White Paper

We make sure

FUJITSU COMPUTERS SIEMENS

Server Virtualization at Fujitsu Siemens Computers

Issue August 2008

Pages 20

Contents

Abstract

This Whitepaper examines the different virtualization technologies and explains to which IT infrastructure layers they apply. It also shows how Fujitsu Siemens Computers incorporates these technologies into end-to-end solutions as an essential part of its Dynamic Data Center™ strategy.

What is virtualization?

Broadly speaking, virtualization is about breaking the dependencies between the individual layers. Introducing a virtualization layer between two layers leads to an abstraction from the complexity of the layer underneath. The virtualization layer hides the complexity of the layer underneath, mediates access to it, creates virtual resources and maps them to physical resources.

This definition can be applied to all the levels where virtualization is carried out.

Virtualization – a technology tried-and-tested over decades

Virtualization was introduced more than 30 years ago on mainframe systems. Basically, two objectives were pursued with it:

1. Server hardware was extremely expensive at the time and so one aim was to use it optimally. Server virtualization enabled the various OS instances to be grouped on one physical hardware system. That significantly improved flexibility and capacity utilization.
2. The applications were to be made independent of the underlying hardware layers. Virtualization thus became a basic technology that enabled raid further development of new technologies, without any impact on the existing software infrastructure.

At the time, a technology was developed that corresponds very closely to the current approach of server virtualization with hypervisor technology.

Server virtualization – a mature mainframe technology: Fujitsu Siemens Computers BS2000 OSD / VM2000 mainframes

Fujitsu Siemens Computers was one of the pioneers, and has correspondingly decades of experience in providing and managing virtual machine infrastructures. VM2000 for BS2000 mainframes is the perfect example of continuity and long-term investment protection in mission-critical data center environments. Resource allocation happens dynamically on a very granular level. Energy consumption, space and administration efforts are considerably reduced. High availability is achieved in a cost-effective manner, and services can be made ready very fast. Except to this, the virtualization concept enables parallel operation of production systems, test and development as well as other pre-production systems on the same hardware.

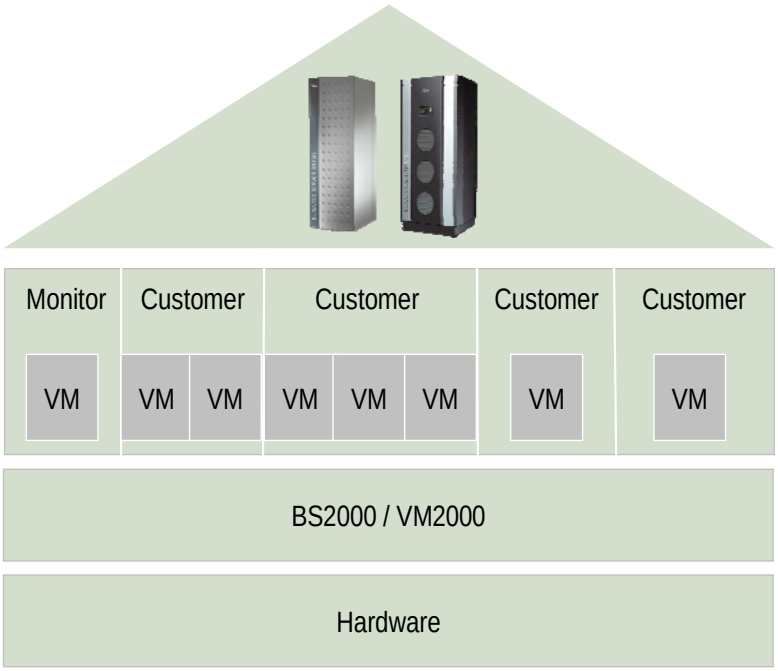

The reasons why enterprises set store by virtualization, the virtualization approaches and the associated levels at which virtualization is carried out are more diverse today. The graphic below shows the levels at which virtualization is now deployed and the solutions that are frequently used for it.

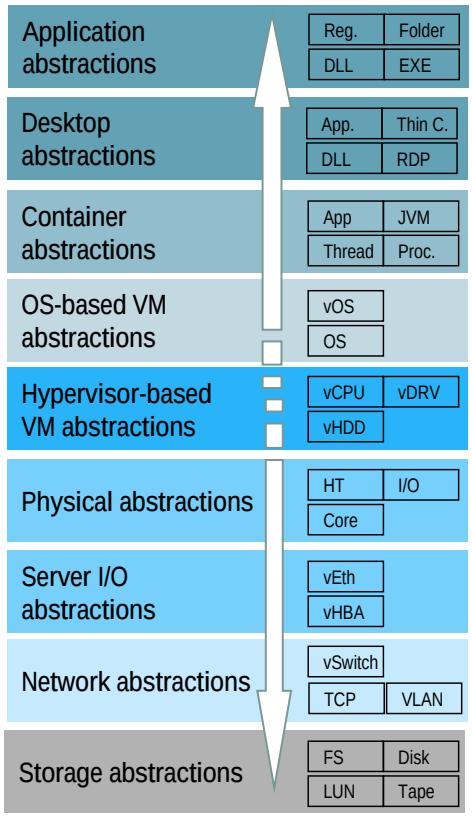

Microsoft SoftGrid, Altiris Software Virtualization Solution, Ardence, SAP, Oracle…

VMware VDI, CITRIX XenDesktop Server, …

Solaris Zones, SWSoft Virtuozzo/OpenVZ, Java EE…

VMware (GSX) Server, Microsoft Virtual Server, KVM…

VMware ESX Server, Xen, Microsoft “Viridian”, PAN Mgr.

Intel Virtualization Technology (VT), AMD Virtualization (AMD-V)…

FSC SV-VIOM, Egenera PAN Mgr.

Cisco, Brocade, Egenera PAN Mgr…

FSC CentricStor, EMC, NetApp, Veritas, Brocade, Microsoft…

The wide range of different virtualization solutions means that one of the key tasks of any IT manager is to identify the right technology for the specific requirements and integrate it in the IT infrastructure. However, the graphic above also shows clearly that there is no one virtualization solution that meets all the requirements for a state-of-the-art Dynamic Data Center. Instead, a combination of the various technologies is needed to create an ideal customer-specific solution.

This White Paper looks at the most important technologies, shows their advantages and describes their connection with other layers and solutions. It also presents which technologies Fujitsu Siemens Computers uses or has developed itself so as to offer its customers consistent, end-to-end virtualization solutions.

Server virtualization with x86-based industry-standard servers (ISSs)

No other server architecture has probably gained as much in importance over the past years than the x86-based IS server platform. The complete PRIMERGY product family from Fujitsu Siemens Computers embodies this form of architecture, which is essentially based on Intel and AMD processor technology. The rapid pace at which this technology has developed means that these servers now have a price/performance ratio that is unprecedented in IT. Other components cannot keep up with the speed at which the hardware is developing. That results in special requirements for this platform, which can be tackled by virtualization.

Traditional x86 platform approach: Challenges

Let us start with something everybody is familiar with: The architecture of a traditional industry standard server. It consists of three important layers. There is first of all the hardware layer, which is a set of different devices and computing resources. Then there is the operation system, which has the overall control over the hardware. It is responsible to manage the hardware resources in terms of sharing them among different applications, based on priorities. And finally there are the applications, which use the operating system as an interface in order to get access to the hardware.

In a traditional platform approach, the configuration of each layer is tightly coupled to the layer underneath. As a result, hardware and software lifecycles are strongly dependent from each other. A legacy application which runs on an operating system, which in turn does not run on new hardware, cannot be moved to such a new hardware. As hardware changes regularly, this is a serious issue.

Due to historical reasons the resource management in the operating system is not very efficient and quite often the applications interfere with each other regarding the joint usage of resources. Hence, most of the servers installed and sized for peak loads run only one application. This leads to under-utilized system resources and correspondingly poor sever utilization. In turn, the result of poor server utilization is an explosion of servers, a high energy consumption, much heat generated, a high amount of space to place all these servers, and of course high expenses in terms of maintenance and support. Likewise there is no dynamics when it comes to scaling, and it is extremely difficult to implement high availability.

The impact is management complexity and low efficiency. It takes considerable efforts to deploy new services, i.e. time to production and response to changing business needs is rather poor.

All things considered, the traditional platform approach implies a high TCO (Total Cost of Ownership).

Server virtualization with virtual machines for x86-based industry-standard servers (ISSs)

For the above reasons, server virtualization with virtual machines (also termed hypervisor technology) is now regarded as a highly promising approach in server virtualization. The technology illustrated in the graphic below is generally termed server virtualization.

By definition, server virtualization means masking of server resources, including the number and identity of individual physical servers, processors, and operating system, from server users. In particular, applications or services are separated from dedicated hardware resources. This finally enables more flexible control of both hardware and software resources, which can deliver value across a wide range of IT requirements.

According to various studies, virtualization is the pervasive and important trend in servers through 2012.

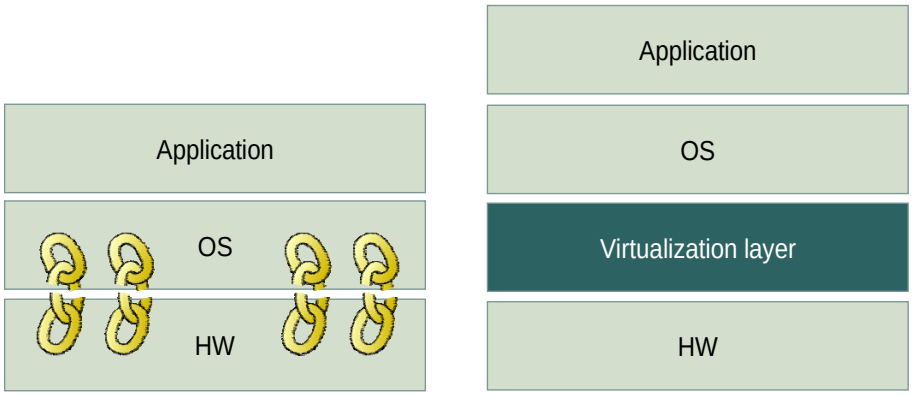

Traditional platform approach **Virtualization approach**

This hypervisor technology enables multiple instances of an operating system and the associated applications to be run on one physical hardware system. That means a large number of underutilized physical servers can be consolidated on a few, well-utilized servers. This creates independence between the software and hardware and so the foundation for new operating processes.

