



Use Workload Placement to **Optimize** **Hybrid IT**

A Strategy Framework that Aligns Cloud
Decisions with Business Priorities





Organizations that approach hybrid IT decision-making from a workload perspective and successfully place those workloads in the optimum environments can **achieve greater flexibility, scalability, security and lower IT costs.**

The cloud computing revolution has been underway for more than a decade, and today, 93% of enterprises have moved at least some of their data and workloads into cloud environments for increased efficiency and flexibility.¹

This rise in cloud computing accelerated hybrid IT as organizations with legacy on-premises or colocation environments added cloud-based environments and infrastructure services. As these third-party services grew, so did the view of hybrid IT as an approach to embrace the advantages of combining on-premises and off-premises environments (colocated and cloud-based). In a recent study, 57% of organizations said they are onboard with *cloud-centric* hybrid IT environments, but the same study showed that 43% have no formal plans or strategies in place for hybrid IT.²

43% of organizations have no formal plans or strategies in place for hybrid IT.²

Despite the many advantages of hybrid IT, a lack of strategy and forethought is ill-advised: Creating or moving into cloud environments without a thorough, upfront analysis for making

the decision can result in systems and services that are more fragile, non-transformable and 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.

A workload-focused framework provides a structure for sound placement decisions and is key to realizing a future-state hybrid IT vision where optimized IT aligns with business goals while remaining agile.

How, then, 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 key to realizing a future-state hybrid IT vision where optimized IT aligns with business goals while remaining agile.

¹ "Cloud Computing Trends: 2020 State of the Cloud Report," Flexera Blog. <https://www.flexera.com/blog/industry-trends/trend-of-cloud-computing-2020>

² <https://451research.com/services/customer-insight/voice-of-the-enterprise>

Getting Started

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

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

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

- Do we have a documented cloud strategy or a 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?
- Do we have an accurate view of where workloads reside today, their use, and what their inter-relationships are?
- Do all our workloads currently reside in locations that deliver the needed performance, security, reliability and user experience?
- Have we faced challenges in adapting infrastructure or applications to different needs or business priorities?

The Critical Decision Point: Workload Placement

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

Identifying the best placement for a given workload is what a successful hybrid IT strategy needs to accomplish. The challenge is prioritizing multiple individual workload needs and matching them to the many placement and service options available. To illustrate those choices, each placement option below represents a unique ownership and management model with potential business advantages and disadvantages. This type of structured approach to decision-making yields smart investments of the available dollars and personnel.

Note that this placement options chart is simplified because its purpose is not to be exhaustive but rather to illustrate considerations when making workload placement decisions. A full spectrum of options would include more granularity, for instance, private cloud implemented on-premises or in a colocation facility. As organizations consider their specific options with greater granularity, care should be exercised to minimize sprawl. It is acceptable to consider many options, but once an organization's core workloads have found a home, companies should carefully consider introducing new locations by weighing the benefits against the incremental costs and associated overhead.

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

	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 MANAGEMENT RESPONSIBILITY	None	Space, power, cooling, connectivity	Dedicated, managed infrastructure	Shared, managed infrastructure	Managed infrastructure, application components & 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 & server-less components	Managed up to the application

Using a Hybrid IT Strategy Framework to Guide Workload Placement and Cloud Decisions

Many technical, operational and financial aspects can sway an individual workload placement decision. There are so many of these “decision drivers” to consider that IT teams need a clear and structured framework to increase the speed and success of workload placement decisions. Analysis utilizing multiple workload placement criteria can identify and prioritize the most critical factors for driving the resulting decision.

Often, what an organization is trying to accomplish with this decision is to solve a problem they have or create an outcome that they want but don’t have, so the framework must include the criteria that act as levers to drive those solutions and outcomes.

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

Flexential Professional Services has created a Hybrid IT Strategy Framework to guide workload placement decisions using 14 technical, operational and financial criteria:

TECHNICAL	OPERATIONAL	FINANCIAL
Latency	Operational Responsibilities	Cost Structure
Performance	Control and Visibility	Cost Predictability
Scalability	IT Skill Set	
Availability	Portability	
Technical Requirements	Security	
Dependencies	Compliance	

Technical Workload Placement Criteria

LATENCY

Network connectivity and latency—the time taken from request to response between workloads and users—is critical to productivity and the user experience. Latency impact is a result of the users’ location and distribution compared to the workload location. Latency-sensitive workloads should be placed as close as possible to users to meet defined latency targets. On-premises solutions are best suited to cases where users are local and centralized, while Software-as-a-Service (SaaS) solutions are well-suited for widely dispersed users. 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.

On-Premises	Colocation	Hosted Private Cloud	Public Cloud	PaaS	SaaS
Centralized User Base ← → Distributed User Base					

PERFORMANCE

Compute and network performance is a combination of computing capacity, storage speeds and bandwidth. While high performance can be obtained in multiple ways, it is most reliable when using purchased purpose-built equipment that resides on-premises or in colocation, or dedicated hardware and reserved compute power from cloud providers. Compute, storage and bandwidth resources will be more cost-effective the more they can be allocated and reserved in advance. For applications and data that grow from distributed sources, such as mobile and IoT, more distributed workload models provide greater performance benefits than a centralized workload model.

On-Premises	Colocation	Hosted Private Cloud	Public Cloud	PaaS	SaaS
Higher performance at lower cost ← → More expensive to achieve					

SCALABILITY

Cloud-based implementations allow computing and storage resources to scale dynamically, and with metered pricing, organizations only pay for what they use. Public cloud platforms are well suited to workloads that “scale-out” through microservices or automated server deployment. Workloads that require “scale-up” by the addition of computing resources may be constrained by hardware limitations and procurement lead times in dedicated hardware models. Platform-as-a-Service (PaaS) solutions offer the most elastic scalability, but such capabilities can also create runaway costs if not properly governed. Note that in a SaaS model, the responsibility to scale effectively is outsourced to the software provider and outside direct control.

On-Premises	Colocation	Hosted Private Cloud	Public Cloud	PaaS	SaaS
Lowest ← → Highest					

³ 541 Research Voice of the Enterprise: Cloud, Hosting & Managed Services, Workloads & Key Projects, 2019

Technical Workload Placement Criteria Continued

AVAILABILITY

A differentiator between workload placement options is how availability is ensured. When an application is designed, developed and built for a particular platform, it can leverage that platform's native features for high availability (HA), such as load balancing or clustering. This often requires additional compute resources to facilitate HA at increased costs. Much commercial off-the-shelf software, however, relies on the infrastructure layer for HA. Clouds can provide infrastructure HA at the hypervisor layer by leveraging redundant compute, network and storage, usually at a lower cost than dedicated hardware. SaaS requires the least direct IT effort to achieve HA, but as with all cloud options, the vendor must be able to meet the required SLAs.

On-Premises	Colocation	Hosted Private Cloud	Public Cloud	PaaS	SaaS
HA provided by infrastructure ← → HA built into application ← → SLA based/ lowest effort					

TECHNICAL REQUIREMENTS

Technical requirements include computing resources, performance requirements, licensing considerations, and needs for specialized appliances or physical equipment. For instance, specific legacy applications, such as those designed to run on mainframes or IBM® AS/400 systems, can't be moved to x86 platforms and must remain on dedicated hardware. Likewise, network equipment may require colocation or on-premises deployment, as cloud placement may not be possible and acceptable use agreements may limit functionality.

On-Premises	Colocation	Hosted Private Cloud	Public Cloud	PaaS	SaaS
Most Options & Control ← → Most Constricted & Dictated					

DEPENDENCIES

Sufficient connectivity to dependent components must be considered in workload placement decisions. In one study, 63% of IT respondents said understanding application dependencies was their top migration challenge.³ Interconnectivity complexity increases when dependencies span different environments and locations. Some workloads require that dependent components be placed within the same data center. Some workloads may have a high 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.

On-Premises	Colocation	Hosted Private Cloud	Public Cloud	PaaS	SaaS
Easier, local dependencies ← → More complex cross-location dependencies					

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.

On-Premises	Colocation	Hosted Private Cloud	Public Cloud	PaaS	SaaS
Highest ← → Lowest					

CONTROL AND VISIBILITY

Keeping workloads highly accessible to IT staff for reliable, transparent and consistent service management is one reason organizations invest in on-premises or colocated solutions. Familiarity and knowledge of platforms, operational processes and system lifecycles increases confidence and predictability during planned maintenance or unplanned disruptions. Control and visibility of components decreases 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 only include application configuration, user access controls and application data.

On-Premises	Colocation	Hosted Private Cloud	Public Cloud	PaaS	SaaS
Highest ← → Lowest					

IT SKILL SET

Eighty-five percent of organizations report having a shortage of expertise for managing their cloud environments.¹ An honest assessment of internal skill sets is necessary to identify training needs, and a lack of internal IT skills is also an opportunity to consider leveraging external vendors to fill team or department knowledge gaps. Once an assessment is done, IT can prioritize the skill gaps to fill in-house, what 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.

On-Premises	Colocation	Hosted Private Cloud	Public Cloud	PaaS	SaaS
Most demanding breadth ← → Specialty Skill Set ← → Least demanding					

Operational Workload Placement Criteria

PORTABILITY

The ease of moving a workload from one environment to another becomes a critical criterion when performing lifecycle management activities such as hardware or software refreshes, changing service providers or moving workload placement closer to users. Containerization can make portability more straightforward, but not every workload is a good container candidate. Container management also requires a specialized skill set that may not currently be found in the IT staff. Additionally, licensing models based on the number of processors introduce cost considerations if moving to different hardware or a new environment, and licenses tied to MAC addresses or hardware keys may require relicensing. 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.

On-Premises	Colocation	Hosted Private Cloud	Public Cloud	PaaS	SaaS
Limited due to rigidity ← → Easiest ← → Limited					

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 controls for a SaaS application. Providers build security into their services, from ISO-certified colocation data centers to fully managed application security from SaaS vendors. Security visibility, auditability and testability are important 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.

On-Premises	Colocation	Hosted Private Cloud	Public Cloud	PaaS	SaaS
Fully Responsible ← → Most Fully Managed					

COMPLIANCE

Organizations must have sufficient controls, auditability and attestations to maintain compliance with standards, rules and regulations. These can be regulatory (governmental) requirements, industry rules or voluntary standards such as those published by ISO and NIST. While cloud providers maintain compliance with many regulations, each customer is ultimately responsible for their compliance program and status and must confirm a chosen vendor meets 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 or access to controls. SaaS and PaaS tend to result in the lowest overall compliance burden since these are highly managed solutions.

On-Premises	Colocation	Hosted Private Cloud	Public Cloud	PaaS	SaaS
Highest Burden ← → Lowest Burden					

Financial Workload Placement Criteria

COST STRUCTURE

Cloud computing provides an opportunity to shift budgets toward OPEX with third-party services and subscriptions and 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, including SaaS. The exercise of comparing OPEX and CAPEX by workload makes expected costs transparent.

On-Premises	Colocation	Hosted Private Cloud	Public Cloud	PaaS	SaaS
Highest CAPEX ← → No CAPEX					

COST PREDICTABILITY

It is essential to understand the cost variability of an individual workload within a given billing format. 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 deprovisioning crucial to avoid expenses for unused resources and services. Moreover, de-provisioning usually entails removing data from the environment, and data egress charges can be quite significant for large data sets. It becomes easy to see how hyperscale clouds, in particular, require cost management to a greater degree than traditional physical server-based environments or more closely controlled private clouds.

On-Premises	Colocation	Hosted Private Cloud	Public Cloud	PaaS	SaaS
← → High Predictability ← → Low Predictability ← → High Predictability					

The Benefits of Optimum Workload Placements

There is no one-size-fits-all path to navigating workload placements or hybrid IT. Analyzing workloads individually identifies the primary technical, operational and financial drivers for a placement location decision and the broader infrastructure and services landscape required. The resulting placement decisions can vary significantly by individual workload. 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, cost predictability or performance considerations.

Organizations that approach hybrid IT decision-making from a workload perspective and successfully place those workloads in the optimum environments can achieve greater flexibility, scalability, security and lower IT costs, putting them on a path to meet business goals.

A hybrid IT strategy framework navigates technology choices and organizational priorities

Today's vendors provide a plethora of choices for off-premises locations, infrastructure and services. No longer is an IT department constrained by a lack of options to on-premises, internal capabilities or a small group of cloud-based services. There are now so many hybrid IT choices that a new challenge is how to choose wisely. Using a hybrid IT strategy framework generates decisions that consistently optimize technical, operational and financial considerations while keeping those decisions aligned with broader IT and business strategy.

Adding resource capacity to support business growth is the most common reason organizations plan to increase 2021 spending on cloud and hosting environments.⁴

Flexential Professional Services 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 reside, Flexential Professional Services partners with IT teams, providing actionable guidance, expert execution and advancing IT strategy, infrastructure and security.

⁴ 451 Research, Voice of the Enterprise: Cloud, Hosting & Managed Services, Budgets & Outlook, April 6, 2021



ABOUT FLEXENTIAL

Flexential empowers the IT journey of the nation's most complex businesses by offering flexible and tailored hybrid IT solutions comprised of colocation, cloud, connectivity, data protection, managed, and professional services. The company builds on a platform of three million square feet of data center space in 20 highly connected markets, and on the FlexAnywhere™ 100GB private backbone to meet the most stringent challenges in security, compliance, and resiliency. See how Flexential goes beyond the four walls of the data center to empower IT through an interactive map found on www.flexential.com.

Flexential is a registered trademark of the Flexential Corp.

