IC Manage Holodeck Enables Google Cloud Platform for Semiconductor Design

Overview

The great expansion of data for semiconductor design poses great challenges for management and IT Infrastructure. Design teams will need to leverage cloud compute resources to meet project schedules and optimize infrastructure costs.

This paper covers how to use IC Manage Holodeck for accelerating EDA jobs on the Google Cloud Platform (GCP) and reducing IT infrastructure and cloud computing costs with instant hybrid cloud bursting.

Challenges Created by Data Explosion (Source Data, Derived Data and Metadata)

Managing size of databases approaching petabyte

Increasing cost of IT infrastructure for compute and storage

Technology Trends Providing Opportunities for Addressing Broken EDA Paradigm

HPC architectures with 100+ threaded CPUs and massive CPU cycles/second

Infinite number of cloud CPUs across the globe

High speed commodity NVME storage with reduced latency and I/O scale out

Introduction

Over the past decade, IC and system designs have grown multifold in complexity, from size of the chips and impact of physical phenomena on their performance to the sheer number of engineers involved in bringing the design to production. It is not uncommon to find hardware design teams, each with 100's of engineers spread across multiple companies and continents working on a complex system project. Source data, derived data and metadata (process, status, analytics) have caused data explosions to petabyte size as designs move toward 7nm and below.

Limitations of EDA workflows based on legacy tools, proprietary flows and shared storage with complex dependencies are creating major challenges to improve or evolve the IT infrastructure. The high cost of capital acquisition and insufficient IT infrastructure are creating major challenges for semiconductor companies to progress.

These trends will continue to have a significant impact on how design teams effectively manage their projects.

HPC architectures with massive CPU cycles/second, access to nearly infinite number of Google Cloud Platform CPUs, and the availability of commodity high speed storage such as NVMe with reduced latency combined with IC Manage's Holodeck software provides an opportunity to improve the EDA paradigm centered on shared NFS-based environments. Augmenting the nearly infinite compute capacity provided by GCP, Holodeck provides an integrated data orchestration platform that efficiently brings data closer to compute across GCP clusters, regions and countries.

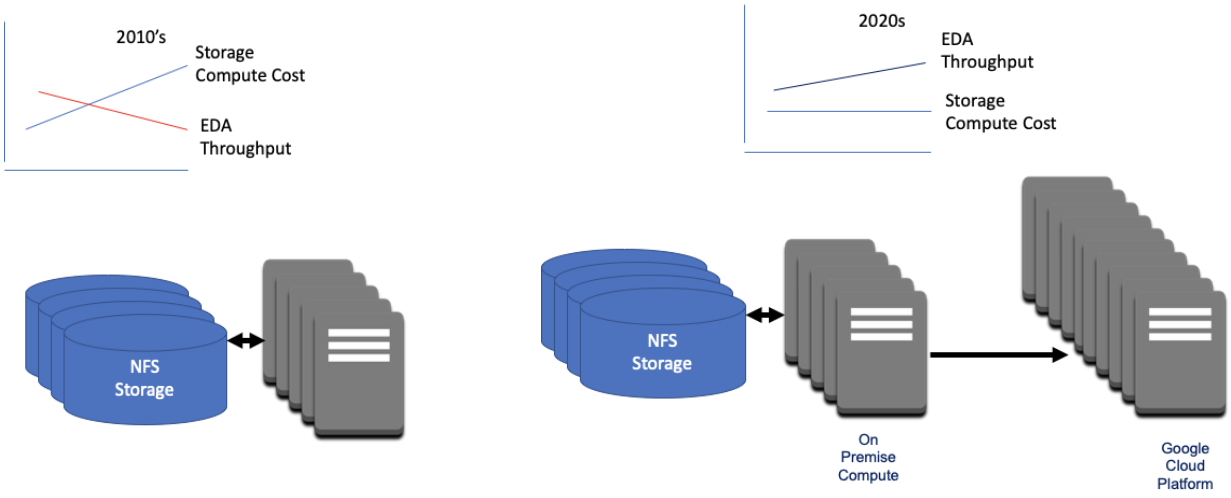

EDA Job Acceleration – Extending IT Infrastructure with Google Cloud

Traditionally, IC design methodology uses compute farms to access shared data on large NFS file servers. However, this approach imposes high capital costs for fixed storage assets and data I/O bottlenecks that dramatically reduce compute efficiency. Storage I/O bottlenecks also result in higher cost of running EDA tools.

