




Inside Track
Research Note



How hybrid IT changes your data journey

Hybrid IT drives a shift in storage thinking
from capacity to service

in association with



Introduction

Storage has long been a diverse world. In the past, it has often been acquired piecemeal as new applications were added, and because workloads have different usage patterns and resource needs, their storage requirements differed too.

Yet in a world where IT as a whole is also becoming more diverse than ever, storage and storage decisions have a key part to play in helping pull it all back together. In this paper, we look at how new approaches to planning and consuming storage - to your data journey - can improve interoperability by supporting an evolution to hybrid IT.

What is hybrid IT?

Few organizations can escape having a variety of IT services and delivery models. While some older applications cannot be virtualized or modernized, newer ones are likely to be 'cloud-native' or to take the software-as-a-service (SaaS) route. In addition, converged and hyper-converged systems are winning fans, as is software-defined infrastructure.

More and more organizations are therefore adding new systems or services. It might be a public cloud user adding on-prem resources such as local storage, a user of traditional IT systems wanting access to flexible and innovative public cloud services, or an organization seeking to bridge on-prem to a public cloud via a private cloud, creating in the process a hybrid cloud.

The risk, of course, is that we may end up with a complex jumble of platforms, each with its own way of doing things – and each with its own storage. What we need is to operate them as a single seamless whole, and this is the promise of hybrid IT (Figure 1).

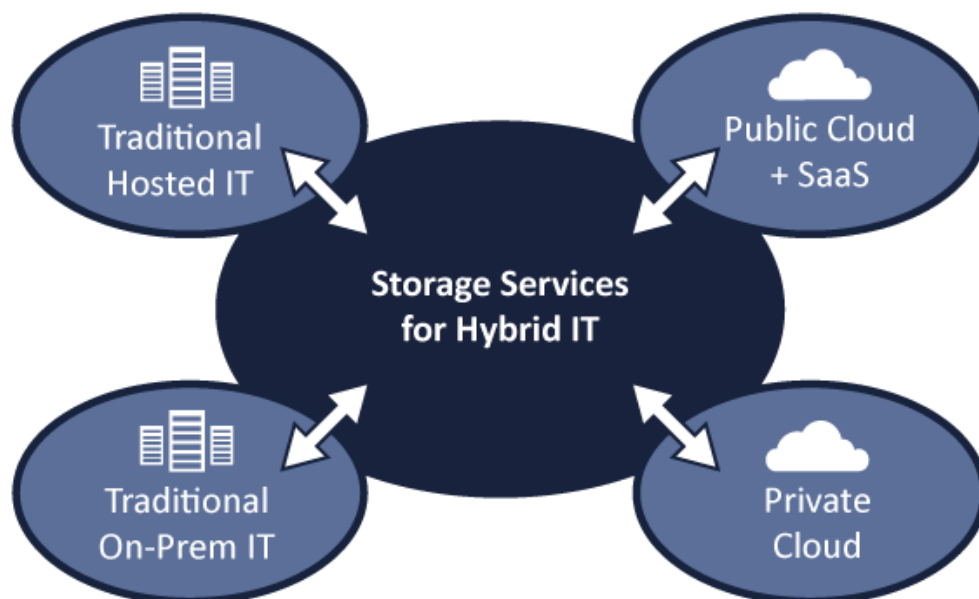


Figure 1: Hybrid IT brings together multiple IT delivery models, and therefore their storage too.

Why this growing need for diversity? Public cloud and SaaS are undoubtedly good solutions in many use cases, but the public cloud isn't 'better' than the data center, or vice versa. They both have roles to play – and this isn't just about cloud, it's about IT in general. Some workloads will suit a cloud-based delivery model and others will not. Some will be appropriate for hosting remotely, and others will do better on-prem, for reasons such as complexity, performance, compliance, cost, security and control.

