Accelerating Electronic Design Automation Workflows
with Hammerspace

SOLUTIONS BRIEF

The automated design of today's advanced semiconductors is an extremely dynamic and performance-driven industry that is experiencing exponential growth, both in the complexity of the designs and the number of engineers needed to rapidly bring new hardware to market.

Key Customer Benefits

- Lower EDA licensing costs with accelerated workflows
- Parallelize EDA pipelines to increase output
- Enable local access to distributed apps and users
- Burst to Cloud for on-demand scale-up and scale-down

Since Electronic Design Automation (EDA) is an integral foundation to technology innovation in virtually every industry there is tremendous competitive pressure on EDA pipelines to accelerate time to market and reduce costs. Both are needed to satisfy the increased demand for new chip designs that can achieve ever-increasing power efficiency and performance goals.

As a result, it is increasingly necessary to spread workloads across distributed engineering teams and infrastructure, creating challenges of how best to collaborate in real time across large distances. And as designs get more complex with new 7nm chips and below, the datasets that must be accessible across distributed resources are also growing significantly. These issues impede rapid access to users and applications to data for processing, and thus slows time to completion. As a result this increases the licensing, infrastructure and operational costs of bringing new chips to market.

Key issues facing the industry:

A key problem that IC manufacturers face as they provision EDA pipelines is to ensure that their distributed engineering teams, their applications, and most importantly their datasets get immediate access to high bandwidth and low latency storage and compute resources to support EDA workflows. As datasets grow due to more complex chip designs, and as resources become more distributed, the impact of needing to orchestrate data between different sites or between on-premises and cloud resources adds significant delays. This increases costs, as well as adds increased time to market for the design process.

These delays and added costs are not only about getting better utilization of infrastructure. The impact of these delays adds significantly to the licensing costs of the EDA software tools, which make up the lion's share of EDA pipeline expenses. If EDA tools are sitting idle waiting for data and infrastructure resources, not only are designs delayed in coming to market, but those licenses end up burning cash waiting for data to be made available to process.

In addition, almost as critical as pure performance of the infrastructure is the ability for EDA production to rapidly ramp up or down compute and storage resources on demand, and to do so in a way that is non-disruptive to both applications and users. This becomes especially important when workflows need to be distributed across multiple locations to take advantage of Cloud-based infrastructure, and to enable companies to parallelize pipelines to bring more designs to market quicker, and with lower cost.

Problems with traditional EDA pipeline strategies:

Traditionally, chip design pipelines use classic High-Performance Computing (HPC) architectures, with parallel applications running in large compute farms attached to multiple pools of NFS file servers. The problem is this approach is both capital intensive and very difficult to adapt to differing EDA design workflows. A new design project may require different workspace requirements, with specific bandwidth, compute, and storage resources. Fixed infrastructure can't easily shift from one design configuration to another, which results in design workflows running serially. The inability to easily support multiple design workflows in parallel slows down the whole process, which adds costs and negatively impacts time to market.

Storage also becomes a bottleneck in running parallel projects, as large NFS filers routinely experience I/O bottlenecks in large-scale HPC pipelines. Modern IC designs may require 1000s of compute nodes to run in parallel, with enough I/O feeding them to keep pipelines running at full speed, and to ensure applications are not burning money waiting for data to process. So even individual designs will challenge traditional infrastructures. Adding more projects in parallel becomes extremely difficult, or impossible.

Finally, as in every industry that is struggling with a distributed workforce, complex EDA projects that require 10s of millions of files across 100s of TB of storage become an accessibility nightmare when trying to maintain the whole workflow in sync across users/applications in multiple locations. The inevitable explosion of file copies when sharing data from site to site creates a heavy penalty in unnecessary storage costs, but more importantly adds to licensing overheads as applications wait for data to be copied to the correct location for processing. In addition, OPEX for such projects also skyrockets, as IT administrators must wrangle copies of files across multiple storage silos and locations, all while normalizing access to users and applications in different geographies.

Solving Global Access to Data and Dynamic Compute/Storage Workspaces with Hammerspace

Hammerspace with its Global Data Environment is ideally suited to help overcome these dual problems of accessibility to data for a distributed workforce/applications, while also providing the flexibility to dynamically manage EDA production resources between on-premises and cloud resources. In this way, Hammerspace enables customers to parallelize EDA pipelines and achieve rapid time-to-market goals, while leveraging and extending their existing storage and compute resources.

