

## SOLUTION BRIEF

# Accelerate EDA Workflows with Pure Storage

Test and fail fast to accelerate electronic design automation (EDA) production workflow with improved chip design quality enabled by Pure Storage.

The EDA industry has seen a recent resurgence in multiple dimensions—existing chip production pipelines are continually evolving with tighter transistor density, hyperscalers like Meta and cloud providers like Amazon and Microsoft are designing and producing custom silicon or accelerated computing, and other sub-10nm designs are emerging that will continue to increase pressure on silicon designers and manufacturers to streamline and accelerate everything.

## Pressure from AI Is Growing

These pressures are now becoming compounded by the addition of generative AI-augmented production workflows to improve system on chip (SoC) design and packaging quality. Front-end design, back-end design, and manufacturing processes are all being influenced by AI.

Agile and high-performance data center architecture, specifically data storage, is critical to successfully enabling these newer accelerated and dynamic EDA workflows. Legacy storage platforms often lack the agility and performance needed to keep up with shifting production needs, causing bottlenecks in the production cycle.

## Generative AI's Effect on Data Center Infrastructure

Generative AI large language models augmented testing and design is being introduced to existing application specific integrated circuit (ASIC) design cycles to provide valuable insight into previously undiscovered architectural design efficiencies. While AI-augmented design offers great potential into future chip builds, it comes with important data center considerations:



### Purpose-built Hardware

- Built from the ground up to natively leverage the performance, density, reliability, and efficiency of modern flash storage



### Cloud Operating Model

- Combines the best of on-premises and the cloud, delivering cloud-like agility and flexibility.



### Cybersecurity Built In

- Get early detection and avoidance capabilities for modern ransomware and other cybersecurity attacks.

- AI-augmented EDA workflows require variable storage performance needs during design verification. Oversubscribed storage causes unwanted delays in supporting requirements, slowing the dev process.
- AI-based operations and high GPU usage result in increased data center footprint and higher power draws. This means higher costs for rack space power overages and potential service outages.
- AI-based chip design workflows push the limits of traditional data center infrastructure designs. Legacy storage was not designed for the complex and diverse workloads of the AI-driven world we live in today.

## Pure Storage FlashBlade for EDA Workflows

The expression “choosing the right tool for the job” is often applied in IT, and enterprise storage for EDA design processes is no exception. Workloads all have unique and different characteristics that can cause significant resource contention if the underlying storage platform is not designed and equipped to handle it all. Pure Storage® FlashBlade®, driven by the Purity Operating Environment, provides an unequalled ability to deliver seamless scale out storage performance that intelligently adjusts its resource allocations for multimodal workloads like AI augmented EDA design.

The combination of NVMe and flash storage in a FlashBlade system is an incredibly fast hardware foundation that supports petabytes of unstructured data, workload multitenancy and quality of service (QOS), design collaboration across remote and global sites, and data protection technologies like replication and indelible snapshots for cybersecurity protection built in.

## FlashBlade System Key Characteristics

FlashBlade systems offer many advantages over legacy storage solutions, including:

- **Multi-dimensional performance advantage.** FlashBlade//S™ delivers high metadata performance and high throughput for any file or object workload, irrespective of data size
- **Multi-protocol support.** Provides file and object protocol services as equal-status peers that share the same underlying storage resource pool natively. This eliminates the need to predict the future and “earmark” dedicated storage resources for each service that can’t be changed.
- **Global storage data reduction.** Additionally, both services benefit from the same umbrella of compression—what gets reduced in file storage is applied to objects, and vice versa.
- **Simple, scalable, and granular data management.** Simple, granular data management is crucial to platform flexibility—delivering a robust data repository should benefit from the fewest interface clicks possible, allowing for any data center generalist to manage and deliver storage.
- **API-led design principles.** Easy product interoperability is a critical component of modern application design and supporting infrastructure management. A robust availability of platform APIs enables current and future integration needs to other datacenter elements without complicated efforts seen in legacy storage providers.
- **Purpose-built hardware architecture.** Built from the ground up to natively leverage the performance, density, reliability, and efficiency of modern flash storage.
- **Platform management simplicity.** Software that enables the ability to deploy, monitor, manage, and perform non-disruptive upgrades without requiring constant tuning or disruptive changes.
- **Dynamic scalability.** Seamlessly scales capacity, performance, and metadata to independently manage both petabytes of data and billions of files to meet modern application needs. Modern File and Object platforms shouldn’t require admins to manually move storage allocations around to make new requirements fit.
- **Always available.** Delivers platform high availability over multiple years with software and hardware upgrades that avoid unnecessary downtime or data migration headaches.



## Better for You and Our Planet

Using FlashBlade for EDA workloads makes sense both financially and environmentally. Combining a cloud operating model with a simple, flexible as-a-service model enables organizations to have:

- **Cloud operating model.** Combines the best of on-premises and the cloud, delivering cloud-like agility, flexibility, and consumption choices with the control of on-prem, with best-in-class [AIOps for fleet monitoring in Pure1](#)—our control plane app available on both desktop and mobile devices.
- **Storage as a consumable service.** Modern IT budgets are shifting from large CapEx purchases to aligning with an OpEx, consumption-based model for both hardware and software. The Pure Storage [Evergreen® subscription](#) portfolio allows for a wide spectrum of acquisition models.
- **Cybersecurity savvy.** And, as the most important icing on the cake, the platform should provide early detection and avoidance capabilities for modern ransomware and other cybersecurity attacks. [Pure1®](#) anomaly detection and [SafeMode™ snapshots](#) effortlessly enable the ability to mitigate the risks associated with today's threats and the ones we don't know about yet.
- **Reduced carbon footprint.** FlashBlade, along with other products in the Pure Storage platform, helps to reduce e-waste from periodic refreshing and upgrading storage infrastructure due to escalating power and cooling demands in the datacenter. The [ESG](#) savings from FlashBlade systems in the datacenter can help towards the computational (CPU and GPU) inflation needed in the compute farms for designing and testing modern chip designs.

## Say Goodbye to Overcomplicated, Archaic Storage Solutions

Pure Storage delivers the simplicity needed to manage explosive unstructured data growth that is a hallmark of the modern IT estate, and FlashBlade delivers multi-dimensional performance benefits for modern EDA chip design workloads using multi-die technology. Its ability to natively support service level agreements (SLA) manages data into separate cost tiers within a storage cluster, while its modular and scale-out architecture allows for on-demand linear improvement and performance for high speed data access during the critical phases of a silicon lifecycle (SLC). As active projects ramp down and designs are taped out, high density QLC-based [Direct Flash® Modules](#) (DFMs) can also provide high capacity storage to store petabytes of chip design data for long term retention and compliance....all of this in the smallest amount of rackspace and power draw that our competitors can't reach.

Designed from the ground up for flash storage to capitalize on next-generation data storage needs, FlashBlade unified file and object storage (UFFO) uncomplicates data storage for organizations with any kind of workload, such as generative AI, vector databases, and file share-based unstructured data with rapid restore and ransomware protection built in.

## Additional Resources

- [Click here](#) to learn more about EDA customers using FlashBlade for their modern chip design projects.

[purestorage.com](https://purestorage.com)

800.379.PURE

