

A Blueprint for EDA Infrastructure

By

Anthony Galdes, Solution Architect, Global Design Platform, IC Manage, Inc.

Steven Klass, Senior Software Architect, Envision, IC Manage, Inc.

ABSTRACT

The great expansion of data for semiconductor design poses great challenges for management and IT Infrastructure. To be successful, design teams will need to do the following:

- Utilize a data orchestration platform to globally synchronize IP, development processes, and oceans of design data
- Leverage multi-cloud computing while avoiding run-away costs
- Adopt analytics for machine generated design data to make better product and resource decisions

This paper covers methods for improving designers' efficiency and team collaboration, best practices for effective IP development and reuse, tools for accelerating EDA jobs and reducing IT infrastructure and cloud computing costs with hybrid cloud bursting, and how to accurately assess project progress and resource utilization.

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

- Managing size of database approaching petabyte
- Increasing cost of IT infrastructure for compute and storage
- Gaining accurate insight into project progress

Technology Trends Providing Opportunities for Addressing Broken EDA Paradigm

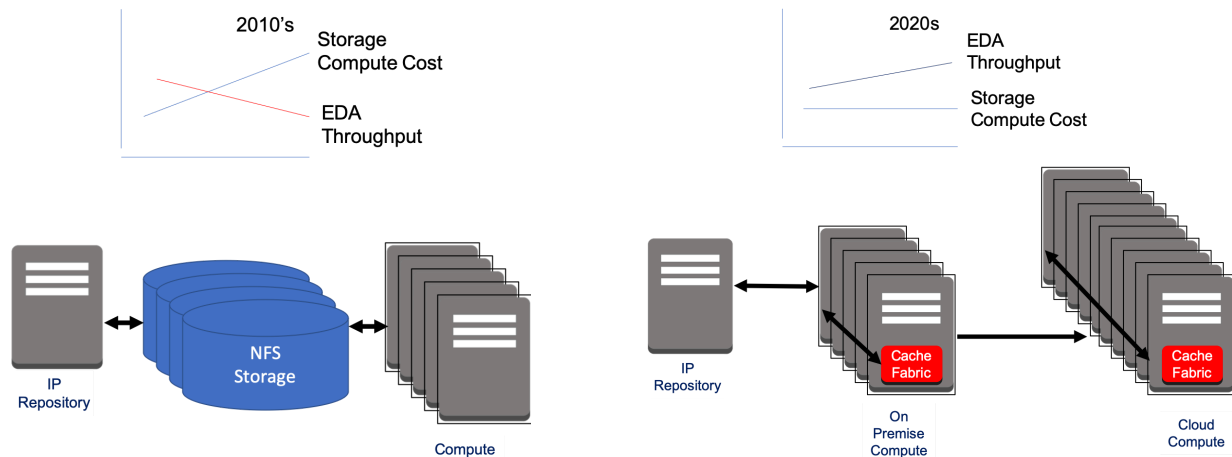
- HPC architectures with 100+ threaded CPUs and massive CPU cycles/second
- Infinite number of cloud CPUs across the globe
- High speed commodity storage such as NVMe with reduced data access latency and I/O scale out

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 or more 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.

Controlling cost and obtaining reuse of IP has become critical in managing complexity of such magnitude. In addition, 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.

As we move into the next decade, these trends will continue to have a significant impact on how design teams effectively manage their projects.



With the latest technology trends, such as HPC architectures with massive CPU cycles/second, access to nearly infinite number of cloud CPUs, and the availability of commodity high speed storage such as NVMe with reduced latency, there is an opportunity to improve the challenged EDA paradigm centered on shared NFS-based environments. In order to seize the opportunity and leverage the availability of these technologies, we need an integrated data orchestration platform that efficiently brings data closer to compute across clusters, regions, clouds and countries.

