cādence®

Accelerating SoC Time to Market with Cloud-Based Verification

By Michael A. Lucente, Product Management Director, Cadence

This paper discusses the growing use of cloud and hybrid cloud environments among semiconductor design and verification teams. The schedule and efficiency benefits seen by verification teams using cloud are specifically highlighted, due to the considerable compute requirements associated with verification of advanced node SoCs, and the significant impact verification has on the overall SoC project schedule. The readiness of public cloud environments for use in semiconductor design and verification workflows is discussed, along with factors to consider when choosing EDA technology for use in the cloud. Cadence® offerings for self-managed and fully managed EDA cloud solutions are also outlined.

Contents

Introduction

Over the past several years, the use of cloud-based technology has become commonplace in our everyday personal and professional lives. At home, we use cloud-based social media applications to communicate and stay updated on the latest events in the lives of our friends and family, while our children use the cloud to play online games with friends that live around the corner—or perhaps around the globe. At work, we use cloud-based applications to help manage customer relationships, employee time and attendance, sales forecasts, and so on.

For years, Cadence has recognized the applicability of cloud-based solutions to EDA, and the value they can bring to customers for system and semiconductor design enablement. At the same time, it was clear there was hesitancy among customers as they contemplated the use of public cloud technology for semiconductor design and verification. Security risks associated with storing design IP in the cloud, suitability of cloud datacenter technology for high-performance compute workloads, and robustness of public cloud offerings for EDA workflows were just a few of the customer concerns surrounding public cloud solutions when they were in their emergent phase.

Cadence has been successfully delivering cloud-based solutions to customers for years, while mitigating concerns associated with public clouds, using Cadence Hosted Design Solutions (HDS). HDS has enabled customers to use Cadence products through a fully managed hosted environment. For nearly a decade, HDS has been hosted on a private cloud architected and managed by Cadence. This capability has been expanded and leverages the use of public cloud resources and Palladium® emulation technology to provide a true EDA-as-a-Service environment.

Growing Demand for Cloud-Based Verification

There has been a significant increase in interest and adoption of cloud-based EDA solutions in recent years, particularly in the functional verification area. Numerous factors are driving this, including:

- The massive number of simulation regressions associated with verification of advanced node designs

- Continued increases in time to market pressure on system-on-chip (SoC) projects

- The need to contain costs and infrastructure demands (space, power, cooling) associated with on-premise datacenters

- Increased focus on the high-performance compute (HPC) market, including EDA, on the part of cloud and cloud ecosystem vendors

- Increased confidence in public cloud solutions among semiconductor companies

A closer look at some of these factors, along with their relationship to one another, will help further explain why the increased attention on the cloud for verification.

Verification and SoC project schedules

Time to market continues to be a critical issue for SoC development projects. The hardware and software development costs associated with a complex SoC can exceed $1B, and failure to have the product ready at the beginning of a market window can have a significant revenue impact. Missing the market window entirely can render a product worthless, and potentially impact the viability of a company. Given time-to-market pressures faced by SoC development teams, it's natural that they are continually looking for ways to shrink schedule. Verification is a logical area to focus on when trying to reduce the schedule, as market research has consistently found that verification consumes the largest percentage of development time on most projects.
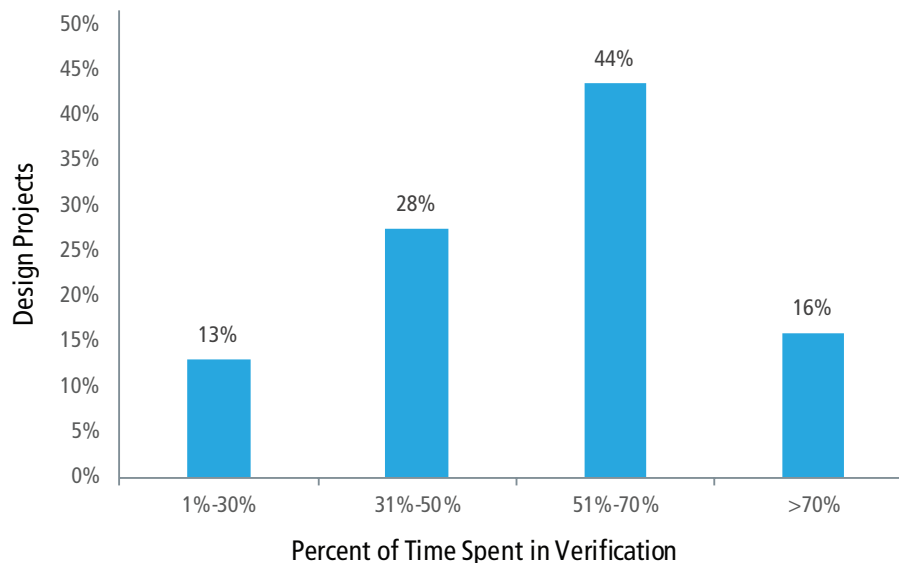

