

PROACTIVE STEPS FOR MANAGING CLOUD COSTS

Proactive Steps for Managing Cloud Costs

Summary

Cloud usage and cloud spend continue to grow. Spend on infrastructure-as-a-service (IaaS) and software-as-a-service (SaaS) was predicted to reach well over \$200 billion by 2020—and that was before COVID-19 accelerated many organisations' moves to the cloud. Factor in the increased number of vendors offering services and applications via the cloud and it becomes imperative for businesses to understand, optimise and manage their cloud spending as much as possible. Furthermore, the economic uncertainty created by the COVID-19 pandemic makes reducing costs more important than ever.

The cloud offers different purchasing methods

For many years, software was purchased in a centralised manner through large and expensive multi-year contracts whose negotiations had a defined sign-off process. The cloud disrupts those traditional methods. Now, anyone with a credit card can purchase enterprise-level cloud services—both IaaS and SaaS—with the corporation having little to no awareness. This lack of central oversight can easily lead to unnecessarily higher levels of spending.


IT asset managers have worked to understand, manage and optimise an evolving range of on-premises technologies, and now they must add cloud management responsibilities.

Impact of COVID-19

The 2020 COVID-19 outbreak caused organisations to take immediate and, in some cases, severe actions to stay afloat. Some of those decisions included:

- Delaying capex investments
- Deferring on-premises implementations/upgrades
- Employee furloughs
- Workforce reductions

Cloud considerations play a large role in the wake of the pandemic. The *2020 Flexera State of the Cloud Report* shows 59 percent of enterprise respondents expect COVID-19 will drive “higher than planned” cloud usage, with 30 percent expecting it to be “significantly higher.”



The data on IaaS usage

The 2020 Flexera State of the Cloud Report shows that 20 percent of respondents spend more than \$12 million a year on public cloud. Thirty-six percent of Microsoft Azure customers and 40 percent of Amazon Web Services (AWS) customers spend at least \$1.2 million.

With survey respondents estimating their average cloud waste to be approximately 30 percent, and Flexera research suggesting it's actually closer to 37 percent, there's plenty of opportunity to reduce public cloud spending.

The always-on-as-a-service nature of public cloud means unnecessary spending is occurring daily. However, it also means any savings you're able to achieve could have almost immediate impact and visibility. Unlike on-premises agreements, where reductions may not be visible until contract renewal in 3-5 years, actions taken to reduce cloud spend can be seen in many cases within 30 days—often within 24 hours.

Optimising IaaS spend

Use available discounts

It's clear that organisations aren't taking advantage of all the discounts available from the public cloud providers, as recent Flexera research shows:

Discount Types Used by Cloud Provider		
AWS	AZURE	GOOGLE
AWS Reserved Instances 53%	Enterprise Agreement 51%	Committed use discounts 43%
AWS Spot Instances 38%	Azure Reserved Instances 43%	Ad hoc negotiated discounts 18%
AWS EDP (Enterprise Discount) 37%	Azure Hybrid Benefit 30%	
AWS Savings Plan 30%	Azure Low Priority VMs 24%	
Ad hoc negotiated discounts 20%		

N=750 Source: Flexera 2020 State of the Cloud Report

With even the most widely utilised discount barely breaking the 50 percent mark, organisations are leaving a lot of money on the table. Many of these options are programmatically available and easy to access without complex negotiations. In some cases, it's simply a matter of checking a box.

Reserved instances

Reserved instances are available across the three major public cloud platforms:

- AWS
- Microsoft Azure
- Google Cloud Platform

Reserved instances (which Google calls Committed Use Discounts) allow you to access reduced rates on certain cloud services such as virtual machines by committing in advance—typically for a period of one or three years. These work best on stable cloud workloads running for long periods of time with little or no changes.

Applying reserved instances to your existing public cloud workloads can bring savings of up to 70 percent, according to the providers. And while the actual discounts you receive may not be as large, they'll still represent a significant savings over the default pay-as-you-go (PAYG) pricing.

Ongoing reserved instance management


Buying reserved instances is only half the battle. You must ensure they're utilised and continue to align with resources being used within your cloud environment. This requires a close relationship with the various teams architecting, deploying and using cloud resources throughout your organisation. Reserved instances are typically purchased for specific types of cloud resources. If cloud usage deviates from those reserved instances, you'll find yourself paying twice—once for the already committed reserved instance and again for the new resources, which will be at the full PAYG price.

