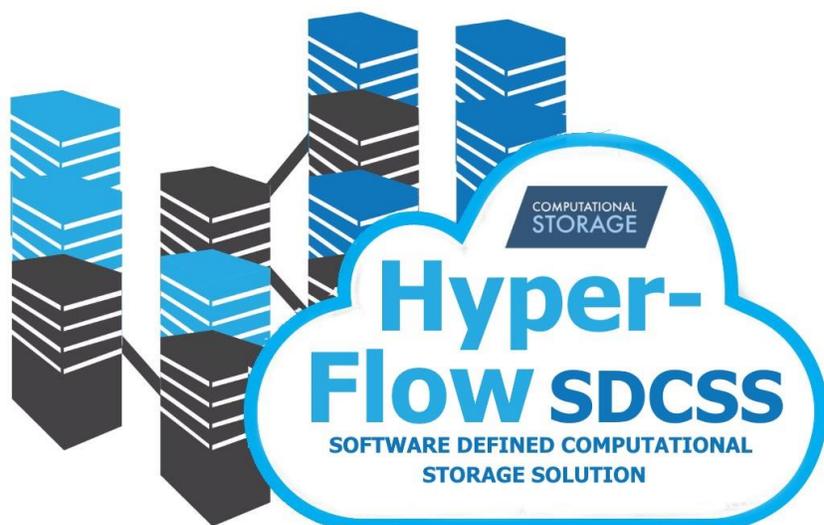


HyperFlow Software Defined Computational Storage Solution (SDCSS)

Based on Ceph using Computational Storage NVMe by ScaleFlux

Introduction

HYPERSCALERS



Thursday, 1 September 2022

INTRODUCTION

HyperFlow is a software-defined computational storage solution (SDCSS) that has been engineered to redefine performance, price, versatility, scalability and availability expectations in the modern storage solutions marketplace.

Hyperscalers [1] and ScaleFlux [2] have partnered as hardware providers to jointly co-develop and qualify a high performance computational storage solution (CSS) against the industry standard Ceph software defined storage (SDS) platform.

The resultant HyperFlow SDCSS platform delivers all of the well-known benefits of Ceph in the context of blazingly fast (NVMe) Computational Storage containing high performance storage processors able to optimise dynamically for both performance and capacity.

Hyperscalers have developed this platform with an all-flash Non-Volatile Memory express (NVMe) Ceph storage cluster using CSD (Computational Storage Drive) technology from Scaleflux [3] in 1U servers provided by Hyperscalers [4].

This unique combination of well-known, trusted hardware elements and the Ceph platform is the perfect combination for delivery of highly available, mature, and flexible Block, File and/or Object-storage services provided by Ceph.

Computational storage drive expansion presents a larger drive capacity to the operating system than the physical size of the drive itself. This ground breaking technology enables customers to realise cost saving multiples typically of 2.5 to 3.5 times (and potentially beyond) for NVMe class drives. This in turn brings the price point of NVMe class drives much closer to the cost per terabyte normally associated with hard disk drive (HDD) equipment.

HyperFlow SDCSS is highly configurable with our minimum base implementation able to support 40Gbps cluster throughput against 691.2 TB total capacity, all in 3RU. If density is a requirement, this base design can easily double in capacity or throughput (or both) with no additional server hardware required. Conversely, capacity and/or throughput can be expanded across additional cluster nodes for greater availability – the choice is yours.

HyperFlow is a fast, reliable, low cost and turnkey SDCSS platform that has been rigorously qualified through Hyperscalers' Hardware Appliance Design process to ensure full-stack compatibility across the Ceph, operating system and device driver layers.

Taken in conjunction with Hyperscalers trusted ability to support enterprise class customers across multiple localities, HyperFlow SDCSS should be viewed as a highly engineered, industry leading storage platform solution that can be trusted to support storage requirements across any customer context.

