10 ways PowerScale surpasses Pure FlashBlade for modern data

By Andrew Glinka, Vice President Competitive Intelligence | May 2022

The modern vs. legacy architecture argument that Pure Storage uses in promotional blogs and write-ups is a smokescreen that obscures a more straightforward, time-proven perspective on how to handle the advancements shaping current and evolving unstructured data storage requirements. You do not need market speak like "born-digital," "multi-modal" and "multi-dimensional" to describe what's happening with data.

It's this simple: Data is exploding. It's being created everywhere. It's diverse, essential and valuable. The fact is that modern data is bigger, more challenging and more critical to business than a minimalist architecture with bare-bones software can handle. This perspective comes from seasoned experts who have been actively solving data challenges throughout all the technology eras. This is Dell's perspective, and it is the perspective that underpins PowerScale.

Pure's modern vs. legacy argument fails to account for operational challenges inherent with FlashBlade's software gaps; its monolithic hardware design that doesn't scale prescriptively or massively enough; and its resiliency deficiencies. What's more, Pure ignores the fact that modern data has a lifecycle; spinning media still offers significant economic advantages; snapshots don't offer full data protection; and cloud interoperability has value. Finally, Pure's modern vs. legacy argument masks the risks surrounding FlashBlade's minimalist design and Pure's narrow solution scope.

For those who are both excited by the opportunities of modern data and sobered by its magnitude, PowerScale offers a more capable, reliable and secure solution than FlashBlade.

Here are 10 ways that PowerScale surpasses FlashBlade for modern data:

1. **PowerScale is proven to feed more GPUs per cluster, so you can produce AI and ML outcomes more efficiently.**

   PowerScale OneFS is uniquely designed to enable concurrent file read access into the millions to support the efficient parallel IO profile demanded by many iterative ML workloads. For a variety of ML benchmarks, including Resnet50 and Inception-

V4, PowerScale demonstrated linear scale to up to 72 NVIDIA GPUs[1]. Compare this to just 32 NVIDIA GPUs in published ResNet 50 results for FlashBlade. If you are buying more than 32 GPUs, why would you invest in a solution that isn't proven to scale linearly past 32 GPUs?

2. **PowerScale lets you leverage different media classes to contain storage costs and merge in advancements.**

PowerScale's SmartPool technology manages data across different media classes, tracking data lifecycle with media economics and enabling seamless merging of media advancements into existing clusters. Compare this to FlashBlade's tier-less monolithic architecture, where you are prohibited from leveraging the $/GB economics of archival-class media within your cluster.

3. **PowerScale enables greater data efficiency to help you achieve better storage economics.**

At the most fundamental level, PowerScale enables up to 85% or higher effective utilization of raw capacity. On top of that, PowerScale both compresses and globally dedupes data for a reduction ratio of up to 3:1. Compare this to around 60% raw capacity utilization for FlashBlade, which also has no dedupe capability.

4. **PowerScale stands in sharp contrast to FlashBlade around protocol support.**

To begin, PowerScale can run AI in place on data using multi-protocol access. PowerScale can support all protocols from a single, unified namespace. OneFS uses a Unified Permissions Model to enable consistent permissions across not just NFS and SMB, but also HTTP, HDFS, S3 and FTP. So, for example, content and metadata can be ingested using S3 (or any supported protocol) and concurrently accessed through other protocols that are configured on the PowerScale cluster. What's more, PowerScale multi-tenancy support enables storage-as-a-service abilities through any supported protocol. Compare this to FlashBlade, which has no direct support for HDFS and requires discrete namespace for object and file storage.

5. **PowerScale enables higher redundancy, so you can better protect business-critical assets.**

Data protection on PowerScale can provide redundancy up to N+4 and is configurable at the file, pool or bucket levels. Also, the redundancy level can be changed on the fly. Compare this to FlashBlade, which has a maximum redundancy of N+2, and it is non-configurable, so business-critical data is protected at the same level as scratch space.

---

[1] Based on Dell testing, Nov. 2020, using Resnet50 and Inception-V3 benchmark test comparing PowerScale vs. previous generation. Actual performance will vary based on usage, configuration, and manufacturing variability.

6. **PowerScale is more secure from both internal and external threats.**

Feature-rich role-based access controls (RBAC), comprehensive and configurable encryption and advanced multi-tenancy enable tight security within PowerScale clusters. Dell Technologies PowerProtect Cyber Recovery, with full air gap capabilities, helps to operationalize protection and speed recovery from cyber-attacks. Compare this to FlashBlade, which lacks full-featured, robust RBAC and multi-tenancy, as well as configurable encryption.

7. **PowerScale is more flexible and adaptable, so you can more easily and cost-effectively handle growth and change.**

You can directly access PowerScale from more types of applications. You can scale PowerScale more prescriptively and economically with different node types (and support for mixed-node clusters) and media choices. You can seamlessly retire nodes or downsize clusters or reclaim capacity. PowerScale clusters can start at 3 nodes and incrementally scale to up to 252 nodes. Compare this to FlashBlade, which requires an initial commitment of 7 blades and >$100K[2], yet it scales to just one-seventh the capacity of PowerScale.

8. **PowerScale is more cloud-capable and available with more consumption choices.**

You can tier to the cloud. You can also consume PowerScale from public cloud, on-prem cloud, and multi-cloud. Compare this to FlashBlade, which has no cloud tiering capabilities and no direct access for applications running in the cloud. With its service-based consumption options, Pure has complex pricing with different rates for base and on-demand storage. PowerScale offers simple, easy-to-understand pricing with no premium for on-demand storage usage.

9. **PowerScale has smarter management, so you can handle advancing demands and unbridled growth with less staff.**

DataIQ lets you visualize and manage all your unstructured data—in the datacenter, at the edge, in the co-lo, in the public cloud—from a single pane of glass. And it gives you full data mobility functions, so you can easily control where data resides. What does Pure have to offer here? Nothing.

10. **PowerScale is part of a bigger solution ecosystem, so you can streamline operations and manage with greater excellence.**

Dell backs PowerScale with data protection, cybersecurity, archival, multi-cloud, computing systems, networking, applications and more. With unstructured data, especially, you must effectively interoperate with all these adjacent solutions to achieve success. Narrow solutions, like FlashBlade, offload broader solution

---

[2] Based on publicly available pricing list, February 2021.

responsibilities on customers. This increases operations workload and carries greater risk.

As you look to seize unstructured data opportunities ahead, listen to advice from veterans who have both the experience and the full-featured solutions you can trust to solve the data challenges of modern applications. Don't be misled by flawed arguments that understate the requirements and challenges of modern data. Beware the snares of minimalist architectures that can derail your success and impact your profitability.