Before purchasing reserved instances, it's important to understand the potential for future changes to cloud resources throughout the lifecycle, work to ensure the purchased reserved instances are being utilised and determine whether any additional purchases or changes will need to be made. The ability to modify and cancel reserved instances varies between public cloud providers but as a general rule, doing either will mean you fail to access the largest available discounts and, in the event of cancellations, forfeit a portion of the money already committed.


Azure Hybrid Benefit

Azure Hybrid Benefit is available only with Microsoft Azure. It currently applies to Windows Server and SQL Server and allows you to use on-premises licences with active Software Assurance (SA) to reduce cloud spend in Microsoft Azure.

When you turn on a Windows Server virtual machine (VM) in Microsoft Azure, you're charged for two components—the underlying server hardware that physically runs your VM and the Windows Server Operating System (OS) that runs on top of that hardware.



With the Azure Hybrid Benefit, you're instead charged only for the underlying server hardware.



This benefit is subject to Microsoft licencing rules, and some elements differ between the products for which the benefit is available. The on-premises software entitlement and licencing rules make IT asset management (ITAM) well positioned to manage the correct and compliant application of the Azure Hybrid Benefit across the hybrid (on-premises and cloud) estate.

Windows Server

Windows Server Standard licences may be used on-premises or in Microsoft Azure, but not simultaneously. However, 180 days of simultaneous use is permitted while workloads are being migrated from on-premises to Microsoft Azure.

Windows Server Datacenter licences may be used on-premises and in Microsoft Azure simultaneously. But the Windows Server Datacenter unlimited virtualisation benefit does not apply in Microsoft Azure.

For both types of Windows licences, every 16 core licence with SA allows up to 16 cores to be used in Microsoft Azure.

SQL Server

The Azure Hybrid Benefit can be applied both to SQL Server virtual machines and SQL database options—the latter being the platform-as-a-service (PaaS) hosted SQL offerings. The licencing rules differ between SQL Server Standard and Enterprise editions. Rules that apply to both include:

- Only pertains to the per-core licencing method
- Four core licencing minimum applies in Microsoft Azure
- No simultaneous use—on-premises and in Microsoft Azure—is allowed*

*Simultaneous use for 180 days is permitted while workloads are being migrated from on-premises to Microsoft Azure.

SQL Server Standard

Every on-premises core licence with SA allows one of the following:

- One core in SQL Server Standard Edition virtual machine
- One core in General Purpose SQL database tier
- One core in Hyperscale SQL database tier
- Every four on-premises core licences with SA allow one core in SQL Server Enterprise Edition virtual machine

SQL Server Enterprise

Every on-premises core licence with SA allows one of the following:

- One core in SQL Server Enterprise Edition virtual machine
- Four cores in SQL Server Standard Edition virtual machine
- One core in Business Critical SQL database tier
- Four cores in General Purpose SQL database tier
- Four cores in Hyperscale SQL database tier

Taking advantage of the Azure Hybrid Benefit

To ensure you're taking full advantage of the Azure Hybrid Benefit, you must:

- 1) Review your Microsoft licencing position to understand your current on-premises licence entitlement, with SA, for Windows Server and SQL Server
- 2) Get an overview of your use of these products in the cloud, not just in Microsoft Azure but other cloud providers as well

The second point offers great potential for finding additional savings and to show the proactive possibilities of ITAM, as many organisations use Microsoft products on non-Microsoft cloud platforms. Analyse your relevant spend on other cloud platforms and calculate what it would cost if it were moved to Microsoft Azure with the Hybrid Benefit applied. Showing internal stakeholders the potential savings will help ensure cloud platform decisions are financially well informed.

The Microsoft Azure Hybrid Benefit calculator can help you estimate savings that could be available: [Microsoft Azure Hybrid Benefit calculator](#)

Turn resources off

Turning off resources is one of the simplest and most impactful ways to save money in the cloud. Compared to on-premises, a significant shift in thinking must occur for an organisation to have a truly effective cloud cost management policy. For on-premises, the bulk of the spend is upfront on the physical server and the software licences running on it. Once that spend has occurred, it's better for those resources to remain on and be used rather than decommissioned and stored away. It's the complete opposite in the cloud. Ideally, your cloud resources should be turned on only when you need them. For internal stakeholders and team members, this shift can be difficult to master at first and cause significant wasted cloud spend.

Test and development

While it's understandable that cloud resources in a production environment—such as a website, ordering portal or customer service system—would need to be running 24/7, this is probably not the case for the development systems that support them. But it is the way public cloud vendors typically price their resources online.

