FEBRUARY 2026

Transforming EDA for the AI Era, With Google Cloud NetApp Volumes

Simon Robinson, Principal Analyst

Abstract: The AI revolution is driving unprecedented demands on semiconductor design for ever more capable silicon, pushing traditional electronic design automation (EDA) processes to their limits. Cloud-based EDA offers transformative potential, but migrating such a large, complex, and multi-step operation requires substantial planning, effort, and support from committed partners, especially given the enormous data volumes involved. In response, NetApp and Google Cloud have strategically partnered to deliver Google Cloud NetApp Volumes (GCNV), a fully managed, high-performance storage solution optimized for cloud-based EDA environments, enabling seamless data mobility, robust lifecycle management, and enhanced cost efficiency.

The AI revolution pushes traditional EDA past its limits

With Enterprise Strategy Group (now Omdia) research finding that 84% of organizations view AI as critical to their future strategies,[1] it's clear that the AI revolution is poised to transform almost every aspect of our lives. It's also clear that this future is dependent on one critical component: the semiconductor. Semiconductors are reaching unprecedented levels of capability; however, the march toward trillion-transistor systems and angstrom-scale devices presents designers with massive scale and complexity challenges.

Designing next-generation chips for AI workloads presents particular issues: AI models are being trained on enormous and varied datasets, which means the capabilities of chip design must increase exponentially, both in terms of capability (e.g., now stretching to billions of transistors) and raw computing power (from petaflops to exaflops).

Many of these challenges are falling chiefly on the shoulders of the EDA process. Specific issues include:

- **Skyrocketing design costs.** As the semiconductor industry moves from 5nm to 3nm chips, costs are set to increase significantly. Current estimates suggest a threefold increase over 5nm designs, pushing the total cost of ownership into the billions.
- **Lengthening development cycles.** Additional costs and complexities are pushing chip design cadences from 18 months to 30 months or more, just as the market is demanding faster cycle times.
- **Growing design complexity.** Current transistor gate technologies, such as FinFET, are reaching their limit and are being replaced by newer, more advancedand more complicated alternatives, such as gate-all-around.

[1] Source: Enterprise Strategy Group (now Omdia) Research Report, *The Critical Role of Storage in Building an Enterprise AI Infrastructure*, September 2025.

These challenges are prompting chip designers to rethink many aspects of their EDA process. Historically, they have centralized their EDA processes and workflows, typically running the entire process on premises and on dedicated infrastructure. This was often the simplest and most effective way of managing a complex, multi-step EDA workflow.

However, as the computational and data requirements for advanced semiconductors grow exponentially, the limitations of traditional on-premises infrastructure are being exposed. The vast compute power required for modern simulations has reached a point where it is no longer economically viable to maintain that peak capacity permanently on premises.

Cloud-based EDA offers a clear path forward for chip manufacturers

One alternative to addressing these challenges is to migrate parts of the EDA process to public cloud providers that are investing billions in building out unprecedented amounts of computational and related infrastructure capacity. Moving EDA to the cloud would be in step with broader industry trends: according to a recent Enterprise Strategy Group research study, almost two-thirds of organizations (64%) said they are using, or planning to use, hyperscale public cloud providers for their AI initiatives.[2]

The appeal of elastically scalable, cloud-based EDA is immense. Cloud-based approaches provide chip designers with near-instant access to the needed computing, storage, and networking capacity in an on-demand fashion. This elastic access can also be achieved across many more physical locations globally than most organizations could ever deploy themselves, making it particularly attractive to EDA designers with globally distributed teams. Crucially, this could be accessed (and paid for) on demand, meaning chip designers can use it only when they need it, potentially significantly reducing overall costs. Additionally, designers are able to access the innovations—around AI and other areas—that cloud providers are developing, both internally and through extensive partner networks. EDA data stores can be enormous, so there may be significant advantages to colocating data for traditional EDA flows and AI-driven capabilities, without having to move data back and forth.

However, the challenge for chip designers is how to transition to a cloud-based approach without disrupting existing processes. Many designers have invested in significant infrastructure and resources over many years, meaning the bulk of the core EDA system and its associated data, including design files, libraries, and historic simulation results, resides on-premises. Additionally, the semiconductor design is a multi-step process, spanning front-end design and verification, back-end verification, and production and test, each of which requires close cooperation between computing and storage, and orchestration of many diverse tasks with multiple complex dependencies.

