

2018

Top 10 Strategies to Manage Cost and Continuously Optimize AWS

HyperGrid

🖱️ www.hypergrid.com

☎️ +1 855 786 7065

✉️ info@hypergrid.com



Table of Contents



01.

INTRODUCTION

p. 1 Introduction

02.

STRATEGIES

p. 2 **1** Start smart: provision the instance you need—no more, no less
p. 4 **2** Stay smart—use continuous, data-driven instance optimization
p. 5 **3** Look for life beyond EC2—embrace serverless computing
p. 7 **4** Leverage reservations to reduce costs—start with EC2
p. 9 **5** Model applications—not instances
p. 10 **6** Centralize governance of Reserved Instances
p. 12 **7** Anonymization is the enemy of accountability—tagging to the rescue
p. 14 **8** Cost governance starts with proactive cost-management
p. 16 **9** Autonomy does not have to mean unrestricted budgets
p. 18 **10** Build an organization-wide, shared responsibility security model

03.

SUMMARY

p. 19 Conclusion

Top 10 Strategies to Manage Cost and Continuously Optimize AWS

The public cloud has unleashed an unprecedented wave of creativity and agility for the modern enterprise. A great cloud migration has upended decades of established architecture patterns, operating principles, and governance models. However, without any replacement for these traditional controls in place, cloud spend inevitably rises faster than anticipated. If not addressed early in the cycle, this is often overlooked until it gets out of control.

Over the course of a few decades, we have created a well-established model of IT spending; to arrive at economies of scale, procurement is centralized and typically happens at three- to five-year intervals, with all internal customers forecasting and pooling their needs. This created a natural tendency for individual project owners to overprovision resources as insurance against unexpected demand. As a result, the corporate data center today is where the two famous laws of technology meet: Moore's Law ensures that capacity increases to meet demand; Parkinson's Law ensures that demand rises to meet capacity.

With its granular and short-term billing cycles, the cloud requires a degree of financial discipline that is unfamiliar to most traditional IT departments. Faced with having to provide short-term forecasts and justify them against actual spend, they need to evolve their governance models to support these new patterns.

Establishing and Maintaining Cloud Governance

Good cloud governance always starts with a well-architected framework and an architecture precisely designed to meet workload requirements. This is just a baseline for optimization, but many organizations stop here. The concept of governance needs to expand to ensure that accountability is as distributed as decision-making. End-users should be empowered to monitor the financial impact of their decisions and given the tools required to optimize for cost.

With a large menu of services and options that can be combined in hundreds of millions of ways, AWS offers a rich array of options—from replicating your existing data center in the cloud, all the way to complete refactoring to achieve a cloud-native stack.

This eBook provides 10 strategies to help you choose among alternatives, streamline your operations in AWS, and continuously optimize your deployment to take advantage of best practices and available arbitrage opportunities within these hundreds of millions of combinations.



1. START SMART: PROVISION THE INSTANCE YOU NEED—NO MORE, NO LESS

SITUATION

Amazon Elastic Compute Cloud (Amazon EC2) web service provides secure, resizable compute capacity in the cloud. EC2 is also the primary contributor to an AWS bill and therefore the best candidate for cost-optimization. EC2 is the service with the largest menu of choices, making it difficult to accurately decide on resource type. For example, with nearly 150 instance types to pick from, EC2 offers 10 choices for 4 vCPU 32GB, with a price differential of 3x separated by optimization type: CPU, memory, storage, and so forth.

COMPLICATION

We see many organizations choose an instance type based on generic guidelines that do not take their specific requirements into account. Think your application is CPU-driven? Choose the “C” family. Think memory is your constraint? Choose an “R.” Don’t know? Just choose “T” because you can always change it later.

IMPLICATION

EC2 instance types are sized and priced exponentially, so choosing the wrong type can result in a significant pricing premium—or a performance penalty.

