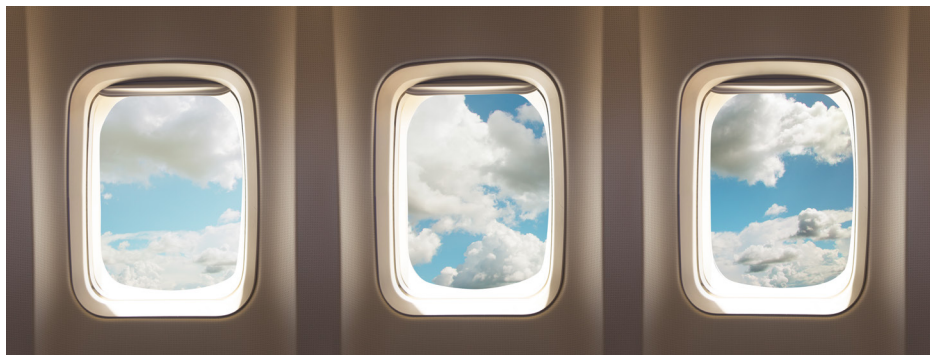FREEFORM DYNAMICS

in association with

NUTANIX™

REBOOT YOUR FUTURE
Shaping a new multi-cloud solution

Are you planning to build upcoming generations of IT infrastructure on top of today's well-known but 10-year-old technologies? Think again. Requirements and expectations have changed dramatically, and meeting them needs a fresher and better starting point. It's time for a reboot.

Converge the clouds to banish shadow IT
The future isn't 'public cloud', but it is public cloud-like. A hybrid system that embraces public cloud and its best aspects can empower users and help them work better.

Stretch, squeeze, step
Do you understand the costs and compromises involved in the various possible routes to the hybrid multi-cloud?

A foundation for the future
Our usage of software tools is both varied and continually evolving, so we need an extensible and open platform that's free of lock-in.

Software-defined and hyperconverged
Both are technically and philosophically essential to hybrid multi-cloud. It cannot exist without the abstraction, agility and flexibility they bring.

IT's cloudy future

Virtualisation, especially of servers but also of storage and networking, has been incredibly useful to the development of enterprise IT. It has helped to vault us out of the Iron Age of dedicated hardware into the Machine Age of interchangeability and mass production.

But it has reached its limits: to go further, to reach the truly modern On-Demand Era of computing, simply adding more of the same won't work. We need to take a step into the future. We need to think differently.
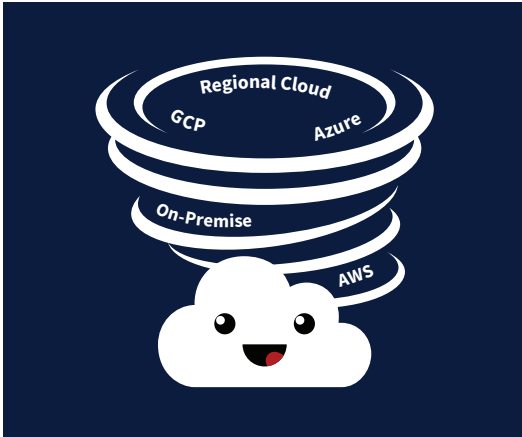
No more shadows

Public perceptions of computing have shifted dramatically in the last 10 years. Users now expect to be able to spin up resources within minutes or even seconds, from applications and software tools to complete systems, but enterprise IT has rarely kept pace.

The result can be shadow IT and loss of control, if users and application teams, frustrated by the corporate options, go instead to the public cloud. They get the flexibility and speed they need, but at the expense of visibility, manageability, security and cost. The problem is that public cloud looks cheap and simple to begin with, but projects quickly get expensive as they grow, plus having many small individual purchases can make security and governance a headache.

Hybrid and multi-cloud

So how can enterprise IT not simply catch up, but get ahead of demand? Supporting public cloud and adding your own private clouds is a good start, but is not enough on its own. The problem is that most hybrid clouds today are one public cloud to one private cloud, while most multi-cloud solutions are loosely-coupled - they are linked more for visibility and cost-control than for integrated operation.

Just as we have converged many of our on-site systems, we need to bring together our public and private clouds too. We need a hypercloud. In this document, we look at the questions to ask and steps to take as we plan to build one.

Remember that

To meet new needs and expectations, enterprise IT has to make fundamental changes, hyperconverging the best aspects of public and private clouds.

Mapping the best route

Once you start investigating, you will most likely discover several potential routes to your converged hypercloud future. Your existing virtualisation provider or supplier of hyperconverged infrastructure products will want to keep you on-board, as will your public cloud provider or providers. Finally there is the option to bring in a new architecture altogether. Let's look at these choices in a little more detail.

Stretch, squeeze, step

The first option is to 'stretch', building from or extending your current platforms to add new capabilities.

