Google Cloud NetApp

Accelerate innovation, streamline your chip design

Purpose-built for the most demanding EDA workloads

Introduction

Chip design moves fast.
Your storage needs to keep up.

The semiconductor industry never sleeps. To stay competitive, Electronic Design Automation (EDA) teams must constantly speed up complex design processes, manage exploding data volumes, accelerate time-to-market, and keep costs predictable — all while maintaining exceptional availability, performance, and reliability.

Google Cloud NetApp Volumes directly addresses the unique challenges and demands of modern chip design workflows. As a fully managed, high-performance file storage solution, it enables EDA teams to seamlessly manage massive data volumes and deliver exceptional throughput. From metadata-intensive front-end tasks to throughput-intensive back-end verification phases, NetApp Volumes provides the scalable capacity, speed, and seamless integration within the Google Cloud ecosystem to accelerate your innovation pipeline and AI-readiness.

In this ebook, we explore why industry-leading EDA companies choose NetApp Volumes to deliver exceptional performance across diverse workloads, simplify storage operations, boost developer productivity, and accelerate time to market.

Exceptional performance for demanding EDA workloads

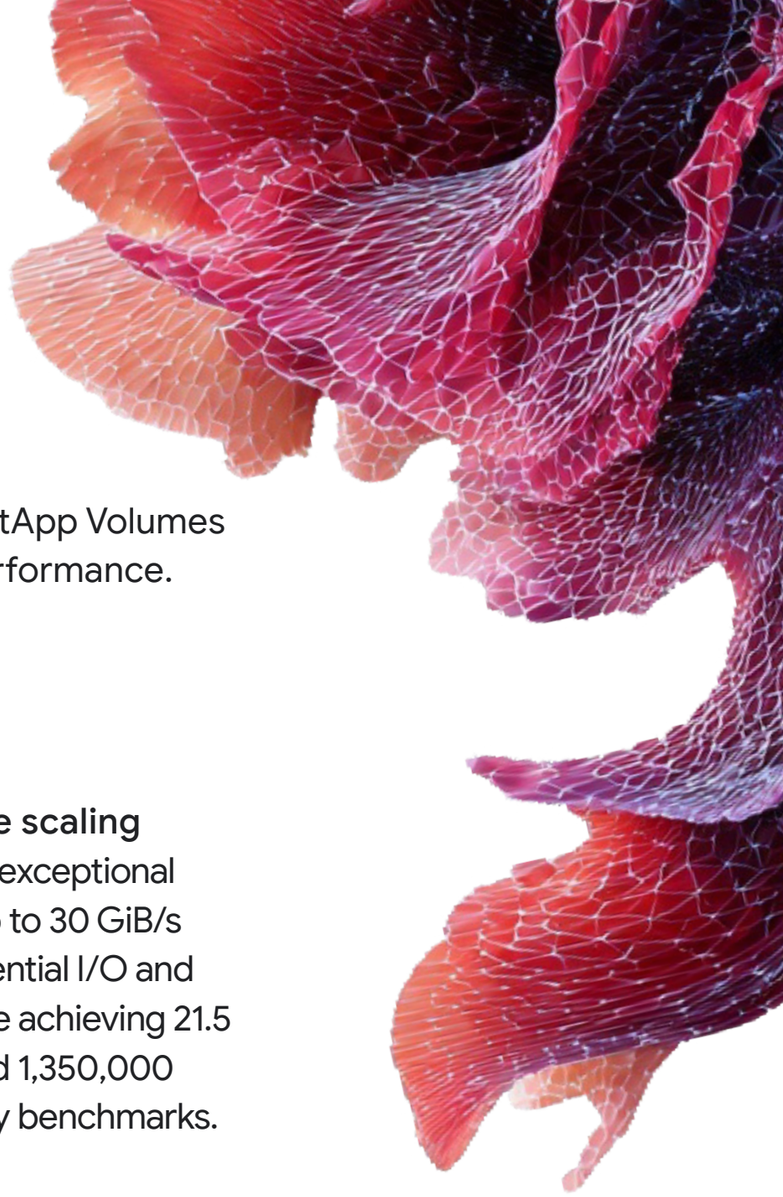

Purpose-built for high-throughput, high-scale engineering workloads like EDA, Google Cloud NetApp Volumes seamlessly supports front-end and back-end design processes and delivers flexible, scalable performance.

