

How to reap the benefits of virtualization in a hybrid data architecture

Virtualization has come a long way in on-premises data centers. Now is the time to extend the use of this technology across edge, core, and cloud to enable real hybrid cloud environments.

Content

Virtualization across edge, core and cloud	2
Server virtualization has become ubiquitous	2
Containers – a compelling alternative to hypervisors	2
Virtual desktop infrastructures (VDI)	3
Storage virtualization	3
Network virtualization	3
Extending virtual infrastructures to the cloud	3
Converged or hyper-converged? – The use case matters	4
Fujitsu Integrated System PRIMEFLEX	6
Flexible service options	9
Summary	10



Virtualization across edge, core and cloud

Virtualization, as a vital step for resource optimization, has become the norm in reducing IT complexity and costs in on-premises data centers. Now, the need of the hour is to build a foundation for a more agile and flexible infrastructure that not only covers on-premises deployments, but also has the potential to span edge, core, and cloud environments.

However, given the myriad of virtualization solutions on the market, the question is how to go about selecting the solution that is right for you. They all have their own place, but lessons learned in the past have shown that there is no one-size-fits-all approach. Any decisions related to your infrastructure always need to consider the individual requirements of your application landscape.

Not so long ago it was common in data centers around the world to run applications directly on physical servers that often sat mostly idle. Within a decade, the use of server virtualization has become common industry practice. Adoption of this technology also increased thanks to additional benefits, such as dramatically improved IT efficiency, application isolation, greater workload portability, improved scalability, and high availability options. Ultimately, these advances have helped many companies to get more out of their hardware and “do more with less”. Although in fewer numbers, many organizations have achieved similar benefits by virtualizing other IT elements, including storage, networks, and desktops. As more companies view the cloud as a viable option for sourcing IT, the integration of virtualized on-premises deployments with those of public cloud providers becomes increasingly important.

Server virtualization has become ubiquitous

Many of the advantages of server virtualization are essential contributions to reducing complexity and administration efforts, as well as operational and capital expenditure. At the same time, server virtualization is an essential step toward making your IT infrastructure more agile. Today, nearly every organization is using some form of hypervisor technology to virtualize physical servers. Hypervisors offer a robust and secure platform to run the vast majority of workloads, whether traditional, cloud-native, or mission-critical. A hypervisor is computer software, firmware or hardware that creates and runs virtual machines. A computer on which a hypervisor runs one or more virtual machines is called a host machine, and each virtual machine is called a guest machine. The hypervisor presents the guest operating systems with a virtual operating platform and manages the execution of the guest operating systems. Multiple instances of a variety of operating systems may share the virtualized hardware resources. For example, instances of Linux and Windows can all run on a single physical x86 machine.

Containers – a compelling alternative to hypervisors

An increasingly important alternative to server virtualization based on traditional hypervisor technology is the use of container technology. These are structures that can encapsulate everything needed to run a particular application: settings, libraries, allocated resources, and other dependencies. Containers isolate applications from other processes running on an operating system, which provides stability and security benefits. And because everything an application needs to function resides inside a container, when a container migrates from one computer to another (for example from a staging to a production server), the application continues to function the same, even if the host operating systems have different settings or are running on slightly dissimilar hardware.

Containers are similar to virtual machines in that they both allow you to isolate applications and divide computer resources, such as memory and disk space. However, a big difference between containers and virtual machines is that each VM encapsulates its own operating system, so a VM is really like an entire virtual computer. In contrast, containers run on top of an operating system, meaning there is only one OS instance responsible for hosting many containers. This makes containers more lightweight than VMs, and remove the need for software licenses for multiple guest OSs. However, there is a lesser degree of isolation between applications because the same OS kernel is shared by multiple containers.

Containers make a lot of sense for organizations that carry out their own in-house software development. They have been seen as a way to mitigate issues that occur during the handoff between development, testing, and operations. Thus, containers have gained a lot of traction in a DevOps environment.



Virtual desktop infrastructures (VDI)

Virtual desktop infrastructures allows desktop operating systems, including desktop applications, to be run as virtual machines (virtual desktops) on a physical server infrastructure in the data center. Instead of a fully equipped PC, a thin client or any other device is sufficient for access to the server. While this form of virtualization can result in IT efficiencies (better use of resources, centralized management and improved resilience, disaster recovery), VDI can also help companies protect sensitive data and support more flexible working conditions. This is increasingly important for users in an evolving workplace where many employees work remotely using a variety of devices.

