

IDC PERSPECTIVE

Define Your Hybrid Infrastructure Needs: A Segmented View of Hybrid Infrastructure Usage and Requirements

Deepak Mohan

Ashish Nadkarni

EXECUTIVE SNAPSHOT

FIGURE 1

Executive Snapshot: A Segmented View of Hybrid Infrastructure Usage and Requirements

Growth of public cloud adoption and its adoption into the enterprise IT environment in a mainstream fashion has made "hybrid infrastructure" the new norm at IT organizations. But hybrid infrastructure is not a homogenous requirement or single capability, and it is important for organizations to understand their specific hybrid usage type, road map, and needs.

Key Takeaways

- Nearly 80% of enterprises are either actively using or planning for a hybrid infrastructure environment.
- Hybrid infrastructure usage can be segmented into six types — best-of-breed infrastructure, tiering and archive, test/dev and staging, backup and disaster recovery, migration, and bursting.
- Hybrid integration can be broken into the following — control plane integration, data infrastructure integration, network layer integration, physical connectivity, and environmental consistency.
- Organizations transition through usage types and needs as they evolve their cloud strategy, understanding the specifics allows organizations to optimally plan hybrid infrastructure investments.

Recommended Actions

- Recognize that terms like "hybrid cloud," "hybrid IT," or "hybrid infrastructure" are broad terms to acknowledge a form of infrastructure use but that these encompass a group of diverse infrastructure usage types with varying needs and priorities.
- Identify the underlying strategic or business drivers for your organization's hybrid infrastructure adoption and the specific integration needs to support it.
- Understand your initial hybrid infrastructure needs and the possible evolution of the cloud strategy over time to better anticipate how your hybrid integration needs and road map will evolve over time.

Source: IDC, 2019

SITUATION OVERVIEW

Adoption of public cloud infrastructure services by enterprise IT organizations to replace or complement their existing infrastructure has evolved at a relatively steady and increasing pace to reach a state well beyond the early stage and exploratory pilots. The majority of organizations are actively using public cloud infrastructure as a service (IaaS) in one form or another today. However, there continues to be demand and investment in on-premises infrastructure systems (on-premises refers to any dedicated infrastructure assets that are owned, leased, or rented by the organization and are deployed as a private cloud or in a traditional noncloud manner). This demand and continued investments are driven by three broad factors:


- First, there are a number of workloads that cannot be immediately moved to public cloud due to latency and colocation restrictions. This could be latency and colocation with respect to physical entities (such as an operations floor) or with respect to other applications and data sets that cannot move to the cloud (such as a legacy data processing application running on a mainframe).
- Second, there are workloads and data sets that are unable to be placed in a public cloud due to regulatory compliance or internal policy restrictions. These are often a result of misperceptions or a lack of understanding of public cloud security models, limited regulatory or compliance certifications offered by the public cloud services provider, or simply the lack of availability of such certifications. These continue to be a major factor slowing the momentum of workload movement into public cloud.
- Last, nearly all large enterprise IT organizations have existing on-premises infrastructure and datacenter investments. The use of public cloud infrastructure does not mean an immediate discontinuation of these assets, and the adoption of public cloud is often a gradual process at these enterprises.

These have resulted in a growing number of organizations adopting a hybrid infrastructure environment, a model that includes both public cloud and on-premises infrastructure. In some scenarios, these are also referred to as multicloud environments, which include the use of public and private clouds within the same organization in an integrated manner. The prevalence of such environments, combined with attempts from the technology provider community to capture and address these needs, has resulted in the terms like "hybrid cloud" and "hybrid IT" being used as a catchall term in the market.

While the terminology does capture the broad theme of mixing public cloud infrastructure with other infrastructure options, it actually encompasses a broad set of varying needs and functionalities, and the interpretation of these terms also varies quite widely among customers. As an example, Figure 2 shows the variance among customer interpretations of the term "hybrid cloud." The variance and lack of consistency make it challenging for IT organizations to accurately defining their hybrid infrastructure related needs and in conveying these needs to the technology provider community. Moreover, the broad categorization of environments as "hybrid" versus "public" versus "private or on-premises" does not emphasize the reality that there are different levels of hybrid integration, and that each of these has specific types of integration requirements.