Direct Local Access to Any Data, on Any Storage, Anywhere

Hammerspace is a software-defined data orchestration and storage solution that unifies all datasets across geographic locations, different storage types from any vendor, and multiple data centers and clouds into a single, global namespace. As such, it is designed to provide organizations with ability to consolidate datacenter and cloud storage while orchestrating global data to decentralized applications and users.

It does so by leveraging the power of metadata to create a shared global filesystem that enables any user or application anywhere direct local access to datasets that may be distributed across not only multiple storage silos in a single datacenter, but also across multiple data centers and Cloud vendors & regions.

When a user or application needs to access files that may physically live on remote storage, Hammerspace presents them with direct multi-protocol access to the files as though they were local with the low latency performance needed for read and write operations. This is not just shuffling copies of data across storage types and locations, which causes delays for applications, adds confusion to users, and creates headaches for IT admins. This is globally accessing the same files, via an advanced metadata control plane that intelligently bridges the underlying physical storage resources of any type.

The Hammerspace Global Data Environment starts from the premise that since data must be accessed by users across multiple locations and stored globally across a myriad of storage choices, so shouldn't those data and storage resources be managed globally as well?

This is an important distinction: There are point solutions that can replicate file copies from Site A to Site B so users on each side can then see a mirror copy of the files. But these are file copies, not the same file. Each of those copies has its own metadata, so they are effectively forked copies. Changes by one user are independent of the other and must be manually reconciled. And the process of moving those copies creates delays in workflows, and adds delays to application access, plus manual processes and/or point solutions that add complexity to IT and data manager workloads

With Hammerspace's Global Data Environment, this complexity is eliminated. Hammerspace assimilates file and storage metadata from existing infrastructure into its global file system, and then presents multi-protocol access via an advanced metadata control plane shared by all users and applications across one or multiple storage types and locations.

At the infrastructure level, IT administrators are now able to create objective-based policies that can pre-stage data close to the compute resources across any location to eliminate application delays. This is completely transparent to both users and applications, who always see all the data online and available regardless of where they are, or the data is. Data orchestration between resources with Hammerspace is an automated background operation since the Hammerspace global file system presents the same file metadata to all users and applications everywhere.

Gone are the days where capacity becomes stranded in storage silos, or IT must manually migrate data copies from one storage type or location to another. With Hammerspace, applications and users simply see all the data they need in the same global file system as though it were on local storage.