A data orchestration platform:

- 1) Maximizes the use of metadata to speed up database and data operations
- 2) Enables efficient, flexible and cloud-compatible design data management
- 3) Creates Effective IP development and reuse
- 4) Accelerates EDA jobs with IT infrastructure and cloud computing cost reduction
- 5) Delivers project progress insight while using existing EDA tools and workflows

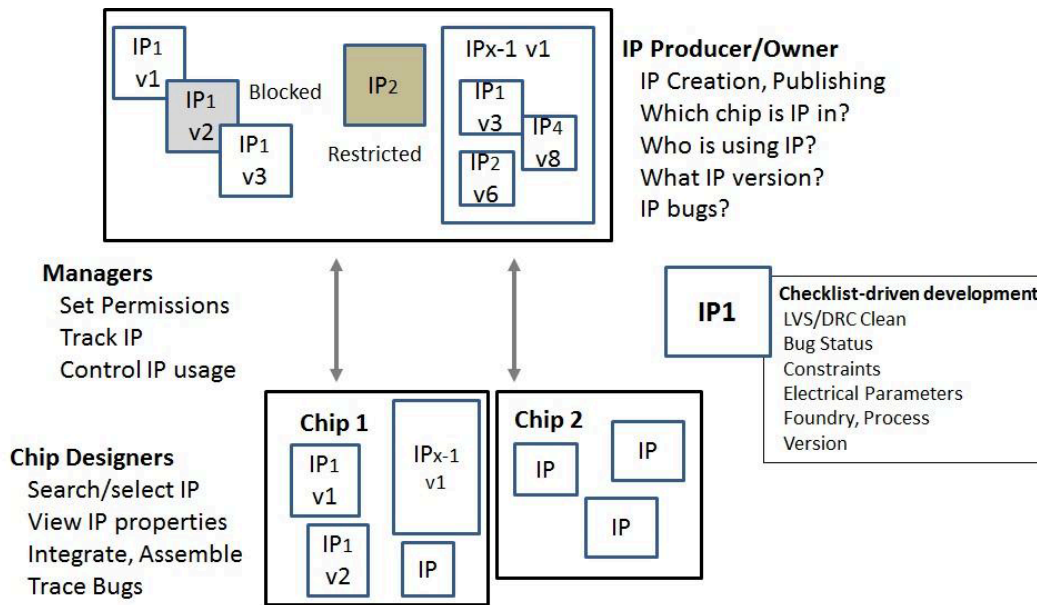
Efficient, Flexible and Cloud-Compatible Design Data and Configuration Management

Design teams are faced with the complexity of managing, verifying and integrating 1000s of design objects with many complex and sometimes hidden interdependencies. They need to make sure that the right data is being delivered to the right engineers to complete their design tasks and successfully tape out. To address this daunting challenge requires a cloud-compatible multi-site design and IP management system that delivers best-in-class performance with flexibility, scalability and solid reliability to enable worldwide design and verification collaboration. The ability to dynamically track, control and distribute design data, IP, foundry data and bug dependencies, while updating configurations with live metadata can significantly improve designer efficiency and team collaboration.

Tracking and viewing design module usage across both linear revision space and derivative space bi-directionally allows designers to easily detect and propagate changes of an object through all the designs using that object. Compared to conventional RCS-style branching, there is no need to diff or manually update all the locations when the object is reused. Change traceability is a must have in mission critical applications such as automotive (ISO 26262) and medical devices (ISO 13485).

Key Requirements for Improving Designer Efficiency and Team Collaboration

- **Traceability:** Dynamically track, control and distribute design data, foundry data, bug dependencies and update configurations with live metadata
- **Revision Management:** Unified system across all data types including RTL, binary data, EDA database and software



Making and organization's intellectual property (IP) accessible for reuse by all design teams is critical for competitiveness. A cohesive system addresses the needs of hardware design and software development that enables management of revisions across all data types including RTL, binary data, EDA databases and software source code. It allows developers to rapidly publish and integrate their IP into existing flows and trace bugs and bug dependencies for effective IP management thus providing access to reusable IP blocks developed at any design site with view of full history in real-time.

Effective IP Development and Reuse

Creating differentiated products within a narrow market window requires effective IP management and reuse during the design process. Semiconductor designs today incorporate 1000s of IP components, developed internally, acquired from third parties or from foundry provided data. These IPs have complex interdependencies, across the original IP, its multiple versions, individual hierarchical sub-modules of the larger IP blocks, and related PDK data.