FIGURE 2

Variance in Interpretation of the Term "Hybrid Cloud" Among Customers



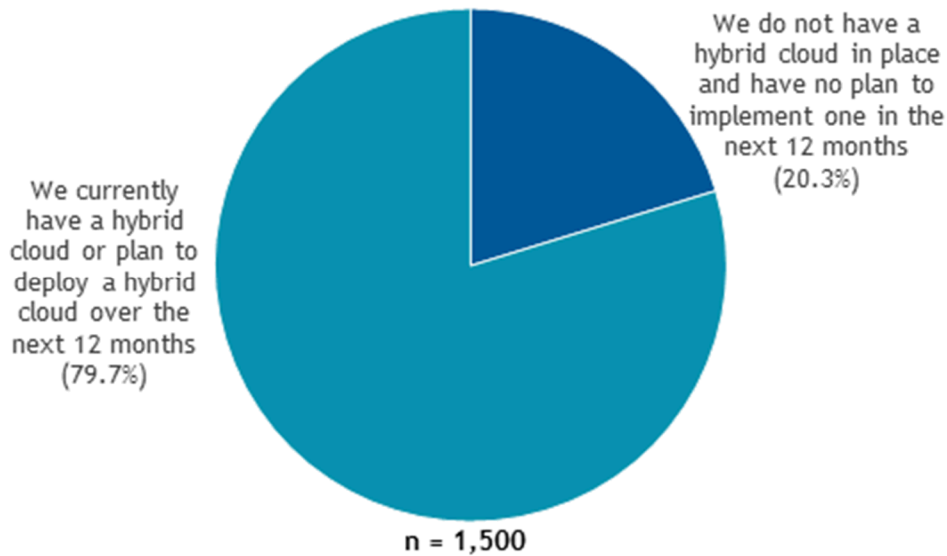
n = 1,500

Source: IDC's *IaaSView Survey*, 2018

As shown in Figure 3, nearly 80% of IT organizations either already have implemented or intend to implement a hybrid infrastructure environment in the near term. Beyond just the use of infrastructure on both public cloud and on-premises, this typically involves some level of integration (custom or off the shelf) of the control workflows or control planes across both infrastructure destinations. This makes it important to get to a common verbiage regarding the descriptions of hybrid infrastructure usage plans.

FIGURE 3

Prevalence of Hybrid Infrastructure on IT Infrastructure Activities and Plans



Source: IDC's *IaaSView Survey*, 2018

This document segments hybrid usage into a few commonly seen patterns of hybrid usage and into specific hybrid integration capabilities that customers seek when using infrastructure in this model.

For the purposes of this document, we will use the term "hybrid infrastructure," which will be interpreted as an organization using a mix of on-premises infrastructure (as mentioned previously, on-premises refers to any dedicated infrastructure assets that are deployed as a private cloud or in a traditional noncloud manner and that are owned, leased, or rented by the organization) and public cloud services for its infrastructure needs. The organizational scope of the discussion is the IT organization, which may be a complete enterprise or a line of business/division within an enterprise with its own IT budget and infrastructure operations. It does not include shadow IT usage of a public cloud, which is the use of public cloud infrastructure within an organization but in a manner that is completely separate from the organization's formal IT processes and budget.

Six Types of Hybrid Infrastructure Usage

IDC research has highlighted six broad usage scenarios within hybrid infrastructure use by IT organizations. The six hybrid usage types are summarized in Table 1 and discussed in the sections that follow.

TABLE 1**Six Hybrid Infrastructure Usage Scenarios**

Usage Scenario	Description
Best-of-breed infrastructure	Use of public cloud or on-premises infrastructure for specific workloads or components, based on what serves the needs best
Tiering and archive	Use of public cloud for tiered placement of older or noncritical data or workloads from on-premises infrastructure
Test/dev and staging	Use of public cloud for development, testing, and staging phases of applications or upgrades, with production being placed into on-premises infrastructure
Backup and disaster recovery	Use of public cloud as a backup and disaster recovery (DR) site to an active on-premises infrastructure site
Migration	Use of hybrid infrastructure as a means to execute a gradual migration from on-premises into public cloud infrastructure
Bursting	Use of public cloud alongside on-premises infrastructure, as a mechanism to provide additional infrastructure capacity when needed for temporary periods of time