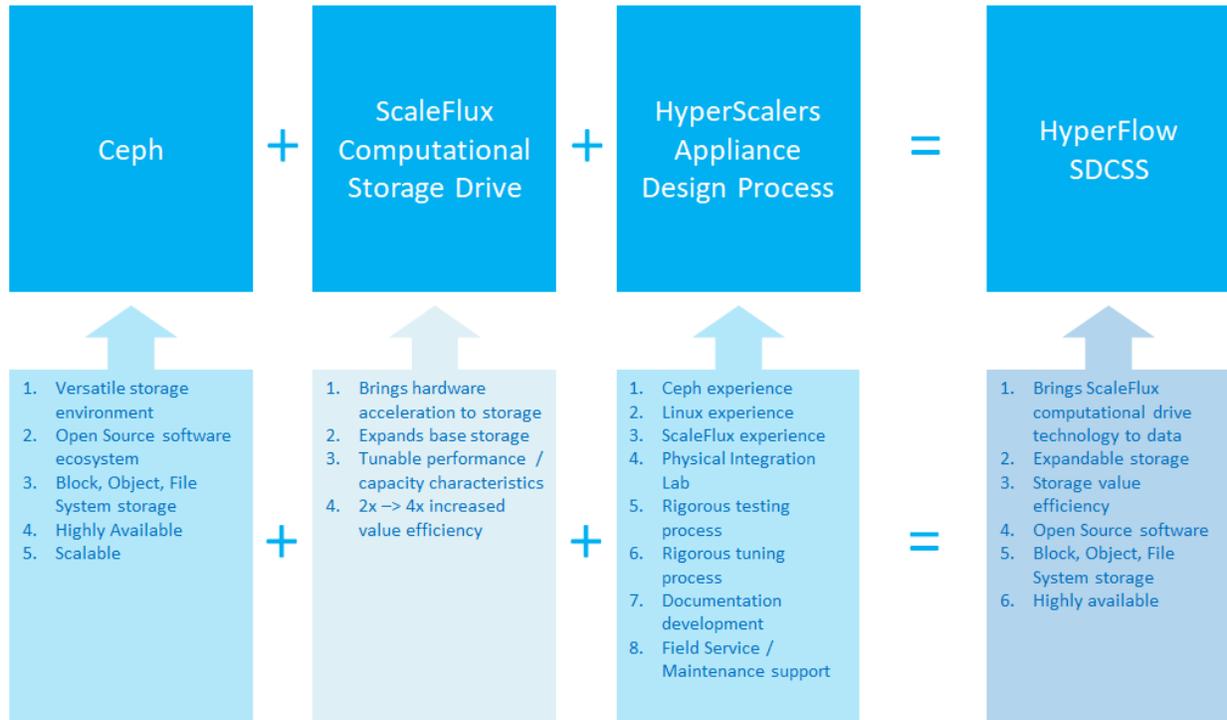


Figure 1 Values of the HyperFlow Software Defined Computational Storage Solution

Ceph with hardware acceleration allows you to accomplish more with your data.

Businesses, academic institutions, global enterprises and others all share a common need for fast access to critical data that in many cases is irreplaceable. Hence we have undertaken the significant engineering effort needed to build the ultimate realisation of software defined storage through the use of a hardware accelerated storage context.

The Ceph platform inherently provides trusted software defined capabilities fulfilling data management requirements across key domains including flexibility (of storage types), configurability, reliability and scalability.

HyperFlow SDCSS takes Ceph capabilities and places them on top of blazingly fast computation storage drives by ScaleFlux. SDCSS speed is derived not only from the use of all-flash NVMe, but in particular by the use of onboard hardware storage processing devices. These hardware accelerators implement key storage processing functions directly into silicon that were once performed in software running on general purpose processors (for example on RAID management cards).

Silicon based (meaning hardware accelerated) functions are what allows HyperFlow SDCSS to deliver a combination of performance and compression metrics that are not otherwise achievable in a software defined storage context (even when using all-flash NVMe).

Ceph, via its advanced CRUSH algorithm, supports automated data redundancy, self-management daemons and much more. This means that HyperFlow SDCSS ensures your data is safely stored, instantly available and optimally distributed for effective disaster recovery.

Engineers within Hyperscalers Canberra based laboratory have worked through an extensive fine tuning regime in order to optimise HyperFlow SDCSS for absolute performance. This has included handling of numerous critical CPU and network layer parameters, in addition to software stack qualification (done in collaboration with ScaleFlux). HyperFlow SDCSS utilises only latest generation Ice Lake Intel CPUs, and very high speed (100gbps) network interfaces.

Data is *invaluable*, so HyperFlow SDCSS keeps it *safe*. Without effective protection, data loss can be detrimental to business, incurring monumental costs and irreparable damage to critical data. All data protection features available under Ceph are equally as reliable in an SDCSS context as for a non-accelerated platform context.

HyperFlow SDCSS provides speed and capacity cost advantages transparently - allowing you to focus on the Ceph storage use-cases that are critical to your organisation.

You can continue to streamline data management and compliance processes, reduce downtime and network bottlenecks, and remove traditional barriers to scale.

HyperFlow SDCSS enables any organisation to build out the Ceph storage use-cases that are critical to them using the most highly optimised hardware / cost solution that is currently available.