Figure 1: Allocation of ASIC project time to verification

Data collected in an electronics industry survey in 2016 (Figure 1) found that, on average, 55% of ASIC development time is spent on verification[1]. Data from similar surveys conducted in prior years shows that this number has remained consistent over time. Moreover, the number of verification engineers working on SoC projects has grown at nearly 3X the rate of design engineers. This significant rise in the allocation of headcount to verification has resulted in simply keeping pace with growing verification complexity. As the complexity of designs continues to grow, and verification challenges along with them, it raises obvious concern among SoC project teams about their ability to maintain schedule through the addition of headcount alone.

SoC verification pushes the limits of on-premise datacenters

Approximately 59%[1] of all ASICs designed today are verified using constrained-random simulation. If we exclude small chips targeting the Internet-of-Things (IoT) segment, and look solely at SoCs, the percentage is significantly higher. Constrained-random simulation, together with metric-driven/coverage-driven verification, has been shown to result in first-pass silicon success in over 70%[1] of projects, which explains its widespread adoption. This approach also delivers a significant productivity boost to engineering teams by leveraging the compute power of server farms to generate vast amounts of stimulus, as opposed to labor-intensive directed test methodologies.

The reliance of state-of-the-art verification on server farm compute power, combined with the significant impact of verification on SoC project schedules as outlined above, results in the following:

• SoC project schedules being accelerated or delayed depending on availability of datacenter resources available to the project

• Verification engineering teams pushing on-premise datacenters to their limits, but still finding their requirements aren't being met

Why cloud versus on-premise only solutions?

The dynamic outlined above doesn't fully explain why cloud use for EDA workflows is growing. In theory, companies could simply ramp spending for on-premise datacenters to stay comfortably ahead of project needs. However, this approach would not have the elasticity and on-demand scalability to achieve the desired outcome in an efficient and effective manner.

Compute resource needs on large development projects are not only difficult to predict, there is also variability in compute requirements throughout the project lifecycle. Early in a project, when verification teams are building their verification environment, compute demands are typically light. Mid-project, once IPs and subsystems have stabilized and simulation regressions are underway, tens of thousands of small or medium size compute nodes may be needed. As the project nears tape-out, full SoC simulations are taking place at the RTL and gate level, likely requiring fewer but much larger compute nodes. Identifying bugs in the latter stages of a project can restart this cycle again and again—hence the interest in cloud and hybrid cloud environments.