Source: IDC, 2019

Best-of-Breed Infrastructure

Public cloud and on-premises infrastructure options both have their respective benefits and challenges. While the specifics of what is perceived as a benefit or a challenge vary by organization, a number of organizations have realized that components of their workload portfolio are best served by the specific benefits available on one or the other. The use of public cloud and on-premises may be either integrated or completely disconnected. An example of integrated use is the separation of layers in a three-tier application architecture, using public cloud for low-latency access at the presentation tier and on-premises infrastructure at the data tier (see *Enabling Hybrid IT – SaaS Provider Q4 Inc. Leverages Rackspace to Build and Manage Its Hybrid Infrastructure Platform*, IDC #US41732916, September 2016). An example of disconnected usage is the use of public cloud for specific customer-facing applications such as digital marketing assets and on-premises infrastructure for internal workloads such as enterprise resource planning (ERP) applications. Hybrid infrastructure usage by new digital economy companies, including those that were born in the cloud, typically falls within this category.

Beyond the initial phase that may involve some level of component or application migration, this type of usage does not typically run the same workload on both infrastructure locations and does not require a high level of consistency in public and on-premises infrastructure environments. But this does often benefit from consistency in the control and operations of the resources (provisioning, auditing, monitoring, etc.) so that all resources can be managed with consistent workflows and tools.

Tiering and Archive

The growth of data volume and the increasing desire to drive value from organizational data have resulted in a growing demand to capture and store data for both operational and business activities.

This growth pushes on the boundaries of available on-premises capacity, and organizations have started looking at public cloud as a destination for placing older or less critical data to optimize the use of on-premises capacity for current and business-critical data. This may be through a built-in tiering capability that is not actively operated by the organization, a capability available in a growing number of appliances and infrastructure stacks or through active movement of data into public cloud.

This type of usage is often one of the first steps taken by larger enterprise IT organizations, as they incorporate public cloud into their infrastructure mix. A pattern of public cloud adoption that is increasingly seen is to start with archive data movement into public cloud, followed by the use of analytics and data warehousing capabilities on public cloud to unlock insights from the tiered data – moving from a "tiering and archive" type usage into a "best-of-breed infrastructure" type usage.

Test/Dev and Staging

Another way in which public cloud is used to address the need to optimize on-premises infrastructure capacity utilization is by moving the comparatively less sensitive phases of application and upgrade development into public cloud – namely, by moving the development, testing, and staging operations into a public cloud environment. This often involves the development of the deployment and upgrade pipelines into the public cloud environment and the creation of an overlay in the public cloud environment to match the one that is on-premises. This type of usage typically sees high value in having consistent operational environments on both public cloud and on-premises, reducing the overhead of deploying and managing the overlay environment.

Backup and Disaster Recovery

A third way in which public cloud infrastructure helps optimize the use of on-premises infrastructure capacity is by serving as an alternative for a physical backup and disaster recovery site. For smaller organizations that do not intrinsically have the scale or geographic spread to have geo-separated infrastructure, public cloud offers a way to build in a level of resilience and availability without the investment and overhead of a separate site. Even for larger organizations, using a public cloud infrastructure option as the backup and DR site for specific workloads can serve as an effective mechanism to gain familiarity and start the integration of public cloud into the organization's IT environment.

This pattern, in fact, using public cloud as a backup site for specific workloads and then eventually moving these workload completely to public clouds, is another pattern of public cloud adoption that is increasingly seen at enterprise IT organizations – starting from a "backup and DR" type usage to a "best-of-breed infrastructure" type usage.

Migration (Often to Public Cloud Infrastructure)

While all in adoption of public cloud are still in their early stage of prevalence, a number of organizations have made an internal strategic call that they eventually want to be completely running the IT infrastructure primarily on public cloud. In most cases, they also realize that this is not a step function transition but a gradual one that may take several months or years, as they build competence, migrate workloads, and transition processes. These organizations build up a hybrid infrastructure with the primary intent of migrating workloads gradually into the public cloud environment as they progress on this adoption path. Organizations on this path typically have a higher affinity to cloud-native stacks and approaches and often look for ways to modernize their applications in place on-premises, alongside the hybrid infrastructure implementation and migration of other parts of the portfolio.