Why Costs of Running EDA Tools and NFS Storage is Rising

- Compute farms accessing shared data on large NFS filers results in data I/O bottlenecks, dramatically reducing compute efficiency
- Workspaces of source and generated data is approaching multiple terabytes to petabyte, resulting in significant capital cost for NFS storage systems
- Sheer size of workspace data prevents parallelizing EDA workflows, forcing EDA jobs to run serially

IC design workspaces contain both source data and data generated from EDA workflows. Nearly 90 percent of the workspace data is derived from multiple stages of the design process such as functional verification, layout, place and route, timing and power analysis, optimization, and regressions. The combined workspace for a full SoC can approach 100's of terabytes and will continue to grow exponentially with each technology node. The sheer size of workspace data prevents us from easily parallelizing EDA jobs, since there is not enough on premise compute resources nor excess storage to enable parallelization. GCP compute resources enable parallelization by making 1000's of compute nodes available for as long as needed and IC Manage Holodeck virtualizes cloud storage for jobs to run in as many parallel workspaces as needed.

Available cloud caching solutions provide some storage efficiency, but limited I/O read and especially I/O write scalability that is needed to keep up as compute capacity is scaled out. IC Manage Holodeck uses the NVMe storage embedded in GCP compute instances as high-speed local file caches along with peer-to-peer network communication to achieve higher read and write performance. The Holodeck Peer-to-Peer networked cache allows all nodes in the GCP region to dynamically share data as compute capacity is added.

Faster I/O Performance and Lower IT Infrastructure Cost with Google Cloud and Holodeck

- Holodeck's scale out storage architecture leverages GCP's large network of inexpensive local NVMe to deliver 100's of GB/sec I/O Performance
- Peer-to-peer network communication along with high-speed local file caches improves aggregate read/write performance and eliminates duplicate NFS storage cost

On-Demand GCP Hybrid Cloud Bursting with IC Manage Holodeck

The capacity of on-premise compute alone is not sufficient in keeping up with the demands of todays' designs and aggressive project schedules. Increasing on-premise capacity can be a very time consuming and expensive proposition, sometimes taking up to 6 months to bring needed compute capacity online. Design teams want to run hundreds to thousands of jobs immediately, creating a pent-up demand for elastic compute power. Taking advantage of the GCP's unlimited compute power, state of the art servers and networking can enable high performance elastic compute to instantly add compute resources. However, moving a design workspace with up to 100s of terabytes to GCP is not a trivial task.

