

Complexity busting, hyper-converged style

The continuing quest for simplified IT

Dale Vile & Bryan Betts, Freeform Dynamics

Complexity is one of the biggest enemies when you are trying to deliver IT services to the business in an effective, efficient and responsive manner. The nature of the problem is clear when you consider all the individual servers, storage devices and pieces of network equipment installed in your data center, and the wide range of systems software necessary to make them work. For each type of component, e.g. x86 servers or storage arrays, you have probably accumulated products from numerous different vendors, often spanning multiple generations of technology.

But it's not just the range of components we need to consider; complexity also arises from the seemingly infinite variety of ways in which equipment and software are brought together to form working systems. As business solutions have been added and updated over the years, each with its own set of requirements, the chances are that you have collected almost as many uniquely configured application stacks as you have individual applications. Then, of course, overlaid across all this, you are probably using myriad monitoring and management tools that may or may not work together as well as you would ideally like.

If much of this sounds familiar, then your IT team will be spending a lot of its time focused on simply keeping everything working, rather than dealing with all those new and changing requirements coming from the business. And when IT staff do try to implement something new, or make a significant modification to an existing system, the complexity we have highlighted acts as a big barrier to progress. Figuring out the infrastructure implications of a change request is hard enough, but working around the dependencies and avoiding the risk of breaking something during implementation is even more difficult and time-consuming.

Against this background, an interesting question to ask is whether you would design your infrastructure in the way it currently looks if you could start again with a clean sheet of paper. The answer is undoubtedly 'no', but sadly this is academic, as you are unlikely to have the luxury of sweeping everything away and replacing it all.

You can, however, begin to break out of the infrastructure complexity escalation cycle and start to make things better incrementally. The trick is to view every requirement for a new system and every request to re-platform an existing application as an opportunity to move forwards in a simpler and more robust manner.

How?

Well, from an infrastructure perspective, one way to simplify things is by buying equipment in the form of pre-integrated blocks of compute, storage and networking, which come with a similarly pre-integrated set of systems software and tools that allow everything to be managed via a single pane of glass. Vendors delivering such 'converged infrastructure' generally provide a single support and maintenance arrangement covering the whole environment. This takes away a lot of the worry, distraction and pain to do with making sure components work together optimally, and ensures that everything continues to work together harmoniously as components are upgraded over time.

If you go down the converged infrastructure route, a big benefit is that because you can simply add more building blocks as additional system resources are required; your environment is extended in a consistent manner. This has the effect of naturally reducing variation at a physical system level, which in turn helps to reduce overall complexity and the related problems we have been discussing.

More recently, the concept of pre-integrated stacks has been taken further through the emergence of so-called 'hyper-

converged' infrastructure (HCI). In this model, convergence is implemented by fusing compute, storage and sometimes even networking capacity together much more tightly to provide a seamless pool of virtual resources with almost infinite scalability. The 'hyper' prefix invites analogies with the industry-standard hypervisor which also pools and allocates compute resources, and is an integral part of most HCI software stacks.

If this is starting to sound a bit cloud-like, that's because it is supposed to. If you read around the topic of hyper-converged, you will often find it mentioned in the context of scale-out storage. The reality, though, is that solutions in this space allow you to create a standardized general-purpose infrastructure upon which complete virtual machines (including storage) can be provisioned and run with the kind of ease normally associated with a public cloud service. The difference, of course, is that the platform is running in your own data center.

Hyper-converged is essentially an example of software defined infrastructure. As such, you can build your own hyper-converged platform by acquiring the relevant software separately, then adding your own hardware devices into the mix (though beware that components usually need to be certified by the software vendor). If you want the full benefits of a one-stop-shop, single integrated solution approach, however, your preferred route would be to buy your hyper-converged platform in the form of one or more all-inclusive, pre-configured and pre-validated systems. This provides all the benefits described earlier in relation to converged infrastructure and the building block approach, but with the architectural and single-source advantages of the cloud-like hyper-converged model.

Zooming back out to the bigger picture, some argue that hyper-converged platforms represent the future of all computing infrastructure. While it would be reasonable to assume that a lot of data centers will move in this direction over time, it's important to bear in mind that the hyper-converged model is only suitable for applications and workloads that can be run effectively in a virtualized environment. The chances are that you still run a lot of software that needs direct access to system resources, and then there are occasions when it makes sense to run specialized integrated platforms, e.g. for particularly demanding applications, or for workloads with specific requirements or legacy licensing terms.

Given the variation in requirements, the likelihood is that you will have to mix and match different types of stack for the foreseeable future. This will increasingly mean a combination of hyper-converged infrastructure for most modern, virtualization-friendly software, and converged infrastructure, which is essentially a more convenient way of delivering a traditional system-stack, for the applications that need it. Hand-crafting dedicated system-stacks may still be required from time to time, but such infrastructure silos are the historical source of troublesome complexity, and will (we hope) become less common over time.

As with everything in IT, it rarely makes good business or practical sense to force wholesale shifts at an infrastructure level. As you evolve your data center, however, the combination of converged and hyper-converged solutions with specialized integrated hardware will allow you to progressively standardize and simplify over time through a building block approach. This drives efficiency and flexibility, and allows you to focus on delivering the only thing that really matters: value to the business.

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, and access to our library of free research, please visit www.freeformdynamics.com.

About Fujitsu

Fujitsu is the leading Japanese information and communication technology (ICT) company offering a full range of technology products, solutions and services. Approximately 162,000 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.

Terms of use

This document is Copyright 2016 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 article for download and/or mass distribution by any means is prohibited unless express permission is obtained from Freeform Dynamics or Fujitsu. The contents contained herein are provided for your general information and use only, and neither Freeform Dynamics nor any third party provide any warranty or guarantee as to its suitability for any particular purpose.