

# Avoiding Sticker Shock in the Cloud Era

*BEST PRACTICES FOR REDUCING COSTS & VM SPRAWL*



## SECTION ONE

# Unrestricted Cloud Resources Lead to Sticker Shock

TODAY, ANYONE CAN PROVISION CLOUD RESOURCES WITH THE SWIPE OF A CREDIT CARD.

*“Many chief financial officers as well as CIOs are gasping at the unexpected expenses seen in their monthly cloud bills.”*

- Joe McKendrick, SOA Manifesto

As public cloud IT resources for the enterprise become increasingly viable, an inevitable tension rises between consumers of these resources and the IT leaders who are ultimately financially responsible. The consumers are the developers, DevOps teams, QA teams, support teams and others. In larger organizations, these consumers might feel less accountable for the expenses and take advantage of whatever they can get to do their specific jobs.

When the time to get the resources is lengthy and the process is cumbersome, some teams will ask for (and hold onto) more than they need just to avoid going through the process again. In the past, going outside central IT departments to get resources to innovate wasn't easy. It required consumers to purchase and configure their own physical IT infrastructure. That's no longer the case with the immediate availability of cloud computing power at the touch of

a button. These resources are readily offered by major cloud providers such as [AWS](#), [Microsoft Azure](#), [Google Cloud Platform](#) and more.

The frustration occurs when there is a disconnect between the promise of agility from the cloud providers with the need to control access from enterprise IT. Consumers see how fast it is to get resources in the cloud compared to the cumbersome process of enterprise IT.



Today, anyone can provision cloud resources with the swipe of a credit card. The public cloud promise is self-service and innovation with agility, and with open access, developers can consume public cloud resources freely. However, without oversight, resource usage goes unchecked, and development teams can charge any number of public cloud resources to specific budgets that suddenly skyrocket.

[According to research from Gartner](#), “through 2020, 80% of organizations will overshoot their cloud infrastructure as a service (IaaS) budgets, due to lack of cost optimization approaches,” and “through 2020, 45% of organizations that perform lift-and-shift to cloud IaaS without optimization will be overprovisioned by as much as 55%, and will overspend by 70% during the first 18 months.”

The monthly bills become the equivalent of [sticker shock](#) for those responsible for cloud spending.

**“Many chief financial officers as well as CIOs are gasping at the unexpected expenses seen in their monthly cloud bills.” -Joe McKendrick, SOA Manifesto**

With too much control, the tension between IT and the developers is palpable. Frustrated developers work with stifled morale or move on to other organizations. The digital business initiatives of an organization become thwarted and it becomes more difficult to disrupt a market or simply keep pace with the competition.

Leading digital innovation from the C-suite requires figuring out how to control IT spending yet allow for more agile provisioning processes for innovation-hungry end users who need to innovate.

Changing the culture to trust end users and developers to spend prudently is one strategy, but it takes time. In the meantime, as provisioning goes unchecked, developers build solutions quickly—spinning up VMs often without following a full lifecycle plan to create and decommission them, which invariably increases the IT spending burden.



## SECTION TWO

# Competing Priorities and Complex Environments

STICKER SHOCK OCCURS WHEN SOLUTIONS HAVE BEEN ARCHITECTED POORLY FOR THE CLOUD AND RUN EITHER UNEXPECTEDLY OR UNNECESSARILY.

*“Different clouds provide different mechanisms for viewing spend. There’s no central point of visibility for users, IT, and management.”*

- Bernard Sanders, Chief Technology Officer, CloudBolt Software

This open access to public cloud resources at anyone's expense leads to a complicated scenario of competing priorities and multiple complex environments for provisioning. The competing priorities of cost control and the ability to innovate with agility becomes even more challenging when considering the vast choices for private and public cloud environments.

Larger enterprises typically have a sprawl of IT resources over several private and public cloud environments. These environments stem from not only individual teams going out on their own to secure resources but also from mergers and acquisitions where technology hasn't been consolidated. With this expanse of resources, it's possible existing and unused infrastructure can be made available on-premises for some workloads and application development, but the environment might not be managed or tracked well enough to implement a good strategy. As a result, this potential set of resources goes untapped as the ease of getting this exact same type of resource in the cloud becomes more attractive.