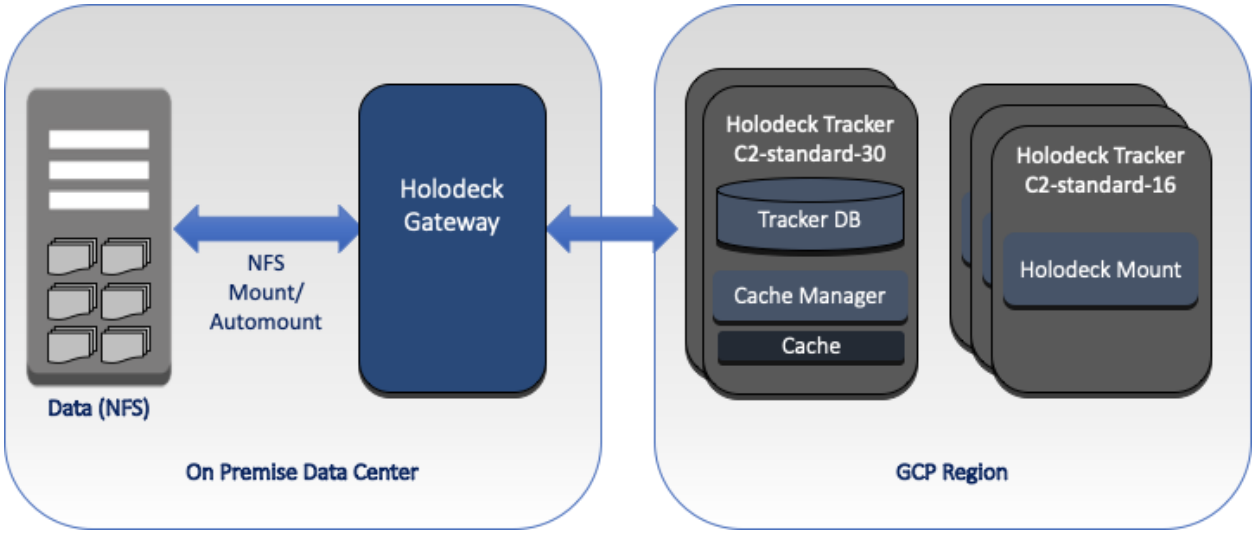

Chip design and verification workflows are very complex with multiple, interdependent NFS environments, usually comprising 10s of millions of files, spanning 100+ terabytes to petabytes of storage when we include all the EDA tools and scripts, foundry PDKs, design data. Enabling GCP to run all the workflows and jobs can be daunting without a means to easily and efficiently synchronize these 10's of millions of interdependent on-premise files across domains. Using software such as rsync or ftp can be very slow and costly to move 100's of Terabytes of data to the cloud, essentially eliminating any ability to gain fast access to GCP compute resources. While very large NFS storage is available in GCP, the cost of storage adds up quickly, especially if duplicate copies of the design are maintained on both on-premise and in GCP. Trying to determine the correct subset of design data to copy to GCP is extremely hard due to all the interdependencies between the data and legacy workflow scripts. More importantly, on-premise EDA tools and workflows are built on an NFS-based shared storage model, where large numbers of compute nodes share the same data. Just like on-premise NFS, scaling shared cloud storage to support 1000's of compute nodes is very challenging.

Challenges of Running EDA Workflows in the Cloud

- Complex workflows with interconnected NFS environments comprising of 10s of millions of files across 100s of terabytes of storage with no means to easily and efficiently synchronize across domains
- Maintaining duplicate copies of the design on both on-premise and cloud environment resulting in significant cloud storage cost
- Infrastructure disparity between on-premise EDA tools and workflow built on NFS-based shared storage model and cloud's local block storage model

The combination of IC Manage Holodeck and GCP compute nodes provide a cost effective and time efficient solution for on-demand hybrid cloud bursting of existing on-premise workflows which run identically in the cloud, enabling elastic high-performance compute by taking advantage of virtually unlimited GCP compute power with local NVMe storage. The ability to automatically run jobs in the cloud using unmodified on-premise workflows helps preserve millions of dollars and person-hours already invested in on-premise EDA tools, scripts, and methodologies. Engineering jobs can 'cloud burst' to GCP during times of peak demand, providing capacity in a transparent fashion.

CUSTOMER EXAMPLE: STATIC TIMING ANALYSIS

Metric	Holodeck	Without Holodeck	Holodeck Reduction
Time to Start 1st Job	10 min	19 days	99.9%
CPU	35 hours*	60 hours	41%
Storage	2TB	200TB	99%

IC Manage Holodeck peer-to-peer cache fabric scales out I/O performance by simply adding more compute nodes as peers. Storing the needed portions of the design workspace in the cache fabric eliminates the need for duplicate NFS storage in the cloud. Once caches in the GCP are hot, they can be fully decoupled from the on-premise environment, reducing storage cost and minimizing NFS filer performance bottlenecks. Since Holodeck works at the file extent level, data is transferred with ultra-fine granularity, delivering low latency bursting. **Only bytes that are accessed by the are forwarded to GCP compute nodes** and selectively filtered design data deltas are written back to on-premise storage for post processing, if necessary. Not having to duplicate all on-premise data significantly reduces the cloud data storage costs, sometimes by as much as 99%. By using temporary cache storage, all data disappears leaving no trace, at the moment that the GCP environment is shut down after jobs are completed. GCP compute nodes can also be shut down upon job completion to ensure that there are no bills for idle CPUs.

On-demand Hybrid Cloud Bursting Delivers True Elastic Compute

- Automatically run jobs in the cloud without any modifications to existing on-premise workflows, preserving millions of dollars and person-hours already invested in on-premise EDA tools, scripts and methodologies
- Transfers and stores only the portions of the design workspace data on demand in a peer-to-peer cache fabric eliminates need for duplicate NFS storage in the cloud, reducing as much as 99% of cloud data storage cost
- Provides immediate and transparent CPU capacity during times of peak demand

SCALE OUT EXAMPLE: VERILOG SIMULATION - 2,015 Regression Tests

Metric	Holodeck (50 Nodes*)	Without Holodeck (1 Node on Premise, Upload Time to Cloud)	Holodeck Reduction
Time to Start 1st Job	8 sec	19 min	99%
Runtime	6.5min	210 minutes	98%
Storage	74MB	1.3GB	94%

Conclusion

A scale out architecture utilizing GCP compute instances with local NVMe as a peer-to-peer cache network enables high performance elastic compute for EDA tool acceleration. IC Manage Holodeck hybrid cloud bursting provides immediate and transparent access to the GCP's virtually unlimited capacity. Virtually projecting on-premise storage to the cloud enables instant bursting, reducing cloud storage cost as well as eliminating NFS filer bottlenecks.

The combination of IC Manage Holodeck and GCP provide a unique solution for semiconductor design teams to extend their compute resources quickly and transparently to overcome the complexities presented by design, verification and integration.