


# White paper

## Why PRIMEFLEX® Attracts Service Providers

Building data center infrastructures is increasingly complex, error-prone, time-consuming, risky and expensive. FUJITSU Integrated System PRIMEFLEX reduces complexity and risk, shortens time to value, reduces costs and increases operational efficiency for service providers.



Content	
What is happening in the service provider market?	2
Building a data center infrastructure is complex	2
FUJITSU Integrated System PRIMEFLEX: Many benefits for service providers	2
Best-in-class components and long experience	3
Delivery models	3
Flexible services options	4
PRIMEFLEX includes hyper-converged systems	4
Classical or hyper-converged – Use cases matter	4
PRIMEFLEX for many important use cases	6
PRIMEFLEX® for SAP HANA®	7
PRIMEFLEX® for SAP® Landscapes	7
PRIMEFLEX for Hadoop	8
PRIMEFLEX for Virtualization and Private Cloud	8
PRIMEFLEX for VMware vSAN	9
PRIMEFLEX for VMware Cloud Foundation	9
PRIMEFLEX for Microsoft Storage Spaces Direct	9
PRIMEFLEX Cluster-in-a-box	<b>Error! Bookmark not defined.</b>
PRIMEFLEX for VMware vSphere	9
NFLEX	10
FUJITSU Enterprise Service Catalog Manager turns your virtual infrastructure into a cloud infrastructure	10
Summary	10

### What is happening in the service provider market?

Let us start with a good message: More and more businesses are looking at IT service providers, because they cannot afford doing everything on their own, and they quite honestly do not have the resources and skills to do so. Some of them use a managed service, turning the operation of their systems over to an expert partner. Others move things to the cloud and allow cloud service providers to manage it – paying for only what they need. The cloud services themselves are hosted in a data center than can be accessed by businesses using network connectivity. The benefit of using a cloud service provider is efficiency and economies of scale. Rather than businesses building their own infrastructure to support internal services and applications, the services can be purchased from the cloud service provider, who provides the services to many customers from a shared infrastructure. Such an approach is absolutely appealing. An IT manager can give his internal users instant access to business applications they may need in the cloud, simply by signing them up. At the same time, the economies of scale decrease the cost of services. All told, many businesses understand that service providers can do certain things much better, faster and even more reliably than they could do themselves. Analysts predict that by 2021 the major share of workloads will be hosted by cloud service providers. This demonstrates a brilliant future for them.

However, there are also some shady sides. Due to competition, the service provider market is tough, while costs and margins are under pressure. To be among the winners, superior services need to be provided for new and existing customers, creating an outstanding value for the end customer's money.

As a foundation for providing IT services, solid data center infrastructures are needed, which have to be absolutely reliable; the investment risk should be low; they should be fast to bill to the service provider's end customers, and finally efficient to operate. That's why service providers are looking for infrastructure vendors who help make them more profitable through higher revenues and / or better margins. At the same time, risk should be taken out of their operation and their business, making them less vulnerable to the impact of market fluctuations, losing customers and other undesired incidents.


### Building a data center infrastructure is complex

Building data center infrastructures can be extremely complex. Servers, storage, network components and software need to be selected from a myriad of options, procured and integrated. As the compatibility of the components is not guaranteed at all, extensive testing is a must. This in turn requires a deep knowledge of all components involved and their interdependencies on each other. Therefore, a DIY (Do-It-Yourself) approach is error-prone, time-consuming and expensive, while presenting businesses with multiple risks. And finally, as every installation is different, maintenance efforts will be high.

