



Use workload placement to optimize hybrid IT strategy

A strategy framework that aligns cloud decisions with business priorities

Author:

Michael Fitzgerald
Sr. Director, Engineering
Flexential




Intelligent workload placement is driving the use of multi-cloud architectures and is the fastest-growing driver for multi-cloud architectures, increasing 20% year over year.¹

This rise in cloud computing has accelerated hybrid IT as organizations with on-premises or colocation environments increasingly adopt cloud-based environments and infrastructure services. As these cloud services grow, so does the view of hybrid IT as an approach for embracing the advantages of combining on-premises and colocation with a multi-cloud architecture. Most organizations are taking a multi-cloud, hybrid IT approach in which private cloud plays an essential role. 87% of organizations have embraced multi-cloud, 72% of which are hybrid cloud (a mix of private and public as opposed to multiple public or multiple private clouds).¹

Organizations that approach hybrid IT strategy development from a workload perspective and have a workload placement framework as part of their hybrid IT decision-making can achieve greater flexibility, scalability, security, and lower IT costs.

Creating or moving into cloud environments without a thorough, upfront decision analysis and a mature and documented hybrid IT strategy can result in systems and services that are challenging to manage, introduce security risk, and are more expensive than necessary. While fast and convenient in the short term, ad hoc or reactive decision-making can result in a landscape that poorly serves users, IT, and business priorities.



87% of organizations have embraced multi-cloud, 72% of which are **hybrid cloud**.¹

A workload-focused framework provides a structure for sound placement decisions and is critical to a successful hybrid IT strategy and realizing a future-state hybrid IT vision where optimized IT aligns with business goals

How can an organization shift from ad hoc or reactive approaches to a strategic approach and plan hybrid IT environments that optimize cloud adoption and avoid poor planning pitfalls? A workload-focused framework provides a structure for sound placement decisions and is critical to a successful hybrid IT strategy and realizing a future-state hybrid IT vision where optimized IT aligns with business goals.

1. "2023 State of the Cloud Report" Flexera

Getting started

Begin with a series of technical, operational, and business questions to understand the current state.

These questions help an organization reflect on current challenges and operating modes regarding cloud and hybrid IT. They also help highlight the ramifications resulting from previous decisions and provide a learning opportunity.

- Do we have a documented cloud or hybrid IT strategy, or are we making ad hoc or reactive decisions?
- Does our current cloud or hybrid IT strategy, if we have one, align with our business priorities?
- Can we accurately view where workloads reside today, their use, and their dependent relationships?
- Do all our workloads currently reside in locations that deliver the needed performance, reliability, agility, scalability, security, compliance, and user experience?
- Have we overcome past technology challenges in adapting to changing needs or business priorities?



An organization's workloads should drive its platform and infrastructure requirements—not the other way around.

The critical decision point: Workload placement

Identifying the best placement for a given workload is what a successful hybrid IT strategy seeks to accomplish. The challenge is prioritizing multiple discrete workload needs and matching them to the many placement and service options available.

The challenge is prioritizing multiple discrete workload needs and matching them to the many placement and service options available

Organizations’ top challenges in migrating workloads to public cloud are understanding application dependencies, assessing technical feasibility, rightsizing, and assessing operating costs.¹

An organization’s workloads should drive its platform and infrastructure requirements—not the other way around. As organizations transition to or further adopt a hybrid IT approach, many choose placements based on what is known, familiar, or most similar to their current environment.

Without considering a broad scope of placement options, and their benefits, this approach can quickly eliminate the best placement for a workload.

To illustrate those choices, each placement option below represents a unique model with potential business advantages and disadvantages. This type of structured approach to decision-making yields intelligent investments of the available dollars and personnel.

Note that this placement options chart is simplified because its purpose is not to be exhaustive but to illustrate considerations when making workload placement decisions. A full spectrum of options would include more granularity, for instance, organizationally owned and operated private cloud deployed on-premises or in a colocation facility.



Multicloud is good; sprawl is no longer the enemy. Lack of strategy and planning is the enemy.