Highly efficient IP creation and reuse requires best practices that include an extensive IP catalog, flexible IP packaging and tracking, linking of the IP repository to other collaboration applications, security policies and automated check-list driven development and acceptance.

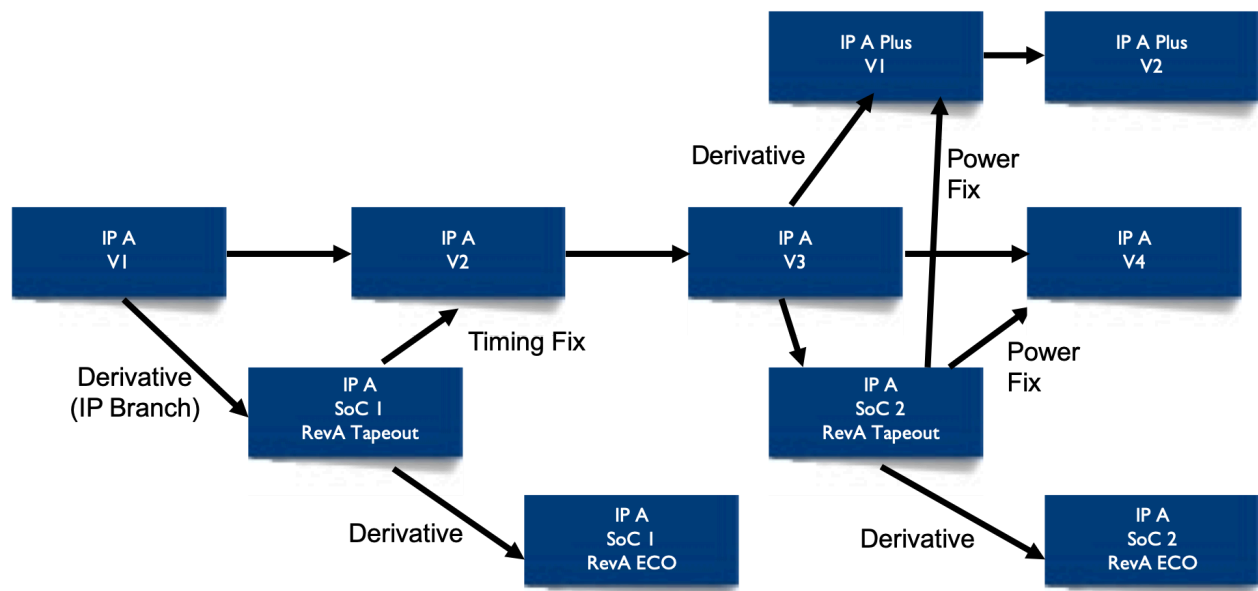
Ideally, each design module can be defined as a reusable IP. Each of these blocks is comprised of mixed data types such as RTL, LEF/DEF, SPF, GDSII, OpenAccess, Liberty, and test benches as well as metadata such as release readiness, coverage, power consumption and target technology. An IP repository organized by data type provides an infrastructure where developers can consistently place their data files rather than using ad-hoc file structure which may be difficult to package, release and make immediately available to remote sites or in the cloud.

Using an assembly methodology to capture IP dependencies from the beginning allows for accurate, real time tracking and reporting of where each IP version is used across the project. Every IP in the repository with its metadata properties is tracked and automatically updated throughout the IP development lifecycle. A simple Bill of Material (BOM) approach containing static information, typically created at the end of project, lacks early visibility and IP dependency and compatibility awareness.

Best Practices for Highly Efficient IP Creation and Reuse

- IP repository organized by data (RTL, layout, testbenches, specifications, etc.) and metadata (release readiness, coverage, power consumption, target technology, etc.) enables immediate access to IP package on remote site and/or in the cloud
- Extensive IP catalog (repository) with rich metadata properties that are automatically tracked and updated throughout the entire IP development lifecycle
- Hierarchical configuration management with virtual cloning of the IP through pointers allows each version of IP to be tracked separately and independently while maintaining the parent-child relationship between original IP and its derivatives
- IP repository linked with other collaboration infrastructure such as bug tracking, requirements management and verification analytics for precise project tracking and resource management
- Continuous Integration (CI) for merging of new and/or changed IP code into the IP repository with automated verification at a sufficient frequency for better ongoing team communications and accelerated conflict resolution.

