

Forbesinsights

AVOID STICKER SHOCK

How to Determine the True Costs of Clouds



IN ASSOCIATION WITH:



TABLE OF CONTENTS

Introduction | Pg 3

STEP 1

Sort Out Cost Complexities | Pg 4

STEP 2

Understand Workload and Optimization Considerations | Pg 8

STEP 3

Pick the Right Cost-Analysis Tool | Pg 10

STEP 4

Turn Insights Into Action | Pg 13

STEP 5

Create a Plan for the Future | Pg 15

Introduction

Think moving to the cloud will save you money?
Maybe it's time to think again.

Unfortunately, the early excitement that accompanies cloud strategies often turns into disappointment for many enterprises, and sticker shock is a big reason why. Despite performing the due diligence of estimating upfront costs based on published pricing from service providers, many companies are hit by unexpected and unnecessary expenses once they make a cloud move.

"Three or four years ago, our research showed that the number one driver for going to public cloud was the perception that it would help rein in costs," says Brian Garrett, vice president of validation services for the IT consulting firm Enterprise Strategy Group. "Cost reduction is dropping as a driver now that people are getting more experience with the true costs of public cloud."

In fact, a study by IDG found that 38% of enterprises have moved workloads from the public cloud back to their on-premise data centers—with costs being among the biggest concerns. Part of the challenge is understanding all the variables that affect cloud pricing, which make it difficult to perform apples-to-apples comparisons to find the right deals.

Cloud veterans also warn that to closely manage cloud costs, enterprises need significant upfront preparation not only for accurate cost comparisons but to determine which workloads will deliver price and performance benefits in the cloud versus keeping them on premise. All this should be done within a larger data center optimization and consolidation effort that's informed by having clear visibility into the entire enterprise infrastructure.

"As part of an ongoing IT simplification, modernization and rationalization program, enterprises must determine the best way to run each application," says Paul Lewis, chief technology officer at Hitachi Vantara, a technology and professional services company. "That could mean either running it in the existing environment or recreating it in a different environment, whether that's a private cloud, a SaaS application, within a new architecture deployed in the public cloud, or managed by a third party."

It's clear that enterprises need a framework for sorting through the complexities of cloud pricing models, which make it difficult to perform accurate comparisons of initial costs and to forecast spending over time. Fortunately, best practices and new tools for analyzing and monitoring IT cloud costs over time are helping IT and financial planners evaluate cloud and on-premise options to find the right balance of cost and performance.



38% of enterprises have moved workloads from the public cloud back to their on-premise data centers—with costs being among the biggest concerns.

PAUL LEWIS

As part of an ongoing IT simplification, modernization and rationalization program, enterprises must determine the best way to run each application. That could mean running it in the existing environment...a private cloud, a SaaS application, in the public cloud or managed by a third party."

A conceptual image of a business meeting. In the foreground, a person's hand holds a pen, pointing at a tablet displaying a line graph. Another person's hand is visible in the background, also interacting with a tablet. The scene is overlaid with various semi-transparent graphics: a world map in the top right, a pie chart and bar chart in the bottom left, and several circular icons with network-like connections scattered across the image. The overall color palette is cool, with blues and greys, accented with warmer tones from the background lighting.

STEP 1 Sort Out Cost Complexities



In an IDG study,

52%

of enterprises moved back from
**PUBLIC CLOUD TO INTERNAL
DATA CENTERS** because of pricing
and cost concerns.

Understanding the true costs of cloud will become even more important in the years ahead. For business agility, hardened security and freeing internal IT staff from time-consuming management tasks, enterprises across all industries are strengthening their commitment to cloud. The timing couldn't be better. As more cloud providers enter the market, pricing becomes a competitive differentiator. "This is creating rate-reduction opportunities for customers," says Vince Campisi, CIO and Chief Digital Officer at United Technologies, a provider of technology and services for building and aerospace industries.

