software built to run on commodity server hardware, aggregate physical storage capacity from many separate systems into a single shared storage pool, and deliver a broad set of storage services similar to the functionality provided by traditional SAN storage arrays.

SDS stands in contrast to the traditional method of enterprise storage management, which has been defined by monolithic expensive storage systems from large vendors that offer limited interoperability with other vendors, and enable no sharing of resources amongst systems. By providing a virtualization layer that abstracts away the underlying physical storage, the promise of SDS is to deliver a new level of flexibility to IT, increasing storage efficiency, enabling unlimited scale of data and operations, clearing away vendor lock-in, solving interoperability issues, and ultimately delivering dramatic cost savings.

Moreover, by enabling a reduction in capital spending and staff time dedicated to infrastructure management, SDS frees up resources for higher order activities. Indeed, the capital and human expenses that are liberated from maintaining the status quo can now be redeployed to the important projects that are intended to drive growth and innovation.

The IT industry at large has already adopted server virtualization and enjoyed its benefits. It is becoming obvious that storage virtualization can bring similar—if not more significant benefits. Indeed, industry experts clearly see that software-controlled storage hardware is now a standard part of the overall IT architecture.

“...by 2024, 50% of the global storage capacity will be deployed as SDS on-premises or on the public cloud (up from less than 15% today).<sup>1</sup>

Gartner

”

The recognition of the promise of SDS is driving a surge in adoption. The market for SDS reached almost \$10 billion in worldwide sales in 2017, and is expected to reach \$16.2 billion by 2021, according to IDC . This rapid growth has drawn in a slew of new offerings in recent years, from both start-ups and traditional vendors.<sup>2</sup>

Within this growing market segment, DataCore™ SDS has emerged as the industry's leading software-defined storage platform. DataCore SDS is the industry's most robust and versatile SDS platform, and brings new levels of performance, availability and agility to the enterprise data center infrastructure.

On the strength of its robust architecture, DataCore™ SDS delivers dramatic cost savings when compared to traditional approaches and various new alternatives. This brief paper demonstrates the many ways in which DataCore SDS delivers significant TCO improvements, in the form of both CAPEX and OPEX savings. It is intended to help infrastructure teams assess the potential economic impact of deploying DataCore SDS within their environment.

## 3. DETERMINING THE COST OF TODAY'S COMPLEX INFRASTRUCTURE

The potential economic impact of software defined storage is best understood in the context of the complexity and cost crisis that characterizes most enterprise IT environments today.

**Hardware and Software Costs:** A peek into the enterprise storage environment today reveals many specialized products, built with proprietary technology, and often using expensive componentry—such as Flash storage, NVMe, or high-end quad-core processors—in order to deliver the reliability and high-performance that enterprise workloads necessitate. To meet all of the enterprise requirements, while accounting for both capacity growth of existing workloads and the addition of new workloads, IT teams have had to devote significant portions of their annual budget to these capital expenditures (CAPEX). Year over year growth in data, applications supported, number of users, and number of sites all drive further CAPEX spending just to maintain the status quo.


**Operating Expenses (OPEX) and the Inability to Innovate:** Infrastructure complexity also consumes significant manpower, as IT staffers must continually monitor each component in the data center, conduct

routine maintenance, manage upgrades, manage patches, etc. The complexity increases with the increase in key variables, including: storage arrays, vendors, locations, applications, and operating systems.

The volume of activities required to keep the existing infrastructure available and working as expected means that the majority of staff time goes to maintaining the status quo infrastructure, leaving a much smaller percentage of IT staff cycles available to dedicate toward innovation or new programs that can enable growth and differentiation for the business.

#### 4. **DATACORE SDS – DELIVERING THE INDUSTRY’S BEST OVERALL TCO (AND FUELING NEW INVESTMENTS)**

DataCore SDS delivers all of the promised value of SDS, while meeting the performance, availability and functional needs of demanding enterprise IT environments. Importantly, DataCore SDS overcomes the limitations found in other SDS products, giving IT teams the ability to address a wide range of use cases and application requirements—including Tier 1 enterprise applications and high-performance databases—with a single, centrally managed platform for both primary and secondary storage. These many strengths all converge to drive dramatic cost savings compared to alternative solutions, giving DataCore SDS the best overall TCO.



Three primary factors allow DataCore SDS to address the full spectrum of enterprise use cases and environments, while reducing spending compared to alternative solutions:

1. Its mature and robust software architecture, built to leverage 3rd party hardware with maximum efficiency
2. A set of software-based patented performance breakthroughs that are incorporated into the platform
3. Several critical storage efficiency features that further reduce the amount of storage hardware required for any workload environment

#### 1. **Enterprise class data architecture:**

DataCore SDS is built upon the robust architecture of the market's first true software-defined storage solution—DataCore SANsymphony™. Originally delivered to the market almost 20 years ago, the software platform is currently in its 10th major release, and is deployed in more than 10,000 customers globally.

