

Innovations in All-Flash Storage Deliver a

# New Approach to Unstructured Data

## Table of Contents

Developing a new approach to unstructured data.....2

Designing a new storage architecture.....2

Understanding the FlashBlade architecture.....3

Benefits of this new model .....4

Conclusion .....5

There have always been certain workloads and applications that stretch the boundaries of storage technologies. While this is nothing new, the number of these workloads is poised to expand dramatically, fueled by the growth of initiatives such as big data analytics, mobility, cloud computing, social media and the Internet of Things (IoT).

The common challenge among these workloads is the requirement for huge data sets to be processed quickly, efficiently and—most important—concurrently. The storage infrastructure must not only deliver extreme performance in terms of bandwidth, IOPS, low latency and high availability, but do so while dealing with potentially huge multi-petabyte data sets.

Examples of these workloads can be found in areas such as genomics, semiconductor design, data analytics, media and entertainment, oil and gas exploration, and computer-aided engineering (CAE), among others. But such applications may be just the tip of the iceberg.

Looking ahead, as more data is being created and data sets are getting larger across *all* applications, IT teams can expect to see many more enterprise business workloads requiring the same combination of all-flash performance, high availability, simple manageability, density and cost-efficiencies as these highly targeted technical workloads. Examples include DevOps, business analytics and SaaS back ends, to name just a few.

Until now, the storage industry has struggled to meet the challenges of these workloads. Traditional disk-based infrastructure and batch-style analytics fall short of delivering the necessary performance. Block-based all-flash arrays also fail to meet performance requirements, primarily because of namespace limitations and lack of protocol support.

There have been stopgaps that have had some success, but only to a certain point. These include scale-up solutions using larger and more powerful controllers, and parallel file systems and hybrid architectures that deploy spinning disk drives with solid-state drives (SSDs). However, each of these solutions comes up short in at least one of the following ways:

- **Expensive to deploy**
- **Highly complex to manage and scale**
- **Unreliable**
- **Unable to deliver the performance required to optimize the targeted workload**

Users in these challenging environments—engineers, designers, animators and others—have long expressed a need for a storage solution that not only can deliver the performance, capacity, availability and density they require in an architecture, but is much simpler to manage, more reliable and more cost efficient to deploy. They know they can be more innovative, creative and productive—and make faster, better and more accurate decisions—if they can leverage real-time analytics, richer queries and interactive simulations.

## Developing a new approach to unstructured data

Such challenging workloads have created an opportunity for an innovative storage vendor to design a new architecture that can meet the concurrent performance and manageability requirements of unstructured data in file- and object-based environments.

In a recent research report, IDC discussed the potential for a new class of all-flash arrays to address these workloads.

It said these arrays can be built around scale-out architectures using custom flash modules instead of SSDs.<sup>1</sup> This new class of arrays, IDC said, should deliver higher storage densities, lower total cost of ownership (TCO) and non-disruptive scalability up into the tens of petabyte range and down again when necessary.

As one of the leading innovators and most successful vendors in the block-based all-flash array environment, Pure Storage was early to recognize the need to address the challenges facing users of file- and object-based environments—and is now one of the first vendors to deliver a viable, cost-efficient solution.

The new Pure Storage solution is called FlashBlade. It is the result of a three-year effort to design an innovative architecture for file- and object-based storage environments. With FlashBlade, Pure Storage is creating a new paradigm in combining performance, simplified manageability, high density and cost efficiencies for the most demanding workloads.

## Designing a new storage architecture

Delivering dramatic improvements to customers required an architectural approach focused on not only IOPS performance, but also metadata performance, massive scale, resiliency and economics. The goal of the FlashBlade design team was to build an architecture that delivers extreme bandwidth, IOPs, resiliency and capacity as well as low latency in a cost-efficient package that is also compact, energy efficient and easy to scale.

To meet these challenges, the FlashBlade design is centered on four core principles:

- **Elastic scale-out:** FlashBlade is built to scale effortlessly and linearly from small deployments to very large deployments. Every dimension of the system can scale linearly with the system as it grows—IOPS performance, bandwidth, metadata performance, NVRAM and client connections.

---

<sup>1</sup> "FlashBlade: Putting File- and Object-Storage Vendors En Garde," IDC, August 2016



- **High performance, efficient TCO:** The design team had to optimize the system for two seemingly opposing goals: game-changing performance and a cost profile that would make all-flash cost efficient for even very big workloads. This led the team down the path of designing unique custom hardware and a minimalist approach.
- **Natively multi-protocol:** Important changes are taking place in the unstructured and semi-structured data storage market. Legacy file access protocols are giving way to newer object protocols and even newer application-specific protocols. FlashBlade is designed with a core object store at the center and the ability to easily add scale-out protocols on top of that. Current versions of the system start with Network File System (NFS) and Amazon S3 object access. Other protocols will be added based on customer needs.
- **Management simplicity:** Simplicity has always been a predominant design consideration for Pure Storage arrays. FlashBlade is no exception. Customers in

large-scale file deployments are typically drowning in volumes, cluster pairs, aggregates and flash caches. The goal with FlashBlade was a design so simple that anyone could manage it at any scale—not just IT personnel and storage administrators, but also developers, data scientists, engineers or anyone else.

## Understanding the FlashBlade architecture

FlashBlade represents a new architectural design for all-flash arrays. It is capable of delivering all-flash performance, simplified manageability and high density in a compact package that can lower TCO for performance-driven workloads. To achieve these capabilities, the architectural model is comprised of three critical components, which are:

- **The blade:** The blade is the scaling unit for FlashBlade. Each blade marries raw NAND flash with an Intel Xeon system-on-a-chip processor, a programmable

processor with integrated ARM cores, DRAM and integrated NVRAM, all connected to the blade via the PCIe protocol. The Flash Translation layer inherent in standard SSDs has been re-architected to eliminate bottlenecks, and the DRAM in SSDs has been re-provisioned to significantly improve parallelism across the system.

- **Elasticity software:** The Elasticity software spans from file system to flash, implementing layers of functionality that require separate code in other systems. Elasticity includes scale-out file and object system software, a core clustered storage system with advanced data and resiliency features, and the software that typically would be found inside an SSD (such as flash management and error correction coding), optimized globally. The software leverages the customized hardware to distribute and provide parallel direct connections to client connections at the network tier. In addition, data access is distributed via the storage blades, which can service any client connection. Finally, control logic parallelism creates efficient access paths to metadata and query updates, removing contention and bottlenecks.
- **Elastic Fabric:** FlashBlade is built upon an embedded 40 Gbps software-defined switch fabric, called the Elastic Fabric. This is a low-latency switching fabric that connects blades, chassis and tens of thousands of clients together on one converged fabric. Each blade contains two 10 Gbps links to the switch fabric, and there are eight 40 Gbps links in the fabric for available bandwidth of 320 Gbps for uplink to client networks. FlashBlade leverages the performance and ubiquity of Ethernet, implemented in a set of redundant fabric modules that slide into the back of a 4U chassis. The Elastic Fabric runs TCP/IP connections to hosts but uses proprietary protocols for internal communication to deliver low-latency communication between blades. This provides integrated networking for super-fast client and data access to the storage unit.

## Benefits of this new model

Early users of FlashBlade are already seeing significant benefits from the performance, density and simplified

manageability of the solution. For example, a semiconductor design company shortened chip simulation operations by more than 20%. Given the typical five-year R&D cycle for chips, this translates to bringing product to market a full year early, presenting a huge opportunity to boost revenues. In another example, an automotive company achieved a fourfold improvement in simulation speed and a threefold increase in concurrent simulations.

These companies and others are achieving such benefits by leveraging features and functions of FlashBlade that are simply not available in other storage solutions. This unique combination of features and functions is manifested in the following ways:

- **Capacity:** Each FlashBlade chassis can support up to 15 blades, and each blade can be configured with either 8.8 TB or 52.8 TB of raw flash storage, supporting tens of billions of files and objects in a remarkably small all-flash footprint.
- **Performance:** FlashBlade supports up to 15 GBps of bandwidth per chassis and 500K NFS operations per second. The solution delivers the consistent performance of all-flash storage with no caching or tiering, as well as fast metadata operations and instant metadata queries.
- **Simplicity and scalability:** FlashBlade has been designed so that anyone can install it. Scale-out is simple, instant and online; to add capacity, you just have to add blades. FlashBlade leverages the Pure1 cloud-based management platform, using high levels of automation to simplify initial deployments and ongoing management.
- **Cost efficiencies:** The FlashBlade system takes up much less physical space than legacy spinning disk drive systems. It also uses much less energy, requires less management time and training, and accelerates time to value. In addition, customers can leverage Pure Storage's Evergreen Storage model to reduce costs and shift to an Opex cost model as opposed to a Capex one. Evergreen Storage also extends the lifecycle of the solution and ensures that customers can stay current with the latest technology upgrades.

The advantages of FlashBlade technology will be obvious for users in the initially targeted workloads of genomics, semiconductor design, data analytics, media and entertainment, oil and gas exploration, and CAE. However, as enterprises strive to leverage big data analytics, social media and the Internet of Things, it is highly likely that other areas of the business will look to FlashBlade technology to solve specific business challenges, particularly as Pure Storage continues to develop the technology and add incremental enterprise-grade features.

## Conclusion

The tremendous growth of unstructured data is creating huge opportunities for organizations. But it is also creating significant challenges for the storage infrastructure.

Many application environments that have the potential to maximize unstructured data have been restricted by the limitations of legacy storage systems. For the past several years—at least—users have expressed a need for storage solutions that can deliver extreme performance along with simple manageability, density, high availability and cost efficiency.

FlashBlade from Pure Storage is the only solution of its kind available today that has been designed to address the storage challenges of these demanding workloads. Based on a new architecture from one of the acknowledged leaders in all-flash arrays, FlashBlade sets new standards for performance, density, resiliency and simple manageability in a package that is extremely efficient in terms of TCO, space utilization and energy consumption.

With FlashBlade, users in industries such as genomics, semiconductor design, analytics, media and entertainment, oil and gas exploration and CAE can achieve design goals that they were never able to achieve before. Beyond that, FlashBlade is a technology that sets the stage for future innovation in areas such as big data analytics and cloud-native applications.

---

**If you would like to learn more about  
how FlashBlade can help your organization,  
please visit [purestorage.com/flashblade](https://purestorage.com/flashblade).**

---