

Media Backgrounder

Integrated Systems and Hyper-converged Integrated Systems (HCIS)

June 2019

The consistent message from the front line is that pressure is intensifying to find and retain IT professionals suitably-skilled for digital transformation. People with the right career experience are in high demand and can pick and choose where and how they want to work. Whole sectors, including medium-sized enterprises and government and related agencies, are experiencing severe difficulty resourcing transformation programs.

The situation is compounded as technology evolves and as options and decisions become increasingly complex. Accordingly, infrastructure deployment projects become error-prone, time-consuming, risky and expensive. The 'grit in the oyster' is the complexity of the infrastructure itself – with its promise of a pearl in terms of performance, agility and efficiency benefits – but composed of highly diverse components: servers, storage, networks, virtualization layers for all these components, databases and other middleware, as well as applications and a management layer to keep all these components under control.

Risk is becoming off-the-scale in do-it-yourself infrastructure roll-outs, where all the options and components are chosen and assembled by an in-house IT team. Unsurprisingly, this tactic has been going rapidly out of favor in recent years. Today, data center availability is make-or-break for organizations in a digital economy that is already desperately short of the necessary IT skills. This has rendered the build-your-own option both unsustainable, and even potentially career-limiting.

Therefore, the alternative pre-integrated approach, where all the options, system integration and testing are handled in the factory, has become the norm and is now accepted as the business-savvy way to handle data center build or overhaul. Integrated Systems reduce complexity through eliminating the high-risk trial-and-error process when introducing new infrastructure in data centers, while the compatibility of all components is guaranteed. At the same time, risk is minimized and the skill-sets required are less demanding. Less time is needed for planning and deployment is massively accelerated, therefore shortening time-to-value.

Due to the optimized design of Integrated Systems, system resource utilization will also be optimized. This can have a positive impact on data center space, cabling, energy consumption and cooling. Furthermore, an Integrated System represents an optimum foundation for efficient operations and reduced maintenance efforts. All these aspects help reduce CAPEX and OPEX costs and refocus IT organizations on the important aspects of their business. Moving away from an operating mode of build and maintain, headlined by break-fix, means improved responsiveness to new business requirements, or even in helping drive business to a new level.

Converged and Hyper-Converged Integrated Systems

The era of Integrated Systems began with the classical converged systems approach, where servers, storage, network connectivity and software were pre-integrated to accelerate deployment, minimize compatibility issues and simplify management. Today, the market has moved beyond that, with CIOs increasingly investigating and selecting hyper-converged infrastructure systems (HCIS). These systems tightly integrate all resources in a commodity server node, meaning it is no longer necessary to have a dedicated physical Storage Area Network (SAN). Whether you choose classical converged or hyper-converged, FUJITSU Integrated System PRIMEFLEX is one of the leading suites, and promises to reduce complexity and risk, shorten time-to-value and reduce costs.

HCIS takes the benefits of converged systems in general to another level. Hyper-converged systems provide customers with an easy path to achieve a fully software-defined data center. HCIS allows data storage to be spread across local disks within the server nodes. Built-in data services – such as data replication, snapshots, deduplication and data tiering – turn hyper-converged systems into a software-defined storage (SDS) platform. This creates unified, single-pane-of-glass management for both compute and storage resources. The new level of simplification this brings, reduces administration requirements and lowers the demand for highly-skilled administrators. Scalability is improved too: while classical converged systems can scale on a component level, hyper-converged systems are best suited to horizontal scaling – through adding more machines to the resource pool, rather than vertical scaling, which involves adding disk capacity or compute performance. As a result, hyper-converged systems are easily aligned to growing business demands, helping always ensure business continuity.

Although software license costs will probably be higher with HCIS, as there is no external storage, the data center footprint will often be reduced, as will energy consumption and cooling requirements and the associated operating costs of staff resources. In addition, the absence of storage arrays also implies savings on SAN switches and licenses for storage management software (in both the arrays and the switches).

Converged or hyper-converged? The use case matters

Hyper-converged systems are not some kind of trump card - a winning gambit to be played at every available opportunity. Fujitsu advocates assessing each use case carefully to select the optimum strategy:

- **Workloads.** When workloads scale horizontally, in terms of needing additional system capacity, then hyper-converged will be a perfect fit, especially if compute and storage resources need to scale in tandem. This applies to workloads which require a fixed amount of CPU performance, main memory, disk space and IOPS. Typical examples are Hosted Virtual Desktops and Hosted Shared Desktops. Where workloads scale vertically, in terms of needing more compute power or more storage capacity, or they require a granular expansion on the component level, then hyper-converged might be less appropriate – as with monolithic applications, for example, which cope with increasing data volumes using a scale-up approach: adding extra computing power or extra storage capacity in a single box.
- **Virtualization.** This is a prerequisite for hyper-converged infrastructures. And this means HCIS is not suitable for workloads which run on “bare-metal” (single tenant physical servers) only, where a virtual environment would be ineffective and just serve to slow things down. And as most hyper-converged implementations are based on a single hypervisor (which creates and runs virtual machines), they won't fit if a mixed operation of multiple hypervisors is needed to run different workloads.
- **Physical location.** Hyper-converged has become an attractive option for remote offices and branch offices. Since no external storage infrastructure needs to be maintained, frequent and costly onsite visits by technicians can be avoided. There are customer cases where travel time has been eliminated simply by replacing a physical SAN with a hyper-converged infrastructure.
- **Storage strategy.** Where workloads benefit from the data services inherent with hyper-converged, those services can be used without any additional investment. However, if they are not needed, you will indirectly pay for something you don't use. If existing storage is to be used for new workloads hyper-converged is not a great strategic fit. On the other hand, if existing storage is to be replaced sooner or later, then choosing hyper-converged may be a good starting point. However, the storage capacity of any hyper-converged infrastructure is limited by the number of server nodes. Therefore, hyper-converged is not a viable option for any workload where it needs to cope with amounts of data larger than the maximum storage capacity of the server cluster – because there's no SAN to store the extra data.
- **Growth.** Another aspect worth considering is the expected growth rate of any organization and its predicted needs. The more frequently infrastructure expands, the more it can benefit from the ease of scalability that hyper-converged confers.
- **Change management.** The unified management of compute and storage resources reduces operational complexity, administration efforts and cost. But bear in mind that going this new way will change existing staff roles and require other organizational changes, perhaps also making headcount reduction viable: enterprises thinking about making the switch should expect resistance from highly-skilled IT staff, especially in the storage area.
- **Network performance.** Although hyper-converged systems promise linear scalability, predictable network performance with larger deployments has been questioned by some users.
- **Software costs.** Software licensing can also be an issue. For instance, a database application may be a perfect fit for hyper-converged, but in cases where license fees are payable per CPU socket or even per CPU core, hyper-converged will quite likely be a no-go for commercial reasons.