Consider the scenario illustrated in Figure 2(a). The on-premise datacenter investment is ramping well ahead of project needs, resulting in under-utilization of compute resources and inefficient use of capital. One could argue that compute requirements will eventually catch up to the capacity of the datacenter, but by then it's likely that the servers purchased are no longer state-of-the-art.
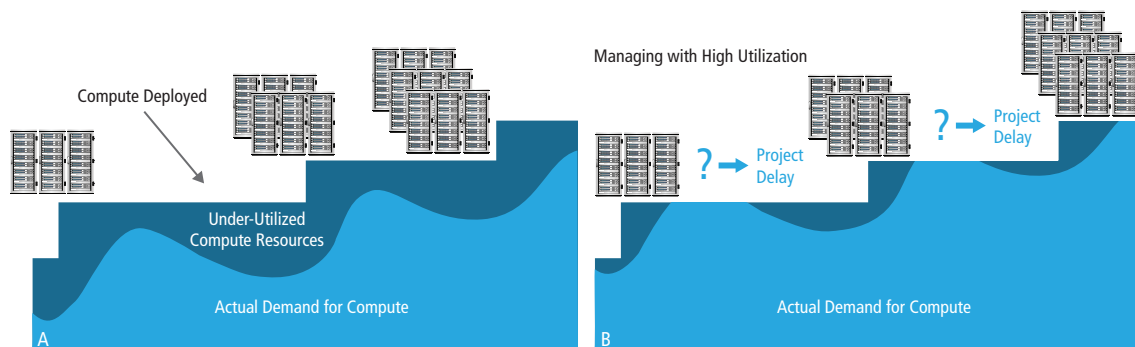

Figure 2: On-premise datacenter efficiency and scalability challenges

A more typical scenario faced by verification teams is illustrated in Figure 2(b). Teams working on verification of advanced node designs using current methodologies can and generally will consume all compute resources made available to them. When using a metric/coverage driven constrained random approach, access to more compute power equates to schedule savings. Conversely, when access to compute becomes a bottleneck, schedule delays result. It's typically very difficult to address compute bottlenecks in a reactive manner on a given project through scaling of on-premise datacenters. Assuming there is sufficient space, power and cooling to handle more servers and storage, it still takes time to order hardware, receive shipment, allocate IT resources, configure and install machines, etc.

Use of cloud and hybrid cloud solutions eliminates the inefficiencies and scalability issues associated with strictly on-premise solutions. Figure 3 illustrates a hybrid cloud approach, where project compute needs are largely handled using an on-premise datacenter.
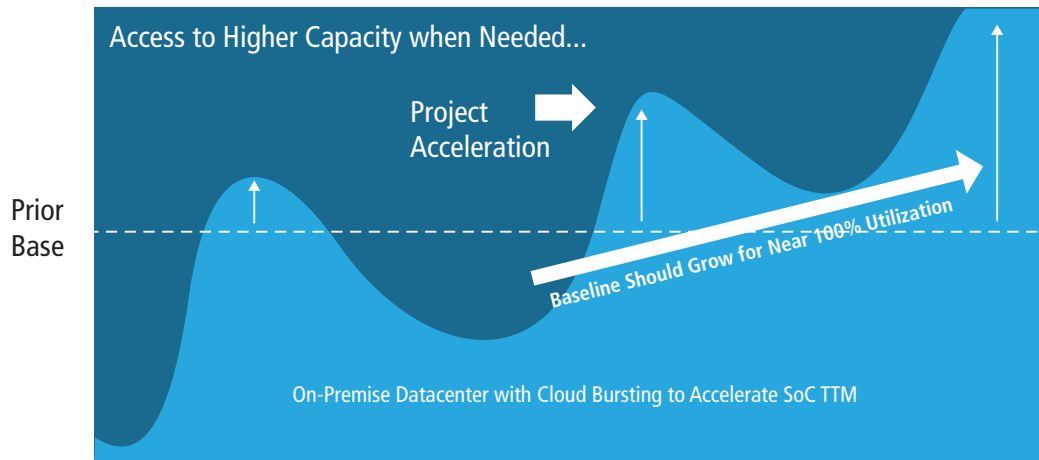
Figure 3: Addressing peak project requirements using hybrid cloud approach

Peak needs are addressed by periodically "bursting" onto the cloud when compute requirements and schedule demands dictate. Using this approach, workloads associated with peak needs are moved to the cloud at a job- or complete project-level, allowing flexible and efficient scalability to meet demanding schedules. If space and power availability allows for it, the on-premise datacenter continues to grow over time to meet steady-state user needs.