Hybrid sourcing models

The result is a growing mixture of platforms or sourcing models. For example, some organizations will be adding public/private cloud to on-prem IT, while others may have containerized existing enterprise applications. Alternatively, they may have gone all-public cloud but are now pulling some workloads back on-prem, after recognizing the implications for control, cost, compliance, etc.

It is useful to remember here that 'cloud' is just the delivery model or mechanism, and that it does not specify any particular deployment or hosting model. For instance, many people assume 'cloud' and 'public cloud' are synonymous, but they aren't – the latter is a specific deployment of the former.

What's new here, and distinguishes the concept of hybrid IT from simply "having multiple platforms and IT delivery models" is the fast growing need for interoperability and data sharing. This means that, rather than the infrastructures and technologies used, organizations must put their data journey and the roles that data plays at the center of their hybrid IT planning. It is the value of data that is key here, and the storage needs to support that data throughout its lifecycle.

The hybrid world revolves around data

Linking multiple delivery platforms under the banner of hybrid IT is not just an opportunity to converge their data storage and rein in its rocketing management costs. It also requires and enables new capabilities within your storage infrastructure, including new ways to leverage your data.

For example, if you are adding private cloud to your existing systems, what storage capabilities and interfaces will it need, and how will you supply them? Or you might already be backing up some on-prem systems and data to a cloud, but can you also backup your SaaS and public cloud data to local storage? What if you need to backup data from one site to another or from one public cloud to another?

Hybrid IT encourages new ways of thinking about data – in particular, it means that at an infrastructure level, all data is potentially both a source and a target, and could even be both simultaneously. It ties in too with the growing awareness that old data can still have considerable value, for example for analyzing customer behavior or business processes, or perhaps for developing predictive maintenance schedules.

Given all of that, let's look at some of the concepts and requirements to consider when investing in storage and data protection solutions to support a hybrid IT strategy.

Service, not capacity

One way to think about this is that we should no longer focus on what physical storage capacity or performance an application requires. Instead, we need to concentrate on what storage services it needs – those services will include a capacity element but they also include performance, availability, security, data protection (backup) and so on.

Insulate your storage decisions

Following on from that shift to thinking in service terms, it was common in the past for storage decisions to derive or spin off from application decisions. Each new software workload needed a certain type of storage or data protection. That was workable in a time when applications were few in number, but that is no longer the case.

Worse, not only are applications far more numerous today, especially once you account for trends such as containerization, microservices and serverless, but data volumes continue to grow. That means we need to insulate our storage decisions from our application decisions, and ensure that the storage infrastructure will be able to accommodate change, growth and varied service demands.

This may, of course, mean adding new elements to the storage infrastructure. For instance, a modern backup appliance might simplify the task of providing hybrid on-prem backup services, or new storage software and systems might be needed to add support for object storage. We might decide to add additional storage tiers - one of NVMe Flash for better performance or a hybrid tier for capacity storage. Alternatively, we might opt to buy in some or all of the necessary services, perhaps from a cloud provider or as a managed service hosted remotely or on-prem.

Plan for tomorrow, not just today

As an extension to the above, almost everything is evolving – business practices, technology models, the communities that use and support them, and so on. As a result, there are fewer and fewer ‘finished products’ in software now. Instead, trends such as DevOps and Agile mean that the application (or rather, the set of services that comprises many modern applications) can be subject to frequent revisions and updates.

Even if the application itself does not change, its service quality and data protection needs may evolve. All this means it is essential for the storage element to be flexible and able to accept change and build on it, as well as being able to grow as freely as possible. There are several ways to achieve this, and we will look at some of them here.

One option is to adopt a software-defined storage (SDS) approach, such as the open-source technology Ceph or similar commercial storage software. These technologies typically make it possible to combine multiple storage systems or devices, whether similar or dissimilar, into an abstracted and virtualized ‘pool’. That pool can readily be expanded by adding more physical storage, and can then be addressed in multiple ways, for example as file, block or object storage, and potentially as all three at the same time.