Whereas, until a few years ago, consolidation and use of older operating system versions on the latest hardware were the main motivation for deploying hypervisor technology, customers now also expect greater agility and high availability from virtualized infrastructures.

What customers expect from virtual machine concepts

When it comes to the question what customers expect from virtual machine concepts, we have to distinguish between large enterprises and midsize companies. For both of them, in an overall perspective the answers are the same, but their priorities are different.

The number 1 topic for large enterprises is consolidation which represents dramatic capital savings in terms of equipment. The number 2 topic is deployment, where virtualization reduces the administrative burden of loading and configuring the endless number of new servers arriving on the loading dock. Number 3 is agility, i.e. moving workloads from one server to another one according to new and changing demands, at reduced administrative burden. The number 4 topic “freedom of choice” allows large enterprises to more easily accommodate a mix of vendors and server models within the data center. And finally the number 5 topic is protection. Using virtualization disaster recovery is enormously simplified, but only for x86 platforms. Due to heterogeneity, most large enterprises have a more complex infrastructure to replicate in the case of a disaster. That’s why protection has lowest priority among the 5 expectations. Quite in general it may be stated that large enterprises tend to focus on the cost savings associated with virtualization.

Midsize businesses have the same expectation, just the significance is partly different. Their number 1 topic is consolidation, too. Of course, here the level of savings is lower, because of the scale of server deployments. Protection, which is number 5 with large enterprises, is number 2 with midsize businesses. They quite often couple virtualization with low-cost SAN (Storage Area Network) solutions, thus tremendously reducing cost and complexity of implementing disaster recovery. Their number 3 topics is deployment. Most midsize businesses have limited administrative resources, virtualization reduces their efforts and increases speed. Number 4 is freedom of choice. Virtualization allows purchases based on competitive pricing without worrying about the overhead of supporting multiple vendors. And last but not least, number 5 is agility, where virtualization helps adapt server resources according to changes in workload demands, although not at the same level as with large enterprises. But it makes it definitely easier to bring up new services in remote branch offices. All considered, midsize businesses are also driven by the promise of consolidation and cost savings, but they see virtualization as a chance to enable solutions which otherwise would be difficult and expensive

Virtual machines

Now we are going to have a closer look at the virtualization of server hardware as such. Here the virtual machine concept has become well accepted in the marketplace. By definition, a virtual machine is a software-driven implementation of a computer, that executes programs like a real machine, and which runs as an isolated environment on a physical computer. A virtual machine is also described as a guest, the physical computer as the host. As quite a series of virtual machines might run on one real server, this server will behave like many servers.

The virtualization layer, in conjunction with virtual machines also denoted as virtual machine monitor (VMM), tricks the guest system into seeing the same virtual HW (CPU, memory and I/O devices) regardless of the underlying physical equipment.

Key features

The virtual machine concept is characterized by 4 key features.

First of all virtual machines are fully isolated from each other. Isolation stands for high availability and security. If for instance any of the virtual machines fails, this will not have any impact on the other virtual machines running on the same physical hardware.

Secondly the whole server hardware is partitioned into several virtual machines, each of them running its own (guest) operating system. The guest operating systems may even be different. Hardware resources, such as CPU, main memory or I/O devices, are shared among the virtual machines.

Due to encapsulation, the status of virtual machines can be saved in files. Just by copying the files, the virtual machines can be simply transferred from one physical environment to another one.

As the VMM is situated between the software stack (containing the operating system OS and the applications) and the server hardware, software and hardware become completely independent from each other. This enables you to run legacy applications on latest hardware technology, although the operating system needed for the application might not support the new hardware, due to the lack of necessary drivers. This means a full investment protection for the legacy applications.

Basically you could even build a failover cluster from virtual machines, which would be much more affordable compared to a cluster with physical servers. Traditional cluster solutions are expensive, complex and difficult to manage. The fact that in a cluster just one server node is active and the second one is passive, does not attest to efficiency. In the event of a problem, rapid solutions are important. Therefore patches are quite often applied to the active system only, while the passive system is forgotten. Hence the active and the passive node diverge, and only after server failure it will turn out that the fail-over does not work. With server virtualization we add complexity in a certain sense, but the administration efforts will definitely be much lower.

Typical usage scenarios for Hypervisors today

The initial motivation behind virtual machine technology was the large number of only partially utilized servers. The importance of Intel/AMD-based servers in enterprise IT infrastructures has grown exponentially in recent years, and with the price / performance ratio of this server architecture being so good, many companies have witnessed an uncontrolled proliferation of these systems. So, workload aggregation, also known as consolidation, evolved as a must have topic.

Besides, operating systems for industry standard servers have a number of deficiencies that make effective coexistence of multiple applications on a single server difficult, if not impossible. Hence, workload isolation is another reason to introduce virtual machines. Without virtual machines, the need to isolate individual applications totally from each other would encourage the uncontrolled server sprawl even more.

Another motivator for virtual machines is workload migration. By workload migration, entire virtual machine environments can flexibly be moved from one physical server hardware to another one on the fly, without interrupting the service itself. E.g. in the event of planned downtime or to improve the load balance of the overall IT infrastructure.

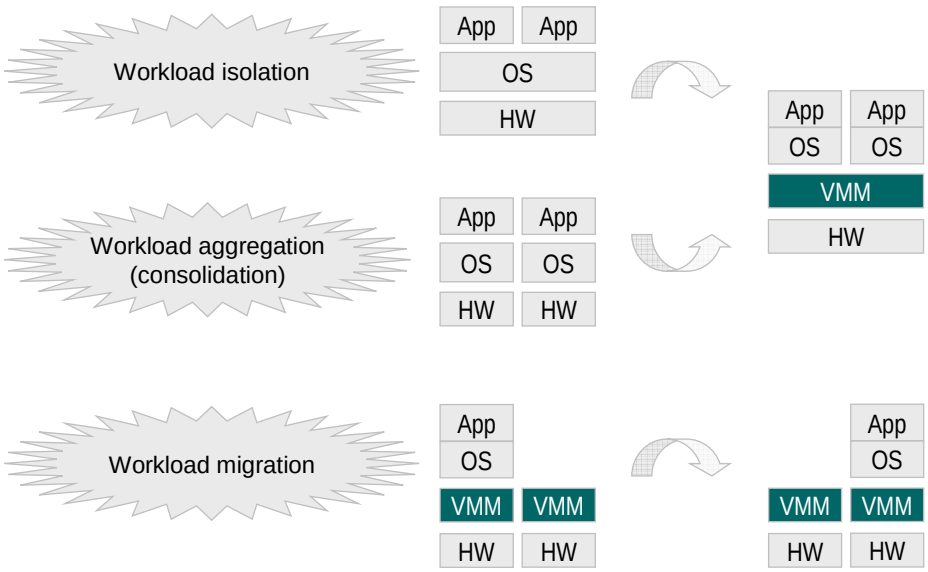

Customer example: Server virtualization at HypoVereinsbank

One memorable example of server virtualization by means of the virtual machine concept is HypoVereinsbank in Munich. They had 650 servers installed and intended to consolidate them in order to reduce floor space and energy cost, optimize resources and create flexibility. Today they run around 1,400 virtual machines on approximately 55 PRIMERGY blade servers with VMware ESX 3.5. As per statement of the customer, they save 2.0M kWh per year which is in accordance with €200,000. The virtualization concept enables them to use computing capacity on demand

Customer example: Server virtualization at Investment Bank in Finland

A Finnish Investment Bank runs self-developed applications in broker offices all over the world. In the broker offices you find a rather heterogeneous server landscape, because they procure hardware they expect to get the best support for in the respecting location. Delivering newly developed or updated applications to all the broker offices, testing them for different server hardware, and deploying them was a huge challenge for the IT department in the past.

Using virtualization technology, quality assurance and tests are done in headquarters, which helps them save the development and test systems in all the other locations. Deployment happens easily and rapidly, just by copying the two binary files of virtual machines set up in headquarters, and moving them to the target systems anywhere in the world, independent from the underlying hardware resources. Thus new services can be provided within minutes. In the event of a disaster, a fast recovery can be achieved with the aid of system backup copies. And everything works with hardly remarkable administration efforts.

Virtual machines in a bare-metal environment

There are several flavors of the virtual machine concept, but certainly the most relevant one is the hypervisor concept, where the VMM (here also denoted as hypervisor) runs in a bare-metal environment. As it is situated directly on the hardware, it has exclusive control over the hardware, meaning that it assumes the role of a real operating system. Therefore, typically the hypervisor includes a fine-grained resource management of all major system resources.

In conjunction with the hypervisor concept, people distinguish between full virtualization and para-virtualization. With para-virtualization, the kernel of the guest operating system is modified for the purpose of improved performance. On the other hand, full virtualization leaves the guest operating system unmodified.

VMware ESX 3.5, Microsoft Hyper-V, Citrix XenServer, Oracle VM, as well as the Xen versions from Novell and RedHat are the most important examples for the hypervisor concept.

Overview of the most important hypervisor products

General characteristics

For a long time, there were detailed debates on what hypervisor architecture is the best and creates a lower overhead and what virtualization method is the most efficient: “Para-virtualization”, “full virtualization”, “trap and emulate” and “binary translation” are some buzzwords in this connection. This debate has been pushed into the background as virtualization has become increasing supported by standard technologies in processors. Basically, all hypervisors now use these interfaces and so achieve comparable performance data. Para-virtualization, for instance, is now only used in exceptional cases, such as where the hypervisor vendor also has unrestricted access to an operating system of its own. Binary translation from VMware is also still used under certain general conditions.

