PURESTORAGE®
Uncomplicate Data Storage, Forever

TECHNICAL WHITE PAPER

UFFO: Unified Fast File and Object Storage

A technical overview of UFFO.

Contents

Background

Unstructured data consists of files, objects, or both. Traditionally, dedicated storage and compute resources were used to store and access unstructured data. Object storage systems were typically low-performing and designed for archived data, while file storage systems can be complex and challenging to scale. The result is an inefficient use of resources with separate storage silos for each type of data.

Some storage systems provide both file and object storage on the same platform via a front end of either object-protocol-on-file or file-protocol-on-object. In this design, either files or objects are prioritized over the other. Therefore, these gateway solutions come at the cost of lower efficiency and performance for either files or objects.

Today's modern applications need to access, analyze, and restore massive amounts of unstructured data at a high performance level. Although all storage vendors have access to the same commodity hardware components, how the software takes advantage of the hardware is critical to meeting the needs of modern applications. In other words: architecture matters.

A unified fast file and object (UFFO) storage platform has independent file and object stores on the back end and native protocols to access the data in each store. NFS and SMB allow for file access, and S3 provides access to object data.

What Is UFFO?

Modern applications need simplicity, performance, and rich data services at scale. Unified fast file and object (UFFO) is a category of scale-out, high-performance storage, and it addresses the needs of modern unstructured data and applications.

Unified refers to a single physical platform that natively stores both file and object data to consolidate critical workloads, resulting in better utilization of resources and a higher return on investment (ROI.)

Fast means exceptionally high throughput and performance regardless of the size of the data sets, the type of I/O pattern (read, write, sequential, random), or the number of files or objects (up to billions.)

File access is a defining characteristic of UFFO. The platform must natively support NFS for Linux/Unix clients and SMB for Windows clients.

Object access is required for a UFFO platform and provides clients an S3 protocol interface to store, access, and manage object data.

Why Introduce a New Category of Storage?

If you are wondering why the storage industry came up with this new category, the short answer is because modern applications and their unstructured data require it.

Legacy storage doesn't work well for these applications because it is:

- Single-purpose and complex with limited scalability
- Too slow and rigid for modern data and applications
- Fast but lacks enterprise-ready data services and robust protocol support

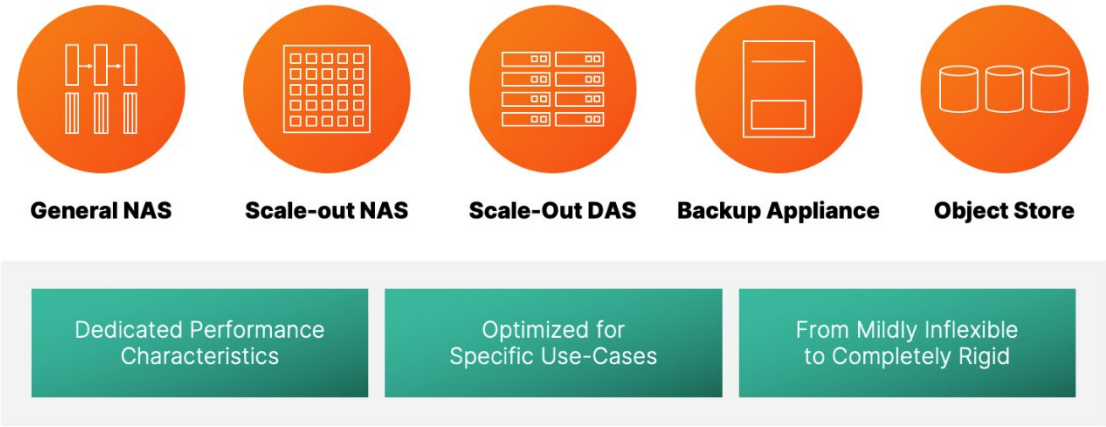

Figure 1. Traditional legacy storage platforms lack the flexibility that modern unstructured data needs.

Traditional legacy storage platforms were built for structured data and older applications such as databases. But the unstructured file and object data from modern applications require higher performance, rich data services like snapshots and replication, and scalability. For example, consider the amount of data generated by sensors, real-time video processing, financial services applications, and predictive analysis.

Architecture Matters

A storage platform can stand out from the others in its performance, scalability, and ease of use. Architecture matters because a solution that meets these criteria not only depends on the software design itself, but it also depends on how the software takes advantage of the hardware. In this sense, hardware does matter, and a better way to say it is that architecture matters because architecture includes both hardware and software.

Requirements for Modern Applications

Modern applications need scalability, performance, and rich data services. Modern data and applications require:

- Multi-dimensional high performance
- Cloud-readiness
- Dynamic scalability
- Intelligent architecture
- High availability
- Accessibility by both file and object protocols

We will discuss how the attributes of a UFFO storage platform meet each of these requirements in greater detail below.