In the final analysis, the choice between converged or hyper-converged is going to come down to considerations of cost, speed and agility for specific use cases. Operational expenditure always tends to be much lower with hyper-converged infrastructures compared to converged. However, when it comes to capital expenditure, it is hard to state a general rule of thumb. Typically, from a hardware cost perspective, hyper-converged systems are more attractive than classic; from a software cost perspective it is the opposite. Hyper-converged often requires a minimum number of server nodes (which may be higher in individual cases than what is actually needed), special, certified hardware components, while license fees must also be paid for the virtualization software. There is no single rule to say which option is best – each use case will be different and therefore, impartial advice from a trusted partner will be extremely valuable.

FUJITSU Integrated System PRIMEFLEX

The Fujitsu PRIMEFLEX brand includes Integrated Systems supporting both converged and hyper-converged infrastructures, making it practical and cost-effective to build infrastructure for general purposes, for heterogeneous usage scenarios, or purpose-built and optimized for a specific use case.

Fujitsu's Integrated Systems are built from best-in-class components, either Fujitsu's own technologies, for example the FUJITSU Server PRIMERGY or FUJITSU Storage ETERNUS, or third-party technologies from leading technology partners, such as VMware, Microsoft, Nutanix and NetApp, all recognized market leaders.

Converged lifecycle management

FUJITSU Infrastructure Manager (ISM) enables a converged lifecycle management across servers, storage, networking, and third-party devices. The impact of unified management for all components is simplified operations. Statistics say that troubleshooting and root cause identification are 23x faster in comparison to silo management, while installation time for operating systems is reduced by 70%. Agility is also increased, as the intuitive software provides actionable insights, enables workloads to easily be re-assigned, and can redirect traffic based on policies. This increases the responsiveness of IT in support of business needs. Another benefit of using ISM is cost reduction. Required time and cost for firmware updates is reduced by 90%, while power consumption is reduced by 50% due to ISM's built-in power cap function.

Delivery options

Fujitsu's integrated systems are either delivered as appliances or reference architectures. Appliances are built and tested in the same way as servers or storage systems, achieving the same quality level. They are not just pre-configured, pre-integrated and pre-tested; they are also fully pre-installed at the factory before being shipped ready-to-run to the customer. Onsite activity is therefore confined to the physical installation and integration into an existing environment. Procurement too is very easy; often there is a single order code (SKU) for the entire system. These advantages mean that there is no design effort at all, eliminating the risk of component clash and providing the fastest path to operational effectiveness.

If more configuration flexibility is required, reference architectures are the better choice. These are pre-tested, validated, design blueprints based on a proof-of-concept. What's more, they can easily be adapted and custom-tailored to customer-specific requirements. With reference architectures, component integration and installation typically happen onsite at a customer. To ensure a smooth set-up, Fujitsu provides detailed configuration and installation guidelines as standard. Depending on the individual offering, onsite deployment by Fujitsu or a certified partner is a mandatory part of the Integrated Systems offering, ensuring that reference architecture-based systems are handed over ready-to-run to the customer – providing similar benefits as with appliances, but with the additional advantage of flexibility. As an option, reference architectures or adjusted reference architectures can even be pre-installed in one of Fujitsu's staging centers, ensuring that the reference architecture is shipped ready-to-run "ex-factory" to the customer, also helping reduce 'time to value'.

+++ ends +++

Media contact

Isabell Horvath

Director of PR

Corporate Communications

Tel.: +49 (89) 62060 4419

E-Mail: isabell.horvath@ts.fujitsu.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 132,000 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. Fujitsu Limited (TSE: 6702) reported consolidated revenues of 4.0 trillion yen (US \$36 billion) for the fiscal year ended March 31, 2019. For more information, please see www.fujitsu.com.

<http://www.fujitsu.com>.

About Fujitsu EMEA

Fujitsu promotes a Human Centric Intelligent Society, in which innovation is driven by the integration of people, information and infrastructure. In the Europe, Middle East, India and Africa region (EMEA), our 28,000-strong workforce is committed to Digital Co-creation, blending business expertise with digital technology and creating new value with ecosystem partners and customers. We enable our customers to digitally transform with connected technology services, focused on Artificial Intelligence, the Internet of Things, and Cloud - all underpinned by Security. For more information, please visit <https://www.fujitsu.com/emeia/about/>