

# FLASHBLADE

## The Cloud-Scale Data Platform of Tomorrow

Built for innovators, FlashBlade™ powers tomorrow's discoveries, insights, and creations — **one blade at a time.**

### Analytics of Everything

FlashBlade helps deliver real-time all-flash analytics to help bring transactional capabilities to your big data repository.

### Cloud-Native Applications

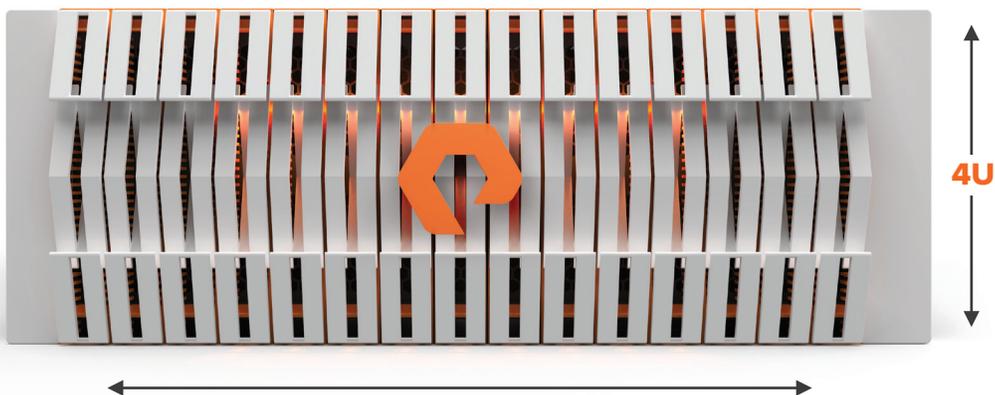
Build your apps to scale from the start, with seamless access to PBs of objects at consistent all-flash latency.

### Digital Science, Engineering & Design

Simulate and test increasingly complex models in shorter time, leveraging 1000s of connected compute nodes.

### 4K/8K Media Workflows

Bring startling new effects to life, and transcode your creations in real-time for every platform.



**15 ALL-FLASH SCALE-OUT BLADES**

1.6PBs usable in 4U

10s of BILLIONS of FILES, OBJECTS, and CONTAINERS

## READY FOR THE BIG + FAST DATA OF TOMORROW

The past decade saw the rise of “big data” – and storage solutions designed to allow Petabytes of data to be centralized and served for everything from data analytics to scientific discovery to movie rendering. These storage solutions were big, but often slow, as they took advantage of bigger and bigger hard disk drives, the last mechanical relics in today's otherwise all-silicon data center. But with the rise of cloud-native applications, needs and expectations are changing. Big but slow isn't good enough anymore, and the dramatic inefficiency of rows and rows of spinning disk in the data center just can't keep up with the efficiency demands of today. What if an all-flash storage platform could deliver big and fast for tomorrow's data, and be simple and efficient enough to be the platform for your cloud-native applications?

### BIG

10s of Petabytes

10s of Billions of files and objects

10s of Thousands of clients connected

All in an unbelievably-small all-flash footprint

### FAST

Up to 15 GB/s of bandwidth per chassis, and 1M NFS ops/sec

Fast metadata operations and instant metadata queries

Consistent all-flash performance with no caching or tiering

### SIMPLE & EFFICIENT

Scale-out everything online and instantly by simply adding blades

NFS, Object/S3, and more to come

Centralized cloud management and support via Pure1™

\*All features and specifications are preliminary until GA. Usable capacity assumes 3:1 data reduction rate.