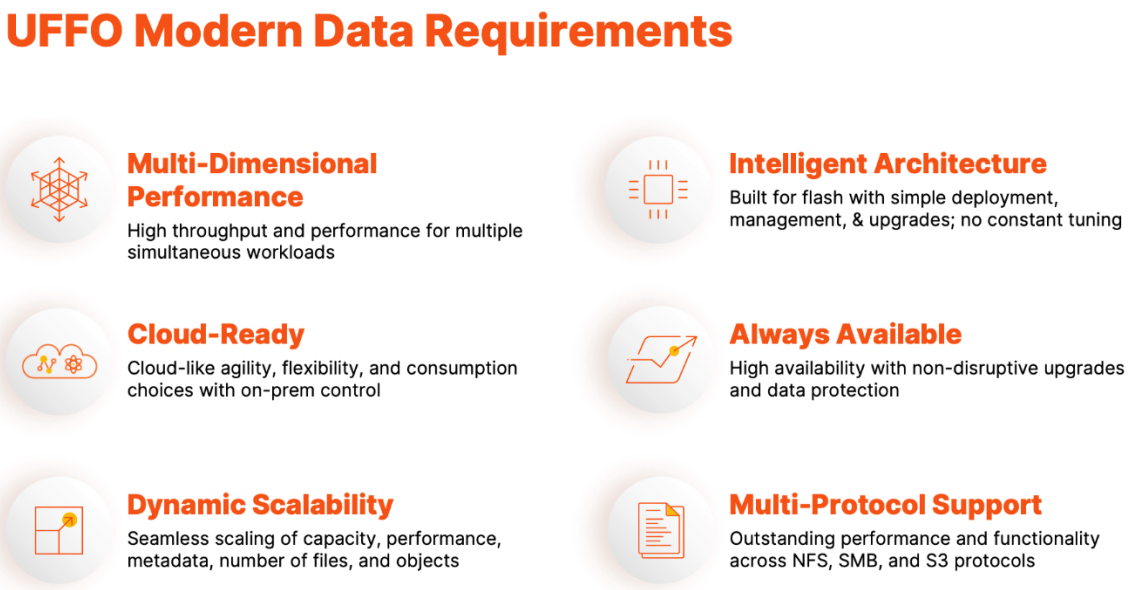

Figure 2. Modern data requirements.

Multi-dimensional Performance

Multi-dimensional performance is the ability to deliver high performance for multiple concurrent file and object workloads, regardless of the characteristics of the workloads. Workload attributes include:

- Data size (large or small files or objects)
- Number of files and objects
- Sequential or random access
- Type of operation (read or write)
- Batch or real-time processes

Traditional architectures can deliver high performance for small or large files and sequential or random file workloads. But modern data workloads require all the attributes listed above at the same time.

A UFFO platform has a core data store and advanced metadata management to natively support file and object protocols and provide multi-dimensional performance. This architecture is fundamental to its ability to take consolidation to the next level.

A UFFO platform delivers scalable, predictable performance with high throughput and can handle tens of billions of files and objects. From small and metadata-heavy to large streaming files, UFFO delivers performance for any access pattern, random or sequential, without the need for manual performance tuning.

Why Fast Object Storage?

Object storage was initially a simple way to store large amounts of archival data at a low cost. But cloud-native applications use object storage for mission-critical data. Modern applications require higher object performance to gain insights from the data, resulting in faster decisions, reduced time to market, and a competitive edge. The public cloud cannot deliver enough

performance for object data, so many organizations need the ability to run fast object storage on-premises or in a hybrid cloud architecture. UFFO architecture should provide higher performance for object data than existing legacy storage platforms.

Cloud-ready

UFFO platforms must be agile and flexible to meet today's performance and capacity requirements and be able to seamlessly adapt as workloads change. There must be different consumption options for the platform, such as subscription services. Modular components and flexible architecture are required so that future expansion and enhancements are simple to implement, while still maintaining predictable performance and high availability.

Cloud strategies often change and evolve. Therefore, a UFFO platform should support the major public cloud providers such as Amazon and Azure, while also supporting a hybrid cloud model. Most importantly, you should have complete control of the platform's operation and security, regardless of its location.

Dynamic Scalability

Scalability is often only associated with capacity, but a UFFO platform must seamlessly scale capacity, performance, metadata, and the number of files and objects. Scaling should be simple and non-disruptive to both availability and performance. When adding more hardware, resources should be immediately available and consumable to the applications.

Storage capacity and compute resources should be modular and scale independently to configure systems that optimize performance, capacity, and cost effectiveness. When requirements change, a dynamically scalable system is simple to reconfigure.

Intelligent Architecture

UFFO architecture must be "intelligent" in its design to take advantage of the performance and efficiencies of flash media. Intelligence also means that the platform is simple to install, manage, and upgrade, and the platform hides its complexity from the end-users.

A UFFO platform should not require constant user intervention for performance tuning and load balancing. Ideally, the architecture should automate network management and storage layout tasks for different workloads, so administrators do not perform these mundane management tasks.

Maintenance—including hardware and software upgrades—and expansions for a UFFO platform must be simple and non-disruptive. As technology continues to change, an intelligent storage platform must future-proof your storage investment.

Always Available

A UFFO platform must provide a high level of availability without compromising performance. Upgrades to the hardware and software must be simple to perform and non-disruptive.

A critical part of availability is data protection, and a UFFO platform must have built-in data protection. "Always available" data means that you have protection from unauthorized access, and you can quickly and easily restore the data after problems like ransomware attacks.

Multi-protocol Support

A UFFO platform must provide fast file and object protocols (NFS, SMB, and S3) for native multi-protocol support of file-based applications and modern cloud-native workloads. Both file and object protocols can run simultaneously or individually, but not impact performance.

A rich set of data services must be available for both file and object protocols, with a full-featured RESTful API stack to enable easy integration and the development of modern applications.