The underlying technology stays the same but you layer additional functionality on top. The challenge here is that you retain all the fundamental limitations of that original architecture. For example, the ability to burst to

the public cloud requires tight coupling, but most virtualisation technology was not designed for that.
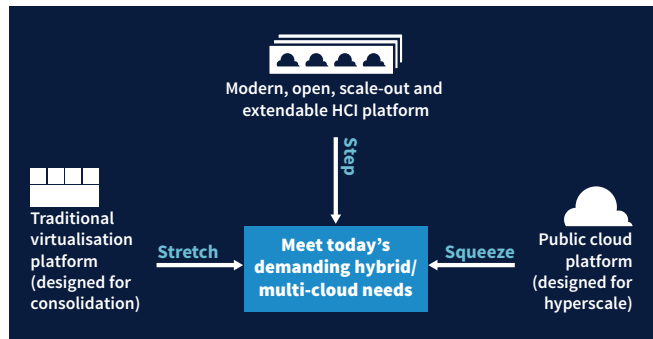
The second option is to take an existing hyperscale cloud platform and 'squeeze' it down to run on local servers – or rather, to take one that's already been repackaged for you to run locally. The big advantage is that it's natively converged and modern, the main disadvantage is that it is specific to one cloud provider. There is also the question of just how well hyperscale systems integrate with existing enterprise tools.

Lastly, you could take a big 'step' and move to a new platform, one that was designed from the start to operate as part of a converged hypercloud. The big advantage here is that you are starting in the right place, not in the old world, using technology the way it was designed to be used. Of course some existing skillsets may become redundant, but this is inevitable, and there will be new ones needed to replace them.

The challenge common to all three options is how much of your existing investment can be carried over without compromising the ability both to tightly-couple public and private resources and to continue evolving into the future.

Remember that

Stretching or squeezing existing platform architectures to create or support hybrid multi-clouds can work, but may require compromises.
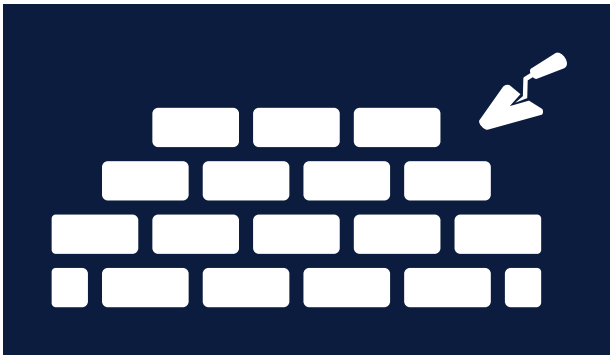
Choosing a platform

Too often in technology we look to solve the immediate problem as quickly and cleanly as possible, but while we consider the solution's future viability we do not always consider the role it will play in our wider systems architecture.

An example might be virtualisation: organisations adopted it to make better use of their physical resources and add flexibility, but few will have foreseen the way it then became the foundation upon which almost everything else now depends. That development 'just happened', with little oversight or direction.

Build a foundation for the future

Anticipating the future is impossible, but we can build a foundation for it. A simple example might be ensuring that the platform we select is open, with APIs and the opportunity to add plug-ins. In this way we leave the way clear for future growth and expansion.

An open platform also avoids the ticksheet situation, where buyers try to future-proof and protect their decisions by making long lists of "necessary" features, only 10% or 20% of which will really be used. Then suppliers in response add layer upon layer of additional functionality, most of which will in reality simply get in the user's way.

The problem, of course, is that while each of us might indeed only use 20% of the functionality available to us from an app, one user's 20% can be very different from another's. In addition, given the high rate of change in technology, everyone's 20% next year could be very different from this year's.

This is where an open and API-enabled modular or platform approach is more useful. Not only can the supplier offer a range of options for different needs, but others, including end user organisations themselves, can develop plug-ins to add or amend functionality.

Remember that

We don't build for the future simply by adding in everything we can think of today; we do it by ensuring that the way is open for additions to meet tomorrow's needs.

Key technologies and terms

The exact mix of technologies required is likely to be different for each installation, and will most probably change with time. However, there are capabilities which will be core in the majority of use cases. There are also key infrastructure technologies which, while not mandatory, will be part of most platform solutions, and which we need to be ready for, such as containers, microservices and serverless computing.

Hyperconverged infrastructure (HCI), is a method of building a system that can be fully software-defined. An HCI system can 'build' infrastructure services automatically and on-demand, drawing upon pools of virtualised physical resources. In effect, it can operate as a private cloud platform.

An equivalent is evolving: the **hyperconverged cloud**. Here, hybrid cloud resources are classified, abstracted and pooled, and can be automatically allocated as needed.

Functional needs

Multi-cloud technology and the new operational and programming models that it supports are both inherently hybrid and abstracted. This means that all the supporting functions and technologies must evolve too. Disaster recovery, file services, object storage, data protection, automation and orchestration, workload migration – all these and more need to be fit for the hybrid multi-cloud.

Also essential, given the trend towards containers, microservices and serverless computing, is a modern storage infrastructure. For example, many containers will be stateful, not stateless, so your storage should be fully container-aware.