If your test and development teams work ten hours a day, five days a week, that's 50 hours per week. So why pay for 168 hours of resource uptime? That's potentially 118 hours per week of unnecessary cloud spend. Turning things off when the team isn't at work can reduce your costs as much as 70 percent. While there are considerations—particularly if teams are spread across different time zones—setting policies and procedures to reduce the amount of wasted uptime will help curb your cloud costs.

Rightsizing


The process of making sure resources deployed in the cloud aren't more than what's needed is called rightsizing. One of the main causes of oversized cloud resources is when typical

on-premises hardware specifications continue to be used for deploying cloud services. For various reasons such as newer hardware, tighter optimisation and more efficient utilisation, the same tasks can often be completed in the cloud with fewer resources than were previously required on-premises.

Defining the correct hardware specifications isn't typically something that can be done by an ITAM team, so work with your infrastructure and cloud architect teams to address this issue. An organisation should first implement a review of existing cloud resources to identify any immediate savings and then work to set guardrails for future deployments to prevent waste.

Process

As we've seen, there are several elements to reducing your IaaS cloud spend, and it's important to realise this is an ongoing process. Just as with on-premises SAM and ITAM, cloud cost management is never finished. As new resources, staff and vendors continue to be added to your environment, you'll find you must regularly revisit each of the steps shown below.



Communication with stakeholders

Key for successful cloud cost management is involving relevant stakeholders across the business. Once processes have been defined, they'll best serve your organisation if applied before resources are created. This is a task for those who directly create and manage the resources daily.

While the ITAM department should regularly review policy adherence and implementation, a looking-backwards approach means waste will continue. Good communication with teams will help ensure waste is prevented before it happens. A great strategy is to outline to the teams the processes and the reasons behind them, define ongoing metrics and regularly meet with department heads to monitor progress.

Optimising on-premises infrastructure

Although it may at first sound counterintuitive, optimising your on-premises environment can help reduce your future cloud costs.

For a variety of reasons, there are often on-premises resources you no longer need much of or at all. This can be caused by underutilised servers, projects being shut down while the hardware remains, a change in overall technology strategy, database consolidation projects and more. Simply moving everything to the cloud as-is will mean higher cloud bills and wasted spend. Instead, keep cloud bills low from day one by understanding the actual utilisation of your on-premises hardware.

Lift and shift versus move and improve

You may have heard the phrase *lift and shift*. This refers to the practice of simply replicating the on-premises hardware estate in the cloud—creating virtual deployments with the same number of servers, processors, cores, etc.

A preferable approach is to move and improve, whereby you move the resources to the cloud while improving them—such as reducing the number of cores due to more modern CPUs in the cloud or consolidating several database servers onto one cloud service.

A good analogy for this process is moving to a new house. Ideally, you take only the necessary things to the new house to start off lean and fresh in your new location. Migrating your data centre into the public cloud is much the same.

SaaS usage

It's not just IaaS that must be considered; SaaS is continuing to grow, too. The *2020 Flexera State of the Cloud Report* showed that 43 percent of organisations cited moving on-premises software to SaaS as a top priority for this year. That's up from 29 percent in 2019. Many software vendors have already made a significant move from on-premises to SaaS—Adobe Creative Cloud and Microsoft Office 365 are two prime examples. And many others are on a similar path. This, coupled with the rise of SaaS-only vendors such as Zoom taking hold within organisations, means SaaS represents a significant portion of their software budgets.

Optimising SaaS spend

Discover your top applications

You can't manage the unknown. SaaS can create hidden problems because purchases are often made at a departmental level. It's important to identify which SaaS applications are in use across your organisation and then categorise them based on those which are most business-critical. They should be further categorised by function, what's most costly and what's most widespread—both in terms of number of users and across different departments/regions.

The sudden and unexpected surge of remote working has led to an increase in SaaS applications being used across organisations. For example, Zoom reports that daily users rose to 200 million (from a previous maximum of 10 million) during March 2020. While these won't all be corporate users, it helps illustrate Zoom's explosive growth. And it's a similar story for many other SaaS companies. As staff within your organisation find ways to continue working, they'll inevitably turn to the myriad SaaS apps in categories such as project management, task management and team collaboration.

Identify zombie users

SaaS applications can often lead to the proliferation of zombie users, i.e., situations where a licence exists without an assigned user. If SaaS purchases aren't tied to your onboarding and offboarding process, the potential exists for licences to sit in limbo rather than being reassigned to a new user. The organisation continues to be charged for the SaaS app on a monthly basis, or it's included in annual renewals. This typically occurs when users leave but may also happen when they transfer between departments and certain role-specific software is no longer required. Wasted spend is then compounded when a new user starts and a new licence is purchased.