Storage virtualization

Just as server virtualization increases the efficiency of server operations, storage virtualization allows greater control and flexibility over physical storage and helps organizations get more out of existing storage resources. For these reasons and more, the adoption of storage virtualization is already fairly common and will soon become mainstream.

In particular, hyper-converged infrastructure deployments using a combination of server and storage virtualization technology based on a cluster of x86 servers are increasing rapidly. According to a user survey conducted by the analyst firm FreeForm Dynamics, more than 70% of the top performing IT organizations already had considerable HCI deployments*. Many organizations view hyper-converged infrastructures or other forms of software-defined storage as a cost-effective replacement for traditional SAN and NAS devices. However, as we will explore later in a separate section, it always depends on the specific use case as to which of the approaches is best for your specific workload landscape.

Network virtualization

Network virtualization is probably the least mature virtualization technology, but it is gaining growing popularity, especially when we look at the increasing use of hybrid cloud environments, where network virtualization is one of the key enabling technologies.

In traditional networks, the control logic and the data flow are implemented in the same network device, typically something proprietary from any network component vendor. Therefore, network configuration has to happen manually at the device level, which is not only complex, costly, and cumbersome, but also inflexible. This often hinders organizations to capitalize on the full efficiency and agility gains of their virtualized server and storage infrastructure. Network virtualization is the process of combining hardware and software network resources and network functionality into a single, software-based administrative entity – a virtual network.

Network virtualization thus decouples the control logic from the data flow. While data is forwarded locally through the network device (which no longer needs to be as intelligent as in traditional networks), management is performed centrally. It is no longer up to the network device to find a route from the source to the destination of a data flow; instead, an optimal routing can be controlled by intelligent, central management software, even in an automated way.

Major benefits include more flexibility with network configuration, improved management capabilities, the ability to reconfigure networks without physical changes, improved security capabilities, more control over network segmentation, automation benefits, and, of course, cost savings.

Extending virtual infrastructures to the cloud

While the use of virtualization technology brings significant efficiency and agility improvements to on-premises data center operations, many organizations look at the cloud as an increasingly important option to gain access to IT resources. And there are good reasons for this. An enterprise may want to add capacity for peak times without investing in additional equipment, which would remain unused most of the time. Or the organization may want to reduce its primary data center footprint in general, or for testing and development in particular. It may need to expand its IT to other geographies (distributing applications closest to the end user) instantly, without going through the lengthy process of making a contract with a local hosting provider. It may want to build disaster recovery (DR) capacities in the public cloud rather than in its own secondary data center, reducing costs drastically. Or the enterprise just might want to achieve the agility needed to react quickly to new business demands, e.g. using DevOps.

A hybrid cloud environment gives you the best of both worlds – the control of your on-premises data center and the cost-effectiveness of the cloud. To make all this a reality, it is essential that on-premises data centers and the public cloud are not separated entities. They should be integrated and complement one another. To achieve a real common management experience and to minimize interoperability issues as far as possible, your on-premises data center systems should be hybrid IT/cloud-enabled. Only this type of infrastructure ensures the management consistency needed to let both sides appear as if they were made the same. One piece of good advice is to implement the same infrastructure foundation in your on-premises location as the one that runs in the data center of your service provider.

* Source: Hyper-converged has come of age - Modern infrastructure options are now integral to service delivery, Freeform Dynamics, 2020



Converged or hyper-converged? – The use case matters

In the area of software-defined storage approaches, hyper-converged infrastructures are increasingly gaining attention in the market. Subsequently, there is an ongoing discussion in organizations as to which storage strategy — converged or hyper-converged — is the best way to move forward over the next couple of years.

