

# The Journey to Design in the Cloud



## INTRODUCTION

For many years, software developers have been building and delivering products in the cloud. The journey continues as semiconductor companies transition to designing chips in the cloud to innovate faster, leaner, and more efficiently.

The reasons to move design to the cloud are compelling but along the way, we will need to shatter some myths and illuminate a new way forward that harnesses the power of the cloud.

# TABLE OF CONTENTS

<b>CHAPTER 1</b> WHY CLOUD? WHY NOW?	<b>/ 03</b>
<b>CHAPTER 2</b> BARRIERS TO ADOPTION	<b>/ 07</b>
<b>CHAPTER 3</b> ADOPTION MODELS	<b>/ 10</b>
<b>CHAPTER 4</b> OPPORTUNITIES TO DO THINGS DIFFERENTLY	<b>/ 12</b>
<b>CHAPTER 5</b> THE WAY FORWARD	<b>/ 13</b>

# WHY CLOUD? WHY NOW?

For a long time, IC design teams have resisted a move to the cloud for hardware development based on several myths, but **three key drivers are now coming together as a perfect storm** to make moving to the cloud a must:

## Design Complexity


The combination of software as a driver for innovation and the use of hardware to bring new ideas to life is driving a new level of systemic complexity. Designs now combine multiple technologies into one highly sophisticated package, requiring a holistic analysis of the entire system. This is the design paradigm that will fuel semiconductor advances for the foreseeable future. This shift requires exponentially more compute and [Electronic Design Automation \(EDA\)](#) resources.

## Artificial Intelligence

AI is [entering design tools and workflows](#). This drives further multiplicative requirements for flexible (i.e., unlimited) access to compute and EDA resources.

## HPC in the Cloud

Cloud service providers (CSPs) have scaled [high-performance computing \(HPC\)-optimized infrastructure](#), availability, affordability, and capacity to handle these workloads.



## The Security Myth

A key issue that originally stalled any serious moves was the semiconductor industry's reluctance to place valuable IP onto someone else's hardware, with limited trust in the security capabilities of CSPs.

This trend to keep all IP in-house and "safe" is understandable; IP is the crown jewel for most semiconductor companies, after all.

The widespread use of the public cloud has spawned a substantial investment in security infrastructure on the part of the major providers. [Microsoft Azure](#) invests about \$1 billion each year in cloud security. [Google](#) will spend \$10 billion over the next five years on cloud security. There are many initiatives like this throughout the industry.

Beyond the security infrastructure delivered by cloud providers, the companies that host cloud applications also provide a layer of application-specific security. Synopsys, for example has a deep domain knowledge in [application security](#) and delivers substantial added protection. Examples include data classification, data segregation, auditing, monitoring and access control.

You can learn more about the [comprehensive security layers Synopsys provides here](#).

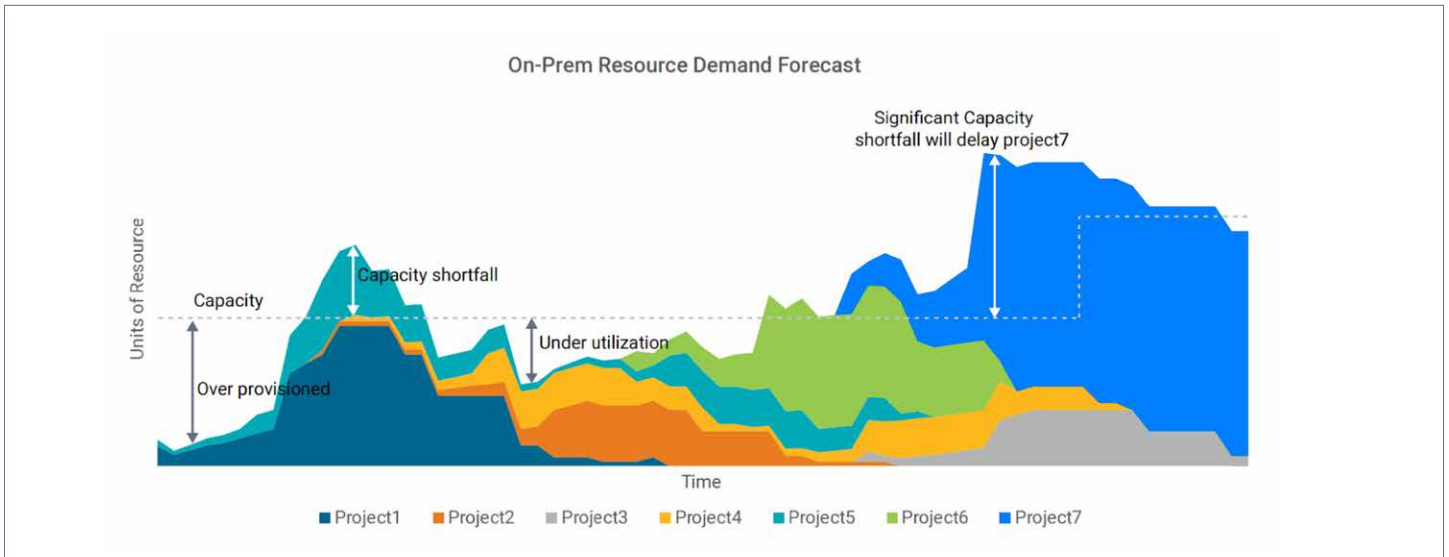
It turns out that CSPs and their partners have invested far more than any single company is able to invest in security, so "safest" usually turns out to be the cloud option.