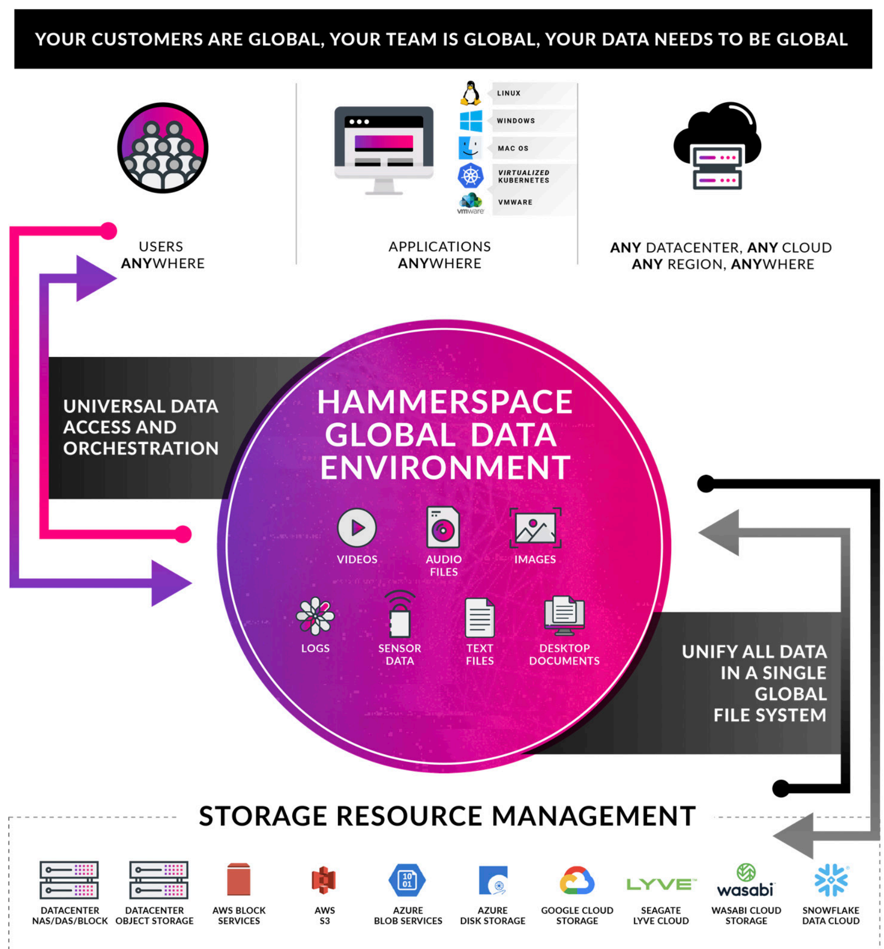

Hammerspace enables a global data environment that provides users and applications seamless global access and control of data across any storage type and cloud. It enables policy-based, cross-platform services for data migration, protection, tiering, and much more at at file-granular level and in a way that is non-disruptive to users or their workflows.

This capability also frees up IT administrators to add or remove storage resources on the back end without disruption to users or applications. For example, when a new EDA workspace is needed to accommodate a new design pipeline, the appropriate storage resources can be provisioned dynamically in the background at any on-premises or cloud location, or combination of both.

Because users and applications are accessing shared metadata via the global file system, increased performance can dynamically be scaled up non-disruptively on the backend across more storage and or compute resources of any type on demand. If extremely high IOPS are required for a task, Hammerspace's automated policy engines enable data placement to be spread across more resources in parallel, eliminating both the traditional NFS file bottlenecks, but also ensuring that user or application access to the data is optimized without interruption.

This capability can also span multiple data centers and Cloud infrastructures as well, which solves the second key problem in managing EDA pipelines: Flexible provisioning.

Flexible Provisioning to Accelerate EDA Pipelines

Because the Hammerspace Global Data Environment spans both on-premises and Cloud storage resources in a single global file system, this also means that EDA workflows can burst into Cloud compute and storage resources on demand, enabling both the flexibility and the efficiency of rapidly provisioning additional resources to parallelize and accelerate design workflows.

In real-world environments, Hammerspace configurations can be finely tuned to the specific EDA workload types, to conserve resources, but also to achieve the needed performance for the specific design requirements.

Hammerspace enables customers to parallelize EDA pipelines and achieve rapid time-to-market goals, while leveraging and extending their existing storage and compute resources.

- Hammerspace nodes can push data up to the maximum throughput of the infrastructure, and performance of the disks.
- Additionally, the parallelism of Hammerspace enables it to be configured to stripe I/O across as many high-performance storage systems in parallel as needed to achieve the throughput required for a specific workload. The limitations of a single disk array are no longer a bottleneck.
- This scale-out capability also ensures load balancing of I/O across the available resources.
- Additionally, the Hammerspace cluster itself can dynamically scale out to concurrently handle any volume of workloads.