# FLASHBLADE COMPONENTS



## BLADE

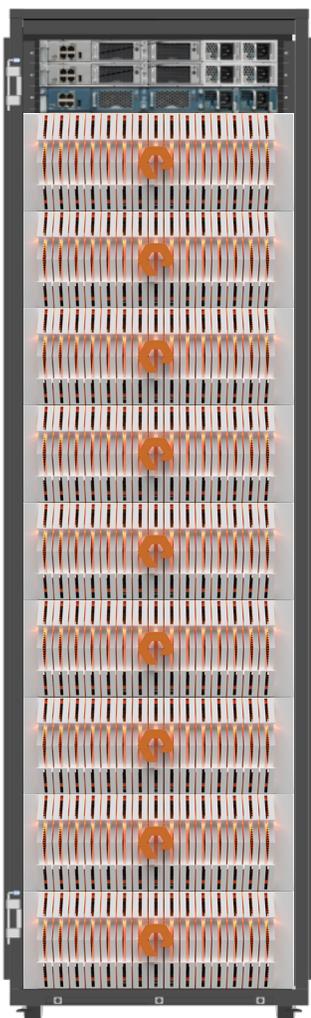
### SCALE-OUT COMPUTE + FLASH

- The Blade is the scaling unit of FlashBlade – adding Blades linearly scales capacity, performance, metadata, and clients
- Each blade contains a low-power Intel XEON System-on-a-Chip processor and FPGA ARM cores to run the Elasticity™ Software
- Blades contain raw flash chips, DRAM, as well as integrated NV-RAM protection to guard write operations from sudden power loss
- Blades are available in 8 and 52TB capacities, and can be both mixed-and-matched and hot-added into the system.
- Blades were designed with ultra-low-latency in mind, leveraging PCIe to connect flash chips to CPUs, and the low-latency Elastic Fabric to communicate with other blades in the system

## ELASTICITY

### SCALE-OUT STORAGE SOFTWARE

- Elasticity software is the heart of FlashBlade, implementing its scale-out storage core, scale-out metadata, scale-out multi-protocol support, and integrated software-defined networking.
- The Elastic Core is the base object store upon which all protocols are built, and implements native data services (data reduction, snapshots, replication, and encryption) as well as flash management services (N+2 erasure coding and LDPC error recovery).
- The Elastic Map is the scale-out metadata service that underlies all layers of Elasticity. It implements an extensible, variable-block metadata engine which allows for instant query of complex operations.
- FlashBlade and Elasticity are managed by Pure1 – Pure Storage's cloud-based management and support platform – making it easy for a single administrator to manage multiple systems of any size globally.



16PB of Usable  
FlashBlade in a Rack

## ELASTIC FABRIC

### LOW-LATENCY, SOFTWARE-DEFINED NETWORKING

- FlashBlade includes a built-in 40Gb Ethernet fabric, used for both communication between Blades and to clients.
- The Elastic Fabric implements low-latency proprietary protocols bypassing the TCP/IP stack, separating client, data, and metadata traffic in the system via QoS.
- Each chassis includes 8x 40 Gb/s Ethernet ports for both client connectivity and direct chassis interconnect, as well as 30x 10 Gb/s internal ports for Blade connectivity via the passive mid-plane.



## FLASHBLADE DIFFERENTIATORS



### ELASTIC SCALE-OUT

Scale-out everything online and instantly by simply adding blades.

Each blade adds capacity, performance, metadata scale, NV-RAM, and client connections.



### CONSISTENT ALL-FLASH LATENCY

100% flash architecture ensures that no operation is ever waiting for disk. FlashBlade accelerates large-block and small-block IO, and both data and metadata are persisted on flash.



### METADATA PERFORMANCE

75% of NFS operations are metadata-only operations, like getAttrs. FlashBlade implements a scale-out Metadata Map – used by all layers of the system, and independently queryable.



### N+2 ERASURE CODING RESILIENCY

Blades are deployed in N+2 redundant fashion to protect against flash or full Blade loss. When Blades fail, the system automatically heals-around them, returning FlashBlade to full resiliency.



### 100% FLASH 0% SSD

FlashBlade uses 100% raw MLC NAND, and all functions of flash management and error recovery are handled by Elasticity in software. This ensures the highest performance and lowest overhead.



### NO STACKED PROTOCOLS

Other arrays often start with a single protocol, then stack others on top – with each inheriting the limitations of the layers. Elasticity implements a common Object Core – and then all protocols are peers on top.

## FLASHBLADE SIMPLICITY

### SIMPLE TO DEPLOY

FlashBlade is deployed as a single 4U appliance, with a minimum of 7 blades. All networking is internal, and all software is included. No tuning or configuration is necessary.

### SIMPLE TO SCALE

FlashBlade is scaled by inserting new blades into the system. No pre-configuration or planning require, and different capacity blades can be mixed-and-matched to ensure future expandability.

### SIMPLE TO MANAGE

FlashBlade includes a built-in GUI and CLI, as well as Pure1, Pure Storage's cloud-based management and support platform. A single administrator can manage a FlashBlade deployment of any size.