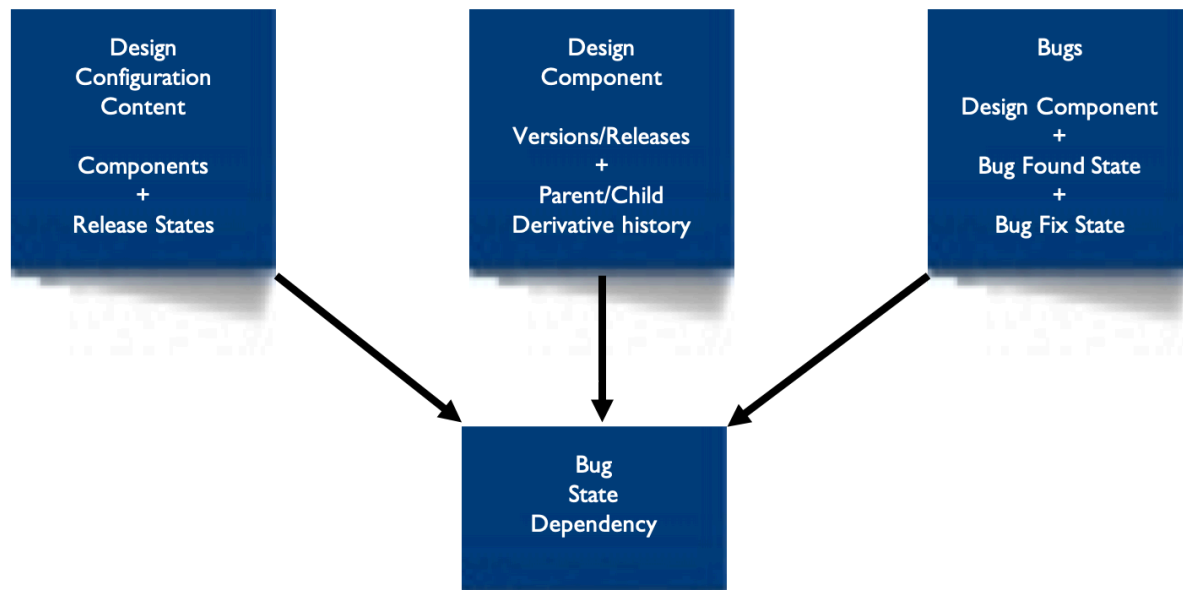
If IP cannot be directly reused, an IP derivative methodology utilizing virtual cloning of IP metadata ensures traceability and management of interdependencies between various IP versions in use across multiple chips and design teams with automated version conflict detection and resolution. Creating virtual copies of the IP as derivatives allows each version of IP to be tracked separately and independently while maintaining the parent-child relationships between the original IP and its derivatives. This hierarchical configuration enables efficient management and propagation of IP updates in both directions. It eliminates the need to create physical copies of different IP versions in multiple repositories with unrecognizable orphans and reduces the associated storage space.



Linking the IP repository with other collaboration infrastructure such as bug tracking systems, requirements management and verification analytics enables precise project tracking and resource assignment. Close linkage of the IP repository and bug tracking system helps identify bugs and bug fixes associated with each IP block, for both original and its derivatives. Design and verification teams can view and trace bug history and IP updates for every IP in the design. Project managers can view rollups of their live bug dependencies for the specific IP releases in their design to be assured that the eventual tapeout doesn't contain known bugs. The ability to mix and match different versions of the same IP in different design blocks eliminates the need to respin all blocks to use the latest version of IP.

Leveraging Automation Framework

Bug or defect tracking is a key part of any IP development lifecycle. When a defect is detected, it is important to be able to accurately and easily replicate the bug and verify that it is fixed in its original configuration. Tightly linking the IP repository with a bug tracking system ensures that all bugs and bug fixes associated with each design element or IP block are automatically recorded. This allows design teams to identify bugs, trace their dependencies and propagate IP fixes across all versions and to consumers of that IP throughout the project lifecycle.