But those opportunities exist only if enterprises understand how to closely manage cloud costs. Those that don't succeed in this area may find themselves reducing their cloud commitments. For example, of the 38% of enterprises cited in the IDG study, more than half—52%—moved back from public cloud to internal data centers because of pricing and cost concerns.

It's easy to see how cloud costs can become an unpleasant surprise. The first question facing decision makers is whether replacing on-premise assets with cloud resources is cost effective. Finding the answer may require a mind shift. "Some people may balk at the cloud, thinking that the same capabilities could be run in existing data centers for free. Well, it's not free," says David A. Bray, executive director for the People Centered Internet, an international coalition working to bring the benefits of the Internet to unserved regions throughout the world. Until recently, Bray was the chief information officer at the U.S. Federal Communications Commission.

Why aren't the existing IT investments that underpin in-house data centers "free?" Enterprises must factor in facility expenses, power and cooling bills, IT updates and maintenance, and salaries. "Those costs often are not brought to the forefront when organizations compare prices," Bray explains.

Next, IT managers must wrestle with the complexities that come with comparing the pricing models published by cloud providers. Decision makers must analyze many variables, which makes direct comparisons challenging and can lead to sticker shock even for enterprises that thought they had pricing locked down. "At first blush, pricing looks simple—just create a workload instance and compare prices of different cloud providers," Garrett says. "But so many factors affect the final costs, including network usage and whether a company needs highly available services or not. One executive I spoke with recently complained, 'Do I really need to be a rocket scientist to understand the pricing?'"

BRIAN GARRETT

“At first blush, pricing looks simple—just create a workload instance and compare prices of different cloud providers. But so many factors affect the final costs....One executive I spoke with recently complained, 'Do I really need to be a rocket scientist to understand the pricing?'"



For example, depending on their service choices, enterprises can capitalize on sometimes significant discounts offered by public cloud vendors. Customers of the Amazon EC2 cloud can get a price break by reserving capacity versus tapping into cloud resources on demand. These Reserve Instances (RIs) benefit organizations with predictable usage levels and offer discounts of 40% to 60%, based on one- and three-year contracts, respectively. "Some large organizations are dedicating people to understanding cloud rates and deciding how many reserved instances to buy based on internal supply and demand," Campisi notes. "It's like a commodities trader hedging investments—is it better to buy a one-year block of services now or wait and bet that there will be a price decline in the future?"

But these are just a few of the choices that impact costs. Amazon also offers two related subcategories with unique pricing: Convertible RIs, which offer 54% discounts off on-demand prices and more flexibility in adjusting usage attributes, and Scheduled RIs, which let enterprises reserve capacity by part of a day, a week or a month on a recurring schedule, according to Amazon's published price sheet.

How do Amazon's discounts compare with what's available on the Google Cloud Platform or Microsoft Azure? Answering that question takes time and analysis since each vendor has its own usage definitions and discounting structures. These are based on numerous factors, such as the number of virtual CPUs, what percentage of the billing month a customer's virtual machines are running in the cloud and, in the case of Microsoft, what terms buyers are able to negotiate in enterprise agreements.

Cloud customers also have options for choosing different underlying storage technologies, and this can impact pricing. For example, cloud providers typically charge a premium to customers who sign up for solid-state drives, which outperform traditional spinning disks. Workload requirements will dictate which technology choice is right—databases associated with financial applications and other core systems may need the highest performance possible—not only because of the criticality of these applications, but also because enterprises inherently take a performance hit when connecting to distant cloud data centers. Other, general purpose, business applications may tolerate slight performance delays, enabling enterprises to choose lower-priced storage options.

The IT staff at the University of Massachusetts Medical School understands how difficult it is to tackle cloud pricing complexities. The medical school uses public clouds to run traditional business applications, such as ticketing and email systems. Over the last two years, it has also moved an extensive array of research capabilities into a hybrid cloud environment, that not only supports big data sets but also satisfies critical compliance needs, and handles large spikes in demand for computing and storage resources. "Our research ecosystem is designed to give faculty members plenty of storage, compute, networking, and software tools—all the foundational elements to do their work," says Greg Wolf, the medical school's CIO.