Apart from this standardization of the virtualization method, all vendors offer functionalities that result in greater high availability and agility and above all in efficient management of the virtualized environments. Every well-known hypervisor vendor now offers a special management suite for virtualization, with the result that the hypervisor itself is becoming more and more of a commodity component. The offerings differ in terms of management and the end-to-end solutions. Consequently, this should be a prime criterion in choosing a suitable hypervisor. In comparing prices of the different products, the product structure and associated functional range should be analyzed closely so as to obtain truly comparable results.

It should also be noted that major ISVs offer their products as appliances, i.e. the application is not shipped in a conventional installation format, but as part of a virtual machine. Consequently, the application is not dependent on the type of server it runs on. As a result, the time and effort ISVs have to spend on Q&As is reduced; however, that may result in end customers using more than one hypervisor. To minimize the complexity this involves, vendors are working on standardized VM formats that will enable the virtual machine of a specific hypervisor to run on another hypervisor at no cost or effort. This feature is also an important criterion in choosing the suitable hypervisor.

VMware VI 3 (Virtual Infrastructure 3)

VMware has been an influential force in shaping server virtualization over the past six years. VMware’s virtualization technology began with PCs – for the first time, multiple operating systems were run on a single PC so as to obtain more flexible test environments. However, this method was very soon applied to servers. As a result, the focus shifted to server consolidation and the possibility of running aged Windows NT4 environments on the latest hardware. With its VirtualCenter, VMware also set new standards for management of virtualized environments. High availability (HA), load balancing within ESX Server farms (DRS) and above all the component VMotion, with which VMs can be moved during operation from one physical instance of the ESX farm to another, were pioneering developments that also had a lasting influence on the planning of other vendors.

The current VMware architecture looks like this:

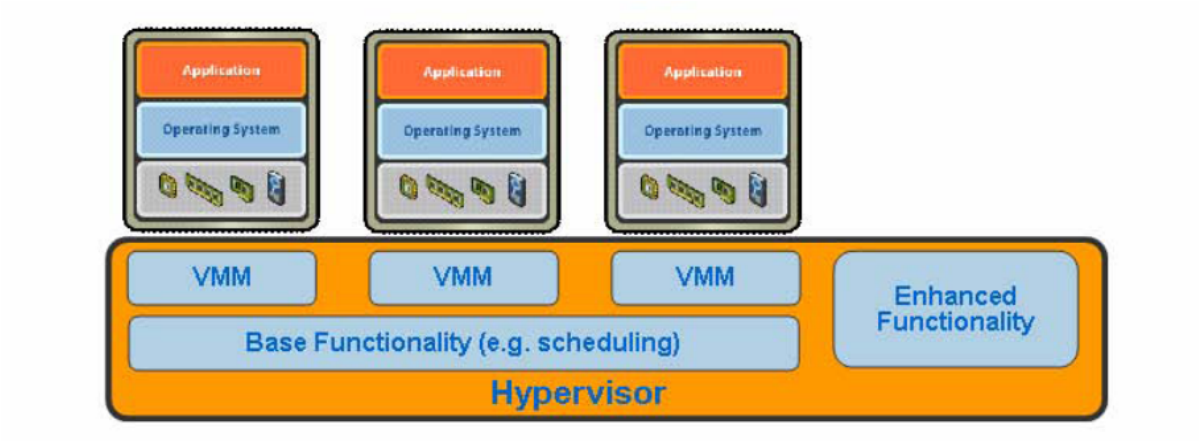

The VMware ESX hypervisor runs directly on the hardware and acts as an independent operating system. This kernel has been developed by VMware itself and impresses with its lean design, requiring just 32 MB. The kernel takes care of all of the following:

- Management of the VMs
- Separation of the VMs from each other
- Efficient distribution of resources between the VMs in line with the administrator's specifications
- Mapping of the machine instructions issued in a VM
- Mapping of the I/O jobs from the VMs.

In particular, it is the last feature that distinguishes VMware ESX from other hypervisor architectures. VMware does not need the support of another operating system instance to handle its tasks. That means there are no dependencies on or restrictions as regards such instances. However, interaction with another operating system instance (see the section on XEN and Hyper-V) also has advantages, as can be seen with regard to I/O handling: the VMware hypervisor is the connecting link between the VMs and the underlying hardware. The VMs "see" only the formally supported hardware (SCSI disks, NICs, processors and storage); however, this does not always correspond to the actually existing hardware the hypervisor has to serve. So that the existing hardware can be addressed from the hypervisor, there are drivers for most of the common hardware components; however, caution is called for if special hardware components are needed for an application (e.g. a security dongle, ISDN cards, special graphics support, etc.). It is then necessary to check on a case-by-case basis whether VMware ensures support for the drivers. These special hardware components must always be certified with VMware by the respective OEMs for the servers in question. If they are not in the VMware HCL (Hardware Compatibility List), no support can be expected from VMware in the event of an error. Details can be found in the list, which can be accessed at:

http://www.vmware.com/pdf/vi35_systems_guide.pdf

As mentioned at the outset, management of the virtualized infrastructures is a core component of every hypervisor scenario. The figure below shows the core functions of VMware VirtualCenter:

VMware VirtualCenter

- Centralized management interface
- VM provisioning with templates
- Performance monitoring
- Secure access control
- SDK for automation

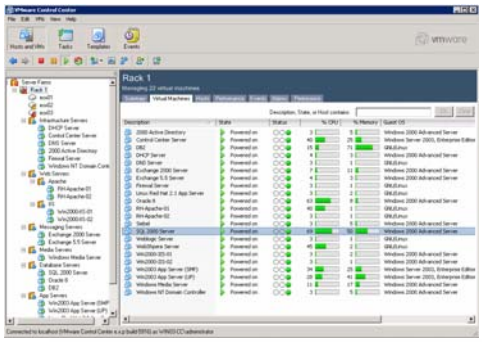

VMotion™ Technology

- Dynamically move VMs
- No downtime or service interruption
- Zero- downtime hardware maintenance

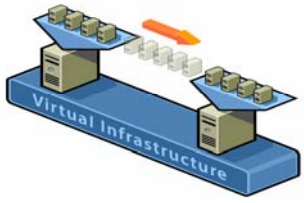

The focus of VirtualCenter from VMware is on managing the virtual instances (VMs). Real server instances play virtually no role any more. That is in line with VMware's philosophy. The real instance, the VMware server, is understood as part of a server farm. The resources the server contributes to the farm are administered comprehensively by the relevant management layers and divided among the VMs as specified. The individual server no longer plays a role in this distribution. VirtualCenter allows a new server to be included in a server farm at very little effort. It also enables servers that have failed as a result of a hardware fault to be replaced very easily. However, VMware VirtualCenter does not take care of providing a physical substitute server. This requires additional automated processes or products, such as x10sure from Fujitsu Siemens Computers. (See the separate section in this White Paper.)

Virtualization with XEN

In the case of XEN, it is certainly worth taking a closer look at how and by whom this technology is currently offered on the market. Basically, XEN is the Linux community's contribution to the subject of virtualization. The basis for this product was developed by the University of Cambridge and made available in accordance with the familiar rules of the open source community. XenSource, a company founded by developers from the University of Cambridge, then took charge of developing this hypervisor further. XenSource not only provides the market with basic developments of the open source community, but also a XEN distribution of its own. Compared with the open source product, this distribution has a far greater functional range and, of course, a fee must be paid for it. All major Linux vendors (and also since recently ISVs like Oracle) develop their own derivatives on the basis of open source XEN and market these using different license models. So that a sensible comparison can be made, the functional range of the individual distribution must be examined precisely. The basic architecture of XEN must be supported by a Linux operating system. However, the various distributions differ in this aspect. The figure below shows the basic mode of operation of a XEN hypervisor:

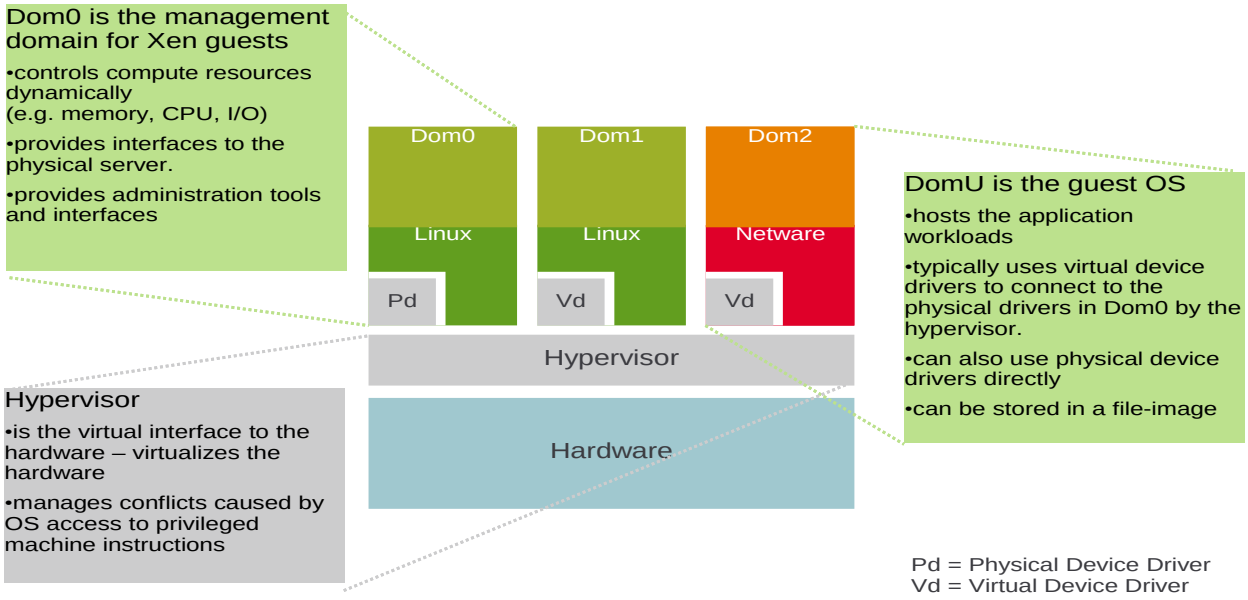

As in the case of VMware, the actual XEN hypervisor runs directly on the hardware. Unlike VMware, however, the important functions are not handled in the hypervisor, but in what is termed a "domain 0." This is a privileged virtual machine that runs the standard operating system of the respective Linux vendor. Linux Enterprise Server, for example, runs in the case of SuSe/Novell, and RedHat Enterprise Linux in the case of RedHat. This operating system controls the virtual machines (called in this case domain 1, … domain_n). The domain 0 is especially important when it comes to I/O handling. Every I/O issued from a virtual machine to a virtual hardware is redirected internally to the domain 0. This I/O is then processed there on the real hardware. Original drivers from the respective Linux distribution are used in this. That means the costs for an OEM of certifying such hypervisors are low. XEN supports para-virtualization and full virtualization.