One point to note is the distinction between migration as a primary driver for the hybrid infrastructure and migration as an activity early in the hybrid infrastructure adoption. Most hybrid infrastructure implementations, especially at existing IT organizations, will start with the migration and transition of certain workloads or activities into the public cloud environment. The discussion in this usage type specifically refers to the former, which is where migration into public cloud is the primary reason for the hybrid infrastructure – not the initial migration as part of the hybrid infrastructure initialization.

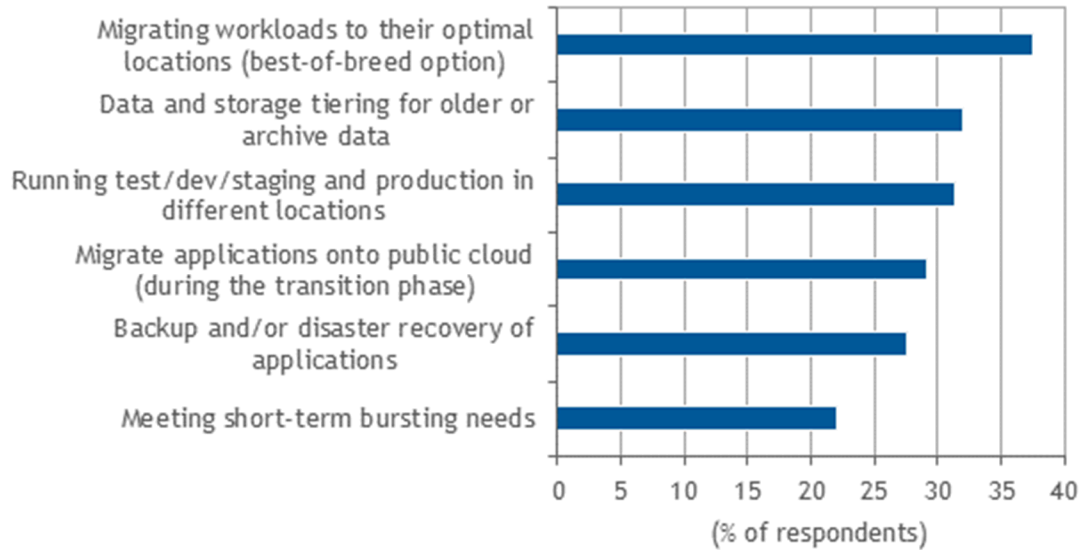
Bursting

Last, there is the infrastructure bursting usage type where public cloud infrastructure is used to meet temporary infrastructure capacity usage needs, such as short-term scaling for a specific need or for periodic expansion of infrastructure capacity. Compared with the other usage types discussed previously, bursting requires a higher level of integration and consistency across the on-premises and public cloud environments. Not all applications or workloads are candidates for this type of usage – especially large, monolithic, vertically scaled, enterprise applications. For these reasons, bursting is not yet a prevalently seen use case among enterprise IT environments, beyond specialized use cases involving infrastructure-intensive operations running custom or semicustom applications.

Thinking through these hybrid infrastructure usage types and using these to identify the initial adoption plan and the long-term path regarding public cloud adoption into the environment will help narrow down the hybrid infrastructure-related functionalities and operational capabilities that need to be acquired or developed. While these are six broad categories of hybrid usage types, they are not hard and mutually exclusive categories. Organizations may find themselves with a mix of needs across types, especially in scenarios where organizations are maturing on their cloud adoption and expanding from the initial usage types to other ones (as discussed in some of the usage types). Even so, thinking of hybrid usage and plans in terms of these types can help organizations be specific and anticipate evolutions in their hybrid infrastructure needs and plans. Figure 4 shows the prevalence of these hybrid infrastructure usage types among enterprise IT organizations.

FIGURE 4

Customers Hybrid Infrastructure Usage Today



n = 1,195

Source: IDC's *IaaSView Survey*, 2018

A consistent understanding of the drivers will also help identify the specific hybrid infrastructure-related capability needed for your organization. The Hybrid Infrastructure Integration Requirements section discusses a segmented view of these hybrid infrastructure requirements.