In other words, moving such a critical, sizable, multi-step process to an entirely different environment requires an evolutionary approach that will take time and careful planning, not to mention substantial commitment and strategic partnerships from both on-premises infrastructure vendors and public cloud providers. For many, the optimal route might be a hybrid approach that enables semiconductor designers to adopt a graduated approach to cloud-based migration that can be executed over time and at a pace that suits their own business requirement and timings.

[2] Ibid.

Hybrid approaches offer an on-ramp to cloud-based EDA without disruption

Recent market developments are set to help significantly with this approach. In particular, key players are partnering to offer sophisticated integration between traditional on-premises data storage infrastructure and related management toolsets and public cloud-based capabilities.

Growing cooperation and development here is seeing capabilities that previously resided only on premises being integrated with public cloud-based services. Enterprise Strategy Group research validates the multiple benefits that enterprises are seeing from such integrations, including superior performance, providing the best set of features and functionality, and providing the best total cost of ownership, as well as strong alignment with an organization's strategic direction (cited by half of respondents, see Figure 1).[3]

Such integrations highlight that semiconductor designers stand to benefit from increasingly rich cloud-based services that offer some of the infrastructure and data tools they are familiar with using on premises.

However, this alone does not address the entirety of the problem. How, for instance, does an organization optimally manage an EDA environment that spans both on-premises environments—which might host the core, high-performance EDA environment—and cloud-based services that might be used for bursting specific simulations or similar capabilities? Ensuring that the cloud-based compute cluster has efficient access to the vast datasets that are often stored on premises is key.

Few technology suppliers are capable of effectively and efficiently handling such a hybrid EDA workload. One notable exception to this is the evolving and strategic partnership between NetApp and Google Cloud. We will explore the substantial capabilities that the two offer chip design companies looking to build a true hybrid EDA environment in the next section.

Figure 1. Key benefits of third-party storage offerings running on public cloud service

You indicated your organization uses third-party storage offerings that run on a public cloud service. Which of the following drivers influenced this decision? (Percent of respondents, N=315, multiple responses accepted)

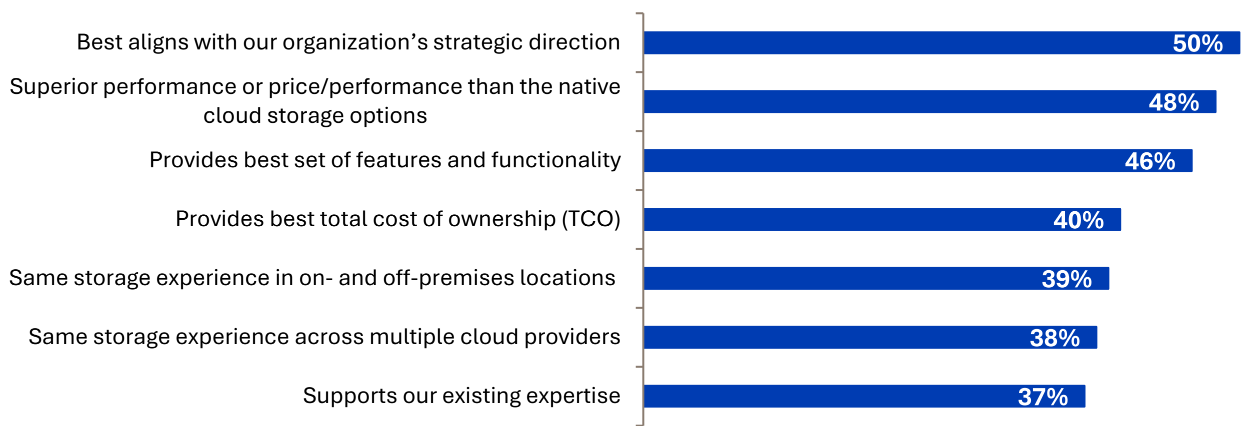

Source: Omdia

[3] Source: Enterprise Strategy Group (now Omdia) Research Report, *Navigating the Cloud and AI Revolution: The State of Enterprise Storage and HCI*, March 2024.