HyperFlow SDCSS is broadly applicable across all verticals

No matter what type of organisation any of us are situated in, we all share a common need for fast access to critical data.

The following examples highlight specific use-case scenarios that are relevant in the context of some well-known environments:

SMEs

SDCSS provides scalable and reliable data storage for object, block and file storage, without the need for the investment in expensive hardware.

Global organizations

Hyperflow SDCSS can be configured to maximise the value of geographically distributed data centres to deliver highly available, resilient data. By placing relevant organisation-wide data physically at the regional locations where it is needed, your people and processes can be empowered in ways that may not have otherwise been easily possible.

Academic institutions

The extensive customisation capabilities of Hyperflow SDCSS enable tailored data curation and management control, in turn supporting the collection of both unstructured and structured data at scale and reliable backups of your research for secure, robust data storage.

Start-ups

Hyperflow SDCSS makes it easy to purchase only what you need initially and also to scale on demand. This means that you can quickly adjust your cluster balance as you add, replace or remove storage media. If your start-up business experiences sudden exponential growth, CSS will grow your cluster right alongside you with no license renegotiation needed.

Developers

The open-source adaptability inherent in Hyperflow SDCSS is perfect for inclusion into your own software environment and architecture, supporting unimpeded experimentation with cloud hosted services and/or operations across large volumes of block, file and object data.

Financial Services

Modern financial services companies are heavily reliant on storage systems to capture, store and process event information in some cases at staggeringly high speeds. This is particularly true within the new breed of Fin-tech style operations where speed and scale flexibility are paramount. Hyperflow SDCSS can help to maximise the cost-performance-reliability equation that is critical within these environments.

Government

Nowhere is the need for cost-optimised, flexible and fast access to large quantities of data more apparent than in the context of government organisations handling critical data relating to an entire population. Hyperflow SDCSS can provide not only the capacity and performance benefits but also the reliability, maintenance and support backup that are always required by government users.

Use HyperFlow SDCSS to achieve...

Intelligent software integration

The technology architecture employed within HyperFlow SDCSS provides a flexible and adaptable foundation for integration with a broad range of services, applications and cloud technologies. From web-scale content repositories to machine learning and artificial intelligence you can gain detailed insights into your data and leverage it to improve

business outcomes. Use inbuilt HyperFlow SDCSS protocols to work with block, object, and file storage, or create your own interface using the LIBRADOS API.

Scalable storage

With HyperFlow SDCSS there is no limit to your data growth. As a virtualised storage system HyperFlow SDCSS can scale as you require without the confines of traditional hardware storage. HyperFlow SDCSS is designed to ensure the reliability, performance and availability of your data - even alongside exponential growth.

Undoubted reliability

Advanced algorithms, intelligent object storage daemons and automated self-management capabilities enable HyperFlow SDCSS to provide businesses with high reliability and quick, detailed insights into cluster health. It is well known that manual data management processes can result in operational delays and overlooked cluster issues. In order to mitigate these problems, HyperFlow SDCSS automates round-the-clock monitoring, backups and data protection.

Business continuity

There are few businesses that can run smoothly without access to key operational data. Regardless of whether the business context is a customer bookings platform or an automated logistical system (for example), downtime always reduces customer confidence and can have a significant impact on revenue generation. HyperFlow SDCSS will help you to guarantee business continuity by enabling you to implement finely tuned and automated data redundancy across devices, racks and geographic locations.

Controlled investment

In contrast to the hefty payments and costly infrastructure requirements imposed by some storage vendors, HyperFlow SDCSS will run on your existing off-the-shelf hardware. This can make transition to HyperFlow SDCSS far more cost-effective than when dealing with any of the incumbent storage marketplace alternatives. Additionally, by balancing your data replication across a cluster, it becomes easy to monitor cluster capacity and limit any purchasing of additional storage only to the circumstances where it is actually needed.

Choose the hardware that matches your needs

HyperFlow SDCSS provides unparalleled flexibility in your choice of hardware. In fact it is able to run on just about anything. By freeing organisations from committing to a single hardware vendor, HyperFlow SDCSS supports constant adaptation and innovation, with all hardware components able to be swapped out as needs change. This enables you to provide access to Object, Block or File storage from one unified cluster, while decoupling

your data management architecture from any specific hardware elements you may choose to use.