Rightsizing SaaS apps

Many SaaS applications are licenced through a variety of tiered options, each including more features than the last, such as Microsoft Office 365 and Microsoft 365 with F1, F3, E1, E3 and E5 variants. Because of this, it's entirely possible for SaaS licences to be overprovisioned in a similar way as IaaS resources. After identifying your top apps, delve into the specific licences assigned to each user, look at their usage of the various components and decide if they're licenced correctly. If they're not using the majority of

the elements, perhaps they can be dropped down a tier or two. This will require good understanding of the licencing and renewal requirements of a potentially wide range of SaaS vendors.

Product overlap

The easy access and decentralised department-led purchasing of SaaS have made product overlap extremely easy. Product overlap occurs when an organisation has multiple products performing the same task. Video conferencing/team collaboration and cloud storage are two common areas where this occurs due to the proliferation of services available in the market.

What to do

Look at your list of SaaS applications and identify feature overlap—some features may overlap with several other apps. Then it's important to ascertain:

- Are users licenced for multiple services?
- If so, why? Is it due to essential feature differences between the different products, because that's how things have been configured internally, or is it simply personal preference?

Cutting the number of services in use will not only reduce your software spend but also help reduce internal soft costs associated with managing multiple services, including training staff to use them, supporting helpdesk tickets and processing renewals.

Overall cloud cost management concepts

Finished projects

One of the primary selling points of the cloud, both IaaS and SaaS, is that it isn't permanent. You can turn it on when you need it and off when you don't. Turning off cloud resources until needed again can be a great way to reduce costs. But what about that spend that won't be needed again?

Most organisations have several time-limited projects running at any one time, and these can be another source of wasted cloud spend. When a project is finished, organisations generally remember to collect hardware such as mobile phones, laptops and monitors. And if the organisation has a strong software harvesting policy, they'll remove and pool software licences where applicable. What many organisations neglect to do is turn off associated IaaS resources in the public cloud and remove SaaS licences that are no longer needed from upcoming renewals. This could be due to lack of visibility—not knowing those resources are in use—or possibly because of a lack of processes.

Ideally, ITAM should have:

- Visibility into all active projects
- A process for understanding each project's use of IaaS and SaaS
- A procedure for shutting down and removing resources at the end of the project

COVID-19 impact

The above method of cost management is particularly relevant during the current COVID-19 pandemic. As many organisations are being forced to furlough or lay off staff, cloud resources can often be forgotten. While looking at which teams within your organisation have been

impacted, you can identify who may have been using IaaS services and which users can be removed from SaaS licence counts and renewals.

The cloud and security

Reducing the amount of cloud resources and services in use will also positively impact your organisation's security posture. Every cloud resource is a potential entry point for attackers, so reducing your footprint is a key element of a modern security policy. Wherever possible, meet with your security teams regularly and share data regarding cloud usage. They may be able to identify certain cloud resources you'd missed and vice versa.

Managing the cloud, moving forward, making progress

In order to fully manage assets and software spend, cloud cost management for both IaaS and SaaS must become a key aspect of any ITAM function. It currently represents a great opportunity to quickly cut costs and improve cashflow to reduce the short- and long-term impact of COVID-19. Take stock of cloud within your organisation, both current use and future plans, and start building your cloud cost management plan. It's important to consider:

- What can you do immediately?
- Do you have the bandwidth to focus on both IaaS and SaaS simultaneously?
- Which stakeholders do you need to engage?

A small win is better than a big win that never materialises, so look at where you can achieve the largest impact in the shortest time frame. And remember: the sooner you start, the sooner you save.

About the ITAM Review

This document was created in a joint effort with Flexera and ITAM Review, an independent global community for worldwide ITAM, SAM and licencing professionals. The ITAM Review exists to champion the business value of the ITAM profession, develop ITAM industry leaders, and enable organisations worldwide to extract best value from their technology investments whilst minimising risk.

To learn more, visit itasassetmanagement.net

NEXT STEPS

Find out how Flexera can help you manage cloud costs

[DISCOVER MORE](#)

ABOUT FLEXERA

Flexera helps business leaders succeed at what once seemed impossible: getting full visibility into, and control of, their company's technology "black hole". From on-premises to the cloud, Flexera helps organisations unravel IT complexity and maximize business value from their technology investments. For more than 30 years, our 1,300+ team members worldwide have been passionate about helping our more than 50,000 customers optimise IT to achieve their business outcomes. To learn more, visit flexera.com