Cloud Readiness for EDA

It has been estimated that the cloud high-performance computing (HPC) market will grow at a compound annual growth rate of approximately 20% and reach USD 10.8 billion by 2020[2]. Cloud and cloud ecosystem providers have taken notice and are placing greater emphasis on building solutions that meet the requirements of the HPC market, including EDA. Their outbound messaging and collateral also reflects their intent to win market share in this space.

Amazon Web Services (AWS) published a white paper targeting EDA users that describes how to optimize semiconductor design and verification workflows on AWS[3]. The paper, available on the AWS website, makes recommendations regarding compute and storage instance choices for use with various EDA tools used in the design and verification process. It also speaks to the security of AWS customer IP in their cloud environment. Dell EMC also recently published a white paper targeting EDA audiences, in which they describe the benefits of their Isilon scale-out network appliance storage (NAS) solution for semiconductor development workflows[4]. Dell EMC also blogged recently about their work with Google to bring the Isilon NAS solution to the Google cloud platform (GCP). In the post, the author states that organizations "…can now deploy performance intensive file-based workloads, such as automotive design simulation, EDA chip design, and genomic sequencing, on Isilon Cloud for GCP.[5]" Microsoft is also ramping their HPC footprint with acquisitions to help with HPC cluster orchestration (Cycle Computing) and high-performance storage (Avere Systems).

The increased attention to HPC on the part of cloud providers has helped ensure that an EDA optimized environment is available to users in the cloud. It has also helped boost confidence among semiconductor design and verification teams that the cloud community is working to understand the requirements and workflows associated with EDA tools and deliver solutions in the cloud that address their needs.

Evaluating SoC Verification Solutions in the Cloud

Verification teams looking to the cloud to help address their compute and time-to-market challenges should consider several important factors as they evaluate EDA cloud-based technologies and vendors.

Before partnering with an EDA vendor, it's important to be mindful of not only the short-term needs of a current project but also potential future technology requirements that may arise as project and company needs evolve. For example, picking a solution because it fills a short-term simulation need, and finding later that you are unable to address a critical verification IP (VIP) or simulation acceleration requirement, can be devastating to your project schedule. Selecting a vendor that provides you with access to a broad range of cloud-based hardware and software technologies provides you flexibility when building your verification environment in the cloud and eliminates the risk of encountering unforeseen technology requirements that can't be fulfilled.

Another important consideration when evaluating an EDA cloud partner is their level of experience addressing variable compute requirements. Selecting a vendor that has experience hosting customer environments in the cloud, and has experience bursting onto the cloud for its own peak operational needs, gives you a partner with insights and processes necessary for a successful cloud-based project. These can—and should—include:

• Knowledge of best practices associated with building a secure environment in the cloud

• Product validation processes to ensure cloud readiness of products

• Experience selecting compute and storage instances in the cloud, based on EDA workload characteristics

• Resources to facilitate cloud orchestration

In addition to the considerations above, one should also be looking at their EDA cloud partner's current capabilities and plans for delivering solutions in the future that are optimized for cloud and hybrid cloud environments. It is critical that they understand the characteristics of these environments, as well as the technology and vision to maximize project efficiency and productivity in this landscape.

Cadence Verification Cloud Solutions

Using expertise gained through 20 years of hosting experience, Cadence offers customers industry-leading, production-proven cloud solutions that are customizable, scalable, and meet rigorous security standards recognized worldwide. Through a broad portfolio of cloud offerings that includes both customer-managed and Cadence-managed alternatives, Cadence helps customers spanning start-ups to Fortune 500 companies address their variable compute, time-to-market, and cost requirements.