Front-end design
Front-end phases are metadata-intensive, demanding fast I/O for hundreds of millions of small files. NetApp Volumes delivers low latency and high IOPS to maximize developer productivity.

Back-end verification
Back-end processes need high sequential throughput. NetApp Volumes effortlessly delivers high-speed data access — up to 30 GiBps for sequential read workflows — enabling rapid verification.

Linear performance scaling
The solution delivers exceptional performance with up to 30 GiB/s throughput for sequential I/O and 1,980,000 IOPS, while achieving 21.5 GiB/s throughput and 1,350,000 IOPS for EDA industry benchmarks.

Optimized for EDA workloads

Up to

30 GiB/s

throughput

Single

3 PiB

namespace

Accelerate innovation, streamline your chip design

High-performance hybrid cloud architecture for EDA

The elasticity of Google Cloud is essential for scaling EDA workloads, but accessing large on-premises datasets from the cloud often introduces unacceptable latency and complexity. Google Cloud NetApp Volumes leverages NetApp's core data services to overcome these hurdles, enabling a high-performance hybrid cloud strategy.

Enabling the hybrid EDA model with NetApp FlexCache

One of the primary challenges for the "burst-to-cloud" EDA model is running simulation or verification jobs in Google Cloud while maintaining low-latency access to authoritative golden datasets and shared tooling residing on-premises. NetApp Volumes, powered by NetApp FlexCache technology, is the strategic solution to this data mobility challenge. FlexCache creates a sparse, high-performance cache volume directly within Google Cloud, intelligently connecting back to your on-premises ONTAP origin volume.

Reduced WAN latency

FlexCache caches only the active, requested data blocks (the working dataset) in the cloud, bringing the data physically closer to your Google Cloud compute instances. It supports up to 100 Cache volumes per origin, with sizes from 100GiB to 1PiB, with support for both SMB/NFS protocols. This dramatically reduces latency, allowing read-intensive EDA tasks like synthesis and verification to run at local-like speeds.

Faster job completion

By serving frequently accessed tool libraries and design data from the NetApp Volumes cache, cloud compute resources spend less time waiting for data, accelerating design cycles and minimizing the run-time and cost of cloud instances.

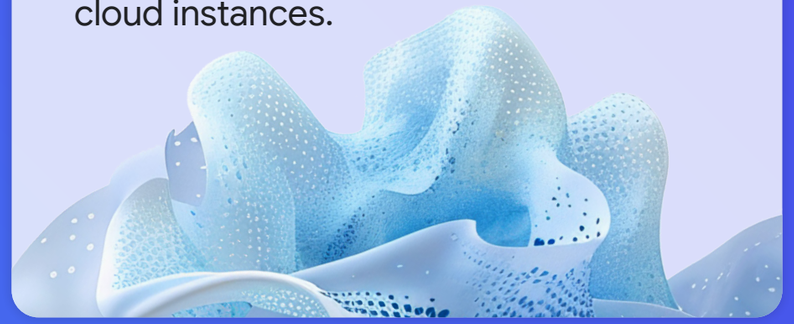

Seamless cloud bursting

FlexCache enables the true “Design Anywhere” model. The cache volume looks exactly like the origin, allowing engineers to burst workloads without needing to manually copy or manage redundant data. Write operations are consistently passed through to the original volume, ensuring data coherence and integrity across the hybrid environment.

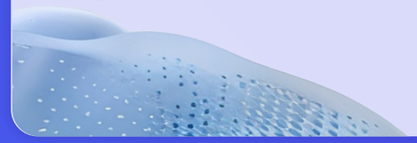

NetApp FlexCache accelerates data access, reduces WAN latency, and lowers WAN bandwidth costs for read-intensive workloads, especially when the same data needs to be repeatedly accessed.

Operational simplicity and seamless Google Cloud integration

Google Cloud NetApp Volumes is fully managed and cloud-native, freeing your engineers to focus solely on design innovation.

Fully managed cloud-native service
No infrastructure upgrades or maintenance required. Seamless integration with familiar Google Cloud tools (Google Cloud Console, Terraform, RESTful APIs) simplifies automation and visibility across your storage environments.

Instant snapshots and data protection
Built-in instant snapshots and clones allow rapid checkpointing and rollback, enhancing your productivity without interrupting ongoing jobs and protecting critical IP.

Built-in security and high availability
With built-in data replication, encryption, IAM integration, and up to 99.99% guaranteed uptime, your data remains secure, compliant, and always available.

