SOFTIRON

Solution Brief

Achieve sustainable efficiency and limitless scale for your HPC cluster with SoftIron

HPC faces many challenges in the decade to come, including reducing the net carbon emissions produced by data centre infrastructure, and managing the diverse requirements for successful edge deployments. HyperDrive provides the future-proof storage solution that enables power-efficient and flexible block, file and object storage.

Performance isn't the only storage challenge in high performance computing. Infrastructure teams must account for flexibility, simplification, and sustainability needs to meet increasingly diverse workloads and deployment scenarios.

Diverse HPC deployments still share commonalities

Imagine the HPC world as the castle between supercomputers in the IO500 list, battling it out for computational supremacy. And while all this is going on, they're processing breakthrough scientific work. Maybe a little quantum physics – the time step analysis of material interactions during the first few seconds of a nuclear explosion, for example.

The labs doing that kind of work tend to buy really big compute clusters, which need a large and wide variety of storage to support each stage of the data life cycle. In the past decade, though, the HPC mindset – having been branded at different periods as "big data", "business intelligence", "analytics", etc. has come into regular use in enterprises of all kinds of sizes and shapes, with mixed rates of success. Typically, the level of success tends to be driven by whether a project's focus is on clearly-defined outcomes and setting realistic and incremental expectations. But across use cases and deployment sizes, there's little variation amongst the tools and processes employed from an architectural and data interaction perspective.

Each project will tend to:

- Use clusters of compute resources (either computers or dedicated ASICS like GPUs) operating in parallel toward a common goal.
- Require a variety of data interactions that have particular characteristics, and have relatively well-understood project start and ends or else cyclical flow and feedback mechanisms.

Whether a single researcher renting time on a cluster or a team of scientists building complex testing and feedback automation systems, there always has to be a point at which the data is generated, shared, and aggregated.

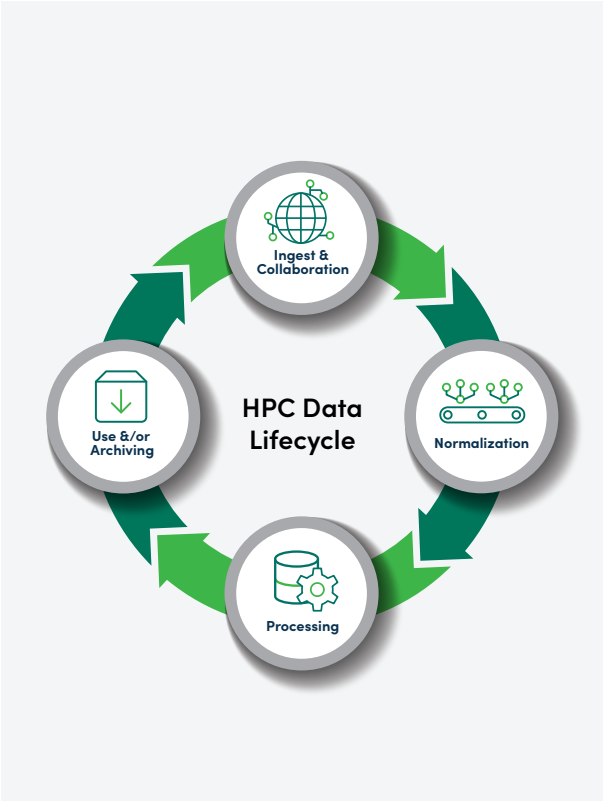

For many, scalable shared file (or increasingly block storage) systems provide the platform for collaboration that reduces the risk of data corruption and protects against human error through rich data management features, such as snapshots. As teams and data sets grow, the need for a scalable file system becomes greater and traditional array-based platforms reach limits in either capacity or throughput. Additionally, teams may want to collaborate using convenient human-friendly interfaces, which often require protocols that work across many desktop platforms – like SMB or NFS.