Furthermore the sprawl of untracked spending makes it difficult to quantify business value that is dependent on how much the actual infrastructure costs to run efficiently. The sticker shock mentioned in the previous section occurs when solutions have been architected poorly for the cloud and run either unexpectedly or unnecessarily. [Repatriation](#) of these runaway cloud resources back on-premises becomes a common practice when the initial promise of lower

In a recent survey by IDC, revenue from the sales of IT infrastructure increased year over year to confirm that on-premises workloads are here to stay.



If development teams are using multiple cloud resources tied only to their individual business units, they are not likely held as accountable as they could be under central IT management. On the one hand, the development teams want agile resources that are readily available to build and innovate in a highly-competitive market for digital services. On the other hand, the cost to do this can spiral out of control without oversight depending on how and where the resources are utilized. There are instances when a private cloud or on-premises environment would work just fine for a given project and save quite a bit on cloud spending.

In addition, there is also the ability to negotiate rates and reserve computing power with public cloud providers. When these public cloud providers work closely with enterprises, they see the potential for usage across more business units within the organization and ultimately provide greater business value at a lower cost so both parties win.

IDC found revenue from sales of IT infrastructure equipment grew 48.4 percent year over year in the second quarter of 2018 to \$15.4 billion



# What's at Stake Without Effective Cloud Controls?

A GOOD COMPARISON OF SCALING THE PROVISIONING OF CLOUD RESOURCES ACROSS AN ENTERPRISE IS THE ONBOARDING PROCESS FOR NEW EMPLOYEES.

*“Cloud is no longer a place to get some cheap servers or storage. It’s not even a single place. Cloud computing is now shorthand for how companies turn amazing ideas into winning software – faster.”*

- Dave Bartoletti, Vice President, Principal Analyst Forrester



With overspending and competing priorities in place, what's the implication of this dynamic? It clearly does not scale very well for enterprises and large organizations. Without effective cloud controls in place, each cloud environment must be accessed separately and managed by at least one person who has the knowledge and expertise across the multiple cloud venues or at least coordinated with an expert for each of the environments.

This makes each environment its own separately managed resource. Without a reporting or standardization platform that consolidates the view, the usage and management of the resources cannot be easily compared from one venue to the next.

Moreover, having one or more teams consuming these resources can duplicate the complexity with each team environment siloed or tied only to a specific business unit. Each team could have a separately managed set of resources without participating in the more holistic

view across the enterprise. Extend this across several teams and it becomes quite a sprawl for central IT to manage. Even worse, it will likely go unmanaged with the hope that everything will work out properly.

The results? The strategic leaders in IT, finance, and operations risk adequate insight for business decisions based on cloud usage and adoption across all the various teams using them. They do not have a comprehensive view of what resources are being used and by whom. Knowing what the costs are for similar workloads in different venues helps them make better decisions that might cost less yet don't sacrifice performance or security standards required by internal end users and/or customers.



In smaller organizations, this multiple cloud issue might not be a problem at all. It's relatively easy to manage a handful of accounts and resource connections to separate public or private cloud resources. The impact, however, is far greater when you're scaling up to provide resources for hundreds or thousands of users. According to a Gartner 2018 study, the breakdown of moving infrastructure and resources to the cloud is trending up (see Table 1).

**Table 1: Cloud Shift Proportion by Category**

	2018	2019	2020	2021	2022
System Infrastructure	11%	13%	16%	19%	22%
Infrastructure Software	13%	15%	17%	18%	20%
Application Software	34%	36%	38%	39%	40%
Business Process Outsourcing	27%	28%	29%	29%	30%
<b>TOTAL</b>	<b>19%</b>	<b>21%</b>	<b>24%</b>	<b>26%</b>	<b>28%</b>

*Source: Gartner (August 2018)*



A good comparison of scaling the provisioning of cloud resources across an enterprise is the onboarding process for new employees. In a smaller company, when the onboarding process takes an extra day or two because of a technical glitch (e.g., getting a laptop or enrolling in benefits), most of us would be relatively unphased. We'd say it's "no big deal." On the other hand, consider ramping up hundreds of employees over a month's time. If all of them have the same kludgy onboarding process, the impact will be much more significant. It delays productivity, reduces efficiency, and ultimately, delivers a pretty big impact on the business goals of an organization.

The point is that when inefficiency scales—like an HR onboarding process or provisioning cloud resources for developer

teams—finding ways to reduce the time it takes to get resources and to improve the processes will have a huge impact that drives significant organizational benefit.

The need for a centralized mechanism that balances agility and cost for these complex scenarios has never been greater. In the next section, we explore the benefits of how a cloud management platform solves these problems.