By consolidating object and file storage services onto a single platform, UFFO:

- Accelerates applications beyond the limits of cloud object stores
- Consolidates file and object workloads, eliminating the need for silos of storage
- Unifies object and file management with a single intuitive interface

What Is Modern Unstructured Data?

To understand unstructured data, you first need to understand structured and semi-structured data. Structured data has a well-defined schema for the information it holds, such as a database or a software program like a spreadsheet. Semi-structured data is self-describing, such as XML or JSON, but it does not reside in tabular form as in a database.

Every piece of unstructured or semi-structured data belongs to the class of unstructured data. Unstructured data includes things like text files, images, audio, and videos. By contrast, *modern* unstructured data is created digitally, often by sensors or software applications.

Modern unstructured data has the following characteristics:

- Born digital
- Unpredictable
- Continually generated
- Blended, multimodal, and interoperable
  - *Data blending* is the process of combining data from multiple sources into a functioning dataset.
  - *Multimodal* data comes from a variety of sources (for example, cameras, wearable sensors, infrared imaging).
  - *Interoperable* data can be reused and processed in different applications, allowing multiple information systems to work together.
- Constantly flowing (billions of files and objects generated, processed, and analyzed in real-time at scale).
- Replicated for better access and data protection.

Modern unstructured data needs high performance and throughput for a variety of applications, along with processing of real-time/streaming data. The high volume of unstructured data requires storage that has massive capacity and the ability to handle billions of files or objects.

Use Cases for UFFO

The use cases for UFFO span across most industries, including science and health, financial services, manufacturing, automotive, oil and gas, food service, and state and local government. There are many use cases for UFFO, and the following are just a few examples.

Rapid recovery/Ransomware mitigation. All storage platforms must include some form of data protection. Backing up data is rarely an issue, but rapidly recovering data can be a big problem, especially in the case of a ransomware attack.

A UFFO platform delivers rich enterprise data services such as immutable snapshots and replication to protect data. UFFO is built for performance on flash technologies, making it the best solution for rapid data recovery. Data replication is a required feature to provide a strategy for disaster recovery.

Modern analytics. Virtually all businesses use data analytics to discover trends and answer questions about their line of business. Without these insights, companies can quickly lose their competitive edge.

Several challenges come with data analytics, including slow search and query, silos of unused data, and the complex operations required to process the data. UFFO is the best solution for data analytics because of its high performance, cloud readiness, simplicity at scale, and high availability with non-disruptive upgrades.

Healthcare: PACS and genomics. The picture archiving and communications system (PACS) used in the healthcare industry stores medical images and reports. Imaging technology is always changing, and data sets grow exponentially over time. Medical professionals must be able to quickly and easily access images to analyze them.

Genomic sequencing platforms can generate up to 2PB of raw data per week. Getting actionable insights from the raw data requires rapid read processing from a high number of concurrent users through the sequencing pipelines.

Artificial intelligence (AI). Artificial intelligence (AI) is a branch of computer science concerned with creating self-learning systems that can perform tasks that normally require human intelligence. From self-driving cars to predicting the future, AI is revolutionizing the ways in which we can use data to shape our world.

AI workloads require a high-performance, seamlessly scalable UFFO platform with a centralized hub to store file and object data and share that data with many concurrent users.

Software development/DevOps. Traditionally, development and operations teams were separate entities. Today, software powers business, and combining and empowering these teams into one DevOps organization is the key to unlocking innovation and productivity.

Because of its simplicity and performance at scale, a UFFO platform can accelerate the entire continuous integration, continuous delivery (CI/CD) pipeline, providing a competitive advantage. Developer productivity determines the pace of innovation, which directly affects business growth.

Other technical computing. Technical computing includes applications in categories such as high-performance computing (HPC), computational modeling, and simulations. These applications usually involve enormous data sets, and many users need access to this data simultaneously.

UFFO platforms solve the challenges of simultaneously supporting an extremely high number of clients while delivering high performance and throughput for both file and object protocols.

Pure Storage FlashBlade//S: A UFFO Platform

Multi-dimensional Performance

Flash technology enables high performance, but the architecture of FlashBlade//S is what enhances performance even further. The integrated architecture in Pure Storage® FlashBlade//S efficiently uses flash technology and makes the solution perfect for many use cases.

The software design for FlashBlade//S is an optimized, transactional key-value store with the following three principles:

1. Distribute everything across the platform (metadata control, large file chunks, protocol handling)
2. Optimize for small and large file and object sizes (variable block encoding, optimize for random access, large files distributed across the platform)
3. Provide direct access to flash media (expose concurrency of flash devices throughout the Purity//FB software)

Each blade runs identical software with three types of processes:

- Endpoint
- Authority
- Storage manager

Endpoints manage client connections and there is one instance of an endpoint per blade. Endpoints forward requests to authority processes and relay responses and data between authorities and clients.

There are multiple instances of *authorities* per blade, and they execute the client requests forwarded by the endpoints. Each authority manages partitions of NVRAM and flash on every blade in the platform.

Storage managers perform reads and writes on the NVRAM and flash on its storage units. The storage managers execute requests from on-blade authorities and authorities on other blades.