The fully managed Cadence Cloud Hosted Design Solution enables customers across the semiconductor industry to achieve their IC design and functional verification goals. Employing patented technology and a finely-tuned architecture, the Cloud Hosted Design Solution delivers an EDA-as-a-Service solution that meets variable compute resource needs without the constraints, overhead, and learning curve associated with going to the cloud alone.
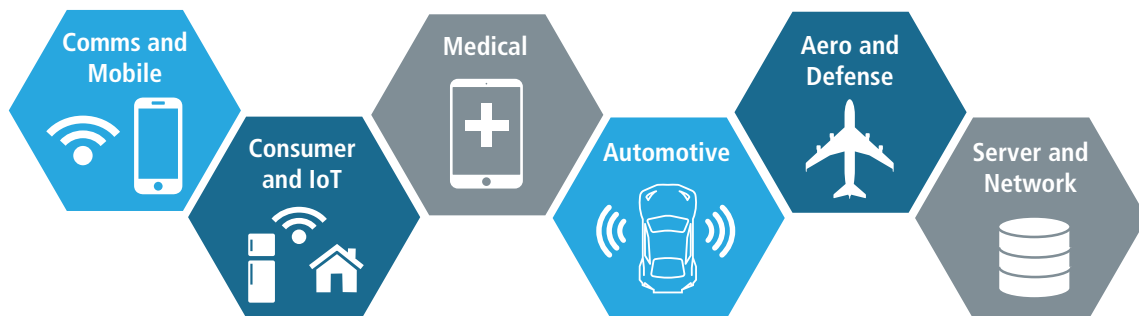

Figure 4: Broad spectrum of customers using the Cadence Cloud Hosted Design Solution

The Cadence Cloud Hosted Design Solution builds upon the native security of infrastructure-as-a-service (IaaS) vendors through additional security applications, best practices, and testing. The IaaS providers we partner with leverage world-class technology that has passed security standards used globally, including:

• FIPS 140-2 Level 2

• ISO 27001, 27017, and 27018

• DoD Security Requirements

• MLPS Level 3

• MTCS

The Cloud Hosted Design Solution follows the Cloud Security Alliance's recommended controls and routinely conducts independent penetration testing to comply with changes in the security landscape proactively.

In addition to offering the full portfolio of Cadence software products, the Cloud Hosted Design Solution offers customers cloud access to the Cadence Palladium series emulation platforms. Palladium Cloud is a unique offering that enables cloud access to Cadence's industry-leading emulation technology. Existing customers with on-premise emulation hardware use it when peak project needs exceed available emulation capacity, thereby avoiding project

schedule delays. New users with no prior emulation experience find it a great way to receive immediate access to emulation technology. Palladium Cloud, powered by the Cloud Hosted Design Solution, supports the full range of emulation use models, including simulation acceleration, in-circuit emulation, and hybrid emulation, enabling customers to accelerate their schedule at multiple points in the project.
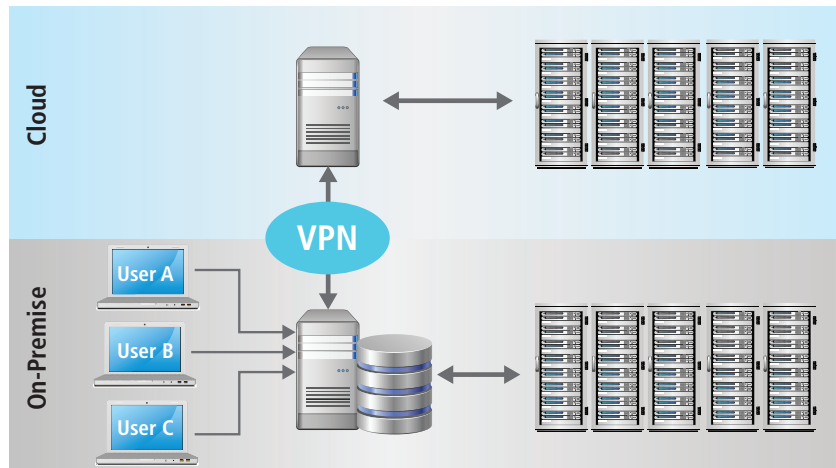

Figure 5: Augmenting on-premise emulation capacity with Palladium Cloud

For customers looking to leverage the scalability and elasticity of the cloud but prefer to manage the environment on their own, Cadence offers robust cloud-ready products. Using best practices gained from experience providing hosting services, as well product development activities taking place in the cloud, Cadence has structured processes to ensure that products are cloud-ready and validated on all major public cloud platforms. For verification teams looking to accelerate their project schedule by leveraging greater compute capacity, the results can be signif-icant.
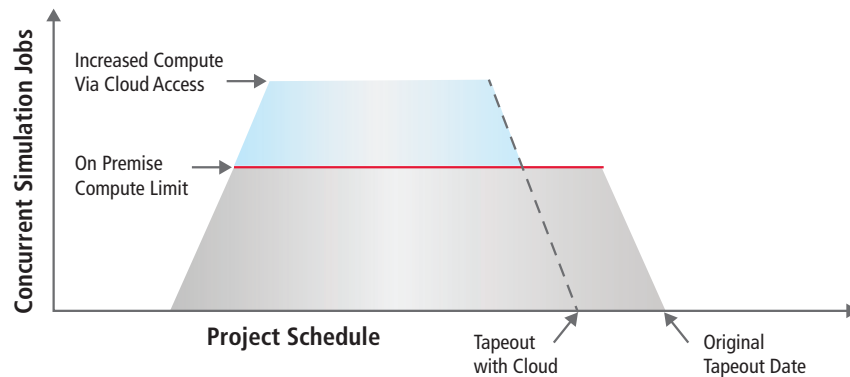

Figure 6: Accelerating project schedule using hybrid cloud simulation approach

Figure 6 illustrates the schedule benefits of having greater access to Cadence Xcelium® best-in-class simulation technology and additional compute capacity during peak project cycles. It is important to remember, however, that executing a project in the cloud requires access to a full suite of cloud-ready verification tools, such as JasperGold® Formal Verification Platform, Indago™ Debug, and Verification IP. These are tools that have not only been validated for use in the cloud but also have been optimized to make project teams more productive when working in a cloud or hybrid cloud environment.