Remove bottlenecks

Bottlenecks that seemed acceptable in smaller deployments may rapidly become costly and unwieldy at scale. As a distributed storage system, HyperFlow SDCSS provides seamless data retrieval by enabling client applications to calculate the location of data within a cluster directly. This capability removes the need for a traditional metadata server. By supporting a direct path for update and retrieval, network traffic is reduced and a critical single point of failure is negated, providing smoother and more reliable service to you and your customers.

A solution you can rely on long-term

Ceph's global community of industry leading developers ensure that Ceph never falls behind the storage technology curve. There are more experts reviewing the Ceph code base on a daily basis than would ever occur for a proprietary solution. This results in quick response to evolving security requirements and new innovations in storage technology. Ceph's core philosophy and open-source model ensures continued development and practicality. Ceph is built to address the needs of its users, and it continues to offer additional features and efficiencies with every new release.

Built to scale

HyperFlow SDCSS offers *high density storage, scalable to petabyte* levels within a smaller footprint (minimum 3 x 1U¹). With typical cost reduction of 250% per TB of NVMe, HyperFlow SDCSS is a fraction of the traditional costs associated with NVMe class storage.

Fast

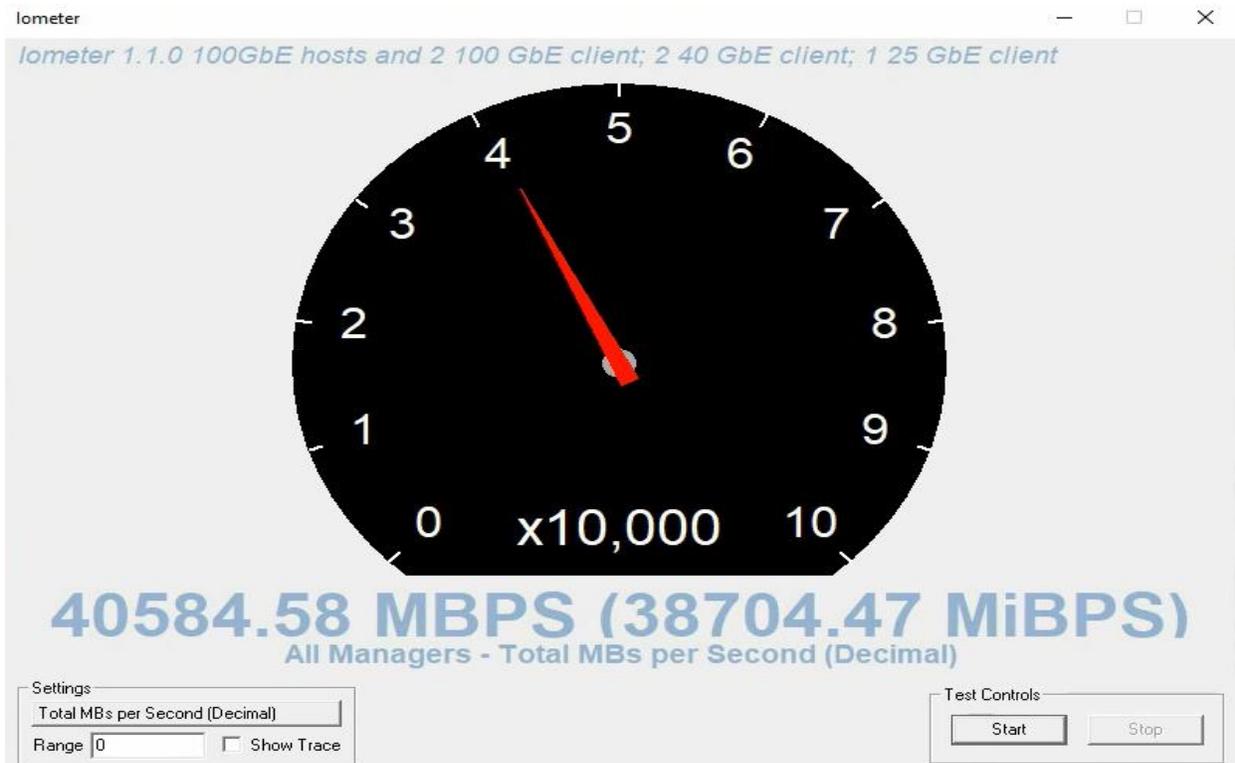
It is possible to *deploy and scale* a POC²/ production ready storage server in no more than few hours³. With ScaleFlux NVMe Computational Storage, HyperFlow SDCSS achieves throughput up to *40 GB/s*⁴.

¹ Refer to CSS Reference guide for details on the hardware.

² Contact Hyperscalers for a demo/ test drive with LaaS.