Continuous Integration (CI) is being used by many software and hardware design teams to enable ongoing communication among team members, reducing the number of merge conflicts, and accelerating conflict resolution. Merging new or changed IP code into the IP repository with automated verification and sufficient frequency allows designers to notice possible errors and start correcting them immediately. Delayed or periodically scheduled check-ins can make design or development conflict resolution much more difficult to achieve. As the IP development and verification process progresses, the IP status is updated automatically, allowing managers to measure progress against specific milestones.

To make all this work across a large, global enterprise, the IT infrastructure needs to deliver reliable, on-demand data availability around the clock, while reducing the cost of storage. It also needs to provide disaster recovery and automatic failover using high availability replicas. To keep up with the volume of data updates and database transactions while supporting up to

1000's of users across the globe and in the cloud, high performance configuration management, scalable revision control, and real-time remote site replication are key factors in managing exponentially growing project data size.

Adopting above practices can dramatically reduce the engineering time spent developing and integrating internal and third party IP into the design flow, as well as minimize the time it takes to achieve confident verification of the IP in the context of the entire SoC, IC and FPGA design and will enable efficient collaboration.

EDA Job Acceleration – IT Infrastructure and Cloud Computing Cost Reduction

Traditionally, IC design methodology uses compute farms to access shared data on large NFS file servers to address the capacity challenges of growing design databases. 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 Cost of Running EDA Tools and NFS Storage is Rising

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


IC design workspaces contain both managed (source data coming from the IP repository) and unmanaged (data generated from EDA workflows) data. 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. As workspaces get larger, more time is spent syncing and re-syncing workspace data. The sheer size of workspace data prevents us from easily parallelizing EDA workflows, pushing EDA jobs to run serially on the original data set. These

serial processing flows present major challenges for effective scaling of compute resources as well as security management.

Scale out storage architectures are needed to deliver significantly faster I/O performance compared to expensive scale up NFS file systems. If we can leverage a large network of inexpensive local NVMe and unused switch bandwidth rather than overloading a small number of expensive storage devices, we can achieve scale out I/O of 100s of GB/sec. Using peer-to-peer network communication along with these high-speed local file caches will improve aggregate read/write performance and eliminate duplicate NFS storage cost. Peer-to-peer networked cache allows all nodes in the grid to dynamically share data, in parallel, for fast cache fills and job execution, either on-premise or in the cloud.

Solutions for Significantly Faster I/O Performance and Lower IT Infrastructure Cost

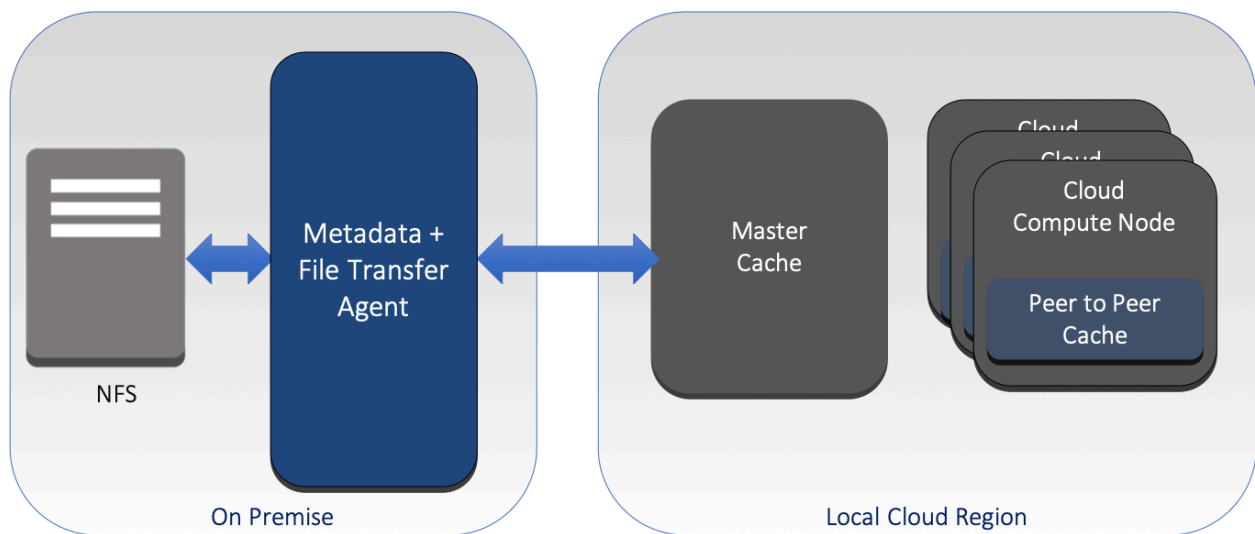
- Scale out storage architectures leverage a large network of inexpensive local NVMe and unused switch bandwidth to deliver 100s of GB/sec I/O performance
- Peer-to-peer network communication along with high-speed local file caches improves aggregate read/write performance and eliminates duplicate NFS storage cost
- On-demand virtual cloning of managed and unmanaged data enables massive parallel job execution with no incremental storage cost