Management is also addressed by the various vendors – and the individual distributions have major differences here.

XEN from SuSe/Novell is very deeply integrated in the platform management product ZENWorks. SuSe/Novell also always refers to its expanded technology partnership with Microsoft. Its aim is to ensure interoperability in virtualization and related management.

RedHat does not have a visible management strategy of its own at present and usually refers to open source management products. There is a strong trend at RedHat to make KVM technology, which is actually part of every standard Linux operating system, the focus of its virtualization strategy. The approach uses kernel-based VMs. In this, the VM is mapped in a normal Linux process that is managed by means of customary Linux mechanisms. A special KVM driver is a firm part of the Linux kernel and, among other things, controls I/O communication between the VMs and the real hardware.

The situation at XenSource is particularly interesting. XenSource was purchased by Citrix at the beginning of 2008. As part of this, Citrix also took over the management product XENCenter along with the actual XEN hypervisor. The functionality of this management product has always set the standard in the XEN arena. However, Citrix has developed a new strategy around this XEN product. XenServer is the elementary component of this strategy. The focus here is without a doubt Citrix' solution for desktop virtualization. Deep integration of the new paradigm in existing Citrix Presentation Server infrastructures (now called XenApp Server) is apparent here. As a result, Citrix is expanding its SBC (server-based computing) portfolio to include dedicated virtualized desktops that, however, can leverage the advantages of the tried-and-tested ICA protocol and all the access paths and management tools of a conventional existing CITRIX infrastructure without restriction. Alongside this, Citrix is positioning its virtualization product XenServer as the ideal solution for virtualizing traditional Citrix Presentation Servers and Microsoft Terminal Servers. This positioning will offer a simple means in future of scaling these servers better to state-of-the-art hardware. A single Presentation/Terminal Server instance will only be able to utilize the capacity of such new servers to a limited extent, due to the restrictions in the Windows kernel. However, multiple instances of such a server can run on one hardware system with the aid of virtualization, resulting in a sharp increase in the number of concurrent users supported. This optimization was possible only in close cooperation with Microsoft, a further major vendor of virtualization technology.

Hyper-V from Microsoft

Microsoft released its hypervisor product in mid-2008 after a lengthy period of development. Up to then, Microsoft had only offered the Virtual Server, a free hosted virtual machine product (see the separate section in this White Paper). Looking at the architecture of Hyper-V, there is a very strong similarity to the XEN architecture. The are differences in the names used for the respective instances of Hyper-V.

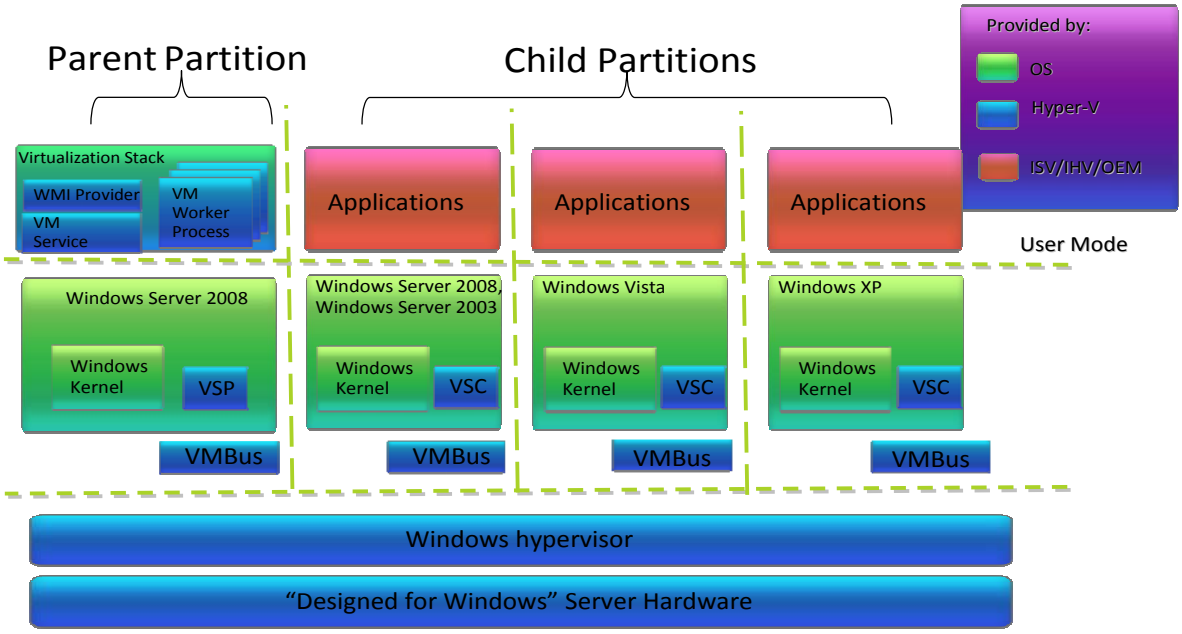

What is domain 0 with XEN is the parent partition with Hyper-V. The child partitions correspond to domain 1 to domain…n of a XEN architecture. A Windows Server 2008 is naturally used in the parent partition. Basically, however, this construct assumes the same tasks as in a comparable Linux environment. If newer versions of Windows or Vista are used in the child partitions, Microsoft uses a para-virtualization approach termed "enlightenment." As part of this, operating system versions that are specially tuned to virtualization and use a specific API for communicating with the hypervisor are deployed.
Hyper-V is a specific role in the new Windows Server 2008 strategy and is shipped as a firm part of a specific Windows Server 2008 distribution. Microsoft's current licensing policy is very restrictive. For example, when they purchase this new operating system variant, customers must decide bindingly whether they wish to acquire this additional functionality for an added charge. The server roles are server applications that can be subsequently installed individually. In earlier versions of the system, applications such as Internet Information Services (IIS) or file release services were configured as standard when Windows Server was newly installed. In contrast, a newly installed Windows Server 2008 cannot yet perform any functions. First of all, the administrator has to assign the required roles explicitly to the Server. Virtualization is also such a role. With its virtualization product, however, Microsoft also aims to address customers who do not use Windows operating systems at all or only do so occasionally. To enable this, Microsoft has provided a specific variant of its Hyper-V, Microsoft Hyper-V Server. This product, which runs in the parent partition, consists of the actual hypervisor and a basic kernel, including a driver model. That means it is possible, for example, to consolidate a pure Linux environment with Hyper-V as well. Hyper-V is excellently suited for virtualizing desktops on which XP or Vista is used. No explicit Windows 2008 CALs (client access licenses) are then required for the individual VMs (which represent a single desktop user). The differences between the two Hyper-V distributions are shown in the figure below:

Microsoft Hyper-V Server

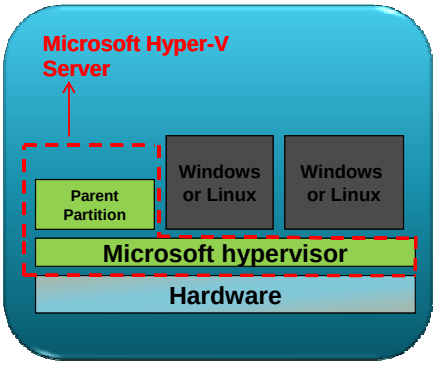

Contains Windows hypervisor and other components, including base kernel and driver technologies.

Hyper-V, role of Windows Server 2008

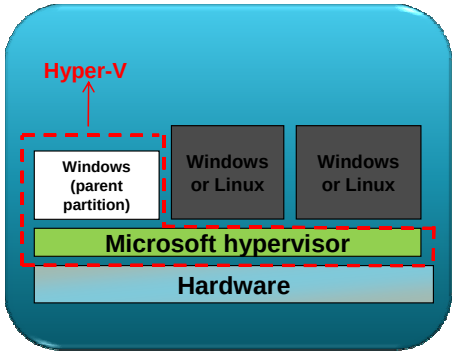

Available as a role in Server Core or full installation of Windows Server 2008

Microsoft's management strategy in connection with Hyper-V is very interesting. Microsoft was very quick to recognize the importance of this functionality and developed a finished product before the actual hypervisor was available. Its mission was not to offer an isolated solution for the virtualized world, but to integrate management as a further strategic pillar in its existing management portfolio.

Microsoft's management strategy covers all IT areas at an enterprise; all the components are administered centrally.