## The Predictability Myth

The second issue is concern for predictability of resources. Hardware engineering teams don't trust that the cloud can deliver the required capacity when they need it. They prefer the predictability of knowing what capacity they have available in their internal compute and storage resources and being assured that it will be there when they need it because it is not being shared with an unknown pool of users.

This would be fine in a world of perfect capacity forecasting where it is certain that all engineering workloads will fit within given capacities. But in practice, internal engineering teams are constantly competing for shared [on-prem](#) resources and someone must make a priority call on which projects win when a crunch arises.





Furthermore, the compute requirements of actual design projects are characterized by significant peaks and valleys, as shown in the figure on the next page. What is the optimal on-prem configuration to support this very real situation?

Of course, you can invest and expand your on-prem resources to meet burst capacity requirements, but capacity expansion is typically non-agile, and in one direction only. So the on-prem estate only grows, it never shrinks when demand wanes. In fact, capacity availability is what the cloud excels at, thanks to the huge expansion in capacities globally by the large CSPs.

## The Affordability Myth

The third myth is around cost. On-prem IT teams always believe they can deliver capacity more cheaply than the cloud, and owning hardware is essentially more cost effective than renting it over time, or at least the costs are better understood and more controllable. They worry that the cloud's infinite capacity will encourage engineering teams to consume compute and storage in an uncontrollable fashion with runaway costs.

However, IT teams will need to provision enough on-prem capacity to cope with peak demands, meaning expensive on-prem hardware may sit idle for a percentage of time, with poor overall utilization. This is an on-prem cost that is often overlooked.

## The Ease of Use Myth

The fourth myth is around migration of IC development workflows from on-prem to the cloud. When workflows have deep dependencies on the architecture of the on-prem estate, lifting and shifting to the cloud seems like an insurmountable effort and cost barrier.

**So, some may conclude that, “If it’s not broken, don’t fix it.”**

In fact, investing effort to migrate existing workflows to the cloud has other benefits. As jobs become better encapsulated and less dependent on the target platform, one can progress to an environment where workflows are portable and can be run both on-prem or in-cloud in a seamless fashion.



# BARRIERS TO ADOPTION

**Let's look at some of the practical issues** preventing hardware engineering teams (in both large and small organizations) from successful cloud adoption.

Existing IC design, implementation, and verification workflows may have evolved with an unknown number of assumptions and dependencies on the compute and storage environment, making it difficult to migrate to the cloud. On top of that, there may be tool licensing assumptions that also prevent using the exact same workflow in an unconstrained cloud environment.

## Can I Use My EDA Tools in the Cloud?

The likelihood is that existing contracts may not have the provision to use EDA tool licenses in the cloud. So, whoever manages EDA contracts is going to have to talk to the vendor to establish the art of the possible. All the main EDA vendors offer multiple models to use licenses in the cloud. However, there may be some contract changes necessary, and you may in fact want to be able to operate some of your licenses on-prem as before, while operating others in a cloud environment.

**After all, there's no point having access to infinite compute in the cloud if you are then limited by EDA licenses.**

## Data and Storage

The cloud is not just about compute, it's also about storage. For some engineering workflows, the storage cost can be a significant factor. As a rule, creating large datasets in the cloud is not an issue in terms of availability, but cloud storage can be expensive. Not everything that is stored in the cloud needs to be in high-performance tier storage.