A classic converged data center infrastructure consisted of servers, external storage systems, network components, and software. While using server virtualization software optimizes the compute part, there are still all the individual components of the storage infrastructure to deal with (storage arrays, SAN network, dedicated storage management tools). This is not the case with hyper-converged infrastructures: These tightly integrate all compute and storage resources in a commodity x86 server node, making a dedicated physical storage area network (SAN) with its management superfluous. Instead, storage is spread across the local disks of the server nodes. As there is no external storage involved, the data center footprint is often reduced, as are energy consumption and cooling requirements. Having compute and storage resources integrated in a single box makes deployment even easier and faster. The built-in data services, such as data replication, snapshots, deduplication, and storage tiering turn hyper-converged systems into a software-defined storage platform. The unified management of both compute and storage resources leads to a higher degree of simplification by reducing administration efforts and skills demands.

While classic converged infrastructures scale on a component level, hyper-converged systems enable scalability on a system level. Compute performance and storage capacity can be scaled by simply adding or removing servers. Thus, hyper-converged infrastructures can be easily aligned to growing demands, while ensuring business continuity. And of course, all these benefits often have a positive impact on capital and operational expenditure. Considering all these advantages, we might conclude that a hyper-converged infrastructure is the only way to go in the future. Nevertheless, let us first consider a number of aspects that might help you make the right choice.

Scalability options

If your workloads scale horizontally, hyper-converged will be a perfect fit, especially if compute and storage resources need to scale in tandem. This applies to those workloads that require a fixed amount of CPU performance, main memory, disk space, and IOPS. Typical examples are hosted virtual desktops and hosted shared desktops. If your workloads scale vertically, or require granular expansion at the component level, hyper-converged might be less appropriate. Another aspect worth considering is expected growth: The more frequently you have to expand your infrastructure, the more you will benefit from the ease of scalability offered by hyper-converged infrastructures.

Storage capacity

The storage capacity of a hyper-converged infrastructure is limited by the maximum number of server nodes. If you have to cope with data volumes larger than the maximum storage capacity of your server cluster, hyper-convergence will not be a feasible option. Even though the hyper-converged approach promises linear scalability, predictable network performance with very large deployments is still an unanswered question due to insufficient data in this area.

Multi-hypervisor support

Since most hyper-converged implementations are based on a single hypervisor, they will be unsuitable in those cases where mixed operation of multiple hypervisors is needed to run different workloads.

Advanced storage management

A hyper-converged software environment will not contain all the sophisticated storage management capabilities you are used to having when operating a dedicated storage management tool from an external storage system vendor. However, after being on the market for about 10 years now, the functional scope of HCI software has significantly improved with every release. Therefore, you could expect that hyper-converged software covers the majority of use cases today.

Remote office and branch office deployments

Hyper-converged concepts have become attractive options for remote offices and branch offices. As an external storage infrastructure is no longer involved, the number of costly onsite visits can be significantly reduced, simply because fewer systems need to be maintained. Some customer cases have cut travel time by 99 percent just by replacing a physical SAN infrastructure with a hyper-converged infrastructure.

Impact on IT organization and roles

The unified management of compute and storage resources reduces operational complexity, administration work and costs. However, it needs to be borne in mind that adopting this new approach will change existing staff roles and require other organizational changes. This could lead to some resistance from your IT staff, especially in the storage area, which would need to be handled. This aspect is also closely related to the storage strategy: If you intend to utilize existing storage, hyper-converged solutions will be a poor fit. However, if you intend to replace your existing storage sooner or later, opting in favor of a hyper-converged solution might be a good start.



Hybrid cloud

A hyper-converged infrastructure may provide you with an easier path to a hybrid cloud, because you will already be using the same software-defined compute and storage foundation that, for example, the HCI software vendors use for their deployments in a public service provider data center.

