Making server decisions for hybrid IT

If you think hybrid IT renders your server knowledge obsolete, think again.

in association with

FUJITSU

Introduction

It's a diverse world, certainly as far as IT is concerned. Using only SaaS and public cloud might be a viable option for start-ups and small organizations, but even there, most will also have, need or want some local systems and storage. Similarly, large organizations with their own IT departments will probably also have requirements for SaaS applications and remote hosting to supplement internal systems or meet new needs.

The challenge is how to make all these systems - and the people who build, run, and rely on them - work together smoothly, without the constant need for integration work and wrestling with mismatches. Perhaps more importantly, it is how to operate them as a single coherent whole. This is the seductive promise of hybrid IT (Figure 1).
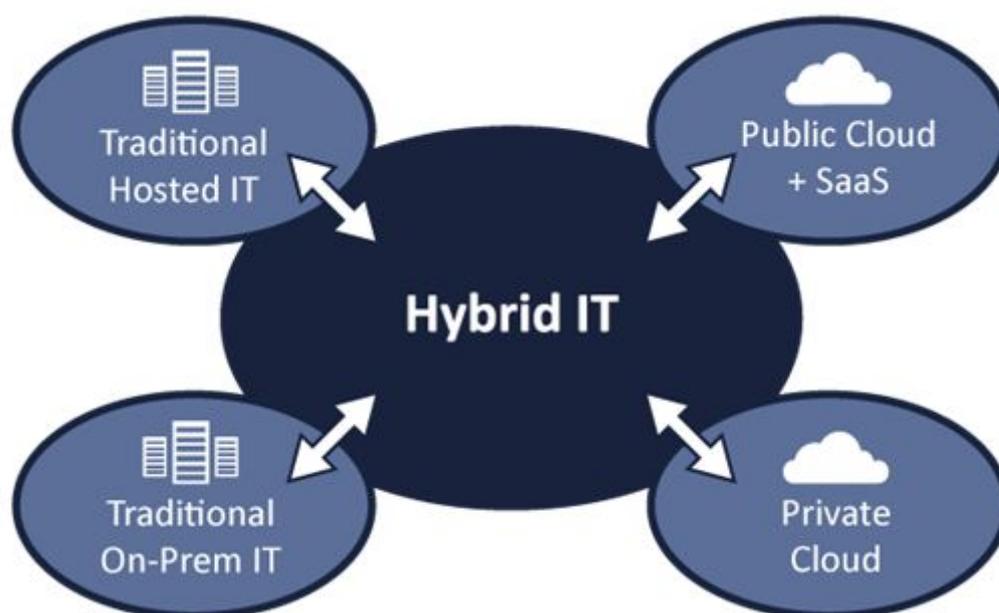

Figure 1: Hybrid IT aims to converge and integrate multiple IT delivery models

How to get there, however? Whether you want to add public/private cloud to on-prem IT, or pull workloads back on-prem from a public cloud for reasons of control, cost, compliance, etc, you will likely need additional server resources to run the new services. So what needs to be considered when investing in server technology to support a hybrid IT strategy, and why does it matter?

Server decisions for a hybrid world

There is a belief that servers are just a commodity now. While this view is correct for some use cases, things can be very different in the enterprise. Here, servers will be the central focus for many of the decisions that you will make in building a hybrid IT strategy. Let's look first at the contexts in which these decisions will be made.

The road to hybrid IT

Businesses and their software workloads have different usage patterns and resource needs, and multiple IT delivery models all have roles to play. In the past, this could have meant mainframes, minicomputers and PCs each playing their part. Today the list might include all of those, plus servers delivering SaaS, private and public cloud, and more.

There are many reasons for this platform variation. For instance, in some cases there will be a sound business rationale for minimizing asset ownership and maximizing short-term flexibility via a public cloud or the SaaS subscription model. In others, longer application lifespans might make the subscription model expensive.

On the software side, as well as workloads designed for traditional IT architectures there are others that are designed for or particularly well suited to the cloud model. And then there is the decision of whether to go public, private or hybrid cloud - this may be directed by a wide range of factors, including how high you think you can push the private system's efficiency and utilization.

Also relevant to the platform choice are your organization's core competencies and a forest of other factors, such as application complexity, data security, regulatory compliance, personal preference, WAN reliability, and so on.

What is hybrid IT and why does it matter?

Simply having multiple platforms and IT delivery models is not the same as having hybrid IT. What distinguishes genuine hybrid IT is that, given today's utilization of multiple platforms and models, it provides seamless management across those platforms, and at least some degree of interoperability between them.

For example, a key subset of hybrid IT is hybrid cloud, where private and public cloud resources operate in concert. If applications are relatively predictable, running them on a private cloud should be cheaper than using public cloud, assuming you can achieve high enough levels of efficiency and utilization. If they are bursty and unpredictable, then a pay-per-use platform such as public cloud may make more sense.

Hybrid IT extends that concept to include traditional IT models, as well as more modern cloud-native ones. For example, a mission-critical application could run on-prem, with a second copy on standby in a public cloud 'virtual data center', ready to take over if disaster strikes.