³ For Hyperscalers tested hardware only.

⁴ Refer to CSS Reference guide for more information.

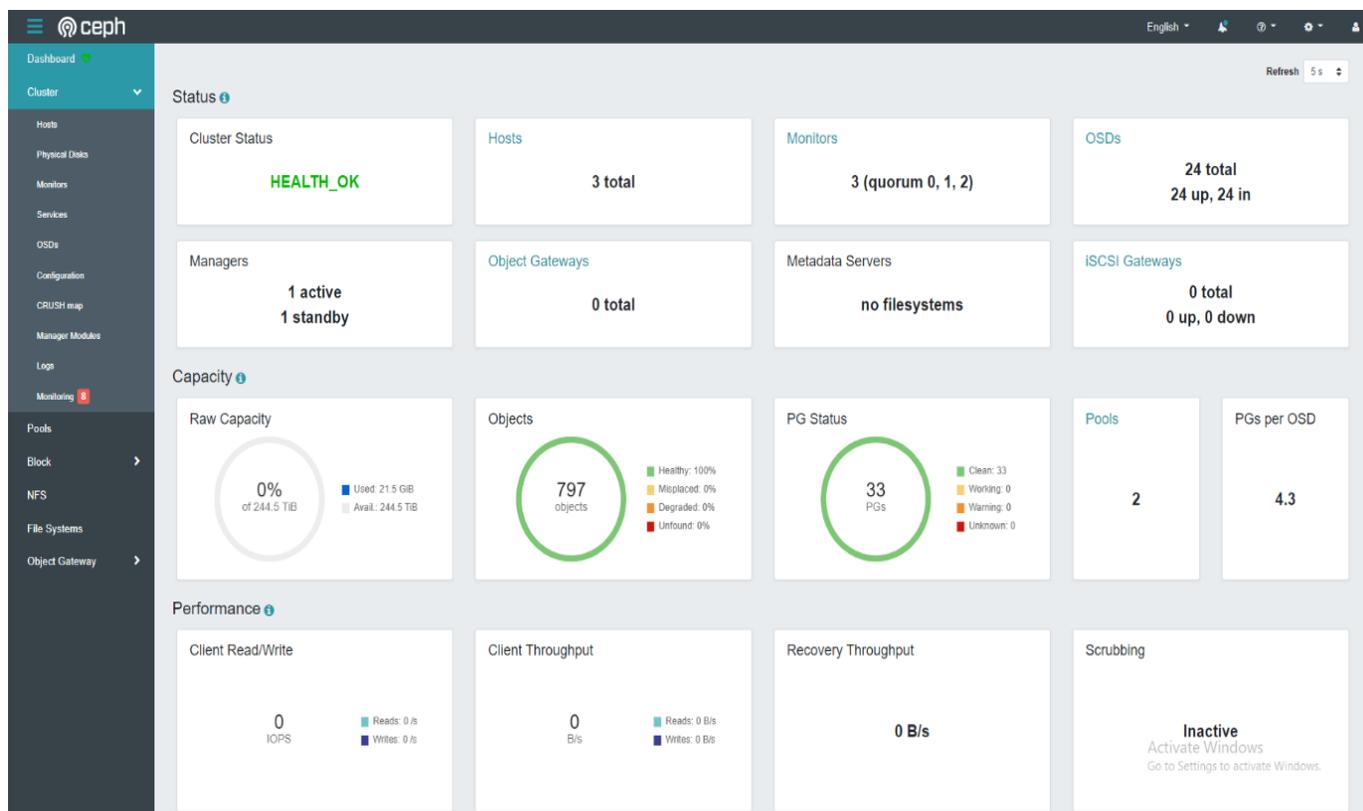


Intuitive Management Interface

HyperFlow SDCSS provides an intuitive and easy to use management user interface / dashboard that enables you to comprehensively manage and monitor all aspects of your environment including:

- Monitoring & audit of cluster logs,
- Addition, configuration and removal of hosts, object storage drives, monitors and managers,
- Creation and configuration of storage pools (with replica/erasure coding) to support Object, Block and File system storage,
- Creation and management of storage volumes and images,
- Creation and management of image mirroring,
- Create and management of object buckets.

The following screenshot of the Ceph management interface depicts high level Ceph environment status including cluster health, number of Ceph system elements and key capacity and performance metrics:



High Availability and Low risk

Courtesy of the highly evolved Ceph architecture, HyperFlow SDCSS offers superior reliability and availability. This in turn allows storage providers to deliver storage services with minimal downtime. *HyperFlow SDCSS mitigates risks and protects data through replication, mirroring, and erasure coding⁵.*

Storage solution for every application

HyperFlow SDCSS offers block, object, and file system storage on demand. This enables it to dynamically serve any combination of hypervisor, container, web/monolithic application and direct clients.