Virtual cloning of managed and derived data can deliver an unlimited number of parallel, independent workspaces in a few seconds with zero storage overhead (zero copy clone). This eliminates the need to create physical copies, lowers the cost and time of continually syncing and re-syncing large workspaces and reduces read, write and filestat performance bottlenecks. The key to virtual cloning is separating file data from descriptive metadata while maintaining one (1) copy of invariant data, even if there are hundreds of cloned workspaces.

On-Demand Hybrid Cloud Bursting

As already shown, scale out I/O based on peer-to-peer network communications with local file caching significantly improves compute I/O performance and reduces NFS disk storage capacity requirements. But the capacity of on-premise compute alone is not sufficient in keeping up with the demands of today's 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 cloud's unlimited compute power, state of the art servers and networking can enable high performance elastic compute. However, moving a design workspace with up to 100s of terabytes to the cloud 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, managed and unmanaged data. Enabling the cloud 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, essentially eliminating any ability to gain fast access to cloud compute resources. The cost of cloud storage can also add up, especially if duplicate copies of the design are maintained on both on-premise and cloud environments. Trying to determine the correct subset of design data to copy to the cloud 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. In contrast, the cloud primarily uses local block storage which is typically accessed only by one host at a time. This infrastructure disparity can become a major challenge to address and overcome.

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

On-demand hybrid cloud bursting enables existing on-premise compute farm workflows to be run identically in the cloud, enabling elastic high-performance compute by taking advantage of virtually unlimited compute power of the cloud and the wide availability of local NVMe storage on cloud compute nodes. 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' during times of peak demand, providing capacity in a transparent fashion.

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 cloud are hot, they can be fully decoupled from the on-premise environment, reducing storage cost and minimizing NFS filer performance bottlenecks. Working at the file extent level allows for ultra-fine granular data

transfer, delivering low latency bursting. Only bytes that are accessed are forwarded into the cloud 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 environment is shut down after jobs are completed. Cloud compute nodes can 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

A scale out architecture utilizing local NVMe as a peer-to-peer cache network enables high performance elastic compute for EDA tool acceleration. Hybrid cloud bursting provides immediate and transparent access to the clouds' 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.