SDS approaches can also enable storage decisions to be made on more practical grounds. For example, for your on-site needs, you might choose hybrid arrays when flexibility and capacity are paramount, or all-flash when performance is key. In a hybrid set-up, the storage layer allocates whatever class of storage is needed according to the service requirements, and if the application's requirements change, the physical data bits can be migrated appropriately and automatically.

Consolidated data management

Another option is to consolidate and insulate your storage at the data management and governance layer instead. At a conceptual level, this is how a technology such as an end-to-end data fabric, shared namespace or enterprise file fabric operates. The aim is to enable unified access and management across a data infrastructure that can combine storage located on-prem, in multiple public and private clouds, and elsewhere.

At its simplest, this approach overlays the different kinds of storage with a single shared file system, thereby eliminating 'data islands'. More sophisticated versions may add features such as data security and auditing, plus automated data migration and tiering, even across multiple geographic locations.

More than just storage

As we unify our storage and insulate it from application-side decisions, we also broaden the variety of services that we can offer. For example, we may still have a backup service – or more likely a range of backup services – but the same hybrid infrastructure could also be used to deliver other services such as archiving, enterprise-wide search tools, data management and governance, and so on, all from the same underlying storage pool or infrastructure.

Of course, while this can bring considerable business value if it is done correctly, it also requires skills that many end-user organizations may not yet possess. That makes it important to work with a trusted partner or supplier that has the relevant expertise and experience, and can supply the training and support needed to build up that skill-base within the organization.

In summary

Creating a hybrid IT infrastructure is not a simple task but it offers considerable opportunities. It will also be difficult to avoid in many organizations: hybrid IT often appears by accident, as new workloads and resource types are added to the organization's technology base. It is a prudent step, therefore, to make the strategic decision to put the necessary frameworks in place to properly support it.

A lot of effort has been put into developing the software side of hybrid IT – containers, APIs, Kubernetes, microservice architectures, service meshes and more, and this is undoubtedly the most glamorous part of hybrid IT. But organizations must also attend to the less obvious but equally essential foundation that underpins all of that, which is where the data is stored, plus its long-term management, security and governance.

The first big step on this new data journey is to think in terms of storage services, not just of physical capacity for an application, if you are not already doing so. This will provide a better fit – philosophical and operational – with the modern service-oriented world of software development, and should also enable you to decouple storage infrastructure decisions from application decisions.

Then it is to eliminate data silos, where possible moving any controls needed for regulatory compliance, security and so on, into the hybrid infrastructure where they can be better managed and audited. And of course, to leverage that unified hybrid storage layer for additional services, such as data governance and enterprise search.

To summarize the summary, getting the storage layer right is near-essential to the success of a hybrid IT strategy. If you can indeed get it right – and that may require working with an experienced partner, as it's a complex task with many dependencies – then there are many gains to be had from making the journey to service-based access to a unified data pool.

About Freeform Dynamics

Freeform Dynamics is an IT industry analyst firm. Through our research and insights, we help busy IT and business professionals get up to speed on the latest technology developments and make better-informed investment decisions.

For more information and access to our library of free research, please visit www.freeformdynamics.com or follow us on Twitter @FreeformCentral.

About Fujitsu

Fujitsu is the leading Japanese information and communication technology (ICT) company offering a full range of technology products, solutions and services. Approximately 140,000 Fujitsu people support customers in more than 100 countries. We use our experience and the power of ICT to shape the future of society with our customers.

For more information, please visit www.fujitsu.com.

Terms of Use

This document is Copyright 2019 Freeform Dynamics Ltd. It may be freely duplicated and distributed in its entirety on an individual one to one basis, either electronically or in hard copy form. It may not, however, be disassembled or modified in any way as part of the duplication process. Hosting of the entire report for download and/or mass distribution by any means is prohibited unless express permission is obtained from Freeform Dynamics Ltd or Fujitsu. The contents contained herein are provided for your general information and use only, and neither Freeform Dynamics Ltd nor any third party provide any warranty or guarantee as to its suitability for any particular purpose.