Wolf says the cloud moves are being made for a combination of agility and cost-savings goals. "It is getting expensive to maintain large research ecosystems on premises, so we would have had to



Gartner says the average enterprise estimates an impact of approximately

\$5,600

for **EVERY MINUTE OF UNPLANNED DOWNTIME** in its primary computing environment.

keep going back to the trustees and saying, 'We need more funding and more real estate,' he adds. "It became harder and harder to keep up with research demands."

But nailing down initial cost estimates required a brute-force exercise. Medical school leaders first tried to weigh of IT services from three major cloud providers. "It was incredibly difficult to do an apples-to-apples comparison across all three," says Paul Langlois, director of research computing. "The providers give you calculators, which are great, but only if you're trying to figure out what one particular service costs—not the entirety of your cloud infrastructure."

Langlois recalls that he spent "a few weeks in the weeds with these calculators," trying to compare the three pricing models. "We put in the effort and ultimately came out with a good comparison to help us choose different vendors for different purposes."

Tracking on-going costs was aided by the IT department's iterative approach to moving workloads to the cloud. This enabled the staff to keep a close eye on the budget and actual spending. The medical school is now using a cost-analysis tool now that the cloud commitment has grown exponentially--Langlois estimates usage has risen 500% in less than two years and will continue to expand as the research community further makes use of the cloud resources.

But there's a flip side to evaluating cloud benefits. Not all cloud-related benefits can be measured with resource utilization equations. Depending on the available internal

PAUL LANGLOIS

[Large public-cloud providers] give you calculators, which are great, but only if you're trying to figure out what one particular service costs—not the entirety of your cloud infrastructure."

resources, enterprises benefit from outside services in the form of higher uptime, which can have a significant financial impact on large organizations. Gartner says the average enterprise estimates an impact of approximately \$5,600 for every minute of unplanned downtime in its primary computing environment. Other pluses include tighter security and the ability to free up time for technical experts to work with business managers to develop new products and services. All of which provide value and must be factored into decision processes.

All the heavy lifting to nail down prices may yield only short-term benefits. The reason: These calculations are snapshots of IT environments that by their nature are inherently dynamic—always changing and usually constantly growing. And that often leads to cost creep, which not only affects budget estimates, but may eventually wipe out any cost advantages anticipated from moving to the cloud.

Thus, while estimating costs is important, that's just one piece of the decision-making puzzle when deciding whether to move to public cloud. Perhaps even more significant are workload requirements.



STEP 2

Understand Workload and Optimization Considerations



\$1,000

Cost of machine that compares
cloud and on-premise costs

Before enterprises can successfully use the public cloud to optimize costs, planners must evaluate each workload to determine which computing platform—on-premise or cloud—will achieve the best balance of cost and performance. For example, public infrastructure-as-a-service models enable organizations to spin up compute power on demand to support sudden increases in demand, such as traffic to a website that results from a new marketing campaign.

But this so-called elasticity comes at a price—as published pricing sheets show, vendors offer discounts for more predictable usage requirements, making on-demand pricing an expensive option. To determine when elasticity pays off, when long-term contracts are a better choice, or when an enterprise should stick with an internal data center requires accurate data and insights about the unique characteristics of each workload. “Important considerations include cost, whether performance allows for certain deployment models, regulatory and security considerations, and architectural restrictions—a mainframe application may be impossible to run anywhere but on premise,” Lewis says.

The value of individual data sets and how they’re used by internal stakeholders also

influences deployment decisions. “If you need to blend and consume data with other information, then the associated workloads probably need to run where everything can be consolidated together,” he adds.

The staff at the University of Massachusetts Medical School created a baseline estimate based on virtual machines to compare cloud and on-premise costs. The \$1,000 per machine, per year benchmark is based on what the medical school charges for resources available from a shared, on-premise private cloud. “The biggest factors we look when gauging total on-premises costs include not just the expenses associated with the shared data center itself, but all salaries, training, time to spin up servers, and depreciation,” Wolf says. “On-premise costs are expensive, and we’re hopeful that as we go further down the cloud road, and take more advantage of discounted pricing options, we’ll be able to fit more workloads into a lower-cost cloud offering.”