POSITION

Your instance choice should be data-driven. A typical organization has already invested fairly significantly in performance management solutions. With a substantial amount of historical data already available, there is no reason to throw that away and make instance type choices based on best practices or generic guidelines.

ACTION

Leverage your existing IT systems of record to help provide data-driven choices of instance type. Yes, the cloud makes it easy to effect changes, but when you’re operating at scale, there’s no substitute for getting it right the first time.

At HyperGrid, we have made instance planning a cornerstone of our HyperCloud™ Analytics solution. By extracting performance and utilization data from your existing performance management systems, and by correlating them with millions of sizing options and pricing points in the cloud, HyperCloud Analytics helps you choose the best instance type for your needs.

BENEFITS

Customers who use data-driven instance choices typically see a cost reduction of about 50

percent, compared to those who use generic guidelines or best practices to choose their instance types.

2. STAY SMART—USE CONTINUOUS, DATA-DRIVEN INSTANCE OPTIMIZATION

SITUATION

Choosing the right instance is only the first step toward cost optimization. The cloud is a very dynamic environment. AWS is expanding choices for instance type regularly, and your business needs are constantly evolving. Making a data-driven choice is a great start, but it is just as important to stay on your feet and constantly evaluate choices available for optimization.

COMPLICATION

Instance Inertia: even though the process of making a change is simple enough, it can be difficult to make a change without having any conclusive evidence of either cost gains or performance improvements.

IMPLICATION

The cloud presents a number of arbitrage opportunities: instance families, generations, types, and regions. The pricing differences between these can be leveraged in the interest of cost-optimization. However, the large number of choices available can make it very difficult to do this manually, and instance-optimization at scale needs a data-driven solution.

POSITION

Continuous re-evaluation of your instance choices will help ensure that you're always staying ahead of the pricing changes in AWS, as well as your business requirements. However, this does not mean you should "change for the sake of change." Change comes at a cost and with its own risks; you should have a mechanism in place to perform a cost-benefit evaluation before making a change.

ACTION

This can be difficult to do at scale, so we've added an Instance Optimization service in our HyperCloud Analytics product. This service allows you to leverage the different arbitrage opportunities available within AWS. By continuously sizing and pricing all alternatives in the EC2 family, HyperCloud Analytics is able to suggest alternative instance types and families (or even generations) that can meet your SLA requirements at reduced cost. By presenting the cost savings, HyperCloud Analytics allows you to use financial justification to drive proactive changes in cloud usage across the organization.

BENEFITS

Continuous instance optimization will help you stay ahead of the pricing curve and leverage cost-optimization without investing in any structural or architectural changes to your applications.

3. LOOK FOR LIFE BEYOND EC2—EMBRACE SERVERLESS COMPUTING

SITUATION

AWS Serverless Computing (Lambda) provides a DevOps-friendly, provision-free, highly scalable, and highly available environment automatically, without having to build the supporting infrastructure in EC2. As a result, AWS Lambda is increasingly becoming a deployment architecture of choice for web services and data processing workflows.

COMPLICATION

Uncertainty around pricing often leads to discomfort in choosing Lambda for large-scale deployments. EC2 may have its complications with overlapping instance types, but once you've chosen, there are no pricing surprises. In contrast, Lambda pricing is multi-dimensional and based on number of transactions, memory allocated to the function, and runtime—each of which can change many times according to business requirements.

IMPLICATION

A poorly architected Lambda deployment may end up costing much more than its EC2 equivalent. For example, running 30 transactions per second—with each transaction having a runtime of one second consuming 1024MB of memory—will cost \$1304.94 on Lambda. A c5.4xlarge instance could most likely run the same workload with similar performance and cost about a third of Lambda (\$248.88 monthly).

POSITION

A good rule-of-thumb for deciding on Lambda is as follows:

If you have a function that is not invoked often, Lambda will almost always be cheaper than EC2.