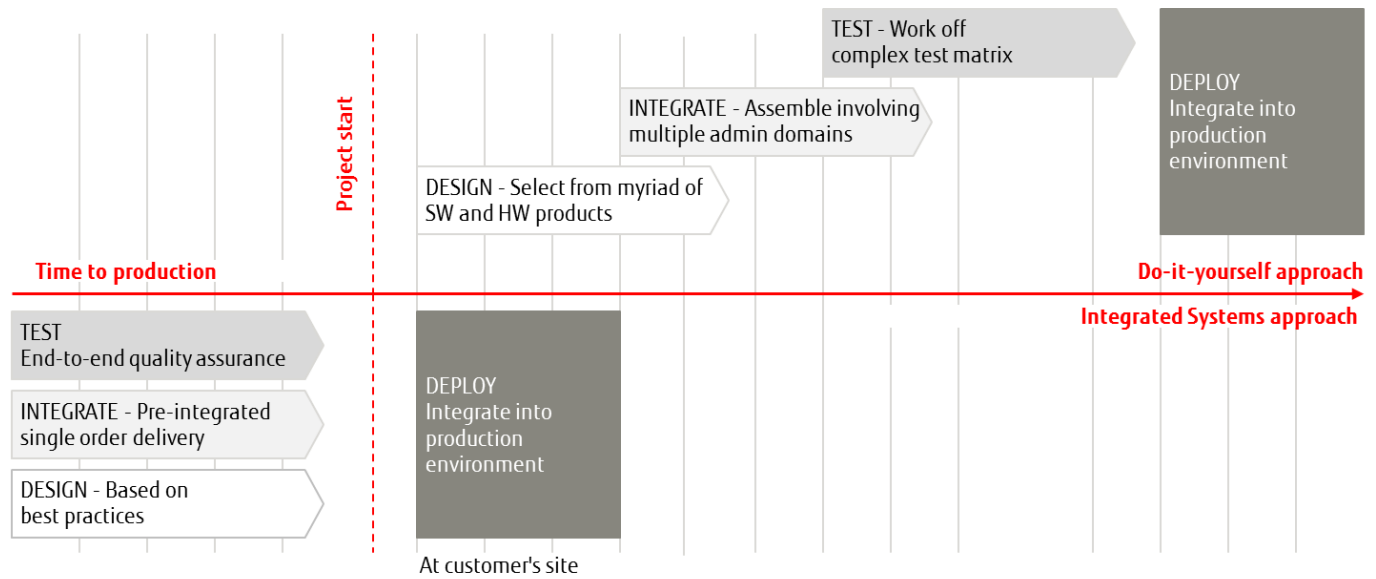
According to analysts, IT organizations spend about 25% of their time and resources on pre-deployment. Hence, more and more organizations are wondering whether building data center infrastructures on their own is really the best way to go.

### FUJITSU Integrated System PRIMEFLEX: Many benefits for service providers

There is of course an answer to this question. FUJITSU Integrated System PRIMEFLEX is a pre-defined, pre-integrated and pre-tested combination of data center components, such as servers, storage, network connectivity and software. All typical activities, such as designing the infrastructure, component integration and testing have been completed before starting the project, guaranteeing the compatibility of all components and minimizing the risk of unforeseen delays. This extremely shortens time to production and the time to bill your end customers. Being freed from all pre-deployment and upcoming lifecycle management activities enables you to focus on superior services which may help you differentiate in the market place and earn money from your end customers instead.



The subsequent figure demonstrates quite impressively the enormous savings in time that can be achieved by choosing the Integrated Systems approach instead of the DIY approach.



### DIY vs. Integrated System

With a DIY approach, all typical activities have to be done after the project has started. You have to design the infrastructure, integrate the individual components and test the integrated combination of selected components before the actual onsite deployment and integration into the production environment.

With an Integrated System, necessary things, such as infrastructure design, integration of components and testing, have been done before project start. The required activities after project start are confined to the deployment onsite and the integration into the production environment.

### Best-in-class components and long experience

PRIMEFLEX systems are built from best-in-class components, either Fujitsu's own technologies, such as FUJITSU Server PRIMERGY, FUJITSU Server PRIMEQUEST and FUJITSU Storage ETERNUS, or from leading technology partners. Fujitsu's proven servers and storage systems have a track record in performance benchmarks, up-time and energy efficiency, and are therefore a solid foundation for highest performance. Providing highest performance translates into a better utilization of your infrastructure, to serve more customers and make more revenue and margin.

Through the introduction of its first system in late 2002, Fujitsu excels by one of the longest track records in terms of integrated systems. Based on best practices, real-life project experience and proven operation in Fujitsu's cloud services, PRIMEFLEX systems are designed in a way that their components work optimally together. This has a positive impact on data center space, cabling, energy consumption, cooling efforts and resource utilization. Moreover, PRIMEFLEX is a perfect foundation for efficient operations, reduced maintenance efforts, and increased service levels. PRIMEFLEX will improve responsiveness to new requirements coming from your end customers, or even driving your business to a new level. This will satisfy your end customers, improve your reputation and attract new end customers. And finally, all these aspects help reduce capital and operational expenditures.