	On-premises	Colocation	Hosted private cloud	Public cloud	PaaS	SaaS
Description	Infrastructure, workloads, and data in the organization's own data center	Third-party-provided data center space, cooling, power, and network connectivity	Third-party-provided dedicated infrastructure with managed hardware and hypervisors	A third-party data center with multi-tenant managed cloud infrastructure	Application and server-less infrastructure components within a third-party managed cloud	Licensed for use applications, hosted on a vendor managed platform, including upgrades
Data center location	Customer facility	Third-party DC, customers may be able to choose a region, country, or city	Third-party DC, customers may be able to choose a region, country, or city	Third-party public cloud DC-customers may be able to choose a region, country, or city	Platform vendor's choice	SW vendor's choice
Services and provider infrastructure management responsibility	None	Space, power, cooling, connectivity	Dedicated, managed infrastructure	Shared, managed infrastructure	Managed infrastructure, application components and server-less components	Managed application and managed platform
Provider security and compliance responsibility	None	Managed facility physical security	Managed up to the hypervisor	Managed up to the hypervisor	Managed up to the application components and server-less components	Managed up to the application

1. "2023 State of the Cloud Report" Flexera

Use a framework to guide workload placement and cloud decisions

Top cloud challenges for organizations are managing cloud spend, security, and lack of resources or expertise.¹

Often, what an organization is trying to accomplish with workload placement decisions is solving a problem or creating a specific outcome. The top initiatives for organizations are optimizing the existing use of cloud, progressing on a cloud-first strategy, migrating more workloads to cloud, and moving from on-premises to SaaS versions of software applications.¹

An analysis utilizing multiple workload placement criteria can identify and prioritize the most critical factors and drive an optimal decision.

Many operational, financial, and technical aspects can influence an individual workload placement decision. There are so many of these workload placement criteria to consider that organizations need a clear and structured framework to increase the speed and success of workload placement decisions.

Identifying the best placement for a given workload is what a **successful hybrid IT strategy** needs to accomplish.

Flexential has created a workload placement framework to guide workload placement decisions using 14 operational, financial, and technical criteria:




Operational

- Operational responsibilities
- Control and visibility
- IT skill set
- Portability
- Security
- Compliance



Financial

- Cost structure
- Cost control
- Cost management



Technical


- Latency
- Performance
- Scalability
- Availability
- Technical Requirements
- Dependencies

1. "2023 State of the Cloud Report" Flexera

Operational workload placement criteria


Operational responsibilities

Outsourcing infrastructure, applications, or services shifts operational responsibilities to a third party and frees up internal IT resources, but when a service impact event occurs, clarity on who is responsible for resolution becomes crucial. When management and availability are the responsibility of a service provider, it is critical to understand the provider’s SLAs and RACIs and identify gaps before a disruption occurs. It is similarly necessary to know how service providers perform lifecycle and patch management activities, as their maintenance procedures and operating policies may conflict with organizational requirements. The availability of outsourced services may be limited for older OSs or applications. Select a provider that can provide support on legacy components if it is required.



Control and visibility


Administrative access to environments for transparent service management is one reason organizations invest in on-premises or colocated solutions. Control and visibility of components decrease with off-premises and cloud environments as infrastructure ownership and management responsibilities are offloaded to providers. SaaS solutions provide the least visibility and control as non-vendor responsibilities are limited to configuration, user access controls, and application data. Control and visibility can be significant factors in highly compliant environments that require regular reporting, auditing, and security testing.



IT skill set

Seventy-eight percent of organizations report that a lack of expertise in managing their cloud environments is a top cloud challenge.¹ Organizations have an opportunity to align and complement provider choices with their existing in-house skill sets. A lack of internal IT skills related to a particular hosting model or vendor is an opportunity to consider leveraging external support to fill team or department knowledge gaps. A strength in certain IT skills will also lead toward or away from certain providers.

An honest assessment of internal skill sets is necessary to understand current IT staff expertise. Once an assessment is done, IT should prioritize the skill gaps to fill in-house, skills to de-prioritize, and the extent of third-party services needed. Managed services can effectively augment staff skill sets and shift IT focus to providing more strategic value to the organization—based on the strategic focus the organization has previously determined.



1. "2023 State of the Cloud Report" Flexera

Portability

The ease of moving a workload from one environment to another becomes a critical criterion when performing activities such as hardware or software refreshes, changing service providers, or moving workloads closer to users. Containerization can make portability more straightforward, but not every workload is well-suited for containers. Container management also requires a specialized skill set that may not currently be found in the IT staff. Additionally, hardware dependent licensing models may introduce cost considerations if moving to different hardware or a new environment, and licenses tied to MAC addresses or hardware keys may be difficult in cloud environments. The use of PaaS also impacts portability as certain functions or specific implementations may only exist on that service provider, requiring code-level changes to re-platform.