If you have a function that is invoked often but is not resource-intensive, then it will most likely be cheaper to run this in Lambda. However, this should be validated during load testing.

If you have a resource-intensive function that runs very frequently, then it is possible that an EC2-based architecture—regardless of its elegance—may be the right choice. Alternatively, it could simply be that the function has been overprovisioned with memory; perhaps the function could be optimized in terms of runtime or memory utilization.

Lambda pricing is a combination of a number of invocations and GB-seconds used by the function. The number of invocations is a function of business success, not an architectural attribute. This leaves GB-seconds as the only lever to optimize pricing: allocate more memory and the runtime will reduce—or vice versa.

Optimizing Lambda costs, then, is all about adjusting the lever between memory allocation and execution time to find the optimal balance.

ACTION

Since Lambda functions read completely headless, understanding their behavior and optimizing cost tends to be more complicated than any other AWS service.

A number of native AWS tools can be used to understand the behavior of Lambda functions: CloudWatch natively supports a number of Lambda metrics.

The Metric Filters function of CloudWatch Logs can be used to convert logging data into additional CloudWatch Metrics. This may require additional instrumentation of your Lambda functions, but the overall benefits are worth the effort.

Once these are in place, you should run your function, at scale, with different memory allocations, and use CloudWatch to measure execution time and calculate cost at scale. When you have this data, you can determine the best cost tradeoff between memory and runtime.

At HyperGrid, we felt there was a better way. We have added a Lambda monitoring and optimization service to our HyperCloud Analytics platform to support Lambda analysis and optimization.

Capabilities include:

Visualize spending on Lambda and detailed metrics on each Lambda function, including key parameters like resource allocations, run time, timeouts, error rates, and so forth.

Easily identify resource wastage—functions with a high error rate, overallocation of memory, under-allocated functions (typically those with a high timeout rate), and so forth.

Make recommendations on right-sizing Lambda functions by adjusting the balance between runtime and memory allocation.

Compare the cost of continuing with Lambda vs. running an equivalent EC2 environment.

BENEFITS

The ability to analyze and optimize Lambda costs will not only result in a more DevOps-friendly and scalable architecture, it will also lower your operational costs by reducing the number of different types of resources that need to be managed.

4. LEVERAGE RESERVATIONS TO REDUCE COSTS—START WITH EC2

SITUATION

Reservations are the easiest way to optimize your AWS costs because using them requires making no technology changes. Reserved Capacity is available for various services; EC2 Reserved Instances are the most popular, due to the relative size of EC2 spend over other services.

COMPLICATION

Accuracy in reservations is key. Reservations are not guaranteed to be applied towards your bill. They apply only if your utilization matches the terms of the reservation. They can only be purchased in 1- and 3-year variants; it can be hard to plan that accurately and that far in advance.

IMPLICATION

Managing your reservations is just as important as purchasing them; unused reservations are one of the easiest ways to lose money in AWS. Not only did you spend money on a resource you did not use, you probably had to spend again on another instance that matches your needs more closely.

POSITION

EC2 Reserved Instances have the most complicated structure of all reservation types due to the abundance of instance options; however, they also provide the greatest flexibility. Understanding and leveraging this flexibility is key to being able to leverage the pricing breaks offered by reservations.

ACTION

Look at your EC2 estate, and find opportunities for saving via reservations. There are typically plenty of low-hanging fruit—long-running instances that haven't changed in a long time. To do this at scale, HyperCloud can help you pick from curated RI options based on deep insights from your historic cloud and instance usage patterns and real time price vs. performance.

Start by purchasing Convertible Reserved Instances; they offer a smaller discount, but provide a trade-in capability that makes a very big difference to the inexperienced user. With experience under your belt, you can venture out to Standard Reserved Instances to save even more.

A simple best practice is to standardize on one or two instance types, which greatly increases the probability that an unused reservation can be applied against an instance. This can be difficult for an organization that isn't building their own services and relying on recommendations from individual vendors.