After data is collected into their lakes, it gets further normalised or transformed to make it ready for HPC compute systems to ingest for processing. This requires manual interaction (i.e. a human has to do it) or automated batch processing systems. Increasingly, this data preparation makes use of object storage as a means to collect and structure metadata – because the "data about the data" is a valuable resource, especially in machine learning or artificial intelligence workloads.

Once ready, the data moves close to the computational system, to prepare it for ingest into the cluster.

> **Traditionally, the migration of data across these three steps in the life cycle has required physical migration from one storage system to another.**

This is because the storage systems for collaboration have not historically been able to deliver the kind of performance needed by the computational system. These systems – especially at larger scales – have historically used the latest storage technologies, high-speed networking, and file systems built for speed (sometimes at the expense of usability and resilience).

The result of that behaviour is that up to 30% of a storage system's processing time/capacity is spent migrating data between tiers rather than completing the work itself, representing immense inefficiency in overall architecture.

The processing of data itself tends to require large amounts of either IOPS or throughput depending on the kind of data and the kind of analysis. Frequently, it requires metadata accessed through specific low-level POSIX calls or through the storage system itself.

The resulting processing does not typically destroy the data used during this process, and it generates more data as output – representing the results of the analysis, scientific model, or inference engine, which gets used in service of the overall business or organisational outcome. Both sets of data get shipped out of the highest tier of storage to their next location – either an archive of some kind, depending on the potential future value the data may provide, or to storage systems accessible by the human or software-based inference systems that derive additional value from that data.

> ### For those seeking a more efficient HPC solution, HyperDrive provides the right combination of usability, scalability, performance, and sustainability

Each point in the data life cycle represents an opportunity for data corruption, either due to human error or systemic failure. Often the human error comes from misunderstanding the fundamental characteristics of the medium they're using to persist data. SoftIron's HyperDrive product line uses the most resilient storage system on the market – Ceph – and provides an interface to it that provides users with context about resiliency and configuration parameters without requiring deep domain expertise in the operation of the storage cluster.

Most HPC processing happens in a Linux environment and is managed by open source tools. HyperDrive provides a consistent interface, using the Linux operating system and open source storage. Open source communities tend to build integrations into other open source community projects, and the robust API and infrastructure and code capabilities available in HyperDrive are no exception.

Whereas historically teams would need to migrate from one storage system to another, HyperDrive provides a degree of flexibility that can remove that burden for many HPC workloads. Because it supports file, block, and object representations of data, and because its file system representation of data spans both user-oriented (NFS and SMB) and clustered computing-oriented (CephFS as a parallel and massively scalable file system), many organisations may find the transition much simpler since it happens behind a single pane of glass, or unnecessary due to the tiering capabilities natively expressed by HyperDrive and Ceph, enabling organisations to tier across both NVMe and HDD based appliances based on how much performance is required at a given point in the data life cycle. With HyperDrive, many organisations can consolidate their storage architecture through the entire data life cycle.