Hybrid Infrastructure Integration Requirements

Hybrid infrastructure integration capabilities can be grouped under the following specific functionalities, based on the specific nature of integration delivered:

- **Control plane integration** refers to the ability to execute control plane functionalities (such as provisioning, monitoring, access control assignment) with resources on both public cloud and on-premises infrastructure using the same APIs or a common "single pane of glass" user interface. This may be implemented and achieved in various ways – including integration within a control tool or API library, alignment across the API on both infrastructure destinations, or overlay stacks (on public cloud and on-premises infrastructure) that deliver this consistency.
- **Data infrastructure integration** refers to the ability to access data sources on public and on-premises environments in a consistent manner. Data sources can include constructs like file systems, data bases, or data warehouses, and consistency in the data infrastructure allows applications to continue interacting with data sources in a consistent manner even as changes are made in the data source location or application location.
- **Network layer integration** is the ability to have infrastructure resources on public cloud and on-premises destinations reside on a common logical network, typically a private layer 3 network with private IPs that can directly communicate with each other. This is particularly valuable in

usage scenarios where components need to seamlessly communicate with each other across premises, such as bursting and best-of-breed infrastructure used in an integrated manner.

- **Physical connectivity** is the availability of a high-throughput low-latency connection across the public cloud infrastructure environment and on-premises infrastructure. This is typically implementable using MPLS and/or other similar site-to-site connectivity services (such as AWS Direct Connect, Azure Express Route, Equinix Cloud Exchange Fabric) between on-premises and public cloud locations.
- **Environment consistency** refers to having a consistent operating environment on both the public cloud and on-premises locations such that applications can move from one to the other in a near-seamless fashion with no changes needed.

Additional hybrid enabler capabilities, such as automated migration and orchestration, can be built on top of these functionalities to further simplify the customer experience around hybrid operations.

ADVICE FOR THE TECHNOLOGY BUYER

Terms like "hybrid cloud," "hybrid IT," and "hybrid infrastructure" have been used extensively in the enterprise IT market over the past five years, but these terms do not refer to a homogenous set of functionalities or underlying needs. Rather, these are broad terms to capture the concept of using public cloud alongside some form on-premises infrastructure. This captures a number of underlying usage types, each of which has varying integration needs and priorities.

Within the IT organization, it is critical to recognize this and have an internal point of view on the organization's specific drivers toward a hybrid infrastructure and its road map with regard to the hybrid infrastructure and the specific hybrid integration capabilities that are needed to meet the hybrid infrastructure usage needs for the organization. This is especially valuable in today's enterprise IT infrastructure market where hybrid infrastructure-related offerings are still rapidly evolving, and new targeted hybrid related offerings continue to enter the market. Actively considering the organization's goals and needs will help plan hybrid infrastructure investments in a phased manner based on the organization's needs and priorities.

LEARN MORE

Related Research

- *VMware Cloud on Dell EMC: Accelerating the Shift to Consumption-Based IT Everywhere* (IDC #US45066119, May 2019)
- *Growth of Public Cloud Stacks in Dedicated and On-Premises Customer Environments* (IDC #US44884219, February 2019)
- *IDC Market Glance: Public Cloud Infrastructure as a Service, 1Q19* (IDC #US44833719, February 2019)
- *Hybrid Cloud Definitions and Hybrid Cloud Adoption Plans by Enterprises* (IDC #US44523318, December 2018)

Synopsis

This IDC Perspective provides a segmented view of hybrid infrastructure usage into a few commonly seen patterns of use. It also lists the specific hybrid integration capabilities that customers need to consider when planning for a hybrid infrastructure environment.

"Hybrid infrastructure will soon be the de facto environment for enterprise IT organizations, and investments to enable this are already a priority at a majority of enterprises today," says Deepak Mohan, research director, Infrastructure Systems, Platforms, and Technologies, IDC. "Building a consistent internal understanding of the planned hybrid environment usage and level of integration will facilitate effective hybrid road map planning and optimized investments in this direction."

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-community.com
www.idc.com

Copyright Notice

This IDC research document was published as part of an IDC continuous intelligence service, providing written research, analyst interactions, telebriefings, and conferences. Visit www.idc.com to learn more about IDC subscription and consulting services. To view a list of IDC offices worldwide, visit www.idc.com/offices. Please contact the IDC Hotline at 800.343.4952, ext. 7988 (or +1.508.988.7988) or sales@idc.com for information on applying the price of this document toward the purchase of an IDC service or for information on additional copies or web rights.

Copyright 2019 IDC. Reproduction is forbidden unless authorized. All rights reserved.