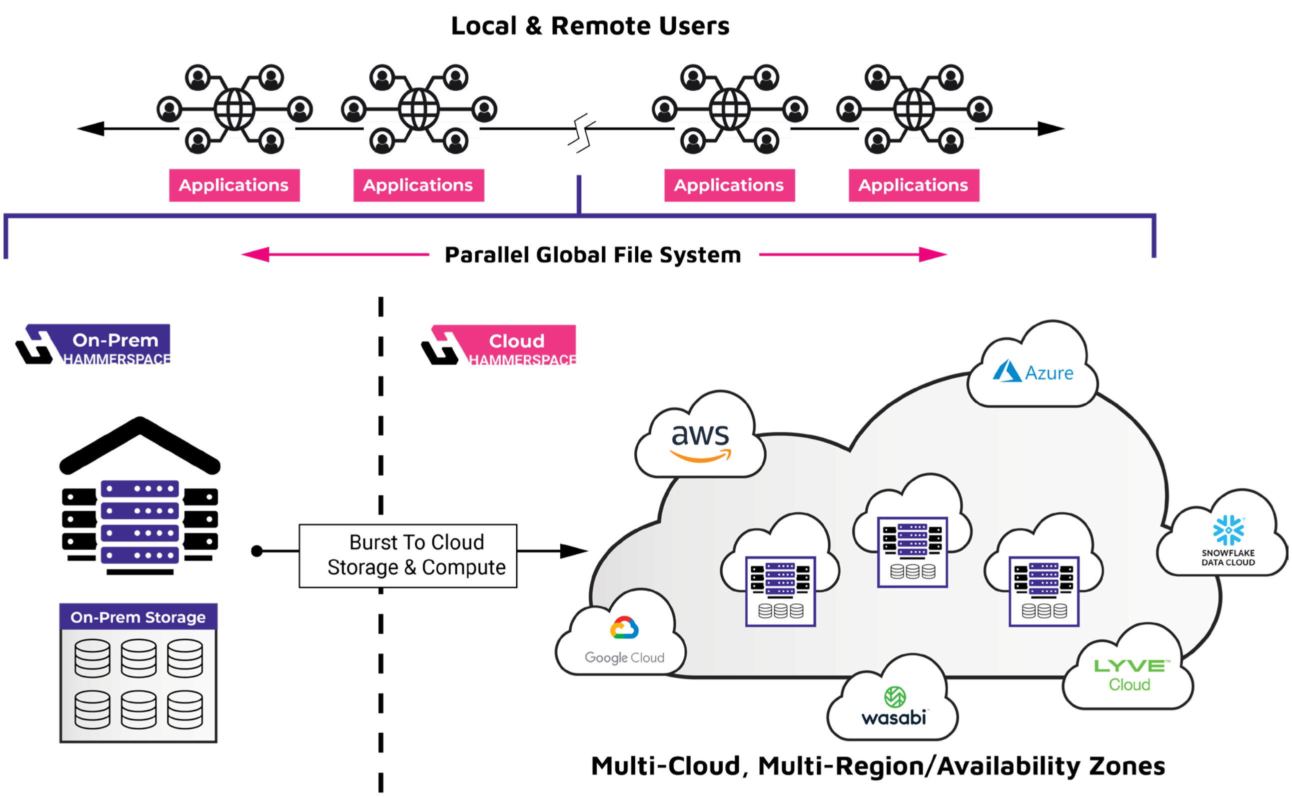

Hammerspace enables EDA pipelines of any performance level to be rapidly provisioned in parallel across any combination of on-premises and cloud resources, providing the experience of local file access for users and applications to the entire global file system.

In this way, the combination of the scale-out capabilities of Hammerspace and its ability to transparently spread workloads out across any storage infrastructure footprint results in maximum flexibility, enabling EDA pipelines to be tuned to achieve the best results with optimal resource utilization for the particular project.

This means that a design workspace can be rapidly provisioned in the cloud to take advantage of additional compute resources when needed, with data orchestration policies established so applications and users have direct local access to data, no matter which storage type it is on, or whether it is on-premises or in one or more Cloud providers or regions.

And unlike with a fixed infrastructure, which is a sunk cost whether it is being used efficiently or not, this rapid provisioning and burst capability into the Cloud without disruption to applications also means resources can be equally rapidly torn down when no longer needed to avoid unnecessary infrastructure and licensing costs. Data access never changes, and users and applications may not even be aware that their underlying infrastructure has changed. But data placement and workloads based upon workflow-driven policy objectives can be tuned to the job at hand, and completely automated.

Pulling It All Together

The Hammerspace Global Data Environment is a solution that is specifically designed to solve the two key problems of enabling uninterrupted global access to datasets anywhere for a distributed workforce and applications, and the need to increase EDA efficiency with rapid provisioning and bursting when needed to Cloud or other distributed compute and storage resources.

Unlike point solutions that may rely on proprietary hooks, gateways, symbolic links and the inevitable delays in pushing file copies across sites or storage types, Hammerspace's unique innovation is the power of its advanced metadata control plane, that presents the same global file system and thus the same files to all users and applications as local, regardless of the underlying storage or where they are. Gone is the need for wrangling copies, as the objective-based policy engines within Hammerspace enable automated data orchestration at a file-granular level across all locations of the entire enterprise.

This central control means that both user access, and administrative management of storage resources is completely unbound by vendor silos or location. EDA manufacturers and similar production workflows in other industries can leverage their existing storage resources, while also dynamically adapting to future requirements without limitation. And no longer do expensive EDA applications need to sit idle, with the meter running on licensing costs unnecessarily, waiting for data to be copied to the correct location.

With the Hammerspace Global Data Environment, all users and applications anywhere have immediate local access globally to all data across any on-premises and cloud compute and storage resources anywhere.