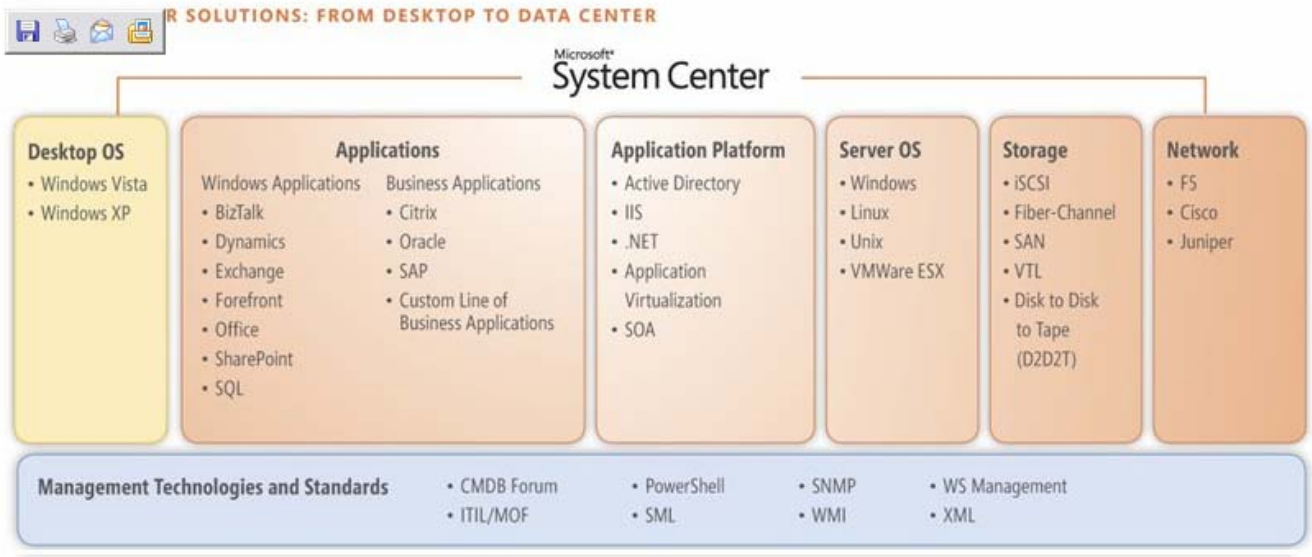

The most important Microsoft products in connection with virtualization are:

- System Center Operation Manager
- System Center Configuration Manager
- System Center Data Protection Manager
- System Center Virtual Machine Manager

More information on these management products can be found on Microsoft's Websites.

The System Center Virtual Machine Manager has the following structure:

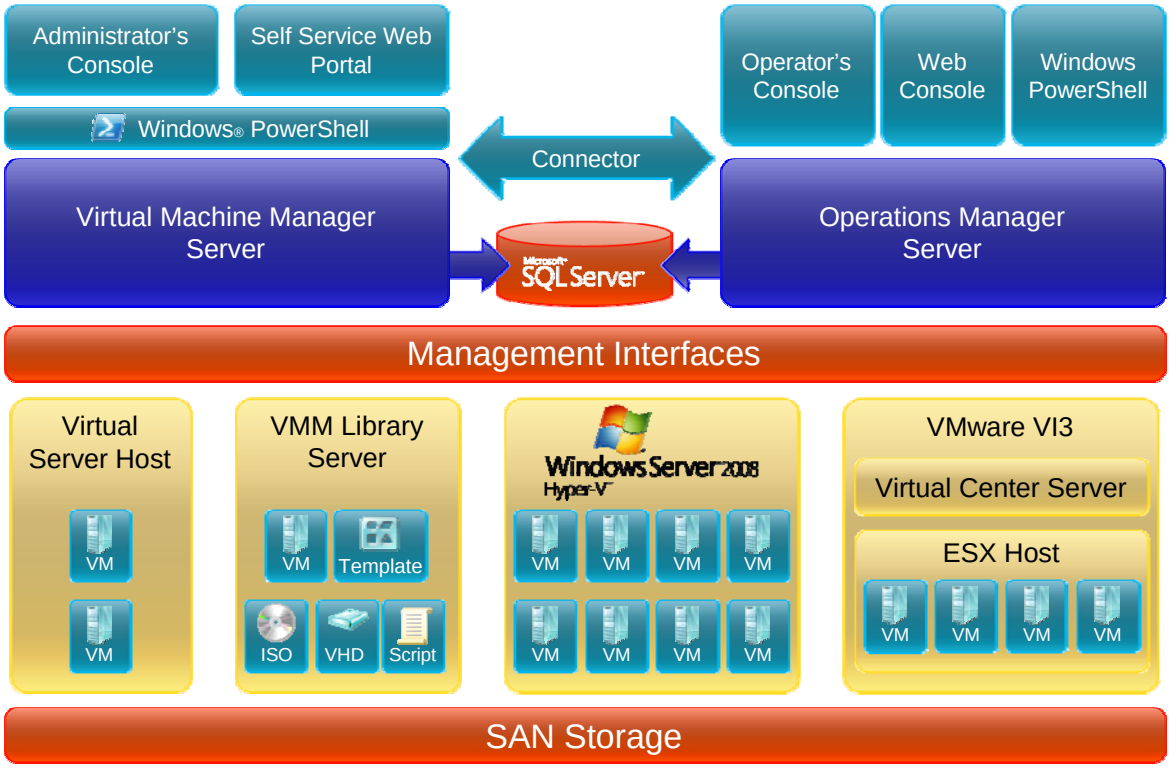

Microsoft also includes hypervisors from other vendors, for example VMware VI3 with ESX hosts, in IT management architecture. In addition, the SCVMM is connected to the Operations Manager, meaning that not only the VMs, but also the real instances (i.e. dedicated real servers and the real instance behind every virtualized server) are integrated in management. The applications, whether they run on a real or a virtual machine, are also covered by IT management. As a result, Microsoft has created a management solution that can administer all the resources in a modern data center. Microsoft sees this as the basis for a Dynamic Data Center.

One obvious strength of Microsoft's concept is that Microsoft boasts development expertise at all levels of an IT infrastructure: the virtualization, operating system and application levels. Microsoft will doubtless leverage this accumulation of expertise to tune its own products for the IT infrastructure optimally to each other, for example by gradually optimizing key applications (Office and others) for its virtualization layer.

How does Fujitsu Siemens Computers support Hyper-V?

Fujitsu Siemens Computers published its Dynamic Data Center strategy some years back. Virtualization is one of the supporting pillars in this concept, alongside automation and integration. Although Fujitsu Siemens Computers has a broader definition of virtualization and does not limit it to hypervisors, hypervisors play a key role in this concept.

Fujitsu Siemens Computers offers and supports various hypervisor technologies from different vendors. It goes without saying that Hyper-V is also supported, which means:

- Certification of relevant PRIMERGY server models
- Provision of specific drivers
- Integration of Hyper-V in the ServerView Management Suite
  - Provision of specific ServerView agents for monitoring the real server instance and the associated VMs
  - Management tools for simple and automated provisioning of guest operating systems in VMs of Hyper-V
- Integration of Fujitsu Siemens Computers' management tools in the Microsoft System Center Management Suite
- Specific support for Hyper-V and the operating systems that run on the VMs
  - Service offerings for all phases of a virtualization project

Summary on Hypervisors

VMware is the benchmark with richest feature set and enterprise-class mgmt. tools for a virtual infrastructure, but lacks mgmt. support for applications, phys. systems and 3rd-party hypervisors

Microsoft currently offers with **Hyper-V** only a „good enough" solution targeted at cost sensitive SME customers; given its limited feature set but aggressive pricing. The management suite is very well aligned with the overall strategy of Microsoft and the related products

Citrix strategy is to integrate **XenServer** with other products from their huge portfolio in order to offer more advanced end-to-end solutions beyond the hypervisor; also focuses on SME, leveraging their existing customer base and huge partner eco system

Novell focuses on good interoperability with other vendors (especially with Microsoft, but also VMware); has a comprehensive systems mgmt. suite; in combination with the acquired PlateSpin technology they are in a good position to provide advanced mgmt. for virtual infrastructures

RedHat has only a comparatively poor mgmt. offering for Xen; currently involved in development of an open source virtualization mgmt. product (oVirt); recent support for a competing virtualization technology KVM shifts focus away from Xen

Oracle and SUN are currently not part of the consideration about server virtualization.

Optimization for Hypervisors through hardware assistants

Depending on the application profiles, the use of hypervisor technology can have a remarkably negative impact on performance. Reasons for this can be simultaneous virtual address translation for multiple virtual machines and I/O emulation.

Nowadays, these problems are addressed by extensions to the x86 instruction set, e.g. with Intel Virtualization Technology (VT, codename Vanderpool) and AMD-V (codename Pacifica), whose support is incorporated in all the hypervisor products. This makes hypervisors simpler, smaller, more efficient and robust, and allows you to move more software onto new industry standard server platforms. Consequently there is no need to modify the guest OS for performance reasons, what makes para-virtualization not so important in the future. But it is worth mentioning that hardware assistants in the CPU chip will not eliminate the Hypervisors at all.

This trend is being rigorously pursued by processor vendors. The new objectives are to reduce the loss in performance when hypervisors are used and to ensure enhanced compatibility of the individual processor generations, which will deliver far greater agility with functions such as VMotion or Live Migration. Two examples may illustrate this:

Reduction in performance loss

The new AMD processor "Barcelona" has a Rapid Virtualization Indexing function that enables optimization of memory management in virtualized environments. The basic idea is that real memory is assigned once to a virtual machine and is then managed as usual by the virtual machine by means of memory tables, without the need for the virtualization software to intervene. That results in a significant reduction in arithmetic operations by the hypervisor, which is usually responsible for managing this. VMware VI3 has already implemented this technology and other hypervisors are to follow. Memory-intensive database applications will benefit from this in particular.