Google Cloud and NetApp optimize hybrid cloud EDA for the AI era

NetApp and Google Cloud have been working together at a strategic level for several years to create a range of solutions that optimize the enterprise storage experience for customers adopting public cloud.

A significant capability here is Google Cloud NetApp Volumes (GCNV), a fully managed, high performance storage service for Google Cloud that supports both file (NAS) and block (SAN) workloads. Built on NetApp's ONTAP technology and co-engineered by both companies, GCNV provides scalable, enterprise-grade file storage with support for NFS, SMB protocols offering concurrent file access across Linux and Windows file systems, along with a rich set of data management services such as snapshots, backup, and replication.

More recently, additional GCNV support for iSCSI block storage enables customers to also seamlessly migrate their SAN workloads to the cloud. This is a particularly good fit for the demanding I/O and throughput needs of EDA workloads, offering consistent low latency and robust throughput for collaborative, multi-user environments.

The GCNV solution, which is sold and supported directly by Google Cloud as a first-party service, allows customers to run enterprise applications such as EDA in the Google Cloud without rearchitecting them, while enjoying close integration into Google Cloud services, billing, and support.

Google Cloud and GCNV – Optimized for cloud-based EDA

Google Cloud and NetApp have worked closely to optimize GCNV for EDA workloads. Key capabilities include:

- **Optimal scale for EDA workloads.** Both scale-out and scale-up performance with file storage volumes ranging from 1GiB to 3PiB, delivering read throughput of up to 30GiBps. Organizations can push over 1.2m IOPS for large numbers of simultaneous EDA jobs.
- **Data mobility for the hybrid cloud.** Replicate tool and library data efficiently to the cloud with FlexCache, offering instant visibility and access, optimizing read performance, and simplifying distribution and seamless data access across multiple locations.
- **Reliable data management and lifecycle solutions.** Snapshots, volume copies, volume replication, and integrated backup enable high availability and disaster recovery while simplifying data protection.
- **Fully managed file storage solution.** 30 years of NetApp storage innovation available in Google Cloud with no need for upgrades, installations, maintenance, or repairs.
- **Data security.** NetApp ensures robust data security with encryption at rest, CMEK, FedRAMP (FIPS 140-2) compliance, and granular access control via role-based access control.
- **Optimized cloud storage economics.** Cost-effective storage with auto-tiering of infrequently accessed data to a less expensive, colder storage tier.

FlexCache – The bridge that enables EDA across the hybrid cloud

Though GCNV can help drive substantial value for running EDA workloads in the cloud, a key issue is initially getting the right data into the Google Cloud environment in a controlled manner. This is where NetApp's ONTAP FlexCache technology is playing a key role in enabling EDA transformation across the hybrid cloud. FlexCache enables customers to leverage the immense elasticity and burst capabilities of Google Cloud compute for EDA workloads while maintaining authoritative data on a customer's premises, without compromising performance or incurring unnecessary costs due to inefficient data access (see Figure 2).

Key benefits include:

- **Accelerated data access.** By caching frequently accessed EDA data in Google Cloud NetApp Volumes, FlexCache drastically reduces latency. This can be on-demand caching, where a cloud compute instance pulls a file from an on-premises ONTAP system into the GCNV cache; while there may be an initial latency for this first access as the data is transferred to the cache, subsequent access from any cloud instance will be served directly from the cloud-based cache. Alternatively, for critical projects or anticipated peak workloads, the cache can also be prewarmed, proactively loading essential datasets into the GCNV FlexCache cache before compute jobs begin, entirely eliminating the first-access latency. Cloud-based compute instances access data at local speeds, eliminating WAN bottlenecks and speeding up simulations, verification, and synthesis.
- **Faster job completion.** Reduced data access times translate directly to faster EDA job completion. Organizations' cloud compute resources spend less time waiting for data and more time processing, which accelerates design cycles.
- **Cost efficiency.** Cost savings are achieved in two key areas:
  - Reduced data transfer time. Faster access to cached data means that compute jobs finish faster, minimizing the time that expensive cloud compute instances are active.
  - Optimized cache footprint. FlexCache enables customers to cache only the essential subset of their vast EDA datasets, optimizing storage costs in Google Cloud NetApp Volumes.
- **Seamless integration.** Leveraging GCNV as the cache target offers a native, high-performance environment for cached EDA data in Google Cloud, extending on-premises storage to the cloud with minimal application changes.
- **Simplified management.** FlexCache maintains data consistency between organizations' on-premises ONTAP source and the cloud cache, streamlining data management in their hybrid setup.