## SECTION FOUR

# Controlling Cloud Costs and VM Sprawl

WHEN YOU OPT FOR A HYBRID CLOUD PLATFORM, YOU CAN TAKE COMFORT IN KNOWING THAT YOUR PROVIDER IS KEEPING YOUR SPENDING GUARDRAIL IN PLACE.

*“Aligning business groups and teams with IT spend is critical to making budgeting and cost recovery easier to implement. If you can see disparity between groups, you can choose to spin servers down to bring budget into alignment.”*

- Aaron Jablonowski, Senior Solutions Engineer at CloudBolt

As enterprises tackle the challenges of adopting new public cloud offerings in concert with delivering digital value to internal and external end users, a solution for comprehensive cloud control that reduces costs and VM sprawl must do the following:

- Empower end users with self-service IT
- Manage complexity and reduce overspending
- Deliver agility and performance
- Comply with security standards

When it comes to self-service IT, there are two options that organizations should consider:

1. Give developers and end users credentials to a particular cloud

environment and let them provision what they want, when they want, and how they want without oversight, hoping they will make the right decisions.

2. Provide self-service IT that is controlled so developers can still get what they need to do their job but based on predefined guidelines set by the organization.

Option 2 empowers end users and developers to deliver results faster without the potential for mismanaging the resources. The cloud management platform used to control cloud access should connect to all public and private cloud resources so an IT administrator can set up and grant access to users, set safeguards to control spending and also



configure the exact set of controls that may otherwise be too complex for the average user. This frees up the time for end users who just need the provisioned resources readily installed instead of having to follow steps to configure some of the complexity themselves. Finally, a central platform for delivering simplified, self-service access to cloud-based resources should harness the value of scalability in the cloud without requiring developers to set up some of the controls and security requirements themselves.

Once a centralized cloud management platform is in place to control cloud resources, it should provide:

- One-click access to private clouds (e.g., VMware vSphere) and public clouds (e.g., AWS, Azure, Google Compute Engine) with all the complexity configured on the backend so that end users get the resources confidently and securely without the risk of overspending or a mismatch in resources from company policy
- Connections to all the best practice cloud environments and DevOps configuration tools (e.g., Ansible, Puppet, Chef, and others) to deploy a distributed architecture of hybrid cloud resources across multiple geographic regions
- Robust role-based access control (RBAC) to integrate with existing Lightweight Directory Access Protocol (LDAP)-based or other user management systems in the enterprise so that each user in the system has “least privilege” user access



## Advancing a cloud control strategy provides for the following outcomes:

- Organizations will have one corporate account for each public cloud provider to get the best rates based on the total volume of resources used across the whole organization
  - Reserved compute power that has been purchased with big discounts can be shared across teams instead of being stuck and unused for a single team
  - IT leaders will have usage and spending visibility across all public cloud provider accounts to make the best decision for which environment matches the specific needs of end users
- IT spending will be controlled from each public cloud environment using:
    - » Best-venue execution strategies based on cost, performance and security
    - » Automated approval processes to speed up provisioning processes
    - » Guardrails for spending using quotas and limits and right-sizing resources
    - » Comprehensive scheduling for powering on and off resources efficiently
    - » Lifecycle management of resources
  - Each end-user will have least privilege access to only the resources needed to complete the requirements of their job
  - Many more end users will be able to access and use the agility of public cloud resources because they will have controlled access that does not violate security policies or exceed IT budgets



Using a cloud management platform, enterprises can finally balance the control that IT wants with the agility and creativity that developers want from competitive public cloud resources. And it's with this control that IT can finally say goodbye to the sticker shock and VM sprawl that would otherwise plague their organization.

To find out how CloudBolt meets all of these objectives and provides a complete cloud control platform, [schedule a demo today](#).





✉ [sales@cloudbolt.io](mailto:sales@cloudbolt.io)

☎ 703.665.1060

🌐 [www.cloudbolt.io](http://www.cloudbolt.io)



©2019 CloudBolt Software

## **ONE VIEW. MANY CLOUDS.** Automation. Flexibility. Control.

In the era of Cloud, CloudBolt helps IT and developers work better together by empowering them with better visibility, control, and self-service. CloudBolt delivers the world's most user-friendly cloud management platform, so enterprises can provision, orchestrate, and consume IT resources across hybrid cloud, multi-cloud, and container environments.