### Delivery models

PRIMEFLEX systems are either pre-installed in the factory, and arrive **ready-to-run** at your destination site; or they are delivered as **reference architectures** providing the flexibility to adapt them to your specific requirements. On demand, the adjusted configuration can be pre-installed and delivered ready-to-run, combining the advantages of both, reference architectures and ready-to-run systems. For PRIMEFLEX reference architectures, installation and configuration guidelines are available by default. It is worth mentioning that PRIMEFLEX reference architectures are treated like a product, i.e. they have a dedicated product manager, they have to undergo the same quality assurance processes as any server or storage system, and there is an overall solution support with a single point of contact. PRIMEFLEX reference architectures are certainly more than just the norm.

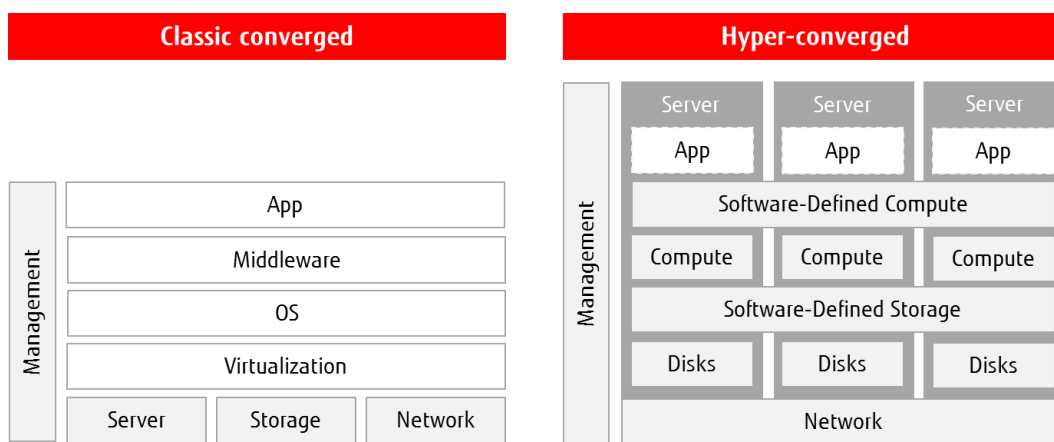
### Flexible services options

PRIMEFLEX is not just about hardware and software bundles; it is supplemented by flexible services options throughout all lifecycle phases. These services options encompass consulting, design onsite deployment, integration of the new infrastructure into the existing environment and migration services, lifecycle management and overall solution support with a single point of contact. If you don't have the resources to operate the data center infrastructure, Fujitsu will provide Managed Data Center and Hosting Services for you. It is all about choice. Just select what you need.

In addition, attractive financing options allow you to shift CAPEX to OPEX which means you may bill your customers first before we bill you.

### PRIMEFLEX includes hyper-converged systems

The era of integrated systems started with the classical converged systems approach, with servers, storage, network connectivity and software being pre-integrated to accelerate deployment, minimize compatibility issues and simplify management. Hyper-converged systems go a step further, as they tightly integrate all resources in a commodity server node, making a dedicated physical Storage Area Network (SAN) with its management superfluous. Instead, data storage spreads across the local disks of the server nodes.



Having compute and storage resources integrated in a single box makes deployment even easier and faster. Data center footprint will be reduced just as energy consumption and cooling requirements. The built-in data services (storage functions), such as data replication, snapshots, deduplication and data tiering turn hyper-converged systems into a software-defined storage platform. High availability is built-in. If a server fails, another one will take over the workload of the failed one, and data will not be lost due to data replication to multiple nodes. Hyper-converged infrastructures may be even spread across sites enabling a cost-effective disaster recovery option. The unified management for both compute and storage resources brings simplification to a new level. It reduces administration efforts and skills demands.

While classical converged systems scale on a component level, hyper-converged systems enable scalability on a system level. Compute performance and storage capacity can be scaled out by just adding or removing servers. New servers will be automatically discovered and can be easily integrated, because everything is software-defined. Hence, hyper-converged systems can be easily aligned to growing business demands, while business continuity is always ensured. All benefits aforementioned have a positive impact on costs and your margin.

### Classical or hyper-converged – Use cases matter

The value proposition of a hyper-converged infrastructure may easily lead to the conclusion that hyper-converged is the silver bullet of IT. But be careful: Whether it is the ideal architecture approach, strongly depends on the use case. It is true that hyper-converged is the right answer to an ever increasing number of use cases, but there are still scenarios where the classic architecture is better suited or sometimes even the only option. In general when it comes to the question "classic converged or hyper-converged", there will not be a black and white answer. However, there are a number of aspects that should be considered before taking the final decision.

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 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. If your workloads scale vertically, or they require a granular expansion on the component level, hyper-converged might be less appropriate. An example is monolithic applications, which cope with increasing data volumes using a scale-up approach.

For hyper-converged infrastructures, virtualization is a prerequisite. Therefore they cannot be used for workloads which run on bare-metal only, maybe because a virtual environment would be ineffective and slow them down. And as most implementations of hyper-converged are based on a single hypervisor, they won't fit if a mixed operation of multiple hypervisors is needed to run different workloads.

Hyper-converged has become an attractive option for remote offices and branch offices. As no external storage infrastructure needs to be maintained, frequent costly onsite visits can be avoided. There are customer cases where travel time could be reduced by 99%, just by replacing a physical SAN by a hyper-converged infrastructure.

If your workloads benefit from the data services coming with hyper-converged, you may use these services without any additional investment. If you don't need them, you will pay indirectly for things you don't use. Another aspect worth considering is the expected growth. The more frequently you have to expand your infrastructure, the more you will benefit from the ease of scalability that features hyper-converged.

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. You may expect resistance from your IT staff, especially in the storage area. Will you counter this resistance? This aspect is also closely related to your storage strategy. If you intend to utilize existing storage, hyper-converged will hardly fit to your strategy. If in contrast you intend to replace your existing storage sooner or later, going for hyper-converged may be a good start.

The storage capacity of a hyper-converged infrastructure is limited by the number of server nodes. If you have to cope with amounts of data which are larger than the maximum storage capacity of your server cluster, hyper-converged will be no option. Though hyper-converged promises linear scalability, predictable network performance with larger deployments is sometime questioned, mainly caused by a lack of experience.

Beside the technical appropriateness of workloads, it is also software licensing aspects which should be taken into account. For instance, a database application may be a perfect fit for hyper-converged, but if you have to pay license fees per CPU socket or even per CPU core, hyper-converged will quite likely be a no-go for commercial reasons.

At the end, it will be all about cost. As mentioned before, operational expenditure always tends to be much lower with hyper-converged infrastructures compared to classical converged ones. But when it comes to capital expenditure, it is hard to make a general statement. Typically, from a hardware cost perspective, hyper-converged is certainly more attractive than classic, from a software cost perspective it is just vice versa. But hyper-converged requires a minimum number of server nodes; it requires special, certified hardware components, and license fees need to be paid for the virtualization software either. You will find lots of examples with cost advantages on either side. Make a simple cost comparison for your concrete project and you will find out which is the more cost-effective option for you.

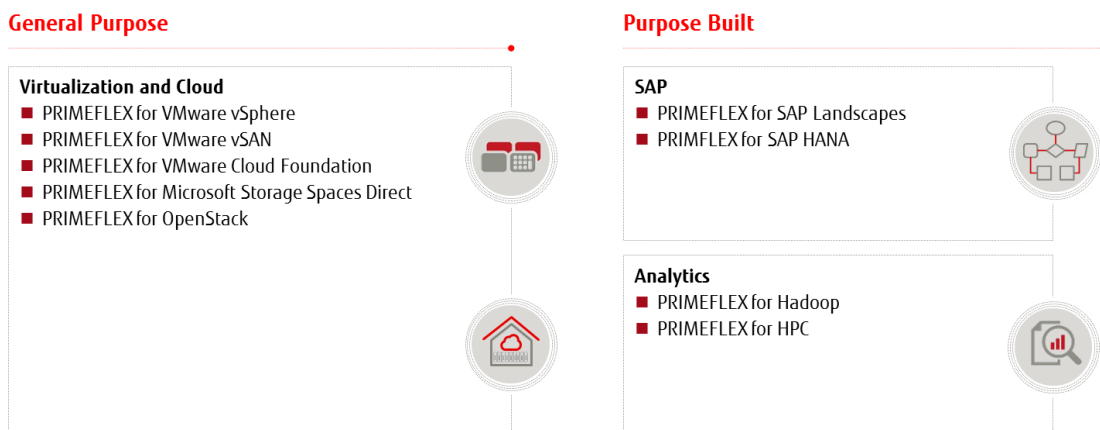
All told: When it comes to the question "classical converged or hyper-converged", the use case matters. There are good reasons to look at both architectural approaches. It is recommendable to take the decision specifically per each use case, and go for hyper-converged systems, if their benefits outweigh the drawbacks. As the PRIMEFLEX family includes both, classic converged and hyper-converged infrastructure systems, you can be sure that Fujitsu will neutrally consult you and figure out jointly with you what the ideal architecture for your use case is. No matter how the result of this joint exercise will look like, we will always have the appropriate integrated system for you, which will help you deploy your infrastructure fast and without any risk.

### PRIMEFLEX for many important use cases

PRIMEFLEX is focused on business-critical areas being of high relevance for service providers too, such as Virtualization, Private Cloud, Big Data and Analytics, High-Performance Computing, as well as SAP environments. The portfolio elements allocated for virtualization and Private Cloud are built for general purposes, while those allocated to Big Data and Analytics, High-Performance Computing, and SAP environments are purpose-built systems, optimized for specific use cases.

Evidenced by an ever increasing number of reference customers who have achieved tangible business benefits, PRIMEFLEX has been implemented across different industry segments and geographies. It is beyond question that numerous service providers have decided to go for PRIMEFLEX.

The subsequent figure shows the PRIMEFLEX line-up of systems and their mapping to the topic areas mentioned above. Some portfolio elements can even be assigned to multiple topic areas.



Let us now have a look at the individual use case areas and their offerings in more detail.

### **PRIMEFLEX® for SAP HANA®**

Today's businesses require speed. However, due to ever increasing data volumes, and the long lasting access times resulting from it, a real-time insight into the business is difficult or even impossible. This means, that existing data cannot be turned into valuable business information.

For this reason, SAP introduced their In-Memory Database platform HANA. All data is stored and processed in the main memory of one or several servers, can therefore be accessed and processed rapidly, and results are instantly available in real-time. Powerful data compression algorithms enable to store huge amounts of data memory-resident already today.

While in the past, HANA was just an option for SAP workloads beside other databases, all today's and future developments of SAP software will be based on HANA. This means that sooner or later a customer who once decided for SAP software will need HANA, if he wants to take advantage from SAP's innovations.

PRIMEFLEX for SAP HANA enables service providers to simplify the deployment and the operation of the SAP HANA platform, whilst maximizing the potential to accelerate and renew your business processes.

PRIMEFLEX for SAP HANA is available as single node and multi-node configurations in various sizes. Hence the requirements of any HANA use case can be met. Due to our flexible reference architecture approach, an easy adjustment to your specific requirements is possible. To protect against data loss, a persistence layer is required which always contains a data copy of the memory-resident data. With multi-node configurations, an external storage system, such as FUJITSU Storage ETERNUS DX, FUJITSU Storage ETERNUS AF or NetApp storage is used as persistence layer. With single node configurations alternatively data may be copied to the local disks of that node.

As an option for mission-critical demands, SAP-certified disaster recovery options are offered, too.

PRIMEFLEX for SAP HANA is not only an important offering for SAP environments; at the same time it may be an important building block of a Big Data infrastructure.

Fujitsu is known for its technical expertise and its deep understanding of the HANA architecture. Moreover, Fujitsu was the first vendor whose multi-node configuration was certified by SAP.

### **PRIMEFLEX® for SAP® Landscapes**

PRIMEFLEX for SAP Landscapes represents an infrastructure for SAP workloads with an extraordinary operational concept to optimize entire SAP environments with or without SAP HANA. For a pool of virtual and / or physical servers, just a single shared operating systems image is needed, which enables you to run any application on any server. This gives you the flexibility to move workloads across servers to achieve the highest performance, even between physical and virtual servers.

There is no single point of failure; this ensures business continuity and dependable mission-critical computing. As the solution is application-aware, workloads concerned by the failure of a server are restarted automatically on any of the other servers of the pool which are at disposal. Using just a few fail-over systems for many productive servers optimizes the utilization of your resources. For mission-critical demands, proven disaster recovery concepts may be applied.

The high degree of automation eases administration and maintenance tasks, and enables better service levels. There are service providers that could keep the number of administrators on a constant level over years, though the number of their end customers was steadily increasing. Applications are provisioned within minutes, enabling you to serve your customers much faster. Having only a single operating systems image makes updates extremely easy. The shared infrastructure can be easily and rapidly expanded to any arbitrary number of validated servers, whenever required by your business. This allows a flexible adaptation to the demands of a growing customer base. There are customer cases with up to 10,000s of SAP users.

The features mentioned before have a positive impact on CAPEX and OPEX. There are customers who even speak of a cost reduction by 90%, though this might be rather the exception than the rule. The cost savings dramatically increase your margins, and you may use the freed resources for generating additional revenue by high value services. The multi-tenancy feature of the infrastructure helps protect your investment.

As the first integrated system ever, PRIMEFLEX for SAP Landscapes was shipped back in 2002, though marketed at that time under a different name. Since its introduction it has evolved as a trusted solution for many SAP service providers. Fujitsu is committed to a future-proof development referring to regular innovations since its introduction into the market.



### PRIMEFLEX for Hadoop

Big Data is about large volumes of data which are extracted and collected from versatile data sources. The data may be of various types (often unstructured), is often generated at a high velocity and often needs to be processed at a high velocity, too. If you correlate this data in an appropriate way, you may create a huge value for the business or the society.

The Hadoop eco-system, a set of open source software products plays an important part in the Big Data game. Meanwhile there are modules in place for basically all disciplines; be it the extraction and collection of large data volumes in predictable time slots; be it processing of continuously generated event streams, for instance by sensors or click streams; be it NoSQL databases which are especially designed for being distributed and scale-out; or be it even in-memory technologies.

PRIMEFLEX for Hadoop is an integrated system simplifying the introduction of Big Data infrastructures for your customers. There are configurations in various sizes for storage-intensive tasks as well as for compute-intensive tasks. Due to our partnerships with the three major Hadoop distributions which are Cloudera®, Hortonworks® and MapR®, we give you and your customers the choice to select your preferred one. The optionally included software from Datameer® enables even business users without deeper IT skills to get in touch with Big Data. It is extremely easy to collect data, analyze it and visualize the results. PRIMEFLEX for Hadoop excels by its easy and basically unlimited scalability which is needed to cope with the steadily increasing data volumes.

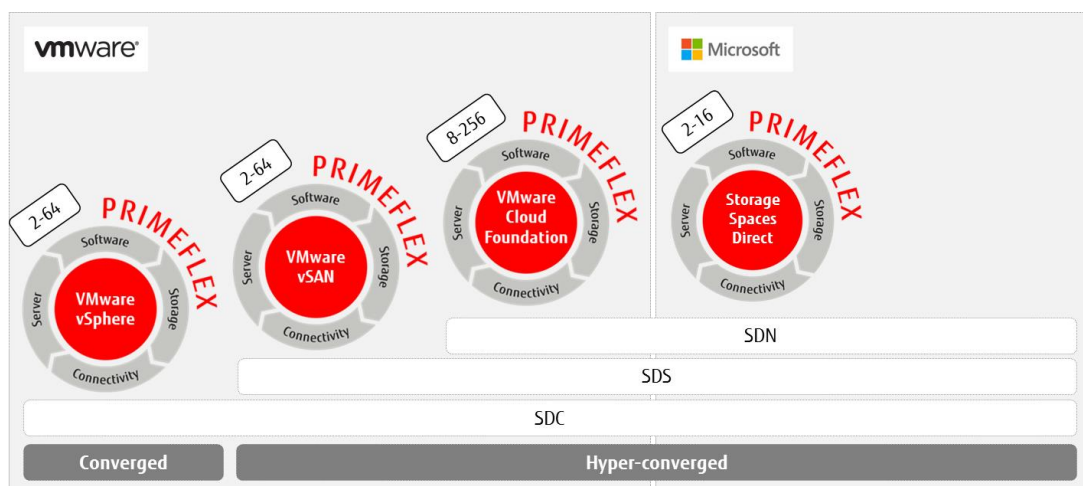
Bridging between traditional Data Warehouses and Data Lakes makes it easier for enterprises to extract valuable business intelligence from multiple sources of data. Therefore, Fujitsu extended its PRIMEFLEX for Hadoop to special configurations especially optimized for SAP Vora®, an in-memory query engine that plugs into the Apache Spark execution framework.

PRIMEFLEX for Hadoop is provided as a ready-to-run Integrated System as well as a Reference Architecture for flexible deployment. In addition, strategic big data consulting, analytics consulting, consulting for Hadoop, and integration and maintenance services supplement the offering.

### PRIMEFLEX for Virtualization and Private Cloud

The virtualization of computing resources has greatly improved data center operations over the past decade. To provide the operational efficiency and agility necessary to cope with the ever increasing demand for more responsive and cost-efficient IT support, service providers are now looking to extend virtualization to storage and network resources by adopting a strategy for establishing the Software Defined Data Center (SDDC).

The subsequent figure shows the PRIMEFLEX offerings for virtualization and their positioning. Beside the separation between classical converged and hyper-converged infrastructure systems it is pointed out which systems included in addition to Software-Defined Compute (SDC) and Software-Defined Storage (SDS) also Software-Defined Networking (SDN), either by default or as an option. Moreover, it is demonstrated which virtualization technology is used (VMware or Microsoft). The numbers with each system represent the number of possible server nodes, giving an impression of scalability.



Virtual infrastructures in turn represent a sound foundation of a cloud infrastructure. Just add a cloud management platform to enable advanced automation, self-service provisioning, accounting and billing. Implementing a private cloud infrastructure is a key stepping stone for service providers intending to gradually modernize their traditional IT infrastructure in order to offer better service levels. Key benefits are improved operational agility and efficiency enabling faster time to production and bill for new business services while reducing operating efforts and costs for increased margins.



No matter what virtualization or private cloud approach you would like to introduce in your data center, Fujitsu provides a broad range of integrated systems based on a converged or hyper-converged architecture, and giving you the choice in terms of virtualization software and configuration size.

Powered by PRIMERGY x86-servers, PRIMEFLEX benefits from all their proven values. Fujitsu x86-servers provide the best performing server platform for running VMware virtual machines. Their long and proven track record is reflected in top rankings across all VMware VMmark benchmark categories. Its patented Cool-Safe® technology drives world records for energy efficiency. Moreover, PRIMERGY is shown as a leading platform in global server reliability reports. All this translates into more revenue and margin opportunities for service providers.

#### **PRIMEFLEX for VMware vSAN**

PRIMEFLEX for VMware vSAN is a hyper-converged system based on the VMware software stack including VMware vSphere and vSAN. The system supports up to 64 server nodes and is available in various vSAN Ready Nodes configurations with pre-installed software for specific use cases, such as All-Flash configurations for write-intensive workloads demanding for low latency (e.g. virtual desktop infrastructure with either linked clones or full clones), hybrid configurations with hard disks and Solid State Disks for mixed workloads, and special high density configurations for use cases with data center footprint being the critical parameter. PRIMEFLEX for VMware vSAN is also a sound foundation for private cloud infrastructures and VDI. The reference architecture approach allows flexible adjustments of the pre-defined configurations with regard to processor type, main memory size and storage capacity.

#### **PRIMEFLEX for VMware Cloud Foundation**

PRIMEFLEX for VMware Cloud Foundation represents a truly Software-Defined Data Center based on a hyper-converged architecture leveraging the VMware Cloud Foundation software stack. The system is available in various configurations from 8 to 192 servers, which are expandable up to 256 nodes. It is based on VMware vSphere, vSAN and NSX for virtualizing compute, storage and network resources. The additional SDDC Manager acts as an automation engine for provisioning, monitoring and lifecycle management, taking care of virtual and physical resources. VMware vRealize is an optional add-on that turns PRIMEFLEX for VMware Cloud Foundation into a full private cloud infrastructure. A further option is VMware Horizon which makes the platform usable as a virtual desktop infrastructure. As the Integrated System is delivered ready-to-run, it will also be ready-to-operate within a couple of hours after arrival. PRIMEFLEX for VMware Cloud Foundation is the ideal general purpose platform for service providers who want to serve multiple customers by a shared infrastructure. It enables service providers to respond quickly as their customers' needs evolve. This ability is a key criterion for competitive success.

#### **PRIMEFLEX for Microsoft Storage Spaces Direct**

PRIMEFLEX for Storage Spaces Direct is a hyper-converged system based on software-defined storage technology (Storage Spaces Direct) integrated in Microsoft's Windows Server 2016 Datacenter Edition. The system supports up to 16 server nodes. Various configurations for a broad range of use cases are in place covering mixed workloads as well as workloads requiring extreme I/O performance. Configuration options include hard disks, Solid State Disks and high-speed NVMe disks that allow for setting up a 2-tier and 3-tier storage infrastructure. The reference architecture approach allows flexible adjustments of the pre-defined configurations.