HyperFlow SDCSS Storage Types

Object storage

Object storage is a data storage architecture that manages data as individual binary objects. Each object includes data, metadata (providing reference information on the data) and a unique ID. Object storage allows customers to connect web applications to uniquely

⁵ Refer to CSS Reference guide to know more on replication and erasure coding.

identified storage objects acting as storage points, and from there to store any form of unstructured data within them. *In HyperFlow SDCSS, object storage is deployed in the form S3 / Swift buckets which can either be attached to by internal consumers or published as a public or private cloud accessible item.*

File System storage

File System (FS) storage is a data storage structure in which data is stored and maintained as files within folders. This approach will be familiar to any desktop computer user navigating their file system folder structure via Explorer (Windows) or the Finder (MacOS). The File System storage structure provides readability and convenience. *File System storage under HyperFlow SDCSS can be deployed as a network file system (NFS) or as CephFS*

Block storage

An abstract form of File System storage is known as Block storage. Block storage supports storage of structured data to be used by VMs, applications and end users. *Block devices provided by HyperFlow SDCSS can be attached to clients as a RADOS block device (RBD) or through iSCSI.*

Support

Hyperscalers offers full hardware and software support to HyperFlow SDCSS customers across multiple geographic localities. Depending on severity⁶, analysis of any specific issue can range between a few business hours to three business days.

⁶ Refer to your Support agreement with Hyperscalers or contact Hyperscalers to know more about support.

Comparison

The following table illustrates key comparison metrics between HyperFlow SDCSS and other storage system products:

Hyperscalers Computational Storage Solution	Purestorage FlashBlade	Purestorage FlashArray	Excelero NVmesh
Minimum 3x 1 RU ⁷	Minimum 4 RU chassis with 7 blades	Minimum 5 RU	Minimum 3x 1 RU
File, Object and Block storage	File, Object storage	Block Storage	File and Block storage
Up to 12 drives per node (1 RU)	Up to 15 blades per chassis (4 RU)	Up to 40 Flash units (5 RU)	Up to 10 drives per node (1 RU)
Up to 1.5 PB ⁸ (for min spec)	Up to 692 TB	Up to 1.48 PB	Up to 230 TB
Up to 40 GB/s ⁹ (for min spec)	Up to 15 GB/s with 15 blades		Up to 35 GB/s
Up to 2 x 100Gb/s NIC per node	Up to 16 x 100Gb/s Ethernet ports per FlashBlade system	2 port 100Gb/s NIC	Up to 2 x 100Gb/s NIC per node

⁷ Refer to CSS Reference guide/ contact Hyperscalers to know more.

⁸ Depends on CSD drive model and compressibility of data.

⁹ Depends on multitude of factors including data composition, network characteristics (and more).

Please refer to the HyperFlow SDCSS Reference guide for more information.

Infrastructure

The following figure shows how the HyperFlow SDCSS architecture relates to some example contemporary storage use-cases, and therefore how it can empower HyperFlow users to store, manage and utilise their data:

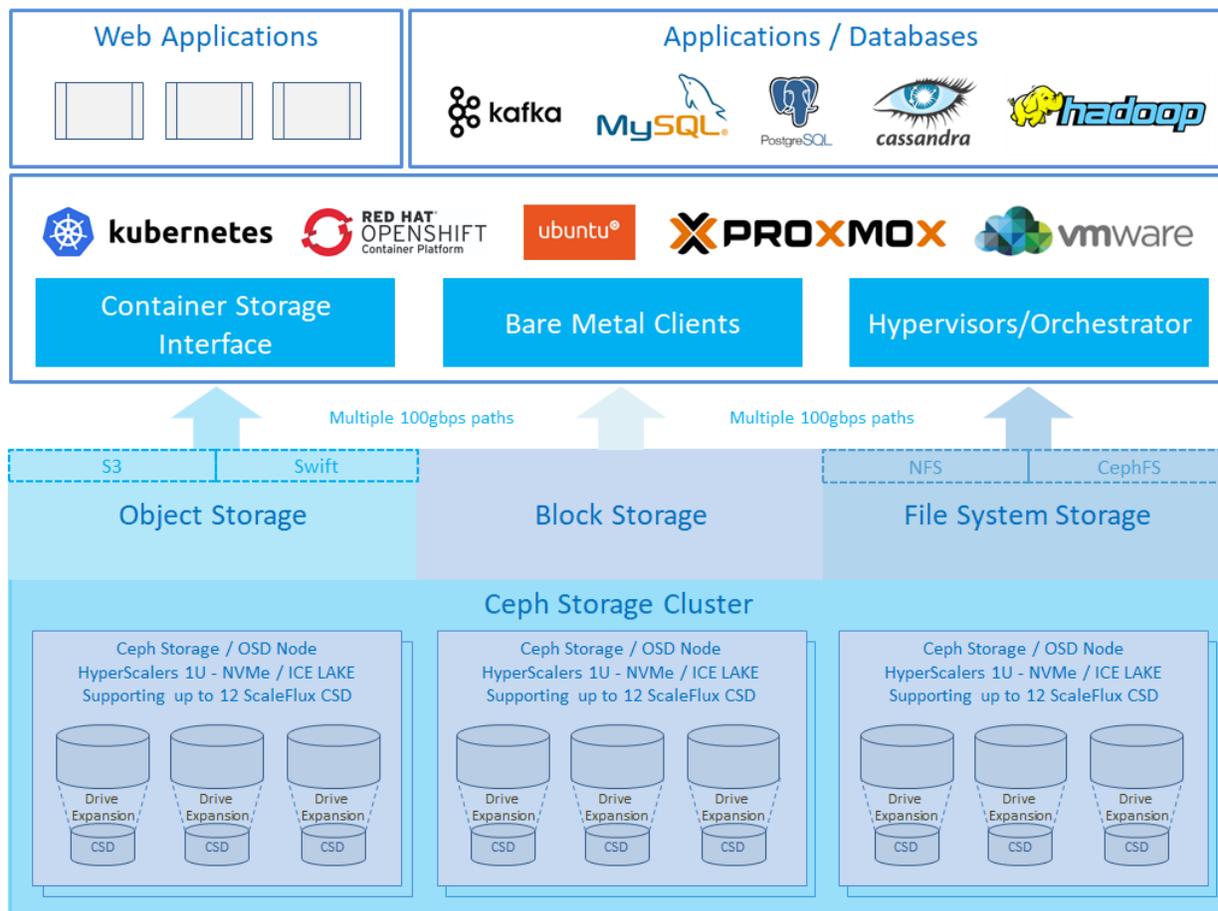


Figure 2 HyperFlow Software Defined Computational Storage Solution generic structure

The requirements of this architecture are mentioned at Page 11 of the HyperFlow SDCSS Reference Guide. Interconnect between HyperFlow storage elements and any applications/clients depends on various factors - please refer to the HyperFlow SDCSS Reference Guide to understand specific technical requirements for storage interconnect.

Why ScaleFlux Computational Storage Drive

ScaleFlux technology transparently minimises the storage footprint with no performance penalty. Advantages of ScaleFlux CSD technology include:

- Faster read / write / mixed speeds
- Extendable capacity – you can extend a 15TB into being a 60TB⁷ drive. i.e. reducing the cost of NVMe by a factor of four.
- Brings “Compute functions to the data layer”
- Offloads compression / decompression of data⁸ inside NVMe drives without overloading motherboard CPUs

