

# Hybrid Cloud Bursting: 5 Principles for Successful Deployment

By Dean Drako, President & CEO, IC Manage

## I. Introduction: Design in Cloud Expected to Double in Two Years

### I-1. Design & Verification in Cloud – Today & in Two Years

Today, 32 percent of companies already deploy cloud for design or verification to some extent today, as their security concerns have been addressed and the EDA vendors have begun adapting their business models. 68 percent still exclusively do on-premise design.

Organizations expect this ratio to flip in two years. **Nearly 70% expect to utilize the cloud by late 2023**; only 30% expect to continue to exclusively do design and verification on-premise.

### I-2. Cloud Workflow Options

Although it's somewhat oversimplified, there are two primary cloud deployments for design workflow today.

The first type is a pure cloud workflow.



The second type is a hybrid cloud workflow, which extends existing on-premise, compute-farm workflows to have identical flows in the cloud.




Let's look at the relative deployment of each type.

### I-2. Twice as Many Organizations Want Cloud Bursting over All Cloud

For companies who are currently doing some design in the cloud today, 10% use an all-cloud solution, while 22% of teams are doing hybrid cloud bursting.

Looking forward 2 years, hybrid cloud bursting is on track to continue to be the preferred method – and by the same factor of 2X.

Nearly half of organizations expect to be doing it. Let's look at why.



### I-3. Why Hybrid Cloud Bursting? On-Demand, Elastic Compute

Our systems & semiconductor industry has billions of dollars invested in their existing on-premise flows which incorporate EDA tools, scripts, methodologies & compute farms.

It's important financially & for schedule stability that they leverage these existing flows. But on-premise compute capacity has not been keeping up. IC design & verification teams spend a lot of time \*waiting\* for compute farm access.

A hybrid cloud implementation can deliver unlimited compute power on-demand. Development teams can continue to run their on-premise compute farms and then temporarily burst into the cloud during peak workloads to instantly run hundreds or even 1000s of jobs.

This enormous need for elastic compute power is especially true for compute intensive simulation, timing analysis, power estimation, DRC/LVS, library characterization, SPICE, and place & route runs.



## II. Hybrid Cloud Bursting Deployment: 5 Principles for Success

Hybrid cloud bursting requires precision execution. There are five key principals to deploying it and making it successful.


### II-1. Cloud Storage & On-Premise Architectures are Different

The first fundamental is that cloud storage architecture is different from On-premise storage. Native cloud storage is object based.

In contrast, on-premise EDA tools & workflows are built on NFS-based shared FILE storage, where large numbers of shared compute nodes and running jobs expect to see the same coherent data.



*On-premise: NFS-based shared file storage*



*Native cloud: object-based storage*


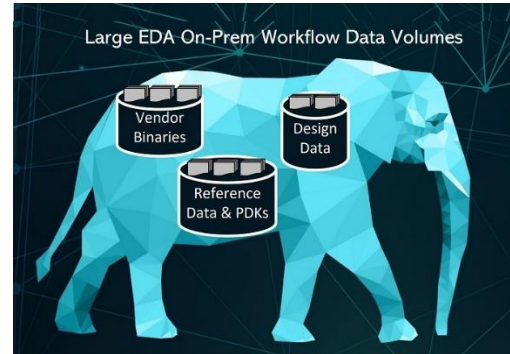
## II-2. Copying from On-Premise to Cloud is Expensive

Principle number two is that copying from on-premise to cloud is expensive.

We are designing enormous chips, with complex workflows made of interconnected NFS environments, usually comprised of 10's of millions of files, easily spanning 50 terabytes.

They include EDA vendor tools and installed trees, PDKs, third party IP, design databases comprised of timing, extracted parasitics, set-up and run scripts tying the data together, and configuration files.

Further, the design flows are extremely complex due to size and data interdependencies. So, it's also hard to isolate the exact data needed.

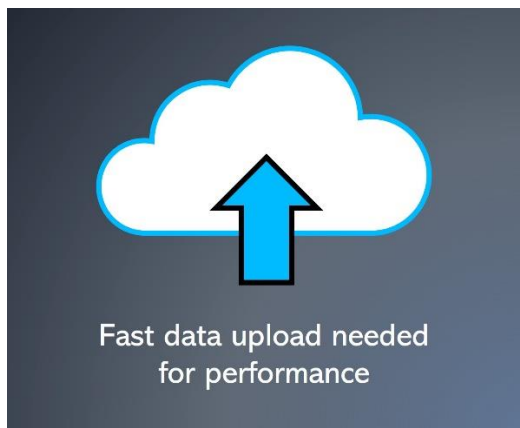


You must determine the correct subset of data to send to the cloud. It's extremely time consuming to determine all the dependencies.

The jobs break due to errors in getting the right data to the cloud. We've seen companies take 3 weeks to get just one block for one process node running VCS. Larger projects take months and months.

On top of this, the design and the workflows are constantly evolving so any cloud transfer becomes out-of-date very quickly. So, you must ensure the data in these 10's of millions of intertwined files is synchronized. RSYNC, TAR, and NFS just don't cut it.