Accurate Insight into Project Progress

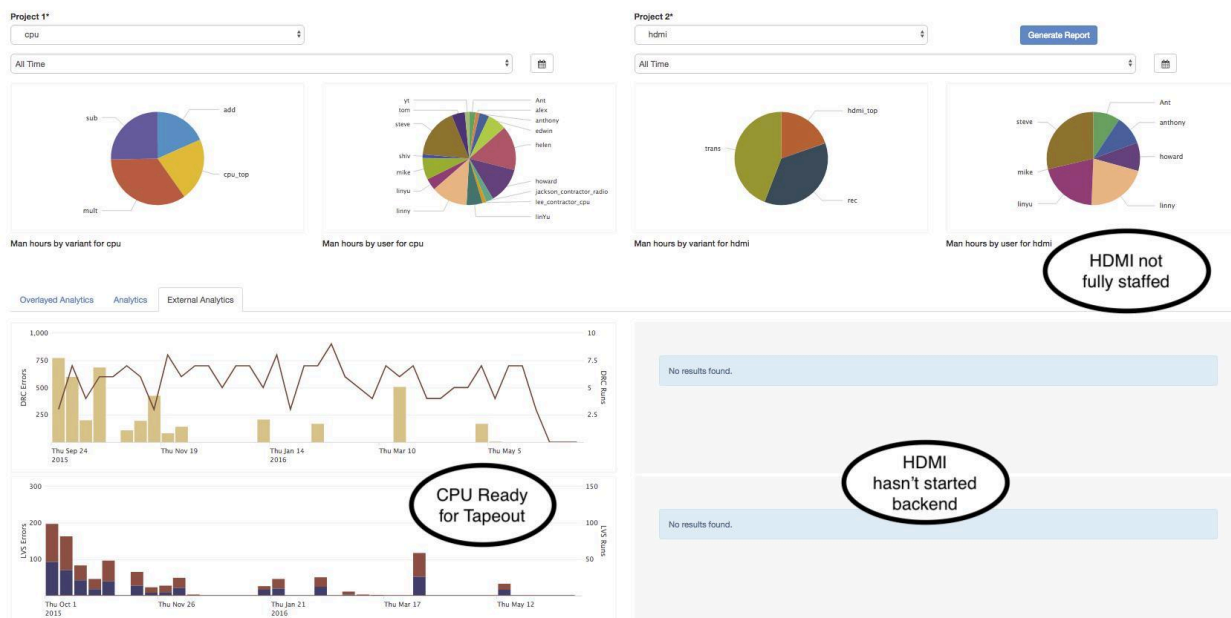
Being able to analyze project progress, identify resource optimization opportunities, predict tapeout schedule and prevent IP theft are critical success factors for today's design projects. The standard methods of top-down project schedule tracking and status reporting are reactive and can only address the problem when it becomes obvious and make resolution much harder and costlier.

Complex IC projects generate terabytes of design change and tool run data per day. Source code edits, verification runs, synthesis, power, timing, place and route and DRC/LVS tasks all generate

large amounts of data that are used by individual engineers in their day to day work, but little to no aggregation or detailed analysis of these results takes place. Aggregating these results and correlating across project components and domain activities can provide an accurate assessment of project progress. Without this large-scale analysis, engineers and project leads spend many hours compiling and analyzing the state of their project progress to determine their chances of meeting the project objective. But the existing, siloed analysis methodologies based on isolated log data can be very slow and hide risks in the process.

Besides tool result logs, design and verification teams today do not have access to file activity logs, leaving them without a clear picture of what is happening with their design data. Being able to instrument a peer-to-peer local cache network will provide detailed log information for every transaction on every file by every tool. This massive amount of activity data can be automatically stored and analyzed to provide accurate design and verification activity analytics, resource utilization and can also be used for tapeout schedule prediction.

By tracking all key resources (compute, storage, EDA license, active users) needed to develop complex IP and SOC and analyzing every facet of design activity in real-time, engineers and managers can see the progress of their project at any level of granularity. Metrics such as release readiness state of IP and SOC based on verification status, QoR, DRC errors or timing and bug rates, and the amount of resources consumed to develop, verify and integrate IP and SOC enables engineers to make real-time data-driven decisions.



For example, applying big data and machine learning techniques to structure this data helps to accelerate verification schedules by providing immediate view of key metrics against customized targets for regression runs, debug convergence and coverage progress. Customizable, interactive analysis provides reports of verification progress versus plans, teams and sub-systems allowing management to focus on project progress and optimize resource allocation. Interconnecting unstructured verification data with structured IP repository design data creates a hybrid database that can be analyzed to identify problem areas and link them to related design activity or changes across SOC/IP hierarchy. This allows engineering teams and management to identify project bottlenecks and understand their root causes.