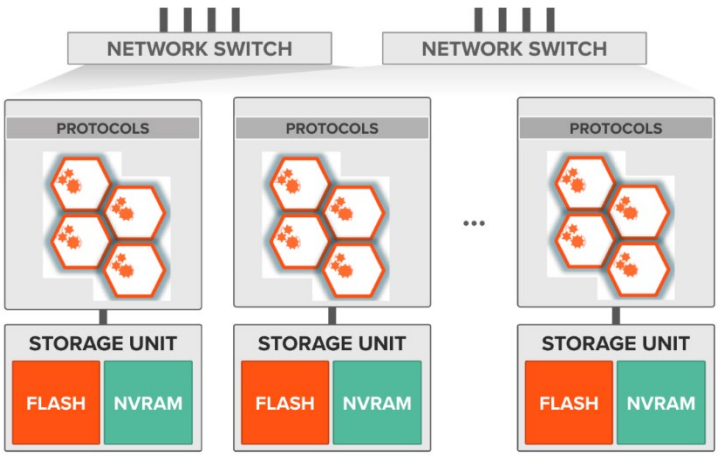

Figure 3. FlashBlade//S design principles.

The design of FlashBlade//S includes a core key-value data store for files and objects and advanced, fine-grained distributed metadata management. It delivers unified fast file and object storage via a highly parallelized architecture for multi-dimensional file and object performance. This architecture is fundamental to its ability to consolidate diverse concurrent workloads. Data is partitioned and distributed across the key-value store and this distribution of data enables parallelizing access and results in higher performance on a larger scale.

A FlashBlade//S system's parallel architecture enables the streaming performance and real-time analytics so that:

- NVRAM stores and (re)organizes I/O to improve small-file and metadata-intensive workload performance and reduce latency. It also eliminates the need to tune for a specific file size or directory depth.
- ECMP (equal-cost multi-path) hashing provides load balancing of client connections to support scale-out performance.
- It supports the consolidation of multiple file and object workloads on a single platform for more significant ROI.

Legacy NAS storage platforms require either human-directed load balancing across NAS controllers or settling for low performance across the board. By employing fine-grained distribution of data and metadata and highly optimized concurrency protocols, FlashBlade//S enables simple, autonomic deployment with performance scalability.

Cloud-ready

FlashBlade//S delivers a cloud-like experience that is agile, flexible, and available through multiple consumption and cloud solution options. The ability to independently add compute resources and capacity allows the system to scale seamlessly while maintaining the control of an on-premises storage platform. FlashBlade//S is a solution that improves over time and because it can accommodate future technology components, there is no need for a complete refresh of the platform in a few years.

Solutions for FlashBlade//S in the cloud include:

- Evergreen//One
- AWS Outposts
- Pure Storage FlashBlade® integration with Azure Stack
- FlashBlade in Equinix with Microsoft Azure
- Pure Storage on Equinix Metal
- FlashBlade Object Replication

Here are the details for each of these solutions.

Evergreen//One is an on-demand subscription model for storage consumption that includes maintenance and support. It is a pay-as-you-go, hybrid cloud, highly efficient storage solution. Evergreen//One provides file and object storage on FlashBlade as a single unified subscription in a physical data center or co-location facility.

AWS Outposts is a fully managed service that extends AWS infrastructure, AWS services, APIs, and tools to virtually any data center, co-location space, or on-premises facility for a truly consistent hybrid experience. As a designated Service Ready Partner, Pure Storage FlashBlade and FlashArray™ have been thoroughly tested and supported with Outposts to deliver simplicity, performance, and consolidation.

Use cases include:

- Modern analytics and AI/ML
- Rapid restore
- Data sovereignty, compliance, or security needs
- Low latency requirements
- Local data processing requirements

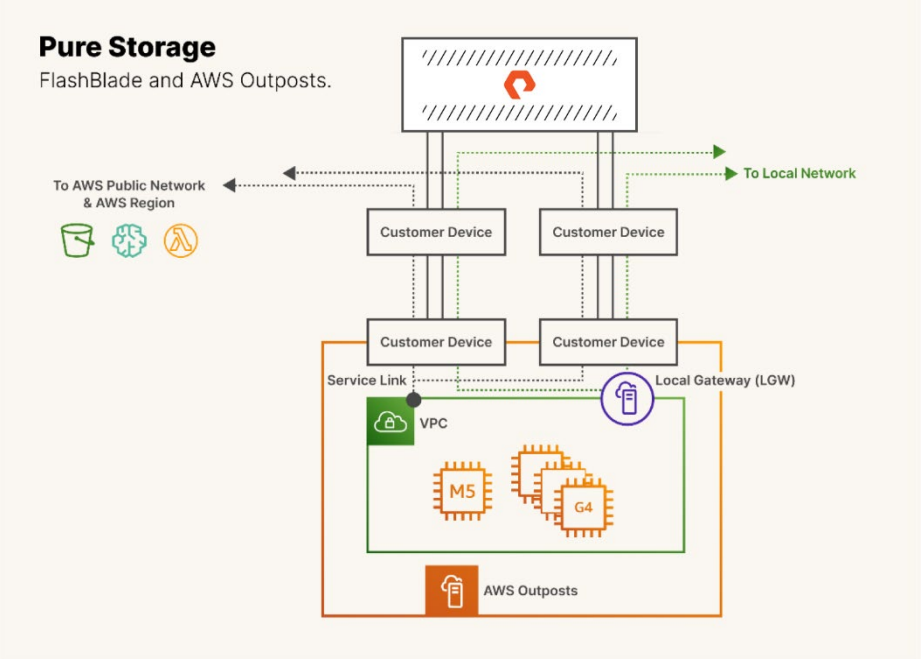

Figure 4. Pure Storage FlashBlade//S and AWS Outposts.