This robust software architecture allows DataCore SDS to drive down OPEX costs by reliably virtualizing and pooling the existing infrastructure resources in the data center. In terms of infrastructure cost, pooling resources pays immediate dividends by freeing up islands of underutilized capacity, thereby allowing application workloads to leverage the entire environment rather than silos of dedicated resources. The impact is a major improvement in resource utilization, which directly reduces the amount of required hardware for the environment. Beyond this, DataCore's core SDS functionality reduces the complexity of heterogeneous storage management, reduces the risk of outages, and improves the overall responsiveness and uptime of applications—all of which feed into direct and indirect cost savings.

The key attributes of DataCore SDS that make it the market leader derive directly from this long, successful run in the market and produce tremendous TCO advantages compared to the alternatives:

- **Widest functional range of data services**, based upon years of responding to the market, and customer requests for new functionality.

By incorporating all required data services into its software, DataCore eliminates the need for a range of expensive dedicated 3rd party products, yielding remarkable cost savings.

- **Highest level of interoperability**, based on close to 20 years of delivering storage software that is built to easily integrate into the existing data center environment, and supports the widest range of 3rd party offerings for each major component of the IT ecosystem. This has the two-pronged benefit of supporting (and extending the life of) the infrastructure already in place, and “future-proofing.” The IT team can easily add any 3rd party hardware offering into the environment in the future, selecting the best price/performance option, and/or allowing them to drive better purchasing terms through the leverage of hardware choice.
- **Most robust**, in terms of Reliability, Availability, Serviceability (RAS), based on billions of system hours running in customer environments. This has a direct impact on TCO through avoidance of application downtime, and elimination of wasted IT labor spent on resolving outages.

## 2. Breakthroughs in Performance Efficiency Drive Huge Infrastructure Cost Savings:

Historically, achieving improved performance (typically measured in I/O throughput, or application response time) required the addition of more expensive hardware to the environment. Indeed, this concept has been the driving force for the widespread adoption of all-flash storage arrays in recent years, and is the motivation behind the recent surge in NVMe popularity which places expensive memory cards inside the application servers in order to boost I/O performance.

In contrast, DataCore SDS incorporates several innovations at the software level which deliver significant performance gains efficiently, without reliance on massive amounts of dedicated expensive hardware:

Parallel I/O:

DataCore’s patented Parallel I/O technology is a standard, “always-on” feature of DataCore SDS. It elegantly eliminates the server I/O bottleneck by processing I/Os in parallel leveraging multi-core processor systems. On the strength of this new technology, DataCore set records for industry benchmarks, with measurements of 459,000 IOPS,

price performance of \$0.10 per IOPS, and application latency of just 0.22 ms—all with a low cost, 2 node HCI configuration using standard Lenovo servers. These measurements were 2-4X better than those of competing solutions, including expensive all flash arrays.<sup>3</sup>

- **High-Speed Caching** is a proprietary caching algorithm that accelerates I/O by leveraging RAM as a read and write cache. DataCore supports up to 8 TBs of high-speed cache per node, creating a true “mega-cache” to turbocharge application performance. Given that RAM is the fastest storage component in the architecture, a RAM-based cache can deliver a 3-5x performance boost to applications, while simultaneously freeing up application servers to perform other tasks. Using RAM-based cache also extends the life of traditional storage components by minimizing the stress experienced from disk thrashing.
- **Quality of Service Controls:** QoS is an optional feature that allows the IT team to ensure that high-priority workloads meet SLAs with predictable I/O performance. This is particularly useful in environments with many applications and workloads of differing priority, all sharing the same storage pool. QoS allows IT to set thresholds and limits on lower priority workloads ensuring that the top tier applications do not suffer in performance based on resource contention with lower priority applications. This feature also helps to reduce or eliminate the need to dedicate expensive silos of resources for high priority applications, contributing to the overall infrastructure cost reduction enabled by DataCore SDS.

Collectively, these capabilities enable SDS to deliver on a given workload’s performance requirements with a fraction of the hardware costs required by competing solutions.

## 3. Storage efficiency drives cost savings even further

In addition to the efficient performance features listed above, DataCore SDS has a number of other features aimed at efficient use of storage resources which drive down the TCO even further. Collectively, all of these features give DataCore SDS a “light footprint” compared to all other alternatives.

Features that allow DataCore SDS to drive up resource utilization, and significantly drop overall TCO include:

- **Auto-Tiering and Load Balancing**

DataCore SDS includes dynamic block-level auto-tiering. This function moves data at a granular level to the storage system that delivers the ideal performance, based on the observed performance profile for each application. The software uses machine learning to assess likely storage bottlenecks, and then automatically moves the “hot” blocks to the fastest media resource, eliminating the bottleneck and bolstering overall application performance, at the same time, migrating warm or colder data to lower cost storage. With this capability, a relatively small amount of flash storage in the DataCore SDS storage pool, can yield superior performance to one consisting of all-flash arrays.

- **Thin Provisioning**

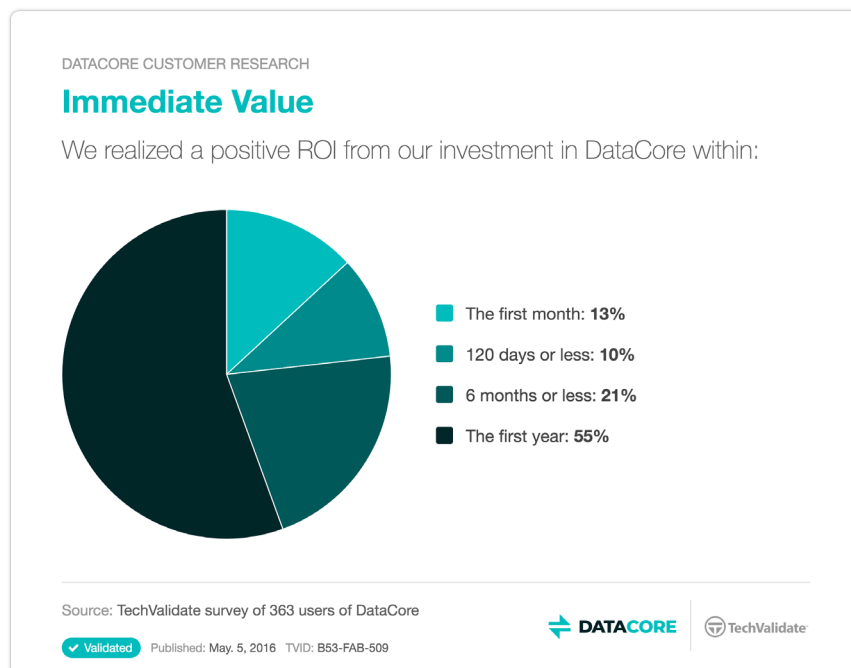
This feature allows IT to define large virtual volumes and assign them to application workloads, without the need to tie-up physical space on disk until it is needed.

- **Deduplication and Compression**

Deduplication and compression are storage efficiency technologies that further increase utilization and reduce the capacity required for a given workload. With DataCore SDS, they are selectable features that IT can choose to utilize for different workloads, and they are implemented as post-process activities, ensuring that the processing required to perform these tasks does not interfere with production application activity.

**Bringing it Together to Demonstrate the Economic Impact of DataCore SDS**

Each of the strengths described above lead directly to remarkable cost savings across a range of capital and operating costs. Collectively, these savings add up quickly, enabling DataCore SDS to achieve a positive return on investment (ROI) far faster than alternative solutions.



Indeed, a survey of 363 DataCore customers found that over half of them (55%) achieved positive ROI within the first year of deployment, and 21% were able to reach positive ROI in less than 6 months, see chart below.

**5. DATACORE SDS TCO SUMMARY**

	<b>DATACORE SDS TCO IMPACT</b>
<b>COST OF ACQUISITION - CAPEX</b>	Requires fewer and less costly resources compared to alternatives while addressing the same use cases
	Defers new storage purchases by unlocking stranded capacity from existing resource silos while extending their useful life
	Reduces dependency on premium priced flash through in-memory caching and auto-tiering
	Enhances negotiating positioning when shopping for new technology since storage hardware and suppliers become largely interchangeable
<b>OPEX- DIRECT SPENDING</b>	
HARDWARE AND SOFTWARE MAINTENANCE	Lowers maintenance expenses through use of fewer and less complex hardware components further
DATA CENTER COSTS	Higher resource efficiency reduces power, cooling and space consumption
LABOR	Less time spent managing and troubleshooting storage-related problems
<b>OPEX- INDIRECT SPENDING</b>	
APPLICATION PERFORMANCE IMPROVEMENTS	Faster app performance improves productivity, customer satisfaction and drives down labor costs
DOWNTIME COSTS	Uninterrupted data access can substantially avoid lost revenue, prevent data loss and mitigate risks

The potential economic impact of software defined storage is best understood in the context of the complexity and cost crisis that characterizes most enterprise IT environments today.