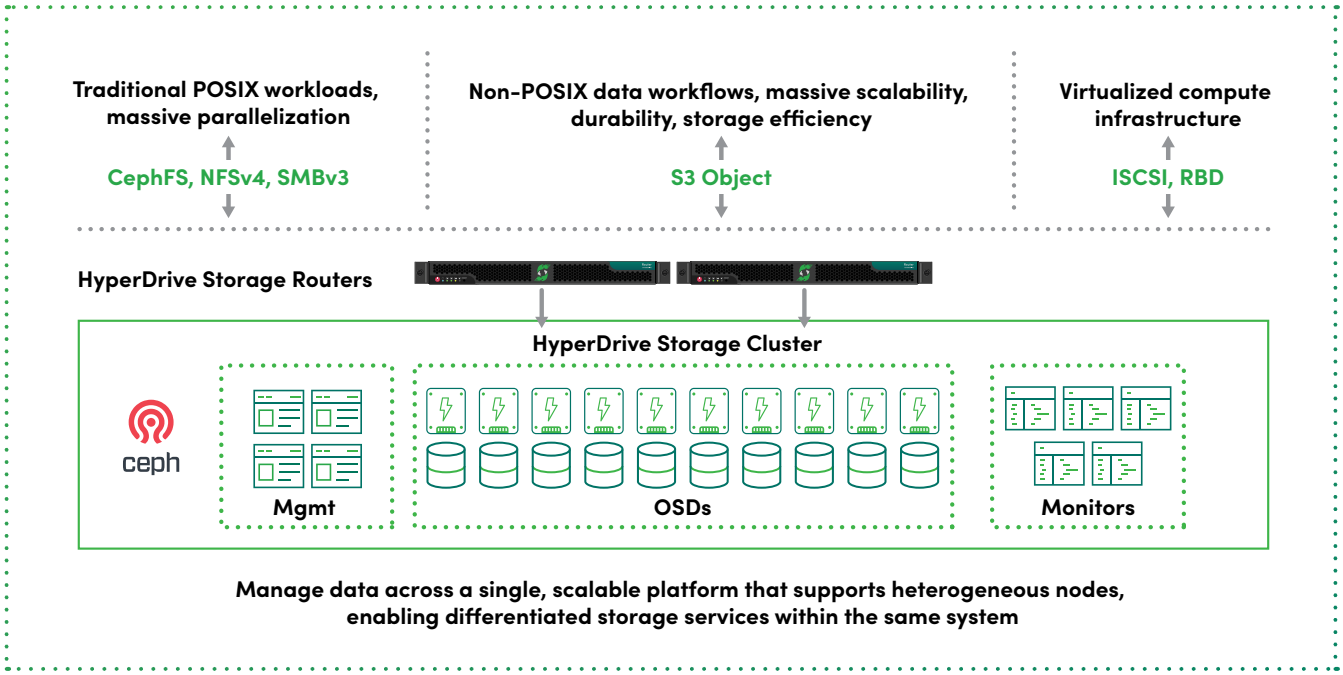

Scalability provided by HyperDrive enables large-scale HPC computing in two key ways. First, Ceph was designed for near-infinite scalability in terms of capacity and performance. CephFS throughput improves semi-linearly as the number of nodes and workloads parallelism increases, as shown in the testing performed by **Sandia National Laboratory**. Secondly – and uniquely to the industry – HyperDrive is designed from the ground up to optimise Ceph performance for a given set of hardware. HyperDrive systems are not based on commodity hardware. As a result, the data path has been optimally streamlined to get maximum performance from the media. This provides the most balanced capacity and performance per node, particularly important during the processing stage.

That task-specific thinking extends not only to performance improvements but also to simplified user experience across the storage system life cycle. Through a holistic approach to running a clustered storage system, we can significantly reduce the operational overhead of routine maintenance and improve the life of components in the system. One such improvement is the Ceph Button, which enables hardware replacement without the person replacing the hardware needing to have a detailed understanding of the storage architecture. As data and data processing move increasingly to the edge, that significantly improves the total cost of storage ownership. In addition to making the cluster simpler to operate, streamlining nodes for a single purpose makes them substantially more energy efficient than commodity hardware. HyperDrive nodes provide up to 80% energy savings vs. systems based on commodity subcomponents. In fact, for every 10 PB of data storage shipped by SoftIron, an estimated 6,656 tonnes of CO_2e are saved by reduced energy consumption alone.*

HyperDrive makes scalable, flexible storage a snap

SoftIron's HyperDrive family of data centre appliances can scale to nearly infinity, smoothly matching your organisation's specific performance and capacity needs. Each appliance in the HyperDrive family is designed to take full advantage of Ceph's software-defined storage architecture and inherent support for object, block and file storage.

Access unified storage, with global support

Utilise file, block and object storage natively in Ceph, allowing for direct backup, archival, and repatriation of your data as you need it.