I/O-intensive applications also prove to be especially critical if they are to be virtualized. Every I/O operation results in activities by the hypervisor or within domain 0 or the parent partition, where the I/O operation is mapped on the real device. The same thing happens when an I/O operation is reported back to the VM as completed. The need for high-performance handling of I/O operations will lead in future to greater cooperation between I/O card vendors and hypervisor providers, with the aim of making VMs able to access the real hardware directly again. That increases the dependency between the hardware and software again, but the I/O frequency is increased significantly. Ultimately, customers must decide where their emphasis lies.

Increase in flexibility

If, for example, a VM writes to a disk, it "sees" only a virtual device, usually a "dumb" SCSI disk. In contrast, the data – which is assigned a LUN – is stored via a physical HBA (host bus adapter) on a real SAN device. The HBA only exists once, but must process the requests from all VMs. A guarantee of QoS (quality of service) for individual VMs is therefore not possible. In future, it is also to be possible to virtualize the HBAs so that every VM can be assigned its own virtual adapter. As a result, the QoS would be ensured for every single VM.

VMotion from VMware or Live Migration from the XEN derivatives permit a new quality of service availability. Planned downtimes, which account for more than 80% of downtimes at a data center, can largely be avoided if running VMs can be moved from a real system to another real system without ongoing services (applications) being interrupted. Up to now, VMs have only been able to be moved within the same type of processors families and generations, a restriction that forces customers to build server landscapes with homogenous processors – an unacceptable requirement given the rapid development of processor technology. This problem has also been solved thanks to cooperation between hypervisor and processor vendors. Once again, it was VMware that was first to respond and develop its Enhanced VMotion Capability technology so that it was ready for serial production. Together with Intel's VT Flex Migration or AMD-V Extended Migration technology, it is possible to move VMs between different processor generations from the same vendor.

Evolution of virtual machine technologies (Example: VMware)

Virtual machine technology has developed rapidly in recent years. VMware ESX provides the best example of how customers can improve the flexibility and efficiency of their IT infrastructures.

In a first phase VMware started with the Hypervisor to address consolidation and isolation. The independence between the hardware and software was also the basis for running old operating systems on the latest hardware.

In a second phase the focus was on managing the virtual infrastructure and so VMware released its VirtualCenter product. The targets of this development were versatile:

- Simple migration of physical to virtual machines
- Fast provisioning (e.g. by replication of virtual machines and flexible distribution)
- Load-balancing (acc to the utilization rate of the servers and specified resource policies of the virtual machines)
- Live migration of virtual machines including allocated disks from one physical hardware to another one while running
- Live migration of related file system entries of VMs (Flexible usage of different storage units)
- Automated high availability for virtual machines (automatic restart on another server in the event of a server failure)
- Enhanced Backup of virtual machines
- Automated patch management for hypervisor and virtual machines (even for suspended ones)

Meanwhile end-to-end solutions are available. The distributed power management which deactivates inactive components in order to reduce energy consumption, considering the fact that most servers consume around 50% of their peak power requirement even when being idle. That permits a totally different type of load balancing. The aim is no longer even utilization of the VMware ESX farm, but to map as few services as possible on physical servers. At the same time, it is always ensured that each service has the resources guaranteed to it.

Another trend-setting end-to-end solution was accomplished with development of SRM (Site Recovery Manager). The aim of this product is for virtualized IT infrastructures to be restarted automatically in a backup data center if a data center fails. Storage products together with SRM deliver a smooth, fully automated process that helps in the event of a disaster to restore orderly operation of the IT infrastructure reliably and with minimum interruption.

The mode of operation and functional range of SRM are summarized in the figure below:

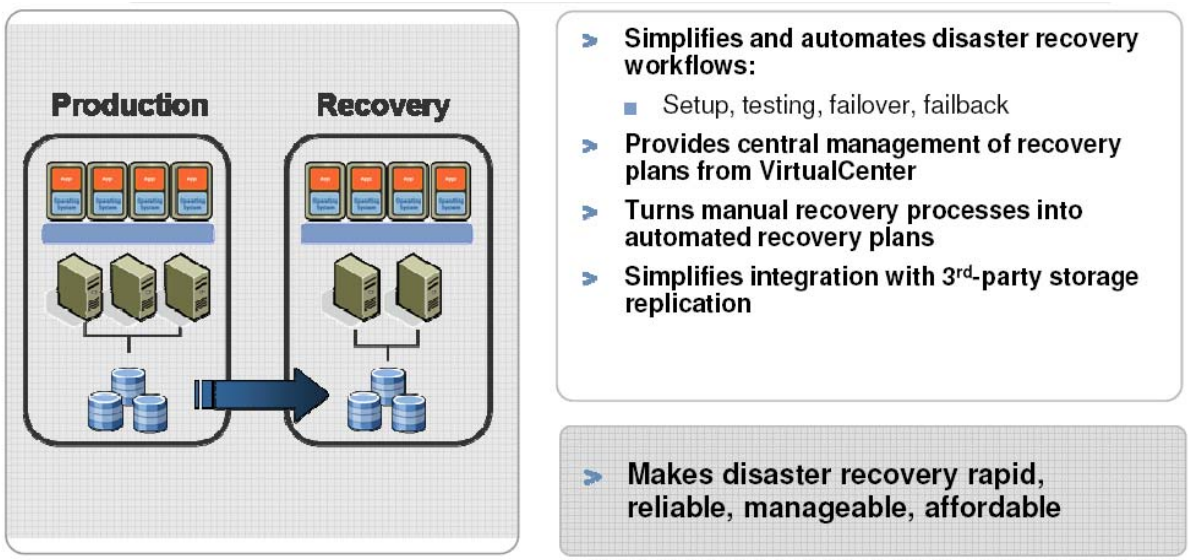

It is important to know that SRM always requires mirroring of the storage system to work.

In connection with virtualization of desktops, new technologies that fulfill the special circumstances of this scenario are also being developed (see the separate section in this White Paper). These essentially relate to the avoidance of redundancy in provisioning operating system images on expensive storage units, new ways of delivering applications (application streaming) and a new method for detecting viruses in centralized desktop environments.

Hypervisor embedded in hardware

First developments of Fujitsu Siemens Computers clearly supporting a trend, that virtualization will become a default feature in the future. Hence it makes absolutely sense to integrate the hypervisor into the server hardware, e.g. by depositing it in a flash ROM attached via internal USB. Doing so, certified server hardware is shipped ready-to-run. From server boot to running virtual machines it will take a couple of minutes only. Thus deployment time and administration cost is dramatically reduced. With a small footprint of code, security exposure is reduced. Likewise patching and updating is very simple. With an embedded hypervisor it is even possible to keep servers diskless.

Embedded Hypervisors are of particular advantage in remote locations without skilled staff. Moreover they are the basis for appliances.

Selected PRIMERGY server models from Fujitsu Siemens Computers are supplied with VMware ESXi 3.5 embedded.

Virtual machines in an OS hosted environment

The technology described in the following is only used now to a very limited extent.
This approach of virtual machine technology is the so-called OS hosted environment approach. Here the virtualization layer resides on top of the host operating system, i.e. the operating system of the physical server, where virtualization is just one of many OS features. Correspondingly, the virtualization layer has the status of any other application, which could run on the same hardware and OS. Of course, it uses the device drivers and the resource management of the host OS.

Examples are VMware Server, Microsoft Virtual Server and VirtualBox from Sun.

The virtual machine concept in an OS hosted environment came into being in the early phase of server virtualization. One of the main areas of application was running different development and test systems on the same hardware. Today it is relevant in small environments only.

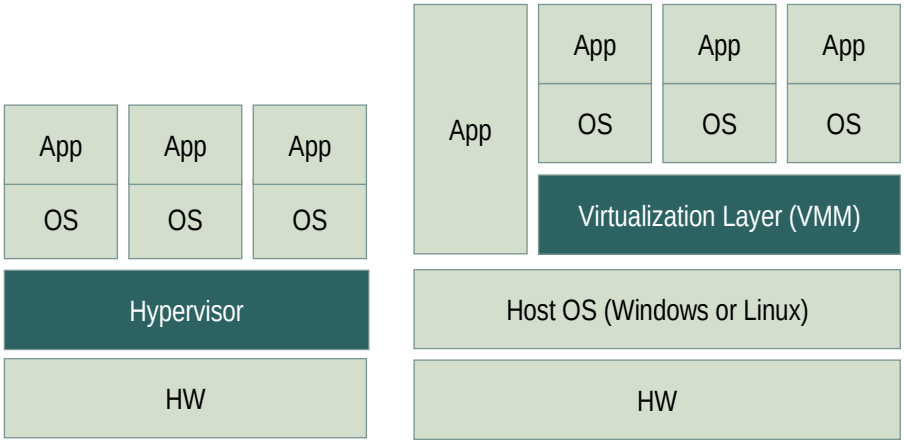

Virtual machines in a bare-metal environment

Virtual machines in an OS hosted environment

Increased availability using x10sure

If you wish to pursue consolidation with virtual machines, you should always be clear about one thing: if a consolidated server fails, not one service is affected, but every VM and the services mapped in it are no longer available. That is why such an infrastructure needs a mature high availability concept that minimizes unplanned downtimes.

It is true that virtualization software from VMware and others includes components that ensure high availability for the individual virtual machines. When a physical server fails, all virtual machines concerned will be automatically restarted on another physical server. But of course the VMware solution will not bring the failed server back. Even if VMware restarts the VM on another server in such a case, this resource is not available to the entire ESX Server farm. That inevitably means that all the servers in the farm must not be fully utilized to start with so that additional loads can be absorbed if a server fails. This is exactly where x10sure from Fujitsu Siemens Computers comes into the game.

Principally x10sure monitors servers and storage systems of an IT infrastructure. Whenever a production server crashes, the application and its operating environment (in our example this would be the hypervisor and the virtual machines) are automatically restarted on the spare server. If the storage system or the storage fabric fails, data or the data access can be automatically recovered.

x10sure ensures n+1 high availability of applications and data in a simple and cost-effective way, without building complex cluster configurations. Applications need not be cluster aware, as they need to be in conjunction with Microsoft Cluster Server, and there is no need to install the expensive Enterprise Edition of Microsoft's Windows Server. For sure, x10sure is also applicable for overloaded server farms.

The figure below shows how x10sure basically operates:

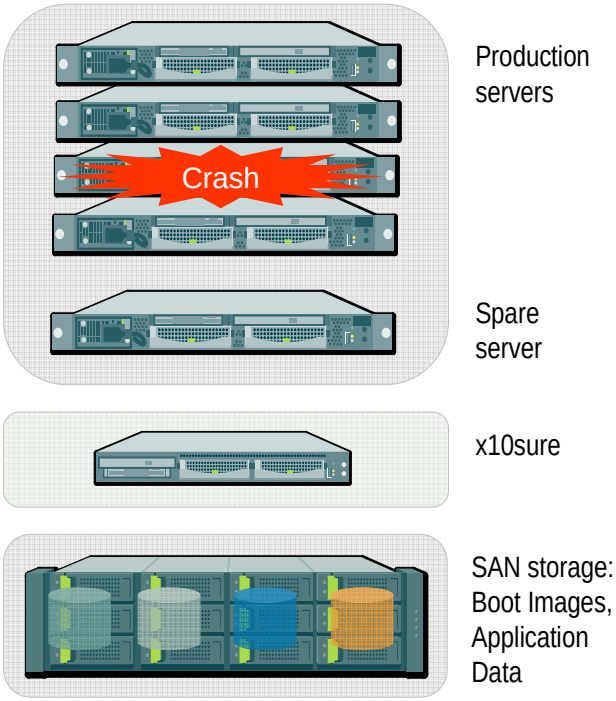

x10sure can be used very expediently in VMware infrastructures. Two constellations where real and virtualized systems are to be backed up may illustrate this.

1. A VMware infrastructure is operated without add-on products from VMware that need to be paid for. Instead of Enterprise Edition from VMware ESX 3i, only the much cheaper Foundation Version is used. If an ESX Server fails, x10sure can restart the substitute server with the aid of the failed server's boot image. The associated VMs are then automatically restarted.
2. A VMware infrastructure is operated with Enterprise Version VMware ESX 3i, which contains functions such as HA and DRS. If a server fails, VMware restarts the VMs of the failed server on the remaining ESX Servers, in accordance with the load balancing rules of the DRS. That remedies interruption to the service very quickly, but the farm has a server instance less. x10sure can supply this missing server in the form of the defined substitute server. After the substitute server has been started, the DRS can again ensure that the entire farm is evenly utilized.

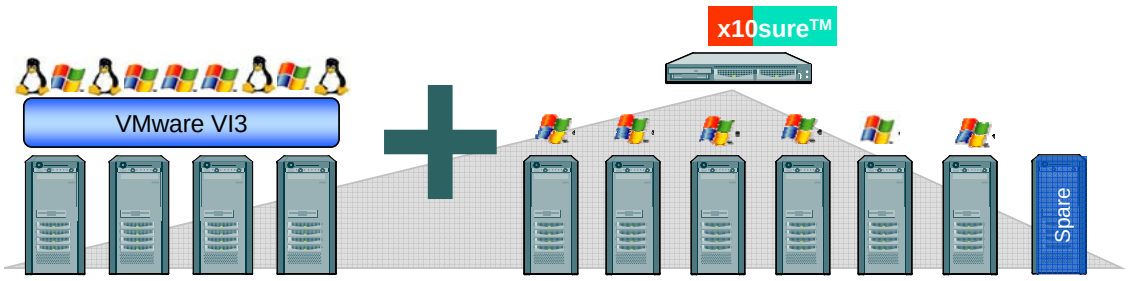

x10sure protects your servers in real **and** virtual environments

It is worth mentioning that x10sure only works with a homogeneous PRIMERGY infrastructure.

Desktop virtualization follows server virtualization

The hypervisor concepts that have been dealt with here in detail have recently gained in importance in a completely different context, namely desktop virtualization. To date, the virtualization concepts have been applied to server environments, i.e. server operating systems and the associated applications have been virtualized. Technically, however, it is perfectly possible to virtualize conventional PC infrastructures as well.
One of the new trends following server virtualization is the virtualization of desktops, which is seen as the most significant PC technology change for years. Virtual desktop infrastructures (VDI) enable desktop operating systems to be run as virtual machines in a hosted datacenter environment.

VDI was influenced by 3 concepts: the traditional desktop, traditional server-based computing and server virtualization. The good thing with VDI is that it gives you the best of all these 3 concepts. Virtual desktops are individualized and isolated, and there are considerable improvements in terms of availability, security, manageability, resource utilization and energy consumption. Virtual desktop infrastructures enable flexible access to your desktop environment from anywhere and its flexible execution, too. And finally, software lifecycles can be widely extended, due to the absolute independence of hardware and software.

All these advantages will considerably improve service levels and reduce total cost of ownership across your workplace infrastructure.

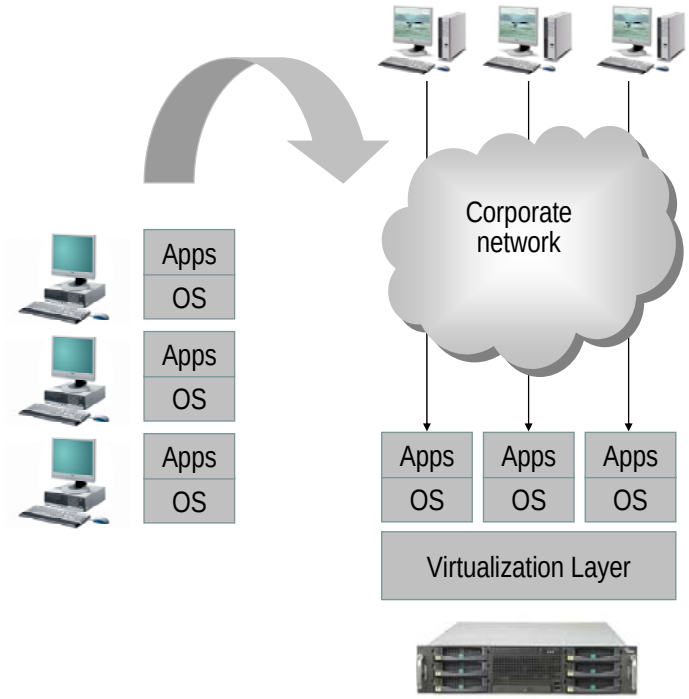

Dynamic IT for Workplaces

“Dynamic IT for Workplaces” from Fujitsu Siemens Computers completes virtual desktop infrastructures and traditional server-based computing by various technologies, such as connection brokering, application virtualization, application streaming, OS streaming and more. A complete solution stack with pre-integrated components can be simply tailored according to individual customer demands. The comprehensive end-to-end service portfolio across the entire project lifecycle even includes ROI services for IT investment decision support that help a customer to determine the financial impact of infrastructure changes already in early project phases. All this helps customers reduce project time and risk.

Only an overview of desktop virtualization can be presented in this White Paper. This is because desktop virtualization results in extensive changes to the traditional PC operating concept and inclusion of many new technologies, such as application streaming, OS streaming and many others. You can find a detailed summary of desktop virtualization with solutions from Fujitsu Siemens Computers at the following Website:

http://www.fujitsu-siemens.com/solutions/it_infrastructure_solutions/dyn_it4worksplaces/index.html

I/O virtualization

I/O virtualization addresses reduction of complexity towards local area networks and storage area networks. Virtualization of server hardware has got a variety of flavors which will be discussed more detailed in this section. Virtualization on the operating systems level originates from UNIX servers; however these UNIX-specific flavors don't play a major part today and will therefore not be regarded in this briefing paper. And finally there is virtualization on the application level.
Servers communicate with local area networks (LAN) via network interface controllers (NIC) and with storage area networks (SAN) via host-based adapters (HBA). Both types of controllers are supplied with static addresses (MAC and WWN respectively). This means that moving the operating system and the applications to another server, e.g. after the failure of the original server, requires a reconfiguration of the LAN or the SAN.

With I/O virtualization you will overcome these issues. MAC addresses and WWN as shown to the network are kept fixed. Virtualizing the real I/O addresses enables you to dynamically move the I/O addresses between servers. All this is fully transparent to the LAN and SAN domain. No need to adapt any network addresses, no interference with network administrators, thus reducing management complexity. Hence I/O virtualization will give you more flexibility when deploying new servers and services, or when replacing failed servers.

An example for I/O virtualization is ServerView VIOM from Fujitsu Siemens Computers.

Let us now have a look at one example demonstrating how I/O virtualization works. Server #1 is supplied with MAC #1 and WWN #1, server #2 with MAC #2 and WWN #2. When server #2 fails, a failover to server #3 will happen which acts as a spare server. The virtual connection is unplugged, and both the addresses MAC #2 and WWN #2 are moved to server #3. This will not require any changes in any of the networks.

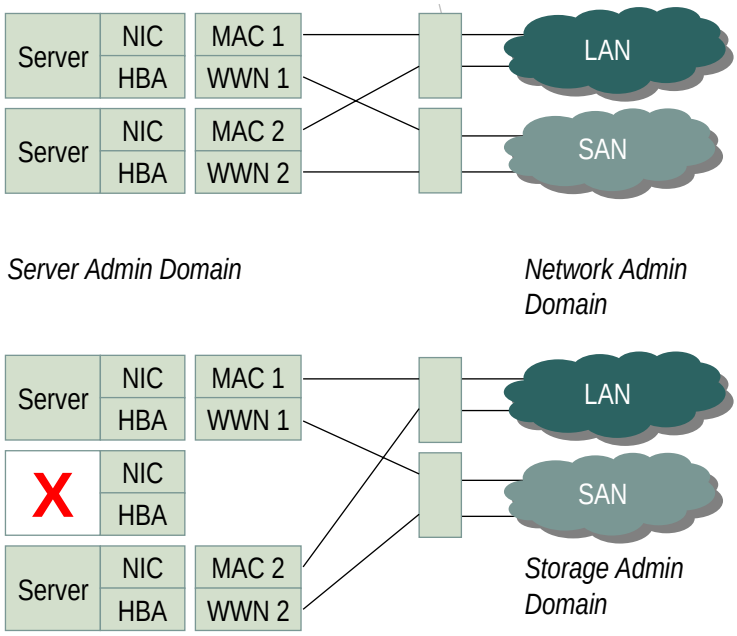

This I/O virtualization technology is provided by Fujitsu Siemens Computers with the *ServerView Virtual-I/O Manager* for the BX600 blade systems. The figure below shows the components that are required:

What it needs - briefly

Hardware

- Latest BX600 S3 Blade chassis
- Special I/O Connection Blades
- Latest BX620 S4 (Intel) or BX630S2 (AMD) Server Blades

Software

- ServerView Virtual-I/O Manager (billable item !)

PRIMERGY BX600 S3

I/O Connection Blades

IBP Panel & FC Access Gateway

ServerView VIOM

The *ServerView Virtual-I/O Manager* does not yet have any automated processes, i.e. all the actions have to be controlled by the administrator. An open API to enable automatic, software-controlled processes is planned for the next version.

Hypervisors are not the answer to every question

A hypervisor reduces complexity in a physical environment, but it also creates a new one at the virtual level. By introducing a hypervisor, you will always take the risk that your virtual server instances will explode, if there is no strict policy how to avoid virtual server sprawl.

In addition, the parallel existence of virtual and real resources requires a management of both; however management for virtual and real resources is of course different.

It is a matter of fact, that certain applications require quite a special software stack. Using a different hypervisor would lead to compatibility or support problems. Examples are applications from SAP, Oracle and Microsoft. There is for instance no certification of Oracle software for VMware. An analysis, which hypervisors collaborate with which operating systems and applications, leads to complex combinatorics.

It should also be stated that a virtual machine concept is not useful, if an application already eats all server resources in a native, real environment. Examples are applications with high I/O needs, such as DB servers or mail servers.

A different approach: Pooling of physical servers and dynamic provisioning

An approach totally different from the virtual machine concept is pooling of physical servers combined with dynamic provisioning. The virtualization of services allows any service to run on any server. This is achieved by storing the complete image of the software stack (operating system, middleware and application) on a central storage system. The allocation to a real server, i.e. provisioning of the server, which is stateless till then, happens dynamically on demand and just-in-time, either when the service is needed, after a server failure or in conjunction with a change in workload. In other words, the applications run on the servers without having been fixed installed there. Here you see the architectural principals of this solution approach

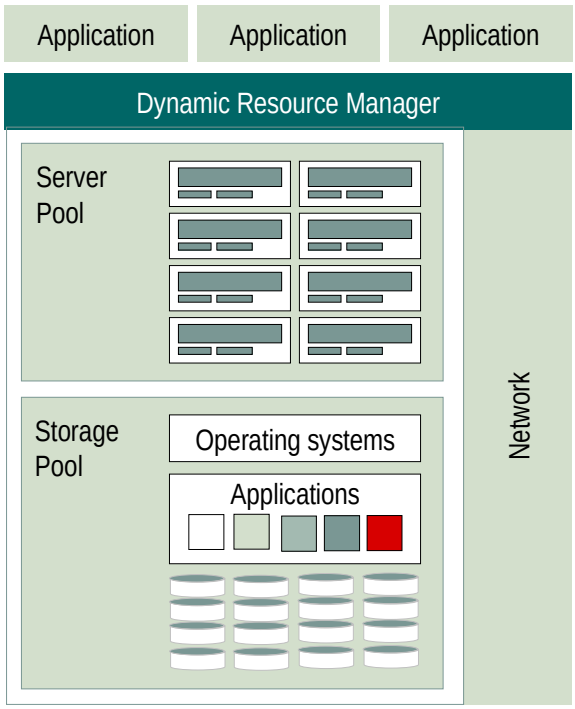

Let us now discuss some of the advantages of this concept. First of all, we do not need an additional layer, which could cause compatibility problems. Existing resources can be flexibly used, and fewer resources are needed. This saves time and cost. High availability is achieved without cluster technology.

The key infrastructure resources for computing, networking, storage and control are separated from each other, and can therefore be scaled independently and assigned dynamically. The result is a highly adaptable infrastructure that responds to changing service needs in real time. All resources are available on demand, giving you greater enterprise-wide agility.

Due to the fact that there is only one shared operating system, installation and deployment can be done very rapidly, and software updates do not require time-consuming software distribution. This reduces complexity and management efforts tremendously.

Examples are FlexFrame for SAP and FlexFrame for Oracle 10g from Fujitsu Siemens Computers. These solutions are the first real adaptive computing solutions in the industry.

FlexFrame for mySAP Business Suite is a joint development of SAP, Network Appliance and Fujitsu Siemens Computers. It is built on SAP's Adaptive Computing, which in turn is an integral part of SAP NetWeaver™, the new business process platform of SAP. It is designed to dramatically increase flexibility, simplify operations and reduce costs as compared with previous infrastructure concepts for SAP.

FlexFrame offers software services on demand. This is achieved by complete virtualization of the software services and by clean separation of the key infrastructure resources:

- Computing
- Networking
- storage
- and control.

A modular architecture replaces rigid hardware and software configurations with virtual bridges. Computing power, storage resources, network components and mySAP™ ERP solutions can all be scaled independently and assigned dynamically. The result is a highly adaptable infrastructure that responds to changing service needs in real time. All resources are available on demand, giving you greater enterprise-wide agility.

Fujitsu Siemens Computers considers the following plus points to be particularly valuable in meeting the requirements for a virtualized application environment, shown in the following table:

Shared operating system	Virtualized SAP services
□ Rapid installation	□ Adaptive computing based on SAP NetWeaver™
□ No need for time-consuming software distribution	□ Flexible capacity allocation
□ Single update for all servers	□ Optimized deployment of server resources
	□ Server changeover in minutes
	□ Simple grouping and pooling
Consolidated storage systems	Automated high availability
□ Innovative backup strategies	□ No need for cluster technology
□ Fail-safe data mirroring	□ Smart autonomous agents
□ Expansion without migrating data	□ Easy system expansions
□ High investment protection	□ Stable business processes

Virtualized resources, automated management – the key benefits of FlexFrame for mySAP

FlexFrame™ organizes all server and storage resources and handles variations in workload by reassigning available capacity in a matter of seconds. This ensures consistently high levels of performance and efficiency with fewer resources. Your investment levels and administration effort drop dramatically.

FlexFrame™ is self-configuring, self-optimizing, self-healing and self-protecting.
FlexFrame for mySAP Business Suite now runs on Linux in 32- and in 64-bit mode and on Solaris it supports mixed environments using NetApp NAS and also support EMC SAN. Since FlexFrame is a grid architecture based on application virtualization, it does not create virtualization-related performance issues and it does not require extra hardware support for non-standard features. It constitutes the best-in-class infrastructure for the SAP NetWeaver™ environments.

While other vendors have been able to gain the ACC certificate of conformity for their respective offerings, FlexFrame has always had a distinct set of features beyond ACC and we are maintaining a continuous dialog with SAP in order to retain our advantage in this area.

FlexFrame and Adaptive Computing Controller have very strong benefits in terms of TCO and ROI. According to a BearingPoint Study the TCO reduction can be as high as 62% with an ROI of 168%. This is a major step in driving down costs for an entire SAP installation and its operation. FlexFrame for Oracle 10g is deeply integrated with Oracle GridControl.

FlexFrame for mySAP Business Suite has a very clear positioning: it is the best-in-class infrastructure solution for supporting SAP-consolidated enterprise applications (EAS) environments based on SAP NetWeaver™ and cannot serve any other purpose.

Customer example: TeliaSonera

TeliaSonera, the leading Nordic and Baltic telecommunications leader, decided in favor of FlexFrame for SAP, in order to set up a flexible and scalable SAP IT infrastructure for procurement, human resources and finances. With the aid of FlexFrame for SAP they enjoy a scalability of up to 20,000 users, fast provisioning of new servers, high availability of SAP applications, and much better resource utilization.

Janne Puhakka, Division Director, TeliaSonera:
“FlexFrame from Fujitsu Siemens Computers gives us the best solution for our purposes. The cost-benefit ratio of the system platform is optimized, and resources can be dynamically balanced to accommodate additional demand at any time.”

Challenges with server virtualization

The coexistence of physical and virtual servers makes unified management difficult. As mentioned earlier, you will rather often find competing hypervisors in parallel in combination with different operating systems, because certain applications require quite a special software stack. This will end up in a complex combinatorics and a complicated support matrix. It should also be taken into consideration that different applications support different virtualization technologies. All this requires a high level of skills and respecting education efforts.

Therefore the objective is a uniform management for virtual and real server resources, as well as for various virtualization technologies and principles.

Our answer: FlexFrame Infrastructure

This is exactly where FlexFrame Infrastructure from Fujitsu Siemens Computers helps. FlexFrame Infrastructure is a unified resource management for physical and virtual servers. Once having created your pools of computing, storage and network resources, the administrator will just define what an application needs in terms of resources, and all the rest will run automatically. Resources are dynamically allocated to applications on demand. Within a few minutes applications can be made available on real or virtual servers. Likewise, due to built in high availability, disaster recovery is carried out automatically by starting applications on a spare system, without the well-known clustering complexity. And there is no need at all to take care of diverse virtualization technