




# Software-Defined Storage For The Enterprise

A strategic guide to the top four approaches to SDS



# Contents

<b>Introduction</b> .....	<b>3</b>
Who can benefit from this guide? .....	3
<b>Where are we now?</b> .....	<b>4</b>
The Data Center Ecosystem .....	4
Before Software-Defined Storage .....	4
AI, IoT, and COVID at The Edge .....	4
The Rise of Open Source .....	5
Challenges & Opportunities .....	5
<b>Where are we going?</b> .....	<b>6</b>
The Data Center Objective .....	6
<b>How do we get there?</b> .....	<b>7</b>
Software-Defined Storage 101 .....	7
<b>Four approaches</b> .....	<b>8</b>
Approach 1: A proprietary, integrated hardware and software solution .....	8
Approach 2: A DIY solution .....	10
Approach 3: A supported solution .....	13
Approach 4: SoftIron's "task specific" solution .....	15
<b>The choice is yours</b> .....	<b>19</b>
<b>Take the next step</b> .....	<b>19</b>



## Introduction

# Who can benefit from this guide?

Today's technology landscape is more complex than ever, particularly when it comes to data storage solutions. The creation, management and value extraction of data underpins virtually every single organization in existence, whether it's immediately obvious or not. A robust, reliable and agile data storage strategy is an essential component of competitive advantage for any organization looking to achieve long term success.

One such strategy is Software-Defined Storage (SDS) which, according to Gartner, is rapidly on the rise as the preferred approach for reasons which we'll cover shortly.

***"By 2024, 50% of the global storage capacity will be deployed as software-defined storage (SDS) on-premises or on the public cloud, up from less than 15% in 2020."***<sup>1</sup>

This guide is designed to help business leaders navigate the world of SDS and better understand some of the core approaches. We'll share the pros and cons of each, and provide a framework for assessing vendor solutions so you can make the best decision for your requirements and available resources.



## Where are we now?

# The Data Center Ecosystem

## Before Software-Defined Storage

In tech, we often talk of technology lifecycles, but sometimes that cycle looks more like a swinging pendulum. IDC reports that **80% of organizations are now undertaking some level of cloud data repatriation activities** <sup>2</sup>, so it's clear that many businesses are deciding it's in their best interests to shift some, if not all, of their data from public cloud providers to local infrastructure environments. This is increasing demand for hybrid-cloud infrastructure solutions and, according to Gartner, **"by 2024, 40% of I&O leaders will implement at least one hybrid cloud storage architecture, up from 10% in 2020."** <sup>3</sup>

## AI, IoT, and COVID at The Edge

Exponential growth in AI and IoT - coupled with the impacts of COVID - has meant staggering changes to traditional data center environments, pushing more and more of it to the edge. "Edge computing" is best defined as a set of characteristics and constraints such as bandwidth, human skill sets, security, environment, space, and power, rather than the traditional concept of a 'single location'.

*"The global Edge Computing market is projected to grow from USD 3.6 Billion in 2020 to USD 15.7 Billion by 2025, at a Compound Annual Growth Rate (CAGR) of 34.1% during the forecast period."* <sup>4</sup>

## The Rise of Open Source


Underpinning all this is the continued, unfaltering rise of enterprise-class open source<sup>4</sup>. This has opened up the market to an array of disruptive innovators and is challenging many of the legacy, proprietary vendor solutions as they try to compete.

## Challenges & Opportunities

Whilst AI, hybrid-cloud, open-source and edge computing solutions are certainly empowering organizations, the influx of new technology players has also created an environment of complexity and confusion.

Collecting, processing and storing more of your data at the 'edge' will likely mean utilizing a data center that is geographically closer to your largest concentration of employees - who will typically live in cities or other densely populated areas. Space will be at a premium, and depending on local conditions, optimizing power efficiency and cooling may become a more complex operation.

Moving to a hybrid-cloud infrastructure also means choosing between a broader variety of vendors, bringing potential headaches with multiple contracts, hardware and/or software compatibility issues, rare skill sets to manage the systems, and spiraling costs.





Where are we going?

## The Data Center Objective

Of course, specific outcomes differ for every organization, but we can all agree on the basic foundations of performance, resiliency, and efficiency when it comes to operation optimization. We're all striving for progress and ROI in the face of relentless change, so where do you begin when it comes to assessing your options? There are infinite combinations of solutions available on the market to address specifics, but software-defined storage as an infrastructure approach is ideal for many reasons.

## How do we get there?

# Software-Defined Storage 101

Software-defined storage (SDS) is an architecture that separates storage software from its hardware.

By separating the functions that software performs from the hardware running it, an SDS approach allows for use-driven customization, scalability and optimization. What's more, independence from storage hardware creates the opportunity for an agile, adaptable data center. As business requirements change and technology evolves, you can modify your setup without expensive forklift upgrades. The ultimate benefit is the flexibility to adapt your capacity as required without needing to rely solely on expensive proprietary hardware.

However, while a well-designed SDS solution will provide you with the tools you need to hit your price, performance and security targets, the very nature of its adaptability can be intimidating. Building a bespoke SDS solution doesn't have to be expensive, but it can be challenging to assess all the options.

To make this process simpler, we've summarized some common approaches to implementing SDS. There is no single correct way to design your software-defined data center - each of the approaches below has some merit - but all have trade-offs you'll need to consider too.

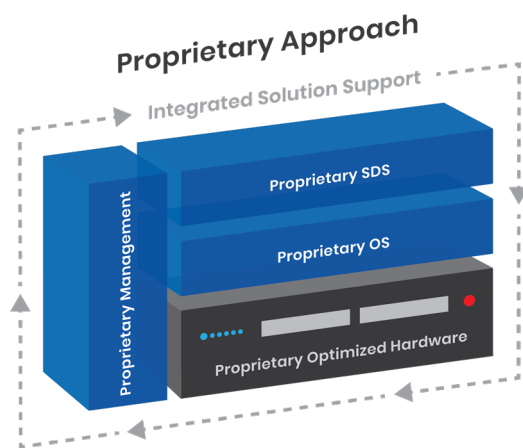
We consider them all in light of the 3 Cs: Cost, Customization, and Contracts.

## Four approaches

### Approach 1

## A proprietary, integrated hardware and software solution

A proprietary storage solution is a combination of integrated hardware and software products and services that are exclusively associated with a specific vendor. In this scenario, everything is provided as a single “package,” including a hardware configuration designed and optimized to run their proprietary software.



### Cost

For most, the appeal of a 100% proprietary solution is the way it essentially simplifies the decisions you need to make. And, from a support point of view, you can generally expect that your vendor is able to provide all the support required to design, deliver and maintain it.

However, financially, this is likely one of the most expensive routes both in initial capital outlay and ongoing support and maintenance contracts. Plus, given the typical nature of these contracts, you are likely to be locked into that vendor, their partners and their product roadmap, including any forklift upgrades they dictate for a good number of years to come.



## Customization

Fewer choices simplifies decision-making, but reduces your ability to adapt. Another pitfall of the proprietary approach is the way it can limit a business' opportunities to be flexible. Vendor lock-in can impede your ability to react as the market evolves, and migrating out of proprietary solutions can at times be costly and time-consuming.

In most cases, a completely proprietary solution will have been designed to fit a specific set of use cases. So if yours doesn't quite match up, you'll find yourself juggling workarounds and business process modifications. Meaning the solution dictates what you are doing, rather than the other way around.

## Contracts

It's easier to hold a single vendor accountable for core solution issues – but be wary of support contract conditions. Your vendor designed, tested and delivered (either directly or through approved partners) the solution, so it's fair to assume they are able to provide all the support required to maintain it.

However, beware of the expiration conditions of support agreements – once support for a product line is discontinued, you may well be on your own, unless you opt for another forklift upgrade.

### Pros

-  Fewer decisions to make
-  Integrated support
-  Single vendors are easier to hold accountable for issues
-  The vendor (should) be able to provide comprehensive support

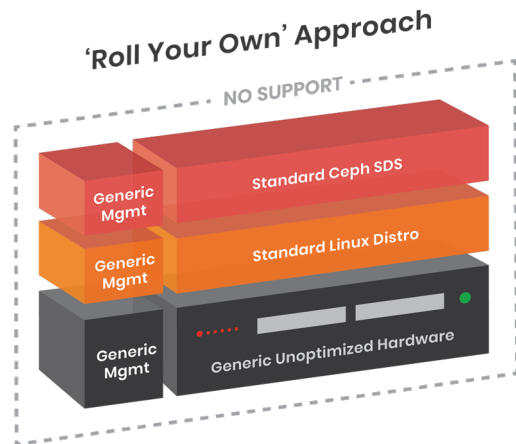
### Cons

-  Typically the most expensive approach
-  Vendor lock-in reduces opportunities to be adaptable and agile
-  If your use cases don't match those the solution was designed for, you'll wind up compromising your business processes to 'fit'
-  If a product is discontinued, you may wind up needing to pay for a complete forklift upgrade

## Approach 2

# A DIY solution

A DIY solution is exactly what it sounds like: you choose your desired open-source software and select white-label, x86 servers from any number of vendors with the best specifications you can afford, and get to work to cobble it together. So how does it compare against the same criteria?



### Cost

For a start, the initial capital outlay is almost certainly the lowest of any option due to license-free, open-source software and cheap commodity hardware. However, an initial cheaper outlay does not always lead to ongoing, long-term savings. There are plenty of hidden costs that might crop up along the way if you're not completely prepared or experienced.

While your initial choices might be broad, your solution will face limitations created due to the lack of a meaningful integration between the layers of hardware and software in the stack. Without careful control, you may end up with an eclectic mix of hardware to manage, which can lead to fragmentation and poor interoperability. The bottom-line costs of poor performance or downtime can quickly escalate...not to mention the costs of expert assistance if you find yourself out of your depth.

Finally, efficiency metrics like power consumption and cooling can rapidly add up. Hot, commodity hardware is expensive to run from a power perspective, and so too is the additional cooling you'll need to manage rack temperatures. And, when it comes to mean time between failures (MTBF), temperature has a direct impact on the lifespan of your hard drives - so while they might run hot and fast, it won't be for long.

## Customization

Another clear benefit of this approach is the level of customization you can achieve when you have complete control over the entire stack. You get to choose the final configuration for everything and it cuts the tie between your hardware and software strategy, giving you the least vendor lock-in and flexibility to seek out great deals, trial new configurations and experiment with software releases.

However, because you're responsible for the integration and optimization of both the box and solution configuration, things are bound to get complicated for novices. If you have an endless supply of "command line warriors," or are perhaps in a test or experimental environment, then maybe this option won't offer as many drawbacks – you will be able to tinker and tweak the deployment to your heart's content, investing all the time you need to coax better performance out of it.

For many however, especially those looking for enterprise-class capability and resilience, this complexity overhead combined with a lack of support and predictable performance might be a stretch too far.

## Contracts

Clearly, contracts in this instance are not an issue, because they most likely won't exist – at least to the same extent as proprietary solutions. So who do you call when your knowledge and time is exhausted? As mentioned above, if you know what you're doing and don't need the support, this is an ideal approach. However there will be no one to turn to for those who reach the point where they need help.

## Pros



Ability to create a completely bespoke solution for your business



Lowest initial capital outlay



Minimal contractual obligations

## Cons



Limitations caused by no meaningful integration between your SDS software and hardware



Managing a completely custom solution can be extremely time consuming



Staff turnover can lead to lost knowledge regarding your bespoke solution



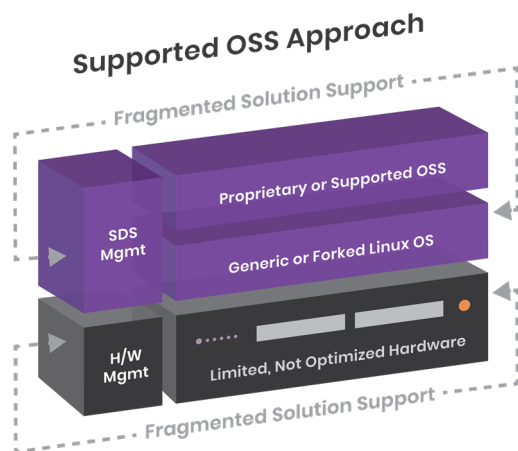
If you lack the depth of knowledge to solve complex integration issues in-house, you may need to rely on external consultants for ongoing support



## Approach 3

# A supported solution

Another popular option is to run a supported software-defined stack on top of generic hardware. Here you might still choose to go with an open-source SDS, but use a supported stack from one of the “red, green or orange” vendors. Alternatively, you could opt for a proprietary SDS, but choose one which relies on third party hardware for their platform to run.



## Cost

Whilst generally not as economical as the pure DIY approach, these types of solutions are typically also not as expensive as proprietary solutions due to the fact you still maintain some flexibility in your choices. Of course, depending on the badge on the hardware you deploy, those costs might well rack up to the same, or even more than a fully integrated proprietary solution.

## Customization

On the plus side, you've still abstracted your hardware platform away from your software architecture which inherently enables better adaptation to your specific requirements. In this instance however, your hardware choice is still going to be limited to approved hardware vendors and configurations dictated by your software vendors which could hamper the outcome of your unique objectives.

## Contracts

With this approach, you have the added complication of at least two vendor relationships, contracts, and support teams. While the on-going management and support of this kind of architecture is almost certainly better than "rolling your own," the separation of the hardware and software layer will likely lead to instances where troubleshooting and support become fragmented between support teams and contracts. You may need to pay the vendor to review, or even design and implement any configuration changes to ensure you do not void either support contract.

### Pros



Fewer decisions and configuration challenges than a pure "rolling your own" solution



Reduces support and management complexity (compared to a completely bespoke solution)



More customizable than a fully proprietary approach

### Cons



Your choices of hardware are limited to those that your SDS provider is willing to support



When problems arise, your hardware vendor might point the finger at your software vendor, and vice versa



You may face limitations due to conflicting hardware and software support contracts



Costs can exceed that of a completely proprietary solution if your support requirements extend into configuration review and design

## Approach 4

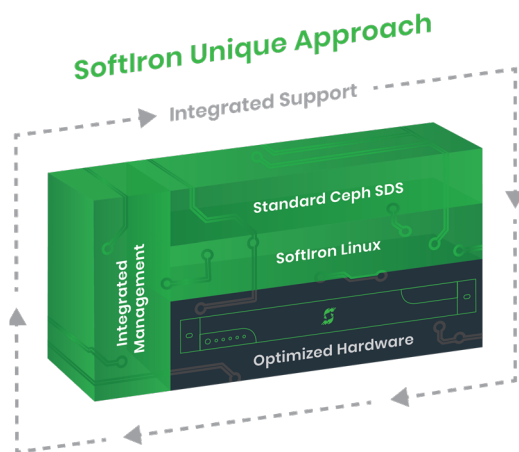
# SoftIron's "task specific" solution

There really is a way to combine the best of all the other approaches while minimizing their drawbacks, and SoftIron® is uniquely, the only company in the world dedicated to pursuing the approach to SDS that we call "Task-Specific".

*We do not optimize a generic unit. We do not configure the cheapest components. We design and build from scratch, specifically for Ceph.*

Early on, we identified the leading, open-source, software-defined storage solution - Ceph - and set out to build a custom hardware platform to optimize its performance. Highly scalable to the exabyte level and beyond, Ceph runs on almost anything. But why run it on anything, when you can run it on appliances that have been purpose-built to optimize its powerful potential?

If you're looking for a solution that combines the best of open-source flexibility with a dedicated, optimized, task-specific platform and integrated support, then SoftIron's HyperDrive® is it.



## But first, why Ceph?

It all starts with Ceph, what we like to call the “Swiss Army Knife” of SDS. Ceph is a unified storage solution, meaning it does block, object, and file-level storage, making it flexible enough to excel across a variety of use cases. Having long been a mainstay within high-performance computing environments, Ceph has been tested under extreme conditions and its adoption within the enterprise is growing quickly. The robust, enthusiastic community brings confidence that this open-source approach is here to stay, and the project has a strategic long term plan.

### Cost

HyperDrive boasts the lowest TCO in its class, a result of several factors. Pure, open-source Ceph means no ongoing software license fees, and, because our appliance is built to do one thing, and one thing only (storage) - you're not paying for unnecessary bells and whistles - whether you're after an economical solution or very high-performance appliance.

HyperDrive is also up to 5x more energy efficient than current market-leading appliances. This delivers savings on power consumption, real estate (due to the fact you can get more nodes in a rack before hitting your power cap), and cooling costs. A cooler appliance naturally needs less air con, and enjoys a longer MTBF, meaning less chance of frequent replacements.



## Customization

By designing and engineering our HyperDrive storage platform specifically for Ceph, we're able to meet enterprises on their own terms, delivering a configuration that exactly matches their unique requirements. We've productized Ceph to fit into various legacy environments, allowing us to connect to virtually everything inside a traditional enterprise.

Our appliances are designed to deliver the best of all worlds – blistering performance and incredible efficiency, manageability, and scalability – underpinned by transparent design and manufacturing, giving you secure provenance no matter your configuration.

And, because we've built the hardware for the software, our solution offers unparalleled, granular levels of control. Our intuitive HyperDrive Storage Manager simplifies the overheads of storage software and hardware management, enabling adoption of Ceph at scale, for the entire organization.

We provide the reliability and support of a fully proprietary solution, while still leveraging the flexibility of SDS and Ceph.

## Contracts


Last but not least, we offer end-to-end support across the stack, but crucially – without the vendor lock-in associated with proprietary solutions. Because you're running standard open-source Ceph, you've still abstracted your hardware and software strategies, so you're free to switch vendors and your software architecture remains untouched.

In the meantime, you enjoy the benefit of the simplicity and performance that comes with integrated hardware and software consulting and support, utilizing clearly defined performance characteristics and integration points. At your fingertips is enterprise-class management and performance, powered by cutting-edge platforms finely tuned to efficiently run open source at scale.

## Pros

-  Secure provenance hardware
-  Task-specific design
-  No vendor lock-in
-  The flexibility of open source combined with the reliability of proprietary solutions
-  Works with mixed-hardware Ceph environments
-  Up-front pricing available through our online [Storage Configurator](#)
-  Our power-efficient hardware offers significant savings for ongoing data center costs

## Cons

-  It's not free. We had to draw a line somewhere!

## In summary

	Proprietary	DIY	Supported	SoftIron
Software lock-in	✓	✗	✓	✗
Hardware lock-in	✗	✗	✓	✗
Integrated support	✓ / ✗	✗	✗	✓
Integrated management	✓	✗	✗	✓
Performance optimized	✓	✗	✗	✓
Cost/TB	\$\$\$	\$	\$\$	\$\$

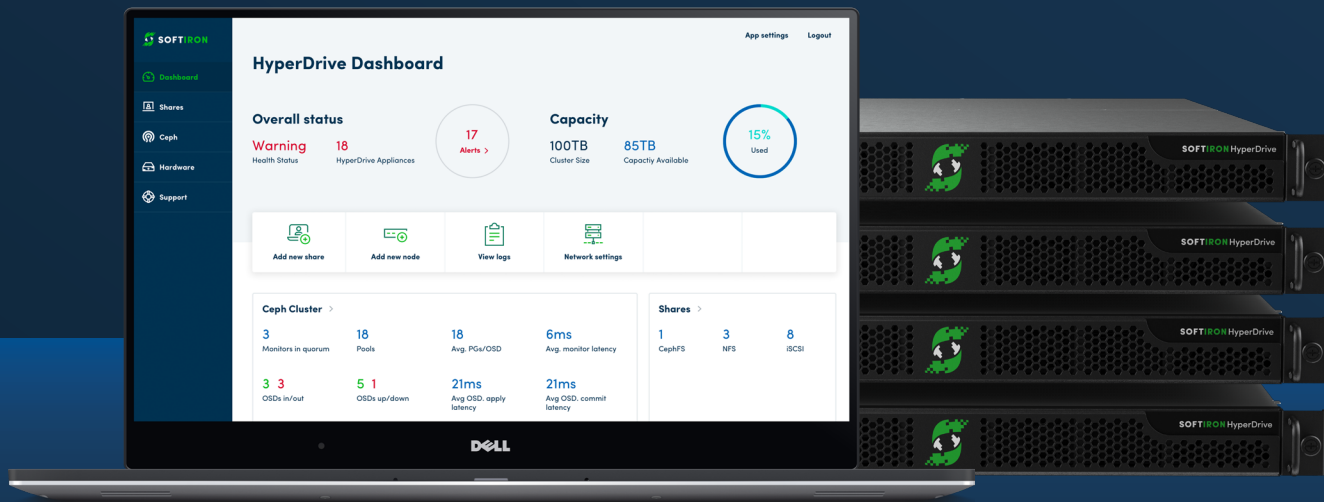
## The choice is yours

There's no one right way to do SDS. But we're confident that our approach is one of the best. By designing and building from the ground up, we are able to provide a highly performant, ultra-efficient and secure platform on which to build a world-class data center. Coupled with our end-to-end support, you really do get the best of everything.

If our task-specific approach sounds like a good fit, **get in touch**, and let's talk about your organization's specific storage requirements.

## Take the next step

Is your data center being impacted by power and space constraints? Is archaic and inflexible infrastructure adding bottlenecks to your processing targets at the edge? Or, is performance being hampered by legacy storage solutions that are not keeping up with the complex and evolving requirements of modern-day workloads? **Take HyperDrive for a test drive** and discover what a secure, task-specific storage solution can do for your organization.



### References

<sup>1</sup> IDC's Cloud and AI Adoption Survey, January 2018

<sup>2</sup> Gartner, 2020 Strategic Roadmap for Storage, August 2020

<sup>3</sup> Research and Markets: Edge Computing Market by Component (Hardware, Platform, and Services), Application (Smart Cities, IIoT, Content Delivery, Remote Monitoring, AR and VR), Organization Size (SMEs and Large Enterprises), Vertical, and Region - Global Forecast to 2025, August 2020

<sup>4</sup> 2020 Red Hat Report: The State of Enterprise Open Source / 2020 TideLift: The Managed Open Source Survey



SoftIron® is the world-leader in task-specific appliances for scale-out data center solutions. Their superior, purpose-built hardware is designed, developed, and assembled in California, and they are the only manufacturer to offer auditable provenance.

SoftIron's HyperDrive® software-defined, enterprise storage portfolio runs at wire-speed and is custom-designed to optimize Ceph. HyperSwitch™ is their line of next-generation, top-of-rack switches built to maximize the performance and flexibility of SONiC. HyperCast™ is their high-density, concurrent 4K transcoding solution, for multi-screen, multi-format delivery.

SoftIron unlocks greater business value for enterprises by delivering best-in-class products, free from software and hardware lock-in.

For more information visit [softiron.com](https://softiron.com)

---

 [info@softiron.com](mailto:info@softiron.com)

 [@SoftIronNews](https://www.facebook.com/SoftIronNews)

 [@SoftIron](https://www.linkedin.com/company/SoftIron)

 [@SoftIron](https://twitter.com/SoftIron)

 [SoftIron](https://www.youtube.com/SoftIron)

[softiron.com](https://softiron.com)

Copyright © SoftIron Limited, 2020. All rights reserved. SoftIron, HyperDrive, HyperCast, HyperSwitch and the SoftIron logo are registered trademarks of SoftIron Limited. ARM is a registered trademark of ARM Limited (or its subsidiaries) in the EU and/or elsewhere. AMD, the AMD arrow logo, and combinations thereof are trademarks of Advanced Micro Devices Inc. Socionext is a registered trademark of Socionext, Inc. SoftIron disclaims proprietary interest in the marks and names of others. This document is for information only. No warranties are given or implied. Contents are subject to change without notice. SoftIron Ltd is registered in England (no. 8172199) at Level 1, Devonshire House, One Mayfair Place, London W1J 8AJ, United Kingdom.