Nearly all Reserved Instances can be split or combined, as long as you stay within the bounds of the overall Reservation. HyperCloud provides reports on RI Utilization and Waste, which

can be used to find unused RIs and discover opportunities to apply them towards existing instances.

Be a bargain hunter. Browse the AWS Reserved Instance Marketplace for an RI that may meet your exact needs better than a brand new 1-year and 3-year RI. Note that the Marketplace only lists Standard Reserved Instances, and not Convertible.

If HyperCloud identifies a Standard Reserved Instance lying unused, sell it on the RI Marketplace to recover cost. Before you list your instance on the RI Marketplace, you can modify it (e.g., split a large instance, combine smaller instances, exchange an RI on an unpopular instance size for a more mainstream one, and so forth) to make it easier to sell.

Purchasing RIs with all payment up-front is the best discount that you can get from AWS. However, when you try to sell an unused RI on the RI Marketplace, AWS will charge you a fee based on the up-front price you charge, which can make up-front RIs more expensive to sell. This makes it even more important to use a predictive analytics tool like HyperCloud to help with your proactive RI selection.

BENEFITS

Reservations allow you to get pricing breaks on AWS instances without making any changes to your deployment. Intelligently purchasing Reserved Instances is important, but managing Reserved Instances and ensuring that they are being applied towards utilization will have an even bigger impact on AWS costs.

5. MODEL APPLICATIONS—NOT INSTANCES

SITUATION

Cloud app planning is a much broader discipline than just instance selection; it requires taking a look at the “whole application”—that is, the application itself and the components that make up the deployment (e.g., load balancers, proxy servers, cache, database servers, log management, performance management, and so forth).

COMPLICATION

Cloud app planning means mapping the data transfer between these and choosing an optimal resource placement that takes into account such factors as network and data transfer costs and data durability.

IMPLICATION

Lack of a holistic view of costs before deploying the “whole application” can often lead to pricing surprises. For example, before you deploy an application in an HA configuration, you should be able to quantify the network characteristics of the application before the deployment and understand their effect on pricing. An application that generates a lot of traffic could generate more in cross-AZ and cross-region transfer costs than the business benefit of High Availability.

POSITION

Applications need to be modeled in their entirety, not just the individual instances. Their interactions and network traffic between components need to be captured and analyzed before making recommendations on placement.

ACTION

When planning your application, make sure that you treat your application as a collection of services that communicate with each other, and not just as a collection of individual instances. For example, before deploying an HA architecture across regions, factor in the cost of having your application communicate across regions. Similarly, before deciding on S3 as a storage mechanism, factor in not just your storage costs, but also your retrieval costs.

Our solution, HyperCloud, has a Cloud App Planning service that allows you to model an entire application, with all its dependencies. The service will recommend a placement for the application (along with costs) that factors in such details as data and network transfer costs and storage tiers. It also allows for easy what-if scenarios that help with optimal placement of test, dev, production and DR instances of the application.

BENEFITS

By modeling applications in their entirety and factoring the bigger picture into not just instance choice but also placement, you can ensure that you have a durable deployment architecture without pricing surprises.

6. CENTRALIZE GOVERNANCE OF RESERVED INSTANCES

SITUATION

A common AWS billing model is to have a centralized account with Consolidated Billing, linked to a number of autonomous accounts. With this model, it is common for individual accounts to purchase Reserved Instances based on their individual usage patterns.

COMPLICATION

Individual accounts may not be linked to each other—the topology is a star, not a mesh. Price breaks due to reservations are isolated in individual accounts, or accounts that are linked to it. This means unused RIs cannot be shared between accounts in this topology.

IMPLICATION

Wastage—RIs purchased in an individual account but not utilized by this account cannot be reused. Also, since the individual accounts are no longer responsible for billing, they may not have the hygiene in place to discover and repurpose or resell unused RIs.