What’s needed to make complete cost analyses like these is an application that takes a comprehensive view of the entire IT infrastructure, from real-time usage patterns and associated costs to idle or underused IT resources and the impact of cloud deployments over time.

A hand holds a smartphone displaying a colorful data visualization. Overlaid on the image are several white, semi-transparent data visualization elements: a bar chart at the top left, a line graph with a peak in the center, and four circular progress indicators arranged in a 2x2 grid. The background is a blurred image of a person's face and a laptop screen.

STEP 3 Pick the Right Cost-Analysis Tool

**VINCE CAMPISI**

Until you sign up for cloud services, it's easy to underestimate the amount of unmet demand that exists throughout the organization. If you're not careful, spending may be higher than anticipated."

Cloud providers offer tracking tools to monitor spending. But larger enterprises run a diverse range of IT resources that include multiple public cloud platforms and private clouds supported by technology from many vendors and running across geographic locations. The inability to manage this environment effectively entails the risk of redundant and overprovisioned services. "Until you sign up for cloud services, it's easy to underestimate the amount of unmet demand that exists throughout the organization," Campisi says. "If you're not careful, spending may be higher than anticipated because it's so easy for people to just sign up for additional capabilities. It's like a faucet—people turn on the tap and use as much as they want."

Understanding these complex operations requires a central view of on-premises and public-cloud infrastructure usage and spending. Fortunately, sophisticated tools are now available that automatically identify underused resources to guard against expensive overprovisioning of cloud services. CIOs should look for five key capabilities as they evaluate analysis applications.

- 1 SIMULATIONS:** These models perform what-if analyses of various cloud and onsite deployment scenarios to find the most cost-effective option for individual workloads, based on their size and service types.
- 2 COST TREND ANALYSES:** Enterprises can quickly monitor daily compute usage and costs, with summaries by application, service type, provider and other factors. In addition, advanced tools can rank tools from most to least expensive and flag underused and idle resources that are candidates for consolidation or elimination. Centralized dashboards give a clear view of on-premises and public-cloud infrastructure expenditures, with data relevant to IT and business users.
- 3 PUBLIC-CLOUD COST COMPARISONS:** These breakdowns by cloud service provider offer insights into critical factors, such as who's using each service, usage rates and overall costs to the enterprise.
- 4 CHARGE-BACK REPORTS:** For budgeting and forecasting, updated summaries of operational expenses show business managers the costs being accrued by their departments for IT services.
- 5 PREDICTIVE ANALYTICS:** Underlying algorithms drill into utilization and cost totals and help identify emerging problems and threats to achieving service level agreements. The best can also generate yearly cost forecasts based on current utilization and future resource needs.



These analyses will help CIOs more carefully monitor the multiyear migration that enterprises are taking from traditional data centers to modern, cloud-optimized infrastructures. For example, the Enterprise Strategy Group found that 44% of IT executives still equally consider both on-premises technology and public-cloud services when deploying new workloads. In comparison, slightly over a third of respondents used a cloud-first policy, while 19% deployed on premises first.

“Most enterprises can’t quickly move hundreds of applications to the cloud to create one operating environment,” says Lewis. “It’s much more likely that they’ll continue to deploy in a combination of traditional data centers, private clouds and public clouds. So for now, clouds may add expenses to their operating models, such as for new people who know how to architect applications for the cloud. But over time, once most applications go to the cloud, we will see cost savings.”



STEP 4
**Turn Insights
Into Action**

In addition to analyzing IT usage and costs, sophisticated cost-management platforms can take action, based on guidelines set by the IT staff. This enables organizations to quickly adjust to changes in demand, usage and costs. For example, in a hybrid cloud environment, the platform can automatically move a select group of on-premise workloads to the cloud if a more economical option surfaces or to relieve performance bottlenecks that threaten service level agreements. “The ability to do apples-to-apples comparisons of cloud providers and internal data centers could enable more a la carte pricing in the future,” according to Bray. “Managers can run what-if analyses by saying, ‘This is what I’m looking for, what are my options?’”