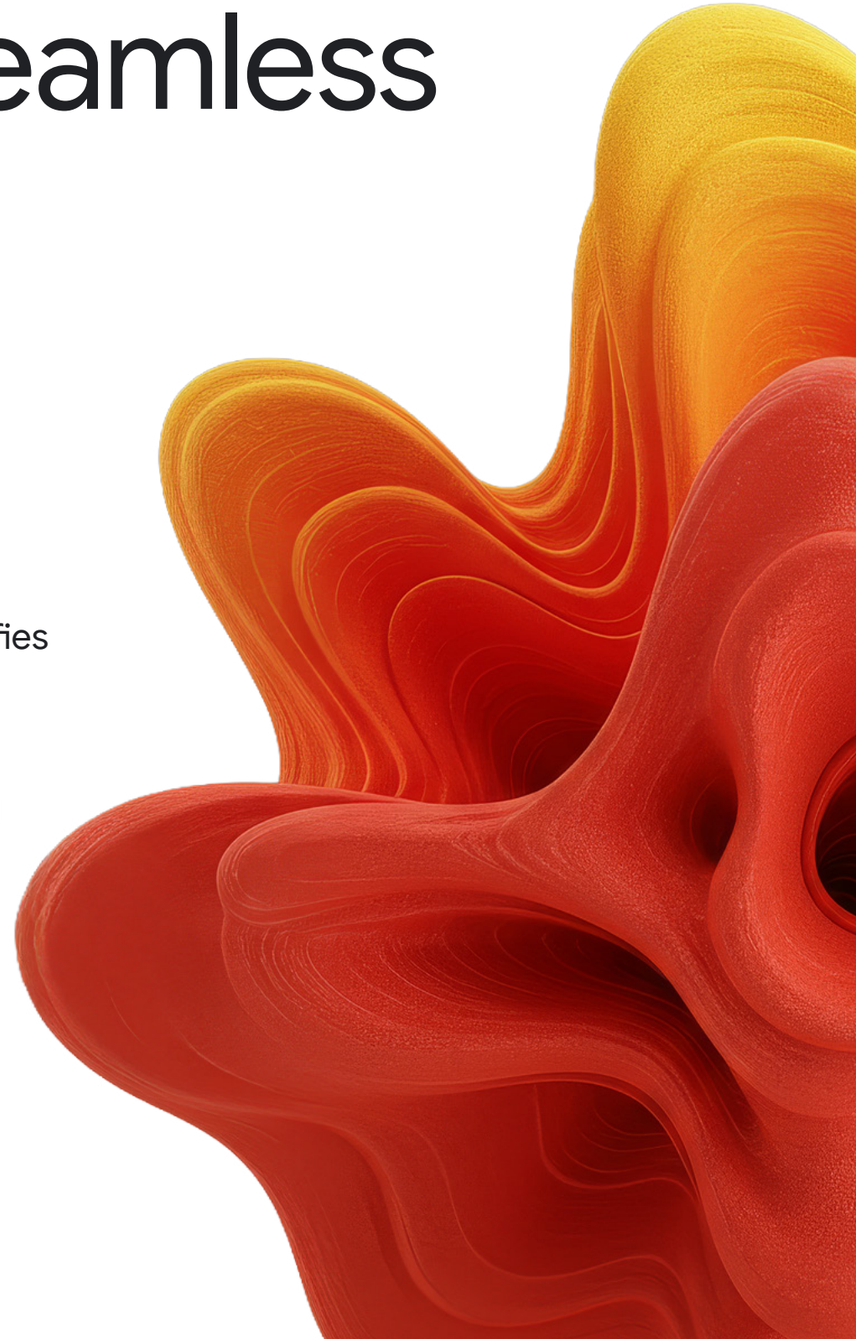

Simplify and speed deployment of EDA workloads on Google Cloud

NetApp Volumes integrates natively with Google Cloud's Cluster Toolkit open-source software to simplify the deployment of complex high-performance computing (HPC), AI, machine learning (ML), and EDA environments. It helps ensure the entire stack — compute, networking, and high-performance storage — is deployed efficiently according to Google Cloud best practices. This integration enables:

Rapid deployment

Quickly deploy validated blueprints for EDA workloads, reducing administrative provisioning time.

Flexible workflows

Seamless support for both run-in-cloud and burst-to-cloud models.