POSITION

Reserved Instances should be procured and managed centrally. Purchasing an RI is only the beginning; you should have a process in place to continuously monitor RI utilization and modify unused RIs (split/join or exchange convertible RIs) to maximize their usage.

ACTION

Use master accounts to manage the lifecycle of Reserved Instances—recommendations, procurement, and monitoring. This way, they can flow to all linked accounts. HyperCloud Analytics is a great place to start for discovering opportunities for Reserved Instances. HyperCloud can easily map all your member and master accounts; plus, it can offer a single consolidated view of all instances across all members, and the analytics to make decision-making across all member accounts easier.

Enroll all stakeholders in RI purchases. While RIs can be procured centrally, the individual account owners should be involved in the process of analyzing and approving the process of RI procurement. Ensuring that all stakeholders are empowered to understand the cost implications of their actions is part of good cloud governance.

There will be exceptions: pricing breaks due to RIs can flow between accounts, but capacity reservations are localized to purchasing account. Therefore, there may be workloads (e.g., mission-critical applications or a DR workflow) that need localized management to ensure capacity reservation.

We also recommend having organizational standards around instance types, increasing the

probability that an unused RI will be picked up and applied to an instance in another account. This is not always feasible, however, because individual vendors and business needs drive instance selection.

BENEFITS

RIs are the best pricing breaks you can get in EC2. By creating a centralized RI governance model, you can ensure that you capture the maximum benefit from an RI investment.

7. ANONYMIZATION IS THE ENEMY OF ACCOUNTABILITY—TAGGING TO THE RESCUE

SITUATION

AWS customers use accounts as units of autonomy and as a security boundary between environments. Often, accounts are linked to VPCs and not to organizational or business units. This makes it hard to create the linkage between spend and business initiatives.

In recognition of this issue, AWS has enabled a new service called AWS Organizations to enable management of accounts on organizational boundaries. However, this is a relatively new tool, and a number of customers already have an account structure in place that is based on previous best practices.

COMPLICATION

AWS bills are very detailed; running into thousands of lines is quite common and the “million line bill” is not unheard of. Famously, Netflix tweeted in 2016 that their AWS bill was over 700 million lines long.

With VPCs—and not organizational boundaries—often defining the operational boundary of an account, linking spend back to account is often not sufficient to create linkage between business initiatives and AWS spending.

IMPLICATION

If you lose the direct linkage between operations and spend, you run the risk of removing cost accountability from the individual account owners. This flies in the face of good cloud governance. Good cloud governance means empowering your users to monitor the financial impact of their decisions, as well as giving them the tools to optimize costs.

POSITION

AWS has enabled tags across nearly all of its products and services, making it the best way to group and identify resources. They make it possible to automate the process of cost-allocation and are often the only way to create a linkage between resource utilization and business initiatives. Proper use of tags is critical to maintaining good cloud governance.

ACTION

A first step to using tags is to set up a tag strategy: define a set of tags that are mandatory (e.g., ProjectID, OwnerContact, ApplicationName, and so forth). These must be attached to any resource.

Since AWS tags are not mandatory, it is important to monitor and enforce their use. Use HyperCloud Analytics to find resources that are untagged and set up a mechanism to enforce their tagging, or else you should disable them on a set schedule.

Enable the “Cost Allocation Tags” feature to ensure that the user-created tags are visible from

the billing console and use Cost Explorer or HyperCloud Analytics to create custom views based on tags. Enforcing a well-defined tag structure is important to ensure that all resources can be accounted for in a cost-usage report.

For organizations that wish to automate this process, a cloud management platform with built-in lifecycle management (such as the App Store in HyperCloud) is necessary. The platform carries out the details of creating and managing tags, assigning quotas to individual accounts, and automating chargeback based on predefined policies. The HyperCloud Security and Compliance service can detect and report all deviations from established tagging policy and allow administrators to take one-click automation actions to remediate these issues.