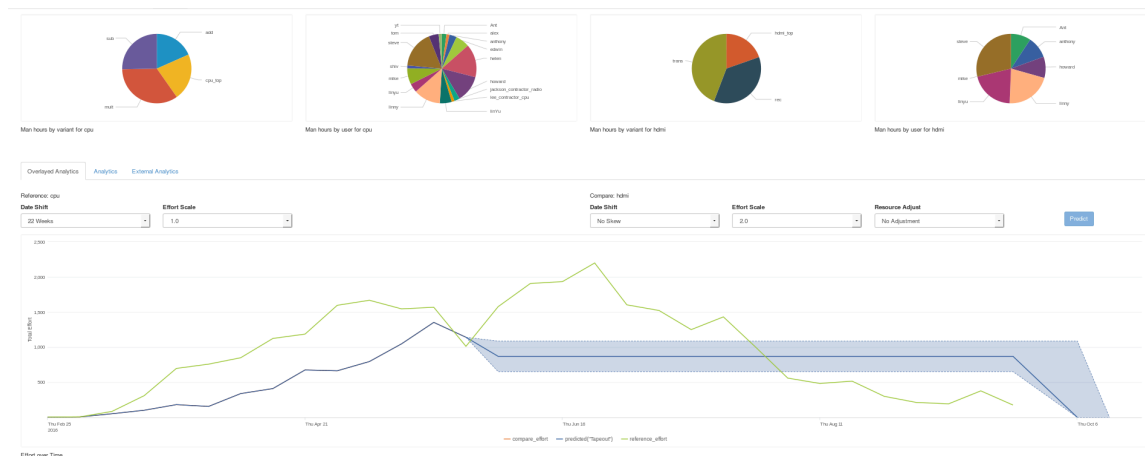
Real-time Insight into Project Progress, Resource Allocation and Tapeout Projection

- Accelerate verification schedule by utilizing big data and machine learning techniques on structured data and providing immediate view of key metrics against customized target
- Identify verification bottlenecks and understand their root causes by interconnecting unstructured verification data with structured design data
- Drill down for resource utilization metrics, helping engineers and managers prioritize compute and human resources
- Predict tapeout date and required compute and human resources by accurately aggregating, filtering and analyzing billions of data points and comparing against historical pattern
- Continuously monitor user login, database access, IP address and amount of data processed to detect potential IP theft in real-time and revoke access

Further drill down of the data can reveal key metrics for measuring resource utilization such as human capital, EDA tool license allocation and the level of effort required to produce and qualify IP. Based on the information, engineers and managers can prioritize compute and human resources for best results. This information can be fed into future IP make versus buy decisions based on cost of development captured with these techniques.

Taking this analysis one step further, predicting tapeout schedules requires accurately aggregating, filtering and analyzing billions of data points to identify which tasks and design components are on time or behind schedule. Every engineering process that consumes or

generates data can be a key performance indicator (KPI) for tapeout and needs to be extracted, compared with historical patterns and tuned throughout the design flow to estimate the completion date well in advance. This allows engineers and managers to not only predict the likely tapeout date but also determine human, license and compute resources that will be required to meet the schedule and adjust resources early enough to reduce schedule slip.



IP is the most valuable asset of today's semiconductor companies and IP theft is a serious and growing problem which leads to a reduction of competitive advantage and lost revenue and profits. Often IP losses occur due to an inability to track internal IP data access, insufficient internal controls and external theft. To stop potential IP theft before it occurs requires continuous monitoring of user logins, database access, network addresses and amount of data being accessed so that any unusual patterns or user credential compromises can be detected in real-time. If attempted IP theft or cyber-attack is detected, the file system access should be immediately revoked for suspicious user accounts without impacting other team members.

Successfully managing design, verification and integration complexities associated with data explosion in the 2020's requires new solutions that can

- 1) Accelerate EDA applications compute cycles and storage I/O operations
- 2) Enable I/O scale out and hybrid cloud bursting low latency and minimal storage footprint
- 3) Perform advanced analytics to accurately measure all important design tasks and resources and enable intelligent project design making

It starts with the adoption of the Blueprint for EDA Infrastructure, an integrated platform that enables efficient design database management, effective IP development and reuse, EDA job acceleration, cloud adoption and accurate project progress insight.