The key on the operational side is that you manage everything as one infrastructure, moving applications between platforms seamlessly and, if need be, automatically.

And on the development side it can help to support the creation of composite applications, perhaps using services from a public cloud while keeping sensitive data on the private cloud, or using APIs to connect a monolithic back-end application to a cloud-native front-end.

Next, let's look at how an understanding of these expectations and opportunities can help us make platform decisions when planning and building business-grade hybrid IT.

What hybrid IT needs from a server

Most of us will already have a mental image – whether imagined or genuine – of what a cloud data center looks like. Rack upon rack of commodity pizza-box servers perhaps, or ranks of black boxes, each with its winking green or blue lights. And for cloud providers that is largely true. The biggest of them even design their own servers, sometimes having several different models for different workloads. These servers then operate as part of a highly virtualized and automated infrastructure that seamlessly absorbs both hardware failures and any new servers that might be added to it.

However, this is a business and operational model that really only makes sense at a very large scale. It relies on there always being a pool of spare hardware capacity, plus the skills needed to build and operate the sophisticated software layers that underpin the apparent simplicity of the public cloud.

The good news is that this is not the only way to deliver cloud-like elasticity, simplicity and power. At enterprise scale, it can make more sense to adopt an operational model based on highly-available and long-lasting hardware, whether in an organization's own data center or on a co-location site. After all, although the software infrastructure and operational processes for private cloud can be very different from traditional IT, the underlying skills are the same. In other words, your data center expertise can be just as relevant in private cloud – and by extension, in hybrid IT – as it was in the past.

As ever, it's the use-case that matters

The shift to hybrid IT can happen by chance or by design, but either way, a common catalyst for it is the addition of private cloud resources. They might be needed because public cloud is too slow or too reliant on imperfect WAN infrastructure, or because the organization needs to run new cloud-native applications alongside traditional IT, but must also ensure data sovereignty. It could even be because the organization has cloud-native applications that need larger servers than are readily available on-demand from a public cloud, or because public cloud costs have become unsustainable.

The important thing is to focus on the use-case and on what really matters to the organization. There is no need to try to emulate the hyperscale cloud giants if that isn't your business model. Conversely, you do not want to perpetuate outdated operational and technology models when adopting cloud-native technology.

For example, the economics of cloud providers and enterprises are very different, with the latter needing much more in the way of service and support. Regular enterprise-grade servers, designed for integration within a wider infrastructure while also being built to have cloud-like capabilities, may therefore make more sense than commodity pizza-boxes that must be custom-built into a larger system.

Hybrid sourcing and financing models

Hybrid IT is not just about how you deliver your applications to end-users, it is also about where you host them and how you finance and source them. On the hosting side, whether the application is cloud-native or traditional, the relevant workloads can be

located on-prem, in a public cloud, in a colocation facility, or in some combination of those three. Even a private cloud could be hosted locally or remotely.

On the financing side, there used to be three main options for an IT project: lease, SaaS or buy. Hybrid IT not only adds a few more, it also allows you to combine them into a hybrid financial model. For instance, the best return on owned or leased resources is when they have high utilization, while SaaS and pay-as-you-go public cloud will typically be better value for lower usage requirements. Models where you reserve additional server resources, but only pay when you use them, might sit somewhere in the middle.

Blending the above, therefore, you might seamlessly and automatically move tier 2 services into a public cloud or your reserved infrastructure pool at peak hours, to allow primary resources to be focused on tier 1 services. Or you might use a service catalog manager to provide subscription-based self-service access to software resources, without the user needing to know or care where those resources are really located or financed (see our paper 'Simplifying Multi-Cloud Service Delivery here).

The reality: converging modern and traditional infrastructures

Modern software platforms or frameworks, such as Kubernetes, are inherently hybrid and continue to evolve rapidly. For example, you can readily run Kubernetes applications across a specified and supported hybrid cloud, and assemble cloud-native applications from services published in many different locations.

However, an important consideration for hybrid IT is that it is more than just hybrid cloud: it may also need to take in traditional application architectures and multiple diverse hardware platforms, whether located on-prem or in a colo. The aim is of course to share and integrate where possible. That could mean using technologies such as APIs, 'wrappers', containers and microservices to enable existing applications to participate in modern composite and cloud-native services and applications. It could also involve moving traditional applications onto abstracted and shared infrastructure, in search of benefits such as better flexibility, consolidation and utilization.

Again, this changes the criteria for platform acquisition, and reminds us that it is not necessary or even desirable to emulate the cloud providers. The aim is to maximize efficiency and utilization for your organization and its diverse systems base.

In summary

Think as an enterprise. Your requirements are likely to be very different from those of a public cloud provider, so it does not make sense to try to imitate one. In particular, hyperscale economics are very different, with enterprises needing more service and support.

Focus instead on what's important to your organization. Of course you should not perpetuate old thinking and old operational models that are no longer relevant. You still need to apply key management and economics criteria, however, and in this your data center experience will be highly relevant.

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.