The cost savings for each line item in the table above can be quantified, allowing for an overall quantitative analysis of the TCO benefit of DataCore SDS in a given data center scenario. The next section provides guidance for assessing the TCO impact of DataCore SDS in your environment.



**6. ASSESSING THE TCO AND ROI BENEFITS OF DATACORE SDS IN YOUR ENVIRONMENT**

DataCore SDS customers typically experience significant OPEX and CAPEX savings, yielding a superior TCO:

- Initial acquisition costs will be much less than alternatives, driven by the customer’s ability to use lower cost storage at a much lower cost per TB. In addition, less overall capacity is needed, as DataCore SDS increases the overall utilization of resources. Additionally, far less expensive Flash storage will be required as the customer can take advantage of various performance and efficiency capabilities with far less expensive hardware.
- 3-Year Operating costs typically yield even more savings than capital costs, as a result of both significant labor cost reductions for normal operations, and the elimination of downtime that carries a cost of lost productivity for both IT and business resources.

Given many characteristics and variables that are unique to each IT environment, a TCO analysis must be customized, and we encourage you to conduct your own analysis using data about your environment. The table below is offered as a way to approach your analysis. It includes the primary cost drivers that should be included in the calculation of the total cost of ownership over a 3-year horizon. Because the TCO of a given solution is a relative metric that is best viewed in contrast to alternatives, the table is structured to allow for a comparison between DataCore SDS and one or more alternatives. Typically, customers compare the TCO of DataCore SDS to their status quo environment and/or other new solutions that they are considering.

The table below exemplifies a simple way to structure the TCO comparison, by looking at the key cost drivers and comparing the alternatives along each dimension.

	3 Year Cost			
	Legacy Approach	Alternative 2	Alternative 3	DataCore SDS
<b>CAPEX</b>				
Acquisition Cost				
Avg. Annual Capital Expansion cost (yrs 2 and 3)				
<b>OPEX- Direct Costs</b>				
HW and SW Maintenance				
Data Center costs				
Labor				
Total Annual Direct Opex				
<b>OPEX - Indirect Costs</b>				
Application Performance Lags / Throttles				
Application Downtime Costs -- Loss of productivity				
Total Annual Indirect Opex				
<b>3-Year TCO</b>				
Total 3 Year CAPEX				
Total 3 Year OPEX- Direct Costs				
Total 3 Year OPEX- Indirect Costs				
3 Year Total TCO				

Once the total 3-year cost of ownership is calculated for each alternative, the cost differences amongst the alternatives can be determined. The savings enabled by DataCore compared to a given alternative can be viewed as the return on the DataCore investment (ROI). Thus, the formula to arrive at the ROI percentage is:

$$\text{DataCore ROI} = (\text{Total Cost Savings Enabled}) / (\text{Acquisition Cost})$$

## 7. CONCLUSIONS

Software-defined storage has experienced a surge in adoption in recent years, driven by the quest of IT teams to fundamentally address the problem of the cost and complexity of their data storage infrastructure. As this paper has highlighted, DataCore SDS offers the industry's best storage TCO on the strength of its architecture which includes a number of innovations that directly reduce the cost of a robust storage environment. With DataCore SDS, customers can avoid or reduce new storage purchases by extending the life of existing storage, purchase less extensive storage without sacrificing performance or functionality, and can make more efficient use of existing capacity. Moreover, DataCore SDS drives significant operational savings based on reduced complexity, improved uptime, lower maintenance cost, lower data center costs, and other indirect factors.

Conducting a TCO evaluation of solutions under consideration is an important step in making any IT purchasing decision. This is particularly true for data storage, one of the most critical pieces of the overall data center infrastructure. This paper has been written to help provide a high level understanding of the potential for dramatic cost savings enabled by DataCore SDS. Readers are encouraged to conduct their own analysis, using inputs from their environment. DataCore professionals are available to assist as needed.

Please contact us at [info@datacore.com](mailto:info@datacore.com) to request assistance in conducting your own **TCO/ROI analysis**.

<sup>1</sup> Gartner, The Future of Software-Defined Storage in Data Center, Edge and Hybrid Cloud, 3 May 2019

<sup>2</sup> IDC, Worldwide Software-Defined Storage Forecast, 2017–2021: SDS Market Growth Significantly Outpaces Enterprise Storage Growth, Led by HCI, September 2017

<sup>3</sup> SPC benchmark 1™ full disclosure report, DataCore Software Corporation, DataCore™ SANsymphony™ 10.0, (Dual node, high availability, hyper-converged), storage performance council, June 2016.

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



© 2019 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.

0519