FlashBlade integration with Azure Stack provides NFS, SMB, and S3 object storage. You can deploy and manage FlashBlade file systems and S3 object accounts directly from the Azure Stack Hub portal and Azure Resource Manager.

With third-party, Microsoft-approved resources, you can directly integrate one or many FlashBlade namespaces into the Azure Stack Hub portal. This integration enables native provisioning and management of FlashBlade file systems and object accounts directly from the Azure Stack Hub portal, CLI, and ARM interface. It extends all the benefits of the FlashBlade high-performance file and object platform to the Azure Stack Hub.

The integration of Pure FlashBlade into Azure Stack Hub delivers a fast file and S3 object storage solution for your cloud strategy.

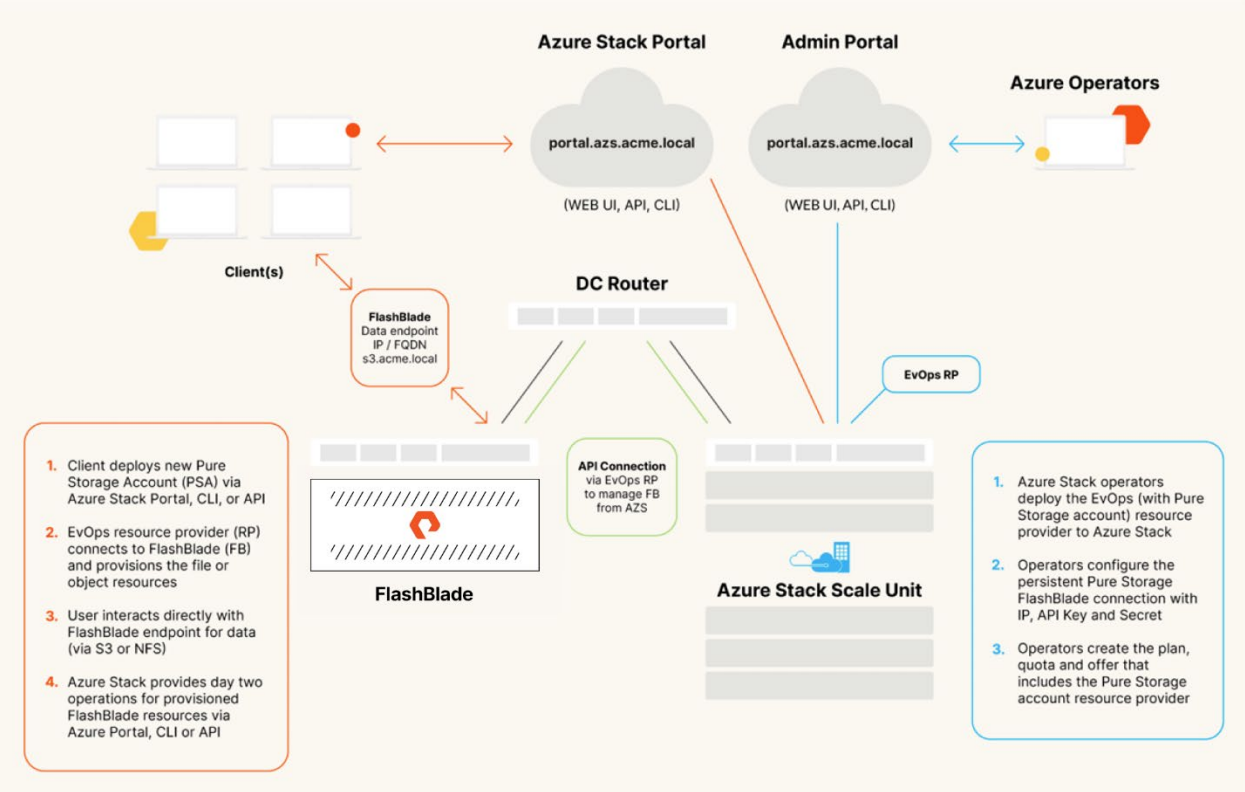

Figure 5. Pure FlashBlade//S and Azure Stack Hub integration.

FlashBlade in Equinix with Microsoft Azure is a cloud-adjacent solution with Pure, Microsoft, and Equinix. Microsoft fully validates this solution for electronic design automation (EDA) simulation workloads. The FlashBlade platform is in an Equinix data center and connects to the Microsoft Azure cloud. The solution delivers predictable high performance and low latency, while giving you complete control over your data. Although this is a Microsoft solution today, you will be able to connect to any primary cloud provider in the future. The main benefit of the connected cloud solution is that you don't have to move your data and can keep security and control.