You can move less frequently accessed data into one of the slower tier 2 or tier 3 storage mediums provided by your CSP. There is also a cost associated with data transfer. It's usually free to upload, but downloads are metered and can contribute to extra fees.

Don't retain large volumes of intermediate data that can be reproduced easily. Perform all your data analytics in the cloud, something that the cloud is increasingly capable of doing, thanks to the emergence of modern big data analytics.

## Controlling Costs

The need for demand forecasting does not go away as you switch to the cloud. What does go away are the technical and physical barriers to capacity expansion, as short-term demands can be more easily met, and you need only pay for the services you consume.

### No need to maintain expensive capacity on stand-by.

However, as mentioned earlier, costs could spiral with runaway consumption if engineering teams no longer feel constrained by fixed capacity and see great opportunities to improve engineering quality and time to market without consideration for the cost ROI calculation. Therefore, built-in utilization analytics and budgetary management are essential parts of any cloud-based IC design workflow.

Cost models for the cloud are also not simple. There are many dimensions to the cost, such as the choice of provider, choice of services (compute and storage), choice of add-on services like [big data analytics](#), pricing models based on pay-as-you-go, pre-purchase plans, or spot-pricing models that can discount services significantly at times of low demand.

**On-prem estates are usually heterogeneous mixes of old and new hardware, fast and slow compute, small memory, and large memory systems. It is often difficult to determine if the right configuration of hardware is deployed based on the unique needs of each step in the design flow.**



## Cost Factors: On-Prem vs. Cloud

There's no one-size-fits-all cost model, as each organization has different needs. Keep these considerations in mind in your calculations.

	On-Prem Cost Factors	Cloud Cost Factors
Facilities & Equipment	Large initial investment and ongoing maintenance of old / new hardware mix	Costs vary based on usage and pricing models
Reliability / Lost Productivity	On-prem may actually result in more downtime for maintenance and upgrades	Cloud typically has more redundancy, meaning less downtime and fewer outages
Staff	Full IT teams required for maintenance and analysis	Fewer human resources required
Business Continuity Risk	Higher risk if all on-prem is in one location or more investment in multiple locations	More redundancy across the cloud means less potential impact

The great thing about the cloud is that all the costs associated with running a job can be accounted for on a single OPEX line. Remember that costs must be considered as the combination of compute, EDA tool licenses, storage, and data transfer.

## Migration Strategy

This "lifting and shifting" of workflows to the cloud necessitates that those workflows are encapsulated in a way that they can be sent to the cloud with all the job dependencies packaged up with the job. For teams looking to port in-house workflows, this may mean some analysis of the existing on-prem workflows to establish the I/O requirements and file dependencies of each job, and then possibly some re-architecting of the workflows to make them cloud-ready.

Further, where to perform the interactive elements of the workflow need to be decided, such as [debug](#). Will you run both batch and interactive workloads in the cloud, or only run batch in the cloud and use on-prem for all interactive debug jobs? Response times are the main consideration for interactive debug activities, such as waveform analysis.



## CHAPTER 3

# ADOPTION MODELS

There are several approaches to adopting a cloud design model. The particular approach will be tempered by **the size of your organization and the current hardware and software investment.**

## BYOL / BYOC

There are two scenarios here that should not be confused. Bring your own license (BYOL) uses existing on-prem licenses with your chosen CSP. That may be the case for larger businesses that already have established tool licensing agreements for their on-prem environment, and they now want to be able to use those same tools in a cloud environment. Users are still license limited according to their license investment level.

Bring your own cloud (BYOC) is a model which may be much more attractive to those users that have already established cloud capabilities with a chosen CSP, and where demands can be very bursty in nature. Licenses are not limited as it is a pay-as-you-go consumption model, similar to the CSP compute model. Metering and analytics are used to bill for usage and to allow for effective consumption budgeting. Other possible pricing models may be available that allow users to pre-plan and pre-purchase license hours at alternative cost points.

## SaaS

Software-as-a-service (SaaS) is another alternative where the user no longer needs to be concerned with the setup and running costs of using a public CSP since the tool vendor is hosting the application on the vendor's chosen cloud service. This is a familiar model for many modern applications that your organization is already using today; think Microsoft Office365 and Salesforce, for example. These are all SaaS solutions that are running on cloud services, where the user is not exposed to the complexities of the cloud service.



The IC development team only cares about running a workflow, such as simulation for example. Effectively this becomes “simulation as a service.” Similar to BYOC, the charging model could be metered/pay-as-you-go, or it could be based on a given capacity of licenses.