Optimized performance

Guaranteed compatibility with Google Cloud compute to leverage the full throughput and IOPS capabilities of NetApp Volumes storage.

Accelerate time to market and lower costs

Your speed determines your market success. Google Cloud NetApp Volumes helps EDA teams accelerate complex workflows and speed time to market.

Accelerate tape-out schedules
NetApp Volumes enables your teams to complete chip design validation faster, shortening tape-out timelines and speeding time to market.

Scale compute instantly
Quickly spin up parallel compute jobs without delays or bottlenecks. Scale seamlessly on demand, keeping your design teams agile and productive.

Lower costs
Faster job completion minimizes EDA license consumption and reduces the engineering hours spent waiting for jobs to complete.

Efficiency that scales with your workloads

NetApp Volumes makes it easy to size and scale storage and compute according to your needs, saving on both costs and time wasted on infrastructure management.

Accelerate your journey from design to market

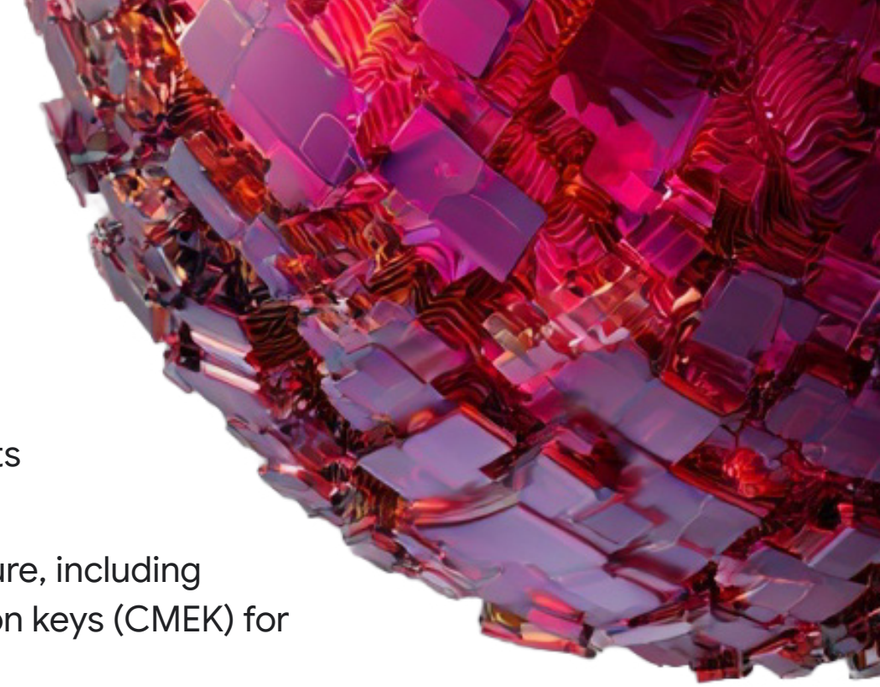

Built to meet your most demanding workload needs, Google Cloud NetApp Volumes supports faster chip design cycles and performance at scale to reduce your time to market.

- A single, scale-out namespace with automatic load balancing, enabling EDA workflows at massive scale without sharding data or rearchitecting pipelines
- Massive capacity and dynamic scaling up to 3 PiB in capacity, up and down, and up to 30 GiBps of throughput
- Scalable throughput and IOPS at low latency for both front-end and back-end jobs
- Cost-effective storage with auto-tiering cold data into lower cost storage; easy consumption with Cloud Billing and eligibility for Google Cloud committed use discounts
- Secure-by-design infrastructure, including customer-managed encryption keys (CMEK) for data encryption at rest
- Secure administration and compliant operations with identity and access management (IAM) and Virtual Private Cloud Service Controls (VPC-SC)
- Integration into Google Cloud console, API, or gcloud to enable provisioning, management, and monitoring of NetApp Volumes as a Google Cloud service
- Integration with Google Cloud Gemini Enterprise accelerates AI-powered design automation without wasteful data movement or replication

Learn more about Google Cloud NetApp Volumes

Innovate faster and go-to-market more efficiently with a fully managed, high-performance, massively scalable, and cost-efficient storage solution, optimized for EDA workflows.

Google Cloud NetApp Volumes trial and documentation
Google Cloud NetApp Volumes for EDA
Talk to a specialist

Google Cloud NetApp