BENEFITS

Enforcing a well-defined tagging policy is the foundation, not just for automating cost-allocation and enabling chargeback, but is also critical to group AWS resources for other governance tasks as well.

8. COST GOVERNANCE STARTS WITH PROACTIVE COST-MANAGEMENT

SITUATION

Cost-management is a challenge for all public cloud users; it is not uncommon for cloud spend to grow significantly past initial projections and still be overlooked until it is too late.

COMPLICATION

AWS Billing is a monthly activity; customers who wait until the end of the monthly billing cycle to see their utilization are likely to be surprised, which affects their ability to address problems as they occur.

IMPLICATION

Customers who rely on the AWS billing cycle are almost guaranteed to be surprised by the bill when it arrives. Proactive cost management is critical to the ability to maintain good cost governance in the cloud.

POSITION

AWS provides numerous mechanisms to manage costs, but unfortunately, many of them are disabled by default and should be enabled as a best practice.

ACTION

Your first step should be to enable Consolidated Billing and link your accounts to a master “paying” account. This will allow you to see the big picture of spend and create reports that take a holistic view across the organization.

You must aggressively use tags (see previous topic) to ensure you have the right level of granularity in the reports. Without this, reporting will be at an account level. In our experience, account-level reporting is not sufficient to draw an accurate picture of spend and trace it back to business initiatives.

The AWS Cost and Usage Report (CUR) is generated multiple times daily and contains real-time raw details on each tracked resource. This report contains a wealth of information, but due to its sheer volume of data, we recommend investing in a bill management service that will help you analyze this data in real time. HyperCloud’s Bill Management service shows a real-time breakdown of our AWS spend by service and by instance, along with recommendations for optimization. This allows you to react to spending spikes in real time.

AWS also provides a native option, AWS Cost Explorer. (This function is disabled by default, so enable it as soon as possible.) AWS Cost Explorer is a visualization tool that helps you view and report on the data contained in the CUR.

BENEFITS

By proactively managing costs, you can start to define and enforce good cloud cost governance. This will increase the accuracy of your short-term spend forecasts and proactively manage any spending spikes in real time.

9. AUTONOMY DOES NOT HAVE TO MEAN UNRESTRICTED BUDGETS

SITUATION

AWS Cost Governance Best Practices state, “Your organization should emphasize the creation of clear, effective policies and governance mechanisms around cloud deployment, usage, and cost responsibility.”

COMPLICATION

This can often be at odds with the high level of autonomy that your individual users expect in the cloud. Users would like to procure, enable, and disable resources at will, without first having to go through a central process for approval.

IMPLICATION

By focusing on cost management and accountability, you are creating and enforcing “back end policy,” which kicks in once the spend has been incurred. This ensures that you have additional accountability within the system, but for good cloud governance, you should be able to enforce proactive “front end policy” as well, to ensure that any deviations from the norm are addressed as early as possible.

POSITION

By letting the centralized team set quotas or budgets for the individual users, IT departments can better distribute the balance of power. End-users can continue to enjoy the autonomy that they expect in the cloud, but they also have the accountability that is needed for good cloud governance.

ACTION

AWS has a built-in mechanism called AWS Budgets, which helps you assign budgets based on a fixed cost, or on usage of a specific resource (e.g., \$100, 50 S3 requests, and so on). However, AWS Budgets is a monitoring solution and not an enforcement tool; it requires that you address the underlying cause of an alert.

Another alternative is to use a cloud management platform—our HyperCloud platform allows you to assign hard quotas to specific resources. Provisioned requests that exceed this quota will fail or, based on policy, can trigger a workflow request for a higher level cloud admin to approve incremental budget. Leases can also be set as part of the proactive policy management, including enforcing default leases for certain types of groups and users who may be engaged in test and validation activities. Finally exception policies like termination protection can ensure that critical production resources are not accidentally shut down because of a default lease policy.