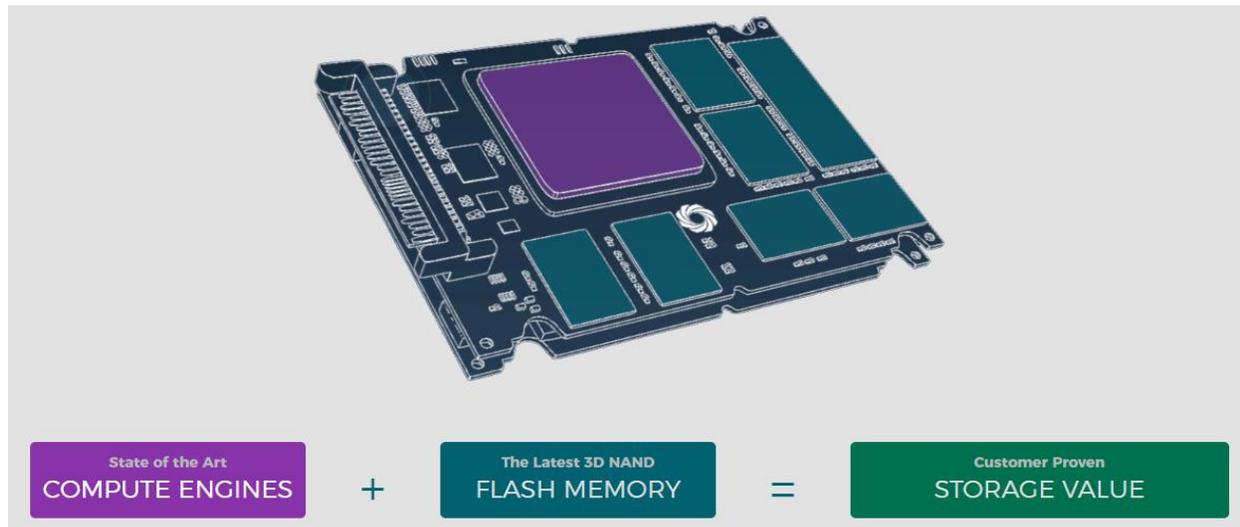


Figure 3 ScaleFlux Computational Storage Drive

Why Ceph

The Ceph platform delivers Object, Block and File System storage in one unified system. Advantages of Ceph technology include:

- Open-source approach and ecosystem
- Block storage (ideal for Host/VM clients) and File-level storage (ideal for web application clients)
- Object storage with Swift and S3 API (Ideal for application development)
- Fault tolerance
- Self-healing and self-managing
- Ceph orchestrator (Cephadm) for easy expansion management of Ceph clusters
- High Availability
- Redundancy via multiple network paths

⁷ Depends on the model of ScaleFlux computational storage drive

⁸ Depends on compressibility of data, Refer to CSS Reference guide for more information.

Why Hyperscalers and ScaleFlux

Hyperscalers [1] is the world's first open supply chain Original Equipment Manufacturer-(OEM), solving Information Technology challenges through standardization of best practices and hyperscale inspired practices and efficiencies. Hyperscalers offers choice across two open hardware architectures:

- Hyperscale - high efficiency open compute equipment as used by macro service providers
- Tier 1 Original – conventional equipment as per established Tier 1 OEM suppliers.

Each of these architectures is complete with network, compute, storage, and converged GP GPU infrastructure elements, and is open / free from vendor lock-in.

Hyperscalers' appliance solutions are packaged complete with hardware, software and pre-built (customisable) configurations. These were all pre-engineered using an in-house IP Appliance Design Process and validated in partnership with associated major software manufacturers. Many can be "test-driven" using Hyperscalers Lab as a Service (LaaS). Hyperscalers appliance solutions are ideally suited to IaaS, PaaS and SaaS providers looking to implement their services from anywhere.

ScaleFlux [2] is the pioneer in deploying Computational Storage technology at scale. Computational Storage is the foundation for modern, data-driven infrastructure that enables responsive performance, affordable scaling, and agile platforms for compute and storage I/O intensive applications. ScaleFlux is a well-funded startup and has leaders with proven experience across deployment of complex computing and solid-state storage solutions at scale.

Computational Storage Drives are integrated into x86/Linux server and storage environments via an easy-to-install ScaleFlux software module. Host-based Flash Translation Layer and Flash Management technologies support consistent latency and performance characteristics. CSD Compute Engines are accessible to applications through APIs exposed by the ScaleFlux software module.

By simultaneously solving compute and storage I/O bottlenecks, CSD technology provides significant and proven run-time improvements to compute and data intensive applications.

ScaleFlux Computational Storage is the ideal foundation for highly scalable, reliable, and low-latency database infrastructure [5].

With data-path compression and decompression that is directly integrated with Flash storage, ScaleFlux delivers the most consistent transactional performance with the smallest Flash storage capacity footprint.

1 REFERENCES

- [1] Hyperscalers, "About HS," [Online]. Available: <https://www.hyperscalers.com/about-us-hyperscalers>.
- [2] Scaleflux, "About-Overview," [Online]. Available: <https://www.scaleflux.com/intro/1>. [Accessed 2022].
- [3] Scaleflux, "CSD 2000," [Online]. Available: <https://www.scaleflux.com/product/item/1002>. [Accessed 2022].
- [4] Hyperscalers, "S5X 2.5" | D53X-1U," [Online]. Available: <https://www.hyperscalers.com/storage/storage-servers/hyperscalers-S5X-D53X-1U-ice-lake-densest-hyperscale-server-nvme-drives-buy>. [Accessed 2022].
- [5] ScaleFlux, "What is Computational Storage," [Online]. Available: <https://www.scaleflux.com/>. [Accessed 2022].
- [6] Ceph, "Ceph Homepage," [Online]. Available: <https://ceph.com/en/>. [Accessed 2022].