Figure 2. GCNV and FlexCache overview

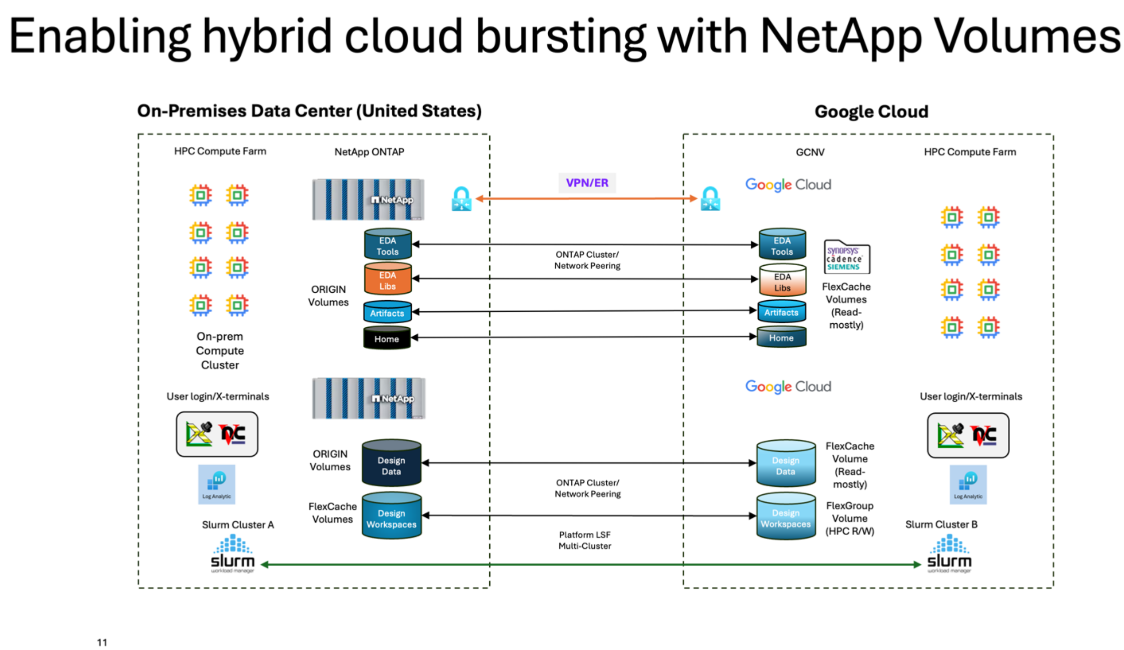

Source: Google and NetApp

Conclusion

The AI revolution is driving unprecedented demands on semiconductor design, pushing traditional EDA processes to their limits. To remain competitive, chip designers should consider cloud-based EDA solutions that offer access to extreme levels of scale, computational power, flexibility, and innovation. But they also need to embrace these cloud-based EDA options at a pace that is right for them, which may make a more graduated approach that leverages cloud in a hybrid fashion more appropriate.

Regardless of the approach, the ongoing partnership between NetApp and Google Cloud—and transformative tools like GCNV and ONTAP FlexCache—offers designers and manufacturers a compelling, strategic path forward, enabling seamless data mobility, faster simulations, and cost-efficient data storage and management. The partnership and deep integration demonstrates strong commitment from both partners, ensuring organizations can meet the growing complexity of AI-driven chip design while minimizing disruption and optimizing workflows over both the short and long term.

By leveraging NetApp and Google Cloud's innovations, semiconductor designers are armed with the capability to confidently take their designs into the AI era and are empowered to accelerate design cycles, reduce latency, and optimize resources. Thus equipped, designers are optimized to tackle next-generation demands effectively, driving innovation and retaining their competitive edge in an era of unprecedented technological evolution.

For more information on Google Cloud NetApp Volumes, please visit this link.

Get in touch: www.omdia.com askananalyst@omdia.com