BENEFITS

Good cloud governance requires that you empower the end-users to be accountable for their own costs and give them the tools required to optimize these. By allowing the central account admin to set proactive policies while allowing the end-users to have the ability to handle the exceptions that are most important, you are ensuring that you have good budget governance

10. BUILD AN ORGANIZATION-WIDE, SHARED RESPONSIBILITY SECURITY MODEL

SITUATION

AWS operates on a shared responsibility model, which is an excellent template for building a shared security model across the organization. By giving end-users the tools they need to monitor their own security and remediate issues, you convert security from a centralized function to one that is distributed across the organization.

COMPLICATION

Lack of awareness is a major roadblock in achieving a shared security model. Very often, end-users are not aware of the security implications of their actions and rely on centralized audits to highlight issues that they will address.

IMPLICATION

By giving your users autonomy to operate in the cloud, but without sharing the responsibility of security with them, you are creating the same security bottlenecks that are found in the on-premise data center. A centralized team, far removed from the details of the applications, will inevitably create a perimeter-based security model—which we know is not applicable to the cloud.

POSITION

By giving end-users visibility into the security implications of their activities, you can create an organization-wide culture of transparency and compliance.

ACTION

AWS has a built-in tool called Trusted Advisor, which can check the status of your resources against AWS Best Practices and can manually remediate.

Third-party tools allow you to expand on this list by checking security against a number of other parameters, in addition to making it easier to pinpoint and remediate the issue. The Cloud Compliance and Security module built into our HyperCloud solution provides a compliance dashboard to visually map the status of all your resources. It pinpoints non-compliant resources, along with providing a link to the specific resource to allow for easy remediation. In many cases it also offers instant remediation of these issues with one-click automated actions.

BENEFITS

By sharing security responsibility with the end-users and empowering them with the right tools, you can afford to give them the autonomy they expect in the cloud, along with the accountability to be responsible for the security that goes with it.

CONCLUSION

Having a well thought out AWS strategy is crucial to your long-term cloud gains. Taking the time to understand and pick the right instances for your apps is well worth the time and effort as it will directly impact your AWS pricing and bill.

We hope these ten strategies help inform and support you as you navigate the sometimes turbulent waters of cloud transition. They are here for you to consult and rely on as best practices and cost-saving opportunities—and of course you are welcome to speak with a HyperGrid professional one-on-one at any time.

Given the virtually uncountable number of combinations, we have tried to identify the most practical and reliable ways to optimize your deployment at all stages and empower your end-users while insulating them from temptations, assumptions and habits that can cost you some unpleasant surprises when the bill arrives.

ABOUT HYPERGRID

HyperGrid™ is helping enterprises and MSPs master the challenge of managing hybrid and multi-cloud environments as requirements and technologies increase in complexity.

Founded in 2016, and headquartered in San Jose, California, with additional offices in the U.S., Europe, and Asia, HyperGrid supports a diverse global customer base across finance, insurance, healthcare, oil & gas, construction, retail, FED/SLED, and telco.

The HyperCloud™ Intelligent Cloud Platform delivers end-to-end visibility, control, and automation, powered by a predictive analytics engine with over 400 million benchmarked data points. This innovative model enables organizations to plan, migrate, optimize, manage, and scale their cloud operations. HyperCloud provides powerful capabilities for cloud cost management, performance optimization, and detailed analytics and reporting, along with security and governance. With HyperCloud, organizations can accelerate cloud decisions, reduce costs, deploy secured workloads, and automate IT and DevOps across hybrid and multi-cloud environments, all in one platform.

Backed by noted venture capital firms, HyperGrid has earned industry recognition for milestone achievements, including Forrester Wave Report Hybrid Cloud Management 2018; CRN Top 10 Hottest Cloud Computing Startup 2018; CRN Datacenter 100 2018; Best of VMworld Finalist 2017; and SVC Cloud Management of the Year 2017.