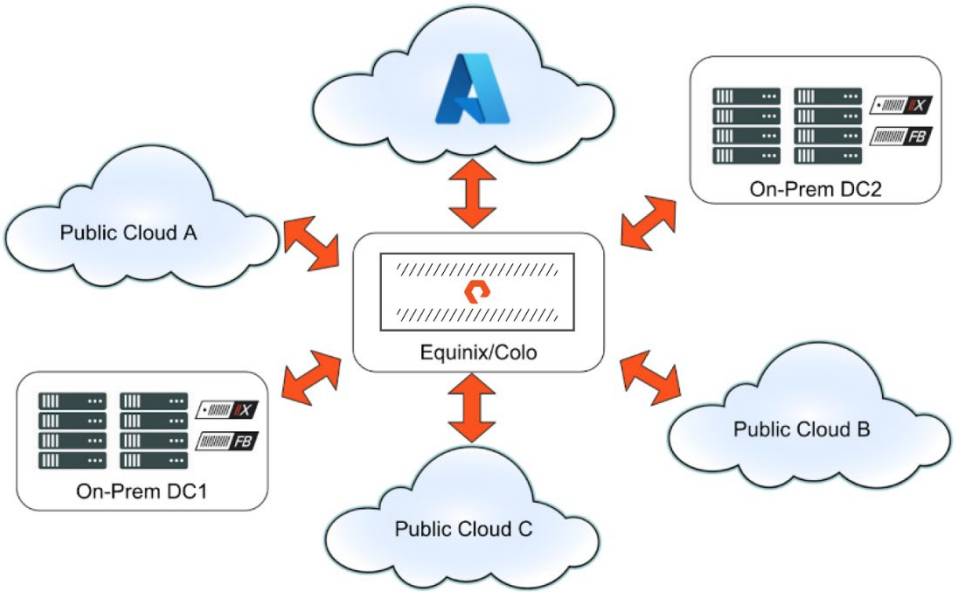

Figure 6. FlashBlade//S in Equinix with Microsoft Azure.

Pure Storage on Equinix Metal offers high-performance, near-edge cloud storage. Combining the FlashBlade unified fast file and object platform with Equinix Metal can consolidate diverse unstructured workloads onto a highly scalable platform with multi-dimensional performance. This consumption-based cloud environment connects to every significant hyperscale provider and reduces the total cost of ownership.

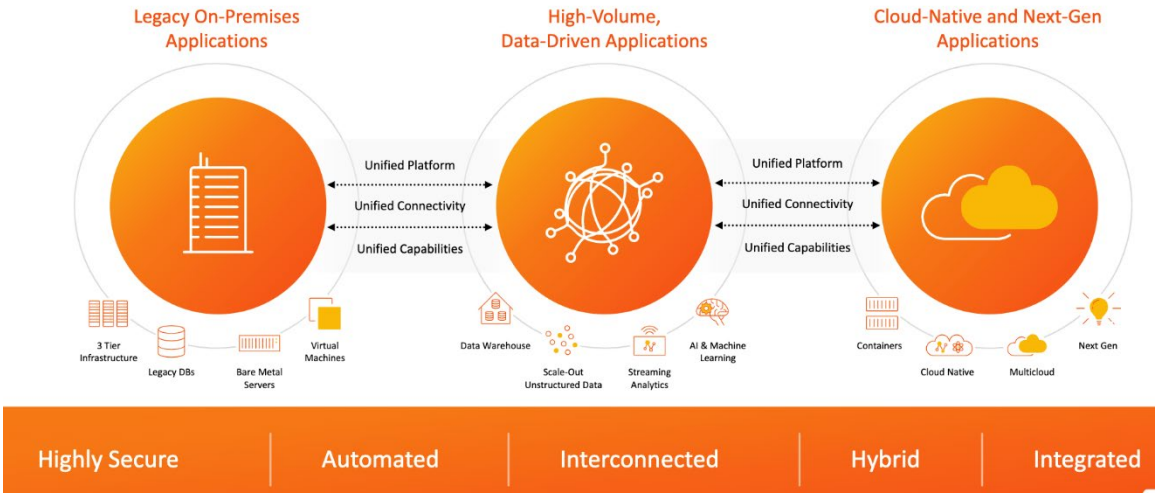

Figure 7. Pure Storage on Equinix Metal offers high-performance, near-edge cloud storage.

FlashBlade object replication delivers quality, simplicity, and efficiency. Asynchronously replicate object data in native format from FlashBlade to FlashBlade and from FlashBlade to Amazon S3.

What makes object replication so compelling?

- **Simplicity:** Easy to set up and manage using the GUI, REST, and CLI interfaces
- **Enhanced performance:** Delivers increased read throughput and lower latency to geographically distributed users.
- **Agility/cloud mobility:** Get the benefits of cloud economics with native S3 replication to the cloud.

- **Secure data-in-transit:** Encrypts data in flight with Transport Layer Security Protocol.
- **No gateways or licenses:** Object replication is part of Purity//FB software.
- **Enterprise monitoring:** Monitor object replication in one central location with Pure1®, the AI-driven cloud-based management platform.