#### **PRIMEFLEX for VMware vSphere**

PRIMEFLEX for VMware vSphere is a converged system based on a classical architecture with external storage and virtualization technology from VMware. You may choose between Fujitsu's ETERNUS DX (hybrid storage) and ETERNUS AF (All-Flash), as well as between iSCSI and Fiber Channel connectivity. Network switches (from Extreme Networks and Brocade), cabling and rack infrastructure are also included. There are multiple configurations varying in size.

### **PRIMEFLEX for OpenStack**

PRIMEFLEX for OpenStack provides an OpenStack cloud infrastructure based on either the Red Hat or SUSE OpenStack platform. It includes a range of validated configurations taking into account different customer engagement levels and varying use cases. No matter where you like to start off, whether you just want to make your first steps into OpenStack, need to evaluate application deployment and backup concepts, different storage, high-availability or enhanced software-defined networking options – we provide a configuration that fits to your needs. As core components, PRIMEFLEX for OpenStack includes FUJITSU Server PRIMERGY, FUJITSU Storage ETERNUS or NetApp storage and networking switches from Extreme Networks or Cisco. In addition, Fujitsu delivers a broad choice of validated extensions covering software-defined storage, software-defined networking, advanced monitoring, cost management, an enterprise portal and an application delivery and migration solution. The reference architecture approach allows flexible adjustments of the pre-defined configurations.

### **NFLEX**

Fujitsu's Integrated Systems portfolio is completed by NFLEX, a factory-integrated, ready-to-run converged infrastructure system based on virtualization technology from VMware and NetApp storage (all-flash or hybrid), jointly developed and marketed with Fujitsu's alliance partner NetApp.

### **FUJITSU Enterprise Service Catalog Manager turns your virtual infrastructure into a cloud infrastructure**

As said, PRIMEFLEX for OpenStack is a private cloud infrastructure, while PRIMEFLEX for VMware Cloud Foundation is optionally delivered with a complete cloud management stack. But what about the other PRIMEFLEX offerings for virtualization, and what about NFLEX? Is there a chance to make them cloud capable, too?

Yes, there is a chance thanks to FUJITSU Enterprise Service Catalog Manager which closes the gap between virtualization and cloud. Service Catalog Manager includes a self-service portal for multi-tenants which may be branded by your corporate design, simplifying service consumption for your end customers. Moreover, it covers automated provisioning and flexible billing, reducing administration efforts and cost, while keeping transparency and control. Finally, it enables you to display your own marketplace, opening up new revenue streams.

Service Catalog Manager can even be used in a hybrid scenario, where you deliver services to your end customers from your cloud and one or multiple 3<sup>rd</sup> party clouds, thus increasing flexibility and speed of service delivery, while keeping governance across all cloud services.

### **Summary**

FUJITSU Integrated System PRIMEFLEX relieves you from all the pain caused by building data center infrastructures on your own. They help reduce complexity, time and risk, while increasing operational efficiency. This is exactly what service providers need to be fast, reliable and attractive for their customers. And this in turn is the prerequisite to be successful in a tough and competitive market. So, if you are tired of the DIY model and want to choose the simple fast track to your data center infrastructure, then choose PRIMEFLEX from Fujitsu as many other service providers have done before. With Fujitsu you have a partner who will help make you more profitable and de-risk your growth.

---

#### **Contact**

FUJITSU Technology Solutions GmbH  
Address: Mies-van-der-Rohe-Strasse 8,  
80807 Munich, Germany  
Website: [www.fujitsu.com/primeflex](http://www.fujitsu.com/primeflex)  
2018-08-12 WW EN

© 2018 Fujitsu, the Fujitsu logo, and other Fujitsu trademarks are trademarks or registered trademarks of Fujitsu Limited in Japan and other countries. PRIMEFLEX is a registered trademark in Europe and other countries. Other company, product and service names may be trademarks or registered trademarks of their respective owners. Technical data subject to modification and delivery subject to availability. Any liability that the data and illustrations are complete, actual or correct is excluded. Designations may be trademarks and/or copyrights of the respective manufacturer, the use of which by third parties for their own purposes may infringe the rights of such owner.