## Hybrid Cloud

For larger organizations with a pre-existing investment in on-prem capabilities, a hybrid approach is a popular strategy. Peak demands can be met by bursting capacity into the cloud for suitable workloads, while less portable workloads continue to be run on-prem with no impact. This more gradual migration to the cloud means that headroom becomes available in the on-prem estate for those workflows that require more time and effort to “cloudify.”

Eventually, you might evolve to a point where all workflows work equally well on-prem or on the cloud, at which point your users don’t need to care where their jobs run. At this point, the on-prem estate can be characterized as an on-prem or private cloud and the hybrid nature of the compute environment is abstracted away from the consumer.

### Which Adoption Model Fits Best?

	On-prem (BYOL)	All Cloud (BYOC)	SaaS	Hybrid Cloud
Scenario Summary	Larger business with significant on-prem investment	Smaller business or startup with established cloud capabilities where demands can vary with bursts and quiet periods	Smaller startups that want to use tools available on a vendor’s CSP	Business established with on-prem investment combined with cloud capabilities
Business Need	To be able to use existing on-prem tool licenses for the same tools in a cloud environment	To handle burst demands without breaking the budget	To run workflows in the cloud without having to manage the cloud	To allow for only specific usage in the cloud for effective consumption budgeting

# OPPORTUNITIES TO DO THINGS DIFFERENTLY

Innovation is the lifeblood of any R&D effort in the semiconductor industry. **The pace of change has been relentless and has been pushed by EDA tool advances** and customers presenting vendors with engineering challenges that would have seemed unimaginable only five years ago.

The advent of the cloud will be one of those major inflection points in history, allowing engineers to improve productivity, performance, and time-to-market.

Essentially, breaking free of the constraints placed on innovation by limited compute can offer opportunities to do things in a different way and it becomes an equalizer for small organizations to compete with larger ones. Many IC design engineering teams are accustomed to the constraints of capacity. This affects the overall quality of the final product and the time to market. What opportunities are missed by not being able to deliver the highest possible quality in a market-winning timeframe because there simply aren't enough on-demand compute resources and EDA licenses available to accelerate the delivery?

At the end of the day, the main resource constraint is the people. Engineer time is the most valuable and the most limiting factor. Engineers being blocked, waiting for lengthy batch compute jobs to complete is not an effective use of engineering talent. With capacity and availability constraints lifted, talented engineers can focus on the things that they do best: innovating.

Earlier we asked, "why cloud, why now?" and characterized the present situation as a perfect storm in which growth in systemic complexity drives compute and EDA demands exponentially, and modern AI-enhanced EDA technologies further increase this demand. Both CSPs and EDA vendors offer usage models that make cloud adoption affordable and scalable for IC hardware developers.

**What's more, if the competition is already exploiting this vast resource to deliver products faster, why shouldn't you do the same?**



# THE WAY FORWARD


Chip development in the cloud represents a way forward for an industry grappling with exploding computational demands along with continued time-to-market pressure. **From established design houses, to system companies, to start-ups, more chipmakers are migrating workloads to the cloud** to take full advantage of the faster time-to-results, enhanced quality-of-results and better cost-of-results that cloud-based design and verification technologies provide.

With [Synopsys Cloud](#), we're taking EDA to new heights, combining the availability of advanced compute and storage infrastructure with **unlimited access** to EDA software licenses on-demand, so you can focus on what you do best — designing complex chips, faster.

The SaaS model provides optimized workflows for every type of IC design, matched to the best hardware for the job. Up to now, only companies with vast flow development resources could provide this capability. Synopsys now makes this same comprehensive support available to all design teams.

The revolutionary [FlexEDA model](#) provides access to the entire catalog of Synopsys software. Reduce time-to-results from days to hours through cloud-scale elasticity using Synopsys' unique FlexEDA pay-per-use model. Get unlimited access to EDA licenses on-demand, enabling the scaling of EDA licenses and compute up or down in real time. You are freed from long procurement cycles and the need to know what tools are needed up-front. FlexEDA is available as both SaaS and BYOC.

Delivering cloud-native EDA tools and pre-optimized hardware platforms, an extremely flexible business model, and a modern customer experience, Synopsys has reimagined the future of chip design on the cloud to facilitate a new world of innovation — without disrupting proven workflows.



## Take Synopsys Cloud for a Test Drive

Synopsys technology transforms how people work and play. Let us power your design journey with cloud-based EDA solutions.

[Sign up to try Synopsys Cloud for free](#)