Security

Cloud solutions include a suite of technical and procedural security controls, but organizations will still have some responsibilities in any solution—even if it's only user access management for a SaaS application. Providers build security into their services, from ISO-certified colocation data centers to fully managed application security from SaaS vendors, yet 79% of organizations report security as a top challenge in cloud adoption.¹

Security visibility and testability are essential for tracking and validating changes and managing risk. Providers will differ in what they will be willing to be responsible for and what tools and processes they can support or integrate.

An essential part of maintaining security is data protection—both backups and disaster recovery. Availability of managed data protection will vary per placement location and provider. In general, an on-premises environment must provide its own data protection, while cloud-based locations offer a suite of managed solutions and SaaS providers include backups and managed DR as part of their standard service subscriptions.



Compliance


Organizations must have sufficient controls, auditability, and attestations to maintain compliance with standards, rules, and regulations. While cloud providers maintain compliance each customer is ultimately responsible for their own compliance program and must confirm that chosen vendors meet their specific requirements.

A business-associated agreement (BAA) holds providers legally accountable for their portion of compliance. BAAs will be the broadest in SaaS because customers have minimal visibility and access to underlying infrastructure and controls. SaaS and PaaS providers can result in the lowest overall compliance burden for customers since these are highly managed solutions.

1. "2023 State of the Cloud Report" Flexera

8 | Use workload placement to optimize hybrid IT strategy


Many U.S. government compliance regulations require that the hosting datacenter or cloud be FedRAMP certified. Not all vendors hold this status which may limit options. Additionally, increased development and required adherence to government instituted privacy laws may require an organization to maintain a high level of access to and management of user and customer data, which may be difficult in SaaS environments.



Financial workload placement criteria

Cost structure


Cloud computing provides an opportunity to shift budgets toward OPEX with third-party services and subscriptions and thereby dramatically decrease CAPEX spend. CAPEX-based approaches frequently involve inflexible, longer-term financial positions that engender investment risk for capital-intensive infrastructure. CAPEX for data center expansions and infrastructure refreshes can be reduced or avoided by moving to colocation or cloud-based solutions. The exercise of comparing OPEX and CAPEX options as part of the workload placement decision can surface optimal cost structures for workloads.



Cost control

It is essential to understand the cost variability of an individual workload for a given billing format. 60% of organizations report cost efficiency and savings as the top metric for assessing progress against cloud goals and report an estimated 28% of wasted spend.¹ Reserved resources can provide more stable cost control and predictability compared to use based or metered pricing structures.

With cloud-based automatic scaling and ease of provisioning comes an increased risk of runaway costs. Rapid provisioning, and empowering developers and larger groups of users to create their own environments, shortens deployment and development cycles but makes managing unused resources and services crucial to avoid extraneous expenses. Moreover, provisioning and deprovisioning may entail moving data to and from the environment, and data ingress or egress charges can be significant for large data sets. It becomes easy to see how hyperscale clouds require cost management to a greater degree than traditional physical server-based environments or more closely controlled private clouds.



1. "2023 State of the Cloud Report" Flexera

Technical workload placement criteria


Latency

Network connectivity and latency—is critical to productivity and user experience. This criticality is compounded in Hybrid IT multi-cloud deployment models. Latency impact is primarily a result of the location and distribution of users and dependent workloads compared to the data sources and workload location.

The analysis for latency includes consideration of user proximity and dependent workloads. Latency-sensitive workloads should be placed as near as possible to users to meet defined latency targets. On-premises solutions are


best suited for local and centralized users, while Software-as-a-Service (SaaS) solutions are well-suited for widely dispersed users. Similarly, it is important to consider the proximity and latency sensitivity of dependent workloads.

Edge computing use is growing as the approach to attain the low latency of locally-based workloads without the high management overhead of owned infrastructure. Knowing the available cloud regions and testing latency ahead of moving into production is recommended.




Performance

Compute and network performance is a combination of computing capacity, storage speeds, and network resources between workloads or sites. While high performance can be obtained in multiple ways, resources are most available, controlled, and monitored when using purpose-built equipment that resides on-premises or in colocation, or dedicated hardware and reserved compute power from private cloud providers. Within the public cloud space, compute, storage and network resources will be more available if they are allocated and reserved. For workloads with data growth from distributed sources, such as mobile and IoT, more distributed workload models provide greater performance benefits than a centralized workload model.



Scalability

Cloud-based implementations allow computing and storage resources to scale dynamically. Workloads that require “scale-up” by the addition of computing resources are best suited to on-premises, colocation, and some private cloud environments and may be constrained by hardware limitations and procurement lead times in placement options that require dedicated hardware. Public cloud platforms are well suited to workloads that “scale out” through microservices, load balancing, and automated server deployment. Public cloud and platform-as-a-Service (PaaS) solutions offer the most elastic scalability, but such capabilities can also create runaway costs if not properly governed. In a SaaS model, the responsibility to scale effectively is outsourced to the software provider and, except for subscription size, is generally outside the direct control of the subscriber.




Availability

When a workload or infrastructure is designed, developed, and built around a particular platform, it can leverage that platform’s native features for high availability (HA), such as load balancing or clustering. Some commercial off-the-shelf applications, however, do not have native HA capabilities or features and rely on the infrastructure layer for HA. This often requires additional compute, storage, and network resources to facilitate. Public cloud providers can deliver availability based on infrastructure HA with redundant compute, network, and storage. For SaaS solutions, availability is the responsibility of the vendor, however customers must ensure the required SLAs are in place.




Technical requirements

Technical requirements include compute and storage minimum requirements, licensing considerations, and needs for specialized platforms, appliances, or physical equipment. For instance, specific legacy applications, such as those designed to run on mainframes or IBM® AS/400 systems, cannot be easily moved to x86 platforms and sometimes must remain on dedicated hardware. Likewise, network equipment may require colocation or on-premises deployment, as public cloud environments may restrict the deployment of hardware devices, and use agreements may limit functionality. Workloads that only operate on specific operating systems may also limit placement choice, especially if the operating system required has reached end-of-life or end-of-support.



Dependencies

The performance and reliable access of dependent components must be considered in workload placement decisions. In one study, 49% of respondents said that understanding application dependencies was the top migration challenge.¹ The possibility of workload performance challenges and interconnectivity complexity increases when dependencies span different environments and locations. Some workloads require that dependent components be placed within the same data center to maintain workload performance. Some workloads may have a dependency on external services and are therefore suited for well-connected colocation and clouds. PaaS and SaaS solutions may need custom integrations or APIs for cross-application data exchange. The principle of data gravity applies to workload placement as part of evaluating a workload’s dependencies. Workloads with the highest data access requirements are placed in the closest proximity to the data.



1. "2023 State of the Cloud Report" Flexera

The benefits of optimum workload placements

There is no one-size-fits-all path to navigating workload placement decisions or hybrid IT strategies. Analyzing individual workloads identifies the primary operational, financial, and technical drivers for a placement decision and the broader infrastructure and services landscape required. Resulting placements can vary significantly by individual workload and their related business goals and success criteria. An organization may prioritize cloud-based scalability or rapid provisioning using infrastructure as code for one workload. At the same time, a second workload needs an off-premises edge placement for improved user experience, and yet another will be kept on-premises for control, visibility, or simplified cost predictability.

Organizations that approach hybrid IT decision-making from a workload perspective and successfully place those workloads in optimum environments can achieve greater flexibility, scalability, security, and lower IT costs, putting themselves in a stronger position to accomplish their business goals.

A hybrid IT strategy framework navigates technology choices and organizational priorities

There are now so many hybrid IT choices that the new challenge is how to choose wisely. Using a hybrid IT strategy that includes a workload placement framework generates decisions that consistently optimize operational, financial, and technical considerations while keeping those decisions aligned with broader IT and business strategy.

There are now so many hybrid IT choices that the new challenge is how to choose wisely.

[Learn more](#)

About us

Flexential experts partner with customers on architecture and strategy, cloud migrations, cloud optimization, DevOps, cybersecurity, compliance, and IT resiliency to solve today's hybrid IT challenges. We take a consultative approach based on industry best practices and standards, combine our expertise, experience, and methodologies, and then tailor each engagement to each customer's needs. Wherever data, applications, or infrastructure resides, Flexential partners with IT teams, providing actionable guidance, expert execution, and advancing IT strategy, infrastructure, and security.