Enterprises may also be able to squeeze more costs out of IT infrastructures in the months ahead. To do this, many IT leaders will avoid layoffs or project cancellations and instead try to leverage their buying power to renegotiate the terms of existing contracts. Detailed reports about costs and usage created through thorough analyses will give these firms ammunition when sitting down with cloud vendors.

DAVID BRAY

“The ability to do apples-to-apples comparisons of cloud providers and internal data centers could enable more a la carte pricing in the future. Managers can run what-if analyses by saying, ‘This is what I’m looking for, what are my options?’”

The background image shows a close-up of a person's hands typing on a laptop keyboard. The image is overlaid with several digital and network-themed graphics. In the top left, there is a line graph with several data points connected by lines. In the center, there are several horizontal bars of varying lengths, some with dots at their ends, and a network of lines connecting these dots. In the bottom left, there are concentric circles and a central dot, resembling a target or a radar screen. The overall color palette is a mix of warm and cool tones, with a blue tint in the lower half and warmer tones in the upper half.

STEP 5 Create a Plan for the Future



With an understanding of true cloud costs, decision makers can plan for the future—based on accurate estimates of expected changes in workload requirements and how this impacts the mix of services. “We’re currently discussing what we can do with these granular details to make good business decisions,” says CIO Wolf.

He and his staff have already decided that all the data associated with the research ecosystem will move to the cloud, so the next step will be to complete that migration process. “We’ll also focus on leveraging the cloud more for its ability to provide elastic compute resources.”

Rather than guesstimates and gut feel, a growing number of IT executives are now updating their cloud road maps based on timely and accurate data, ensuring that their IT infrastructure priorities are aligned with their strategic business goals.

Acknowledgments



Forbes Insights and BMC would like to thank the following individuals for their time and expertise:

- **DAVID BRAY**, Executive Director, People Centered Internet
- **VINCE CAMPISI**, CIO and Chief Digital Officer, United Technologies
- **BRIAN GARRETT**, Vice President of Validation Services, Enterprise Strategy Group
- **PAUL LANGLOIS**, Director of Research Computing, University of Massachusetts Medical Center
- **PAUL LEWIS**, Chief Technology Officer, Hitachi Vantara

ABOUT FORBES INSIGHTS

Forbes Insights is the strategic research and thought leadership practice of Forbes Media, publisher of Forbes magazine and Forbes.com, whose combined media properties reach nearly 94 million business decision makers worldwide on a monthly basis. Taking advantage of a proprietary database of senior-level executives in the Forbes community, Forbes Insights conducts research on a host of topics of interest to C-level executives, senior marketing professionals, small business owners and those who aspire to positions of leadership, as well as providing deep insights into issues and trends surrounding wealth creation and wealth management.

FORBES INSIGHTS

Bruce Rogers
CHIEF INSIGHTS OFFICER

Erika Maguire
PROGRAM DIRECTOR

Casey Clifford
DIRECTOR, ACCOUNT MANAGER

EDITORIAL

Kasia Wandycz Moreno DIRECTOR

Hugo S. Moreno DIRECTOR

Alan Joch REPORT AUTHOR

Zehava Pasternak DESIGNER

RESEARCH

Ross Gagnon DIRECTOR

Kimberly Kurata SENIOR RESEARCH ANALYST

Sara Chin RESEARCH ANALYST

SALES

North America
Brian McLeod EXECUTIVE DIRECTOR
bmcLeod@forbes.com

Matthew Muszala DIRECTOR
mmuszala@forbes.com

William Thompson MANAGER
wthompson@forbes.com

EMEA

Tibor Fuchsel MANAGER
tfuchsel@forbes.com

APAC

Serene Lee EXECUTIVE DIRECTOR
slee@forbesasia.com.sg