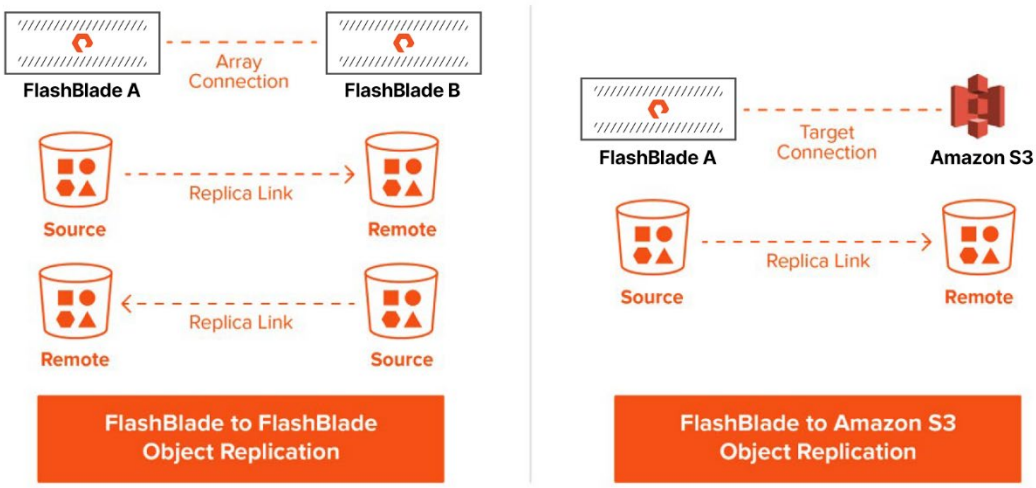

Figure 8. FlashBlade object replication.

Dynamic Scalability

Telemetry from the Internet of Things (IoT), medical imaging, cybersecurity applications, and Application Performance Monitoring (APM) and logging drive real-time/streaming data high-volume small-file operations. Legacy NAS storage platforms are unable to deliver the performance required for this type of data. The variable block size storage engine at the heart of Purity//FB software in FlashBlade//S optimizes the layout of billions of small files.

The architecture of FlashBlade//S architecture enables greater scalability without performance trade-offs. Capacity and performance scale independently. Adding DirectFlash® Modules (DFMs) to the blades expands capacity, while adding more blades increases the number of processors, NICs, and DRAM to scale for higher performance.

FlashBlade//S automatically balances the workload across all blades and multiple blade chassis. Data handling and forwarding via stateless connections mean that each blade can service any client connection. With a remote procedure call (RPC) cache across all blades, each blade can restart any client connection.

Because of these features in the architecture, FlashBlade//S dynamically load balances without manual tuning by administrators and non-disruptively scales out. FlashBlade//S performs dynamic multi-dimensional load balancing at every step:

- Connections between blades
- Front-end connections from clients across all blades
- Where to place the data

FlashBlade//S and Evergreen™

FlashBlade//S hardware modularity simplifies capacity increases and non-disruptive hardware upgrades. This has made it possible for Pure to offer Evergreen//Forever service for the new systems. Evergreen//Forever has several advantages, including the ForeverFlash lifetime media guarantee, Ever Agile upgrades with trade-in credits, "flat and fair" service pricing guarantees, and perhaps most importantly to users, it includes periodic hardware refresh at no incremental cost.

In all facets of IT, new generations of hardware are typically introduced every 3-5 years. For storage, this has historically meant that every 3-5 years users would effectively repurchase capacity they already owned and be forced to migrate hundreds of terabytes of data from old to new hardware. Evergreen//Forever plus the non-disruptive upgrade capabilities of FlashBlade//S enable systems to track hardware evolution. Repurchasing capacity and data migration are never issues.

Intelligent Architecture

Intelligent architecture is simple to deploy and use, scalable, highly available, high performing, and self-healing. The intelligent storage platform hides complexity from users and administrators. Typical features in intelligent storage platforms include predictive analytics, distributed storage and processing, optimal data placement, and high performance.

With other flash storage solutions available today, the components, such as the flash drives themselves, define the storage platform's behavior, performance, and reliability. This type of architecture requires replacing the entire system as component technology changes, usually on a 3–5-year cycle.

The software and modular hardware design of FlashBlade delivers an architecture that works as an entire *system* maximizing the performance and reliability of today's components. It will adapt and scale as components like flash drives and CPUs evolve.

DirectFlash. Pure designed the architecture of FlashBlade//S to take advantage of flash media, resulting in our DirectFlash® technology. Traditional solid-state disk (SSD) technology has many challenges. A traditional flash storage platform cannot overwrite data on SSD; the system must erase data and then write new data. Traditional SSDs perform like disks, resulting in a performance impact, more wear on the SSDs, and less reliability.

DirectFlash solves these problems by directly using the flash modules at the platform level, not at the SSD level. FlashBlade//S has only one tier of wear-leveling, managed across the entire storage platform. DirectFlash gives higher visibility into the flash and more control than SSDs, resulting in performance, reliability, and global flash management advantages.

Another significant advantage of DirectFlash is that it effectively protects against future changes in flash designs. As SSDs have evolved from storing one, two, three, or four bits per cell (SLC, MLC, TLC, or QLC) and beyond, the wear-leveling and performance characteristics have worsened. The level of control over the flash that DirectFlash provides helps mitigate these challenges.

DirectFlash has a file system and object layer above a key-value store that allows for more degrees of freedom when scaling the platform. Direct visibility into the flash allows for more intelligent decisions about performing the mechanics of garbage collection at a larger scale.

Self-healing and autonomic. The architecture of FlashBlade//S is autonomic and self-healing, which means that the platform runs and balances itself automatically without human intervention. Redundant components ensure high availability and erasure coding, tuning, and configuration choices are always on and auto configurable.

Network fabric. The integrated networking in FlashBlade//S simplifies large-scale deployments by collapsing three networks into one high-performance software-defined networking (SDN) fabric. This SDN is shared across the two fabric modules in the platform, and it hides the complexity of networking from the administrator.