In addition, the massive compute scalability available in the cloud makes it imperative that users be able to leverage best in class technology to help with grid management and data collection/analysis, such as the vManager™ platform.
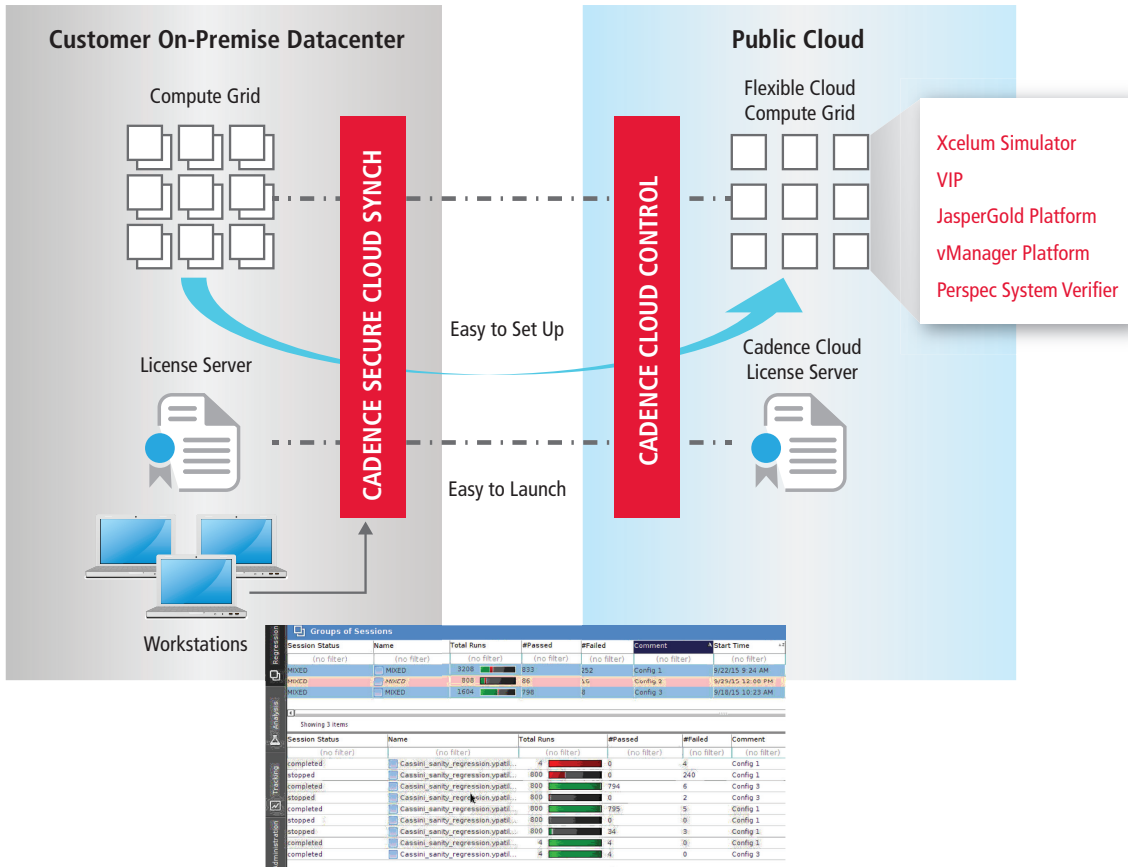
Figure 7: Cadence verification solution cloud

Conclusion

Time to market continues to be a critical issue for SoC development teams, and market research has consistently found that verification consumes the largest percentage of development time on most projects. Shortening the verification cycle by leveraging more compute resources is an effective approach, given the dependency of state-of-the-art verification methodologies on massive compute grids. Cloud and hybrid cloud environments are available today that have the scalability, elasticity, and security to deliver this additional compute capacity in an efficient and effective manner. With years of cloud hosting experience and a broad offering of cloud-ready, cloud-optimized products, Cadence is uniquely positioned to help customers successfully deploy EDA technology in the cloud, through fully managed or self-managed solutions.

References

1. The Wilson Research Group, "*The 2016 Wilson Research Group Functional Verification Study*", 2016

2. MarketsandMarkets.com, "Cloud High-Performance Computing (HPC) Market by Service Type (HPC IAAS, HPC PAAS, Data Organization and Workload Management, Clustering Software and Analytics Tool, Professional Service and Managed Service) - Global Forecast to 2020", November 2015

3. AWS, "*Optimizing Electronic Design Automation (EDA) Workflows on AWS*," October 2017. [Online]. Available: https://d1.awsstatic.com/whitepapers/optimizing-electronic-design-automation-eda-workflows-on-aws.pdf. [Accessed 11 June 2018].

4. Dell EMC, "*Dell EMC Isilon NAS: Performance at Scale for Electronic Design Automation*", May 2017. [Online]. Available: https://www.emc.com/collateral/white-papers/h13233-wp-performance-at-scale-for-eda.pdf. [Accessed 11 June 2018].

5. Dell EMC, "*Dell EMC Isilon & Google Cloud Platform: Bringing Uncompromised Scale-Out NAS to the Cloud,*" 30 April 2018. [Online]. Available: https://community.emc.com/community/connect/anz/blog/2018/04/30/dell-emc-isilon-google-cloud-platform-bringing-uncompromised-scale-out-nas-to-the-cloud. [Accessed 6 June 2018].