Software licensing

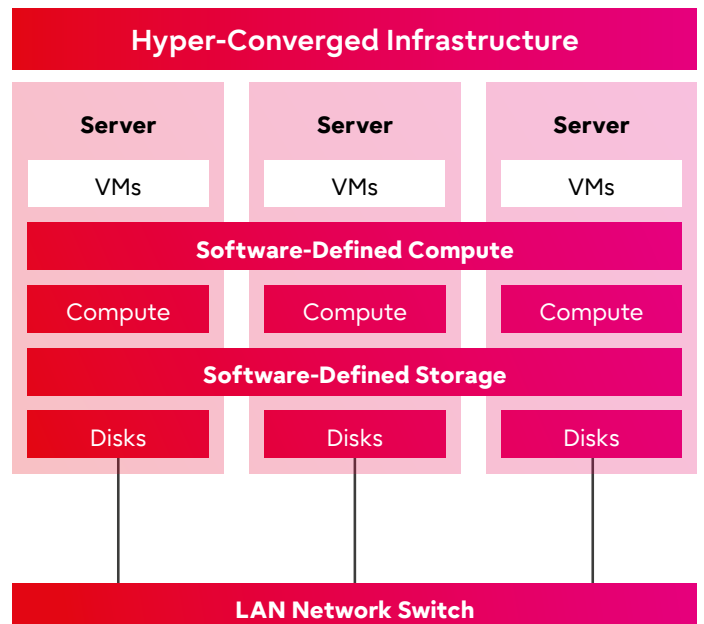
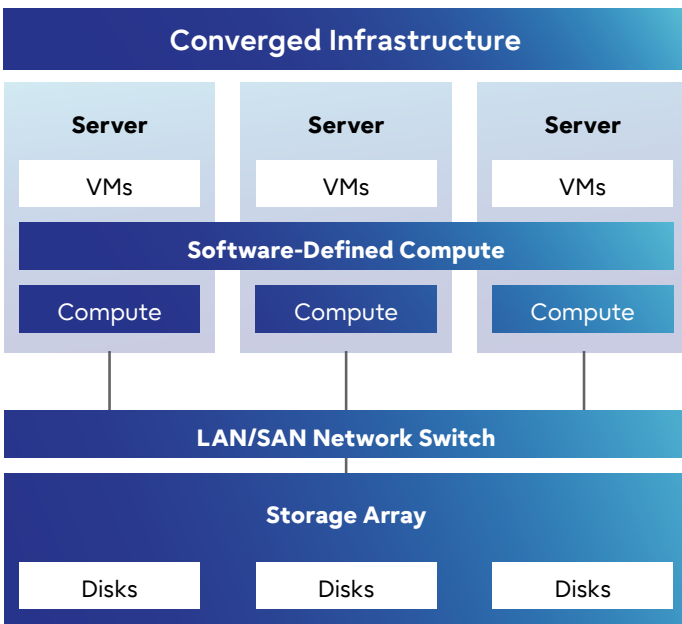
In addition to all the technical aspects, software licensing should also be taken into account. For instance, a database application may be a perfect fit for a hyper-converged scenario. However, if you need to frequently scale the storage part of your database, you would have to pay additional license fees (in the case of per CPU socket or per CPU core licensing) for the compute resources even though they are not required. In this case, a hyper-converged approach would not make sense from a business standpoint.

Ultimately, it is all about cost. As mentioned above, operational expenditures tend to be much lower for hyper-converged infrastructures when compared to classical converged concepts. However, when it comes to capital expenditure, it is hard to make a general statement. Typically, from a hardware cost perspective, hyper-converged is more attractive than classic converged. From a software cost perspective, it the very opposite may be true.

In most cases, hyper-convergence requires additional license fees for the software-defined storage capabilities. Nowadays, some of the HCI software products are also available on a subscription basis, which may allow some flexibility – if the use case is right.

In addition, the functional scope of the different HCI software editions varies greatly. For example, some of the HCI software vendors provide built-in software-defined networking or disaster recovery capabilities. Others require additional licenses for such capabilities. If your business case requires all these capabilities, it may be advantageous to have them built in. If you do not require all of them, you may end up paying for something you do not need. Hence, it is a mandatory exercise to carefully evaluate which of the individual capabilities you really need to cover your requirements.

Whether converged or hyper-converged, you will find many examples with cost benefits on both sides of the equation. There are good reasons to look at both architectural approaches. Make a simple cost comparison for your specific project to find out which option is most effective in terms of cost and functional scope. As previously stated, it is the use case that matters when comparing classical converged and hyper-converged scenarios. You should make the final decision based specifically on each use case. Whenever the benefits outweigh the drawbacks, you should opt for hyper-converged systems.