☑ **Hyperconverged infrastructure**

☑ **Object storage**

☑ **File services**

☑ **Orchestration**

☑ **Public cloud**

☑ **Kubernetes**
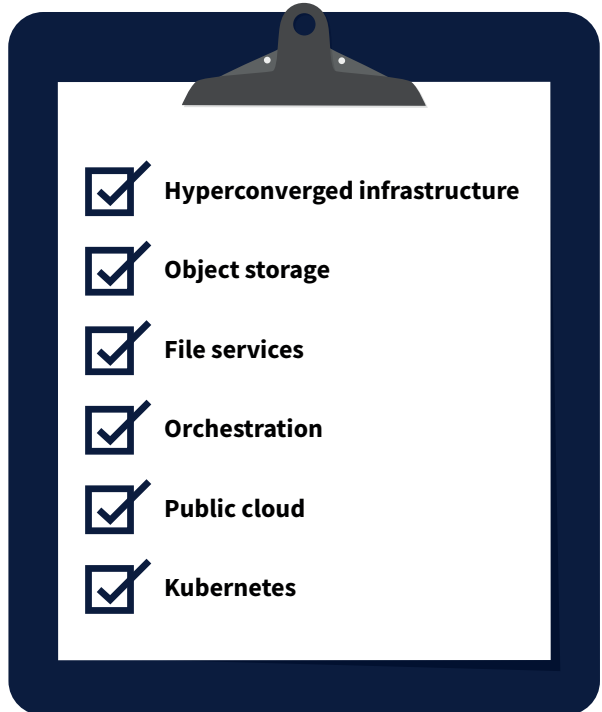
Remember that

Software-defined platforms and infrastructure can provide the necessary flexibility, via the likes of hyperconvergence and abstraction.

The bigger picture

As we move into the hybrid multi-cloud future, service delivery must evolve too. More specifically, to support new distributed application architectures we need significant new capabilities at the delivery level. Examples include:

Application and service discovery; Cloud security; Governance, auditing and compliance; Automated platform right-sizing; One-click software installation; Application lifecycle management; Service dependency monitoring; Edge computing and IoT; and more.

Beware the extremes, balance is best

System architectures designed for the previous generations of service delivery – the locally-managed virtualisation platforms at one end of the spectrum, and the hyperscale clouds at the other – were not designed with all those new capabilities in mind. They must therefore be added, or the system replaced by something more modern which has them built-in.

Two obvious routes are to select individual best-of-breed tools for these various tasks or requirements, or to select a full-spectrum suite or platform designed to handle as many of them as possible.

Bear in mind though that each of these routes presents challenges. In one case there is the work required to integrate tools with your existing platforms and processes, while in the other there is the risk of lock-in, whether it is specifically to that software suite or more generally to that software vendor.

The availability of tools must therefore be evaluated against the degree of lock-in. A broad range of tools is clearly a good thing, but built on an open platform and integrated via APIs.

At the same time, the new system architecture needs to be open to the old ones, able to encompass and interoperate with both traditional virtualisation and hyperscale cloud. It is a tough challenge, but fortunately the tools are there to build it.

Remember that

Building a platform for modern and forward-looking service delivery means finding suppliers with a strong integration story, yet who are committed to openness.

Summary

As demands change and grow, the natural tendency is to adapt and extend our existing technology to keep up. So it is with the drive to construct advanced container and microservices-based application architectures on top of the highly-evolved virtualised platforms that underpin today's enterprise data centres.

There comes a time though when it is no longer efficient or effective to keep stretching what we already have. Stepping instead to the next generation of technology – to a modern, hyperconverged, hybrid and multi-cloud platform, for example – may involve short-term costs as we write off some of our past investments, but once past that hurdle, growth should become faster and smoother.

Integrated thinking

Key aspects of modern application architectures include modularity and openness: rather than being a single off-the-shelf package, they can be composed of multiple services that communicate via shared APIs. It makes sense therefore to select a platform architecture that is also modular and extensible. There are people aspects involved, too. Some technical staff may be reluctant to change with the times, perceiving it as a threat to their hard-won skills. The truth, however, is that those skills are already becoming legacy – they will still have value, but more in the area of maintenance, not at the leading edge. And for most, this is an opportunity not a threat – a chance to learn and work with new technologies.

Opportunity knocks

With digitalisation changing the world and how people think, anyone involved in software needs to act proactively instead of reactively. The current generation of IT has served us well, but it is time to plan for what comes next.

Further Reading

1. Refresh your thinking
   You've got a platform problem

2. The Enterprise Cloud Imperative
   Time to shake things up a bit?

3. Fast-Track to Enterprise Cloud
   Integrated platforms as a route to digital readiness

About Freeform Dynamics

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

For more information, visit www.freeformdynamics.com or follow us on Twitter @FreeformCentral.

About Nutanix

Nutanix is a global leader in cloud software and hyperconverged infrastructure solutions, making infrastructure invisible so that IT can focus on the applications and services that power their business. Companies around the world use Nutanix Enterprise Cloud OS software to bring one-click application management and mobility across public, private and distributed edge clouds so they can run any application at any scale with a dramatically lower total cost of ownership. The result is organizations that can rapidly deliver a high-performance IT environment on demand, giving application owners a true cloud-like experience.

Learn more at www.nutanix.com or follow us on Twitter @nutanix.