## II-3. Fast Data Setup & Upload Times Critical




The next principle is that fast data setup & upload times are critical. Cloud vendors don't charge you for uploading your design data into the cloud. However, the time-consuming upload can negate the entire gain that the cloud might have provided.

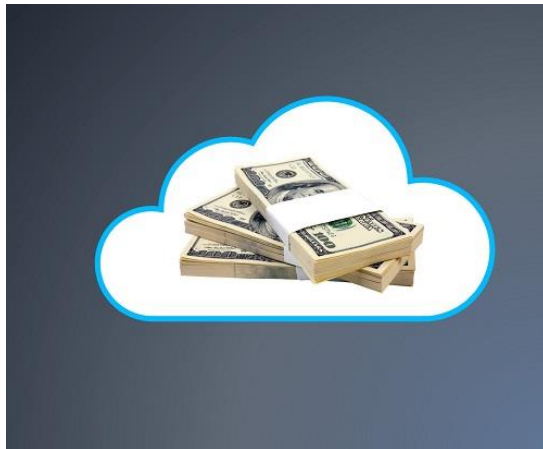
Rsyncing design changes or updating PDKs to the cloud can take hours to days, and lead to data inconsistency.

Multi-cloud environments, give teams more flexibility to take advantage of pricing. However, longer setup times will limit these environments.

The longer it takes to set up and tear down a single cloud environment, the less likely you are to do it.



#### II-4. Cloud Data Storage Adds Costs



Principal four is that cloud data storage adds costs. For cloud bursting to be financially workable, all the cost factors have to be well-understood.

You can buy compute cycles effectively, but if you are inefficient, your costs can easily spiral out of control. Cloud storage pricing may not seem very high until you realize your full environment can be 5 terabytes, which is not cheap to store for a month or longer.

If you leave lots of data in the cloud rather than deleting it, your costs can surprise you.


#### II-5. Cloud application performance is a success metric

The goal of hybrid cloud bursting is to get immediate access to additional resources. The faster your jobs complete, the more you can leverage the cloud's compute power and lower your costs

And yet it's common for application performance to drop substantially in the cloud — for a 50% or lower speed than on-premise. This is particularly true for I/O intensive jobs, because the cloud-compute may not have the I/O you need.

This minimizes the compute power advantage plus raises your compute costs. Thus, it's important to also consider

application performance for your cloud deployment.




### III. On-Demand Hybrid Cloud Bursting

So how do we manage all this? It is not well understood that running a job in the cloud actually requires only a small fraction of the workflow data. You don't need to duplicate all the data in cloud.


This makes the jobs highly suitable for virtual workflows. So, I will share how we can achieve optimal speed and minimum storage for on-demand cloud bursting.

The solution is to get only the data needed to the cloud, transfer it only once, and share it among multiple jobs. Below I review a hybrid cloud implementation method that is in use today. There are 3 short steps:


#### 1. Identify your NFS filer mount points.



#### 2. Replicate your on-premise filer mount points and create a virtual representation on a remote compute node.



3. Then transfer only the files you need, on demand, to the cloud.



By projecting your entire on-premise environment 'virtually' in the cloud, **your cloud workflow is indistinguishable from your existing on-premise workflow**

You have fast setup & data upload times, as there is no data duplication & sync. This approach also minimizes your cloud storage & cloud instance costs and gives you fast compute times.

#### IV. Conclusion: The Cloud – a “Metaverse” for Design & Verification Teams

With the cloud, we have our own “Metaverse” for Design & Verification.

However, instead of putting on goggles and entering fantasy worlds, we enter an exponentially more productive work environment.

And then use that environment to help us build a better future in the real world. It's going to be a great adventure for all of us.



#### ABOUT IC MANAGE HOLODECK

IC Manage Holodeck enables fast access to massive cloud-compute capacity without the overhead of data duplication. IC Manage Holodeck:

1. Virtually projects terabyte/petabyte NFS file systems in minutes.
2. Transfers only specific data segments required by the job running -- in real-time.

3. Delivers high-performance read and write; results can be written back to on-premise storage as file deltas.
4. Reduces cloud storage cost by over 90%; is up to 2000 times faster than Block/NAS/Object storage.
5. Eliminates application performance bottlenecks with I/O scale-out.

## ABOUT THE AUTHOR



**Dean Drako**

Dean Drako is founder and CEO of IC Manage, a design & IP management, data analytics, and hybrid cloud solutions provider. Drako is also founder and CEO of

both Drako Motors, where he launched Drako GTE, a 206-MPH, 1200HP, four-passenger electric luxury supercar, and Eagle Eye Networks, which has ranked on the 2021 “Deloitte Technology Fast 500” for North America for the past 3 years.

Drako was also founder, president, and CEO of Barracuda Networks, where he created the IT security industry’s first spam filter appliance and then grew the company to more than 140 products and 150,000 customers from its inception in 2003 through 2012. Barracuda went public on the NYSE. Drako was named one of Goldman Sachs “100 Most Intriguing Entrepreneurs of 2014.”

Drako received a Bachelor of Science in Electrical Engineering at University of Michigan, and his Master of Science in Electrical Engineering from UC Berkeley.

**BestTech  
Views**

[www.besttechviews.com](http://www.besttechviews.com)

Copyright 2022