Fujitsu Integrated System PRIMEFLEX

With the PRIMEFLEX family of integrated systems, Fujitsu offers you the choice of different platforms to deliver an integrated virtual data architecture across edge, core, and cloud systems. Thanks to the unique combination of a pre-integrated and certified technology stack, new standardized implementation and infrastructure support services providing technical solution support with a single point of contact, these converged and hyper-converged infrastructure solutions remove the complexity to provide you with a quick, cost-effective, and risk-free way to deploy and run the virtual data architecture that exactly meets your workload requirements.

Hybrid-enabled VMware-based infrastructures

Fujitsu addresses VMware environments in on-premises data centers with two integrated systems offering different levels of virtualization.

• PRIMEFLEX for VMware vSphere

PRIMEFLEX for VMware vSphere is a converged system based on a classical architecture with external storage. This integrated system features VMware vSphere including native Kubernetes so that you can run existing enterprise applications alongside containerized applications in a unified manner, while maintaining application portability. You can choose between hybrid storage (Fujitsu ETERNUS DX) and all-flash storage (Fujitsu ETERNUS AF), as well as between iSCSI and Fiber Channel connectivity. Network switches (from Juniper or Extreme Networks for LAN connectivity and Broadcom for Fibre Channel connectivity), cabling, and rack infrastructure are also included. We also offer the option to bring your own network switches. There are multiple configurations varying in size, and Fujitsu's Infrastructure Manager (ISM) serves for converged lifecycle management of all components involved.

• PRIMEFLEX for VMware vSAN

PRIMEFLEX for VMware vSAN is a hyper-converged system based on the VMware HCI software stack, including VMware vSphere and vSAN. The system provides you with a broad choice of certified vSAN ReadyNodes™ giving you the highest degree of flexibility in choosing a VMware HCI environment. PRIMEFLEX for VMware vSAN supports any HCI use case, including general-purpose virtualization, VDI (Virtual Desktop Infrastructure), big data and analytics, remote and branch office, and edge, and has been certified for mission-critical workloads like SAP HANA. For companies who want to extend their VMware HCI deployment to a completely software-defined data center including networking virtualization, Fujitsu offers VMware Cloud Foundation on PRIMEFLEX for VMware vSAN, a solution that provides a complete set of software-defined services for computing, storage, networking, security, and cloud management to run your enterprise apps—traditional or containerized—in private or public environments.



Going hybrid with VMware

Powered by VMware Cloud Foundation, VMware Cloud on Amazon Web Services (AWS), Microsoft Azure or Google Cloud integrates VMware's compute, storage, and network virtualization products (VMware vSphere, VMware vSAN, and VMware NSX) along with VMware vCenter Server management, optimized to run on dedicated, elastic, bare-metal AWS, Azure or Google Cloud infrastructures. The key enabling feature to connect the VMware Cloud with your on-premises data center instance is VMware Hybrid Cloud Extension, an application mobility platform designed to simplify application migration, workload rebalancing, and business continuity across data centers and clouds.

Hybrid-enabled Microsoft-based infrastructures

Fujitsu serves Microsoft environments in on-premises data centers with two hyper-converged integrated systems based on different operating system platforms.

• PRIMEFLEX for Microsoft Storage Spaces Direct

PRIMEFLEX for Microsoft Storage Spaces Direct is a hyper-converged system based on software-defined storage technology (Storage Spaces Direct) integrated in Microsoft's Windows Server Datacenter. Various certified configurations for a broad range of use cases are already in place, covering mixed workloads as well as those requiring extreme I/O performance. With the Azure Hybrid Services available through Windows Admin Center, IT organizations can connect to Azure cloud services, including Azure File Sync, Azure Site Recovery and Backup, Cloud Witness, Azure Monitor and centralized Azure Update Management.

• PRIMEFLEX for Microsoft Azure Stack HCI

PRIMEFLEX for Microsoft Azure Stack HCI is another hyper-converged system, based on Azure Stack HCI – the new purpose-built operating system from Microsoft that offers an easy route to hybrid cloud, in-built Azure connectivity, subscription-based licensing, and advanced features for disaster recovery through stretched clusters, as well as Kubernetes capabilities. Various certified and validated Azure Stack HCI nodes with different form factors are available to cover a broad range of use cases, such as general-purpose virtualization, Kubernetes, SQL server, VDI or ROBO (Remote Office and Branch Office) environments.



Going hybrid with Microsoft

With Microsoft, connecting your data center with the cloud is based on Azure capabilities. The new subscription-based Azure Stack HCI operating system for on-premises HCI deployments released in 2020 has built-in cloud-native capabilities enabling you to create a hybrid cloud environment. Making it simple to connect your virtualized workloads to cloud-based services for backup, monitoring, identity access management, security, and more. In addition, you will automatically have the latest feature and security updates applied to your on-premises HCI deployment. If you are running the Windows Server operating system, your route to Azure cloud services is through Windows Admin Center, along with the agents needed to set up the respective Azure Hybrid Services, such as Azure Backup, Azure Site Recovery, Azure File Sync, and Azure Monitor.

Hybrid-enabled Nutanix-based infrastructures

Fujitsu addresses Nutanix environments in on-premises data centers with the hyper-converged system PRIMEFLEX for Nutanix Enterprise Cloud.

• PRIMEFLEX for Nutanix Enterprise Cloud

PRIMEFLEX for Nutanix Enterprise Cloud is a hyper-converged infrastructure system based on virtualization technology from Nutanix. Supporting any number of nodes, the multi-hypervisor system (Nutanix AHV and VMware vSphere) runs various configurations covering a broad range of use cases. These include special ones for VDI, ROBO (Remote Offices and Branch Offices) and pure usage as storage. Cloud management can be added as an option, and PRIMEFLEX for Nutanix Enterprise Cloud is also certified for SAP HANA.

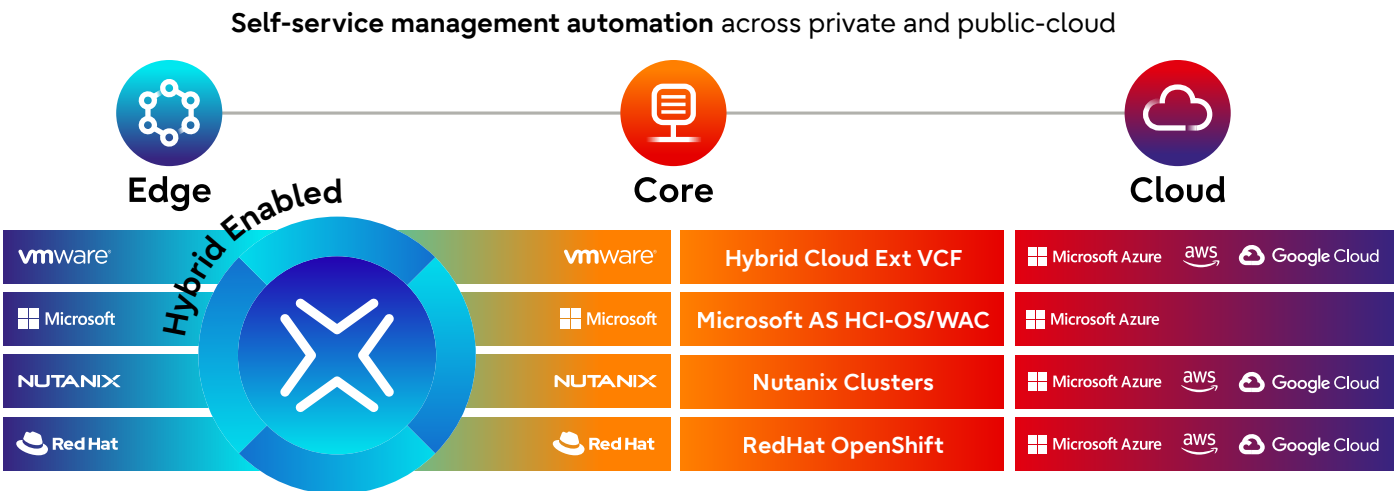
Going hybrid with Nutanix

Similar to VMware, Nutanix supports hybrid cloud environments based on Amazon Web Services (AWS).

The Nutanix hybrid cloud offering is based on Nutanix Clusters, which runs the core Nutanix HCI stack (including Nutanix AOS, AHV, and Prism), along with all Nutanix products and services, on bare-metal cloud instances – allowing you to easily migrate or extend applications from a private to a public cloud infrastructure. You simply connect to AWS while using your existing AWS account, virtual private cloud, virtual private network or Direct Connect configuration, request your virtual machine on a bare-metal Elastic Compute Cloud instance and deploy the Nutanix software in the virtual machine. Prism Central ensures a transparent management across locations. An additional advantage of taking this route is that you have direct access to other AWS cloud services. Nutanix has announced plans to also support Azure cloud, but this was not generally available at the time of writing

Going hybrid with RedHat

RedHat OpenShift has become the de facto standard for container orchestration. Red Hat OpenShift 4 functions like a public cloud infrastructure, building on Kubernetes Operators to automate the critical tasks of updating and managing the entire cloud-native stack. Self-managing and self-updating, OpenShift allows your IT teams to innovate at scale. Being a 100% open-source solution with no lock-in, it serves as a good alternative to extend your heterogeneous converged and hyper-converged infrastructure to cloud. This solution also provides a single set of interfaces for development and operations – both for virtual machines and containers. In addition, the Red Hat DevSecOps framework addresses key security requirements throughout the DevOps life cycle as part of a comprehensive defense-in-depth security strategy. Furthermore, with Ansible, in partnership with RedHat, Fujitsu can offer automated workflow, identification, search, and response to security events – providing integrated, secure, hybrid-managed capabilities across a private, on-premises, and public cloud architecture.



Positioning

Our experts at Fujitsu work closely with you to identify the approach that best fits your needs. To identify this ideal solution, multiple aspects need to be considered.

Among the important questions to be answered are:

- On which operating systems are your applications running?
- Which hypervisors support these operating systems?
- What is your preference if several hypervisors meet your needs?
- Is a converged or hyper-converged architecture more appropriate?
- Do you want to leverage software-defined networking in addition to software-defined compute or storage?
- How many physical servers need to be virtualized?
- What are your compute and storage scalability demands?
- Is a disaster recovery option needed and, if so, which one?
- If you are opting for a hybrid cloud approach, who is your preferred public cloud provider?
- How easy is it to integrate your on-premises virtual infrastructure with the cloud?
- What are your expectations regarding the time needed to have the solution up and running?
- And, finally, what is your budget?

The following table will help you identify the virtualization solution that is best suited to your needs:

	PRIMEFLEX for VMware vSphere	PRIMEFLEX for VMware vSAN	PRIMEFLEX for Microsoft Storage Spaces Direct	PRIMEFLEX for Microsoft Azure Stack HCI	PRIMEFLEX for Nutanix Enterprise Cloud
Architecture	CI	HCI	HCI (Windows Server OS based)	HCI (Azure Stack HCI OS based)	HCI
Server	PRIMERGY RX	PRIMERGY RX	PRIMERGY RX, TX	PRIMERGY RX, TX	PRIMERGY RX
Storage	ETERNUS DX/AF	Local disks	Local disks	Local disks	Local disks
Networking	Juniper, Extreme (LAN), Broadcom (SAN)	Any	Any	Any	Any
Hypervisor	VMware	VMware	Microsoft	Microsoft	Nutanix/VMware
Software-Defined Networking	No	No	Yes	Yes	Yes
Infrastructure Management	ISM (optional)	ISM (optional)	ISM (optional)	ISM (optional)	ISM (optional)
Scalability	2-64 servers	2-64 servers	2-16 servers	2-16 servers	2-any # of servers
Public cloud options	AWS, Azure, GCP	AWS, Azure, GCP	Azure	Azure	AWS, Azure, GCP
Hybrid cloud connectivity	VMware Hybrid Cloud Extension (HCX)	VMware Hybrid Cloud Extension (HCX)	Windows Admin Center	Built-in Azure Stack HCI OS	Nutanix Clusters



Flexible service options

PRIMEFLEX is supplemented by flexible services options throughout all lifecycle phases. These cover consulting, design, onsite deployment, integration of the new infrastructure into your existing environment and migration services, lifecycle management, and maintenance. And for those lacking the resources to operate their own data architecture, Fujitsu can provide Managed Data Center as well as Hosting Services.

Fujitsu Financing Solutions

In addition to the above-mentioned services, Fujitsu offers a complete range of financing solutions, such as IT leasing, trade-in, buy back or even cloud-like financing models. With Fujitsu uSCALE, the pay-per-use consumption model is no longer a domain of the cloud. Fujitsu is now helping organizations to achieve pay-per-use agility across their entire IT landscape, including their on-premises data center. Fujitsu uSCALE pay-per-use is an IT platform consumption service that supports business transformation and IT agility with scalable resources that are measurable, cost transparent, and tailored to your needs. It offers cloud-like delivery for your on-premises infrastructure and is deployed at mutually agreed uSCALE service locations, e.g., customer data center or co-location facility. You pay only for what you use, with simple and predictable monthly costs aligned with the scope of the service—such as per Gb of storage or each instance of VM.*

Fujitsu Implementation Services

To enable a fast time to production of a PRIMEFLEX solution, Fujitsu offers a range of standardized ImplementationPacks with a single order code and price tag covering the creation of a low-level design, the installation including all hardware and software components, the option to deploy workloads provided by the customer, and the handover of all documentation upon successful completion of the project. The new PRIMEFLEX Implementation Desk along with the easy-to-use web-based PRIMEFLEX Deployment Portal makes infrastructure deployment faster, more reliable and secure. Experience from many successful project implementations clearly shows that a well-executed and documented implementation is an important pre-requisite for delivering high-quality infrastructure support services and ultimately for improved customer satisfaction.**

Fujitsu Infrastructure Support

By choosing Fujitsu Infrastructure Support, Fujitsu or one of its certified partners will be your single point of contact for all support matters related to PRIMEFLEX. The Fujitsu SolutionPacks and SolutionContracts for integrated systems provide you with end-to-end, 24/7 infrastructure support covering the complete hardware and software stack including third-party components. Both service options relieve you of headaches caused by unpredicted problems during operation, while ensuring operational efficiency. A unique solution identifier for all PRIMEFLEX systems allows our support teams to make solution-level decisions when working with these often complex systems. For example, this enables us to route support calls to specialized, solution-aware support engineers, resulting in smoother support interactions with shorter resolution times. Beside reactive services based on optimized processes, Fujitsu also offers optional proactive services. These comprise a regular system health check to detect critical system conditions at an early stage such that preventive maintenance measures can be initiated. Just as there are various service level options available, which differ in service scope, response and recovery time, you can also define the frequency of proactive services.**



* Availability depends on configuration size and local terms and conditions. Please contact your local sales representative.

** Availability depends on local terms and conditions. Please contact your local sales representative.

Summary

Applying virtualization technology on various levels is an important step toward reducing complexity in the data center. However, it is no longer sufficient to only apply it to on-premises deployments. In a landscape that is increasingly turning hybrid, with workloads running on-premises and in the public cloud, it is extremely important that all systems in your data architecture are hybrid cloud-enabled. Allowing for an easy integration into a hybrid cloud environment that spans edge, core, and cloud systems.

With PRIMEFLEX Integrated Systems, Fujitsu helps you reduce complexity, time, risks, and costs when building a hybrid data architecture. PRIMEFLEX covers classical converged and hyper-converged architectures and supports virtualization technology from VMware, Microsoft, and Nutanix. This puts us in a unique position to provide you with unbiased advice to jointly identify the solution that perfectly suits your needs.

We help you:

- Understand which virtualization solutions and techniques are available, and how they allow you to think and/or act differently
- Gain insights into the latest developments regarding best practices, technology options, and service delivery
- Devise an effective hybrid data architecture and then select the right mix of infrastructure and cloud services to implement it.



White Paper

How to reap the benefits of virtualization in a hybrid data architecture



Intel® Xeon® Platinum processor

Published by
Fujitsu Technology Solutions GmbH
Mies-van-der-Rohe-Strasse 8
D-80807 Munich

www.fujitsu.com
2022-03-11 WW EN

© 2022 Fujitsu. All rights reserved. Fujitsu and Fujitsu logo are trademarks of Fujitsu Limited registered in many jurisdictions worldwide. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation or its subsidiaries. Other product, service, and company names mentioned herein may be trademarks of Fujitsu or other companies. This document is current as of the initial date of publication and subject to be changed by Fujitsu without notice. This material is provided for information purposes only and Fujitsu assumes no liability related to its use. We reserve the right to change delivery options or make technical modifications.