What about network limitations? A legacy storage platform can quickly become very complex at scale. Traditional platforms have three separate networks:

- **Front-end network**: Handles traffic from clients
- **Control network**: Handles traffic from monitoring and managing the system
- **Back-end network**: Handles the back-end messaging between blades

FlashBlade//S virtualizes the network so that no matter the size of the platform, it appears as one entity. Load balancing and cabling are much simpler than in legacy systems. Each blade can service and restart any client connection and run any protocol, and the platform is stateless because the logic can run anywhere.

Dual Fabric I/O Modules (FIOMs) interconnect blades, connect chassis (in multi-chassis systems), and connect blades to clients. The FIOMs have ethernet switches with eight (8) external ports each capable of 10, 25, 40, or 100 Gbps transmission rates. The switches have a total of 2 Tbps cross-sectional bandwidth. Each FIOM uses 50 Gbps for inter-blade communication in the chassis. Both FIOM switches and blade NICs are capable of 100GB/s for future expansion.

Always Available

Data availability along with traditional platform resiliency are the primary requirements of the design of FlashBlade//S. There are many levels of resiliency in FlashBlade//S. Data is efficiently written in wide stripes across the flash storage.

N+2 erasure coding protects data at rest on flash, and triple replication protects data in flight. There is also a fine-grained resilience in the storage platform that contains faults and proactively repairs at a lower level.

The modular design of FlashBlade//S increases resiliency by separating compute resources from the flash storage. This separation allows for separate failure domains, which means that a DFM can be hot swapped without affecting its blade's ability to function. If a CPU or NIC fails on a blade, there is no need to rebuild the data.

From an availability point of view, the system can tolerate the loss of any blade plus any other DFM (N+1 availability). From a resiliency point of view, the system can tolerate the loss of any two DFMs. Any two power supplies can fail, (2+2 redundancy), assuming independent feeds to the power supplies.

Network modules are highly available active/active pairs and failover is non-disruptive. Metadata is indexed on all flash drives, allowing the platform to rapidly identify the data's location. The FlashBlade//S platform is all flash, enabling rapid rebuilds in case of a failover.

FlashBlade//S delivers high availability over multiple years and upgrades, and it allows for non-disruptive hardware and software upgrades. Software upgrades are non-disruptive, efficient, and happen over time as a background process. Upgrades are a two-step process of loading all the software in the background, and then restarting each blade with the new software. Hardware upgrades are non-disruptive with the simple addition of blades, DFMs, and chassis. Blades and DFMs are field-replaceable units (FRUs) and no data rebuild is required. Removing hardware is also a non-disruptive process built into the architectural design of FlashBlade//S.

Multi-protocol Support

FlashBlade//S supports high-performance file access through the SMB and NFS protocols, including cross-protocol interoperability for Windows and Linux/Unix clients. Access to fast object data is via the S3 protocol.

Protocols can run individually or simultaneously, and data services are available for each protocol, such as replication and immutable snapshots for data protection. The SafeMode™ feature for files and objects provides an extra layer of data protection. Enabling SafeMode disables the ability for an administrator to eradicate snapshots manually, prevents the ability to rollback or restore a snapshot, and introduces an adjustable eradication timer that can be set from 24 hours up to 400 days. With SafeMode enabled, a hidden snapshot schedule creates the snapshots on the platform.

Other multi-protocol features include always-on inline deep data compression for capacity savings, and the RESTful API stack to accelerate the process of developing and deploying modern applications.

Conclusion

The need for a new storage category arose from the requirements of modern applications and their unstructured data. Unified fast file and object (UFFO) storage platforms satisfy these requirements and solve these challenges.

FlashBlade//S is the best-in-class UFFO storage platform for modern data applications. Pure built FlashBlade//S from the ground up to deliver on all the requirements of UFFO, so it provides high performance for both file and object unstructured data, along with unmatched simplicity and scalability.

> "This scale-out platform leverages flash to provide performance and efficiency so customers can consolidate file and object workloads with confidence."
>
> **KERRY DOLAN, SENIOR IT VALIDATION ANALYST, ESG**

There are challenges with modern unstructured data. Advanced workloads are growing, but the public cloud is not always the correct answer and legacy storage platforms are not able to manage all the requirements of massive amounts of data.

The FlashBlade//S architecture addresses these challenges based on three pillars in its design, while keeping the platform simple to use and deploy.

- **Unlimited scalability:** Deliver a modular, flexible platform that seamlessly scales in multiple dimensions to meet the needs of any size problem or business.
- **Enterprise readiness:** Supply the stability, reliability, and performance that customers expect from their storage.
- **Rich data services:** Provide enterprise features like protocols, snapshots, replication, and more so that customers can use their data without sacrificing performance and scale.

To summarize, by design FlashBlade//S is a highly optimized, highly partitioned, highly distributed, autonomically balanced, highly resilient, scale-out, key-value store for both file and object data and metadata.

FlashBlade//S is a proven, best-in-class UFFO solution that is simple to deploy and manage.

Additional Resources

For more information on the FlashBlade//S solution and UFFO, please see:

- ESG Technical Validation: Technical brief on Pure Storage FlashBlade//S Unified Fast File and Object (UFFO) Platform.
- Learn more about FlashBlade//S.
- Gartner on FlashBlade//S: Pure Storage is a leader in the 2021 Gartner Magic Quadrant for Distributed File Systems and Object Storage.
- Discover why FlashBlade//S with Evergreen//Forever is the last UFFO platform you will ever need.

Pure Storage, Inc.
650 Castro Street, #400
Mountain View, CA 94041

PS2181-02 05/2022