Not familiar with Ceph? HyperDrive Storage Manager simplifies day-to-day Ceph operations, and SoftIron offers globally available support services to assist enterprises working with Ceph.

Manage multi-tenancy storage across research teams

Data residency legislation and cybersecurity concerns are pushing organisations to shift the data they've previously stored on a major public cloud provider's servers to ones that are owned, managed and located on unparalleled density make it ideal for creating secure sovereign cloud storage.

Deploy anywhere, even at the edge

HyperDrive offers power-efficient, hyper-dense storage in a 1U form factor allowing you to maximise the storage potential of every rack. Combined with its hyper-efficient cooling features, HyperDrive is ideal for all kinds of edge use cases, whether you're collaborating across geographies or collecting performance data from the back of a moving vehicle.

Flexibly scale up or down as you need

SoftIron offers an OPEX model for HyperDrive – only pay for what you need, with minimal barriers to scaling out when required. And because HyperDrive unites the power of a proprietary system with the flexibility of open source, you modify your cluster to suit your needs, with no limiting contracts.

Enhance your net-zero initiatives with task-specific hardware

Efficient IT infrastructure is a crucial step forward for HPC organisations looking to reduce their carbon impact. HyperDrive fundamentally reduces power and cooling requirements by being engineered specifically for Ceph, achieving efficiency and longevity impossible with generic hardware. And with SoftIron's adoption of an 'edge manufacturing' approach, we're building a more resilient, more sustainable supply chain for the future.

Consolidate your storage footprint across the data life cycle

The HyperDrive range includes a mix of performance (NVMe) and capacity (HDD) appliances that can easily provide multiple storage tiers (hot, cold, archive) within a single name space.

Achieve wire speed performance, without the wires

Performance is nothing without reliability. HyperDrive delivers both, thanks to its high speed backplane, designed to operate at up to 6 terabits per second.

That's perfect data transmission at wire speeds, with wire-free reliability 24/7/365 ++.

Maintain data integrity

Monitor cluster health with HyperDrive Storage Manager, and ensure your backups stay fresh with Ceph's data scrubs, providing a bit-for-bit comparison of your replicated data, protecting against bit-rot.

Greater control, simpler multi-vendor storage management, and a clear path towards managing a globally distributed, locally embedded workforce

SoftIron HyperDrive delivers a solution to support diverse IT initiatives, including:

- Centralised storage control
- Standardised data governance
- Massive scalability
- Award-winning performance
- Energy-efficient ICT infrastructure

Additionally, if your organisation is planning for a cloud repatriation project, this unified solution offers greater cost predictability and control, allowing for a gradual shift of critical resources away from the public cloud and back into your own fully-controlled storage. Similarly, transitioning to an edge computing approach can be handled progressively with minimal disruption to employees using SoftIron's HyperDrive Ceph storage.

Could this solution work for you? No guessing needed - just take HyperDrive for a test drive.

SoftIron's **HyperDrive Test Drive** allows you to try HyperDrive in your environment with our pre-configured storage cluster for up to 90s days. If you decide to go ahead with HyperDrive, we offer flexible contracts without lock-in and pay-as-you-go subscription pricing that can be cancelled at any time.

*Read more in the SoftIron Carbon Footprint Report 2022, Simon Cooks, Earth Capital
https://softiron.ltd/SoftIronCarbonFootprintReport

SOFTIRON

info@softiron.com @SoftIronNews @SoftIron @SoftIron SoftIron softiron.com

SoftIron makes the products that underpin the next evolution of IT infrastructure.

Our blueprint is radical. Taking full control over design and manufacture of platforms optimised for selected open source software, our highly integrated products reduce space and energy footprints while delivering extraordinary performance. Challenging traditional IT manufacturing and organisational strategy, we've developed a model that enables us to create a more resilient and connected business for the customers we serve. A commitment to openness, transparency, and simplicity helps address emerging multi-faceted threats while eliminating the vendor "lock-in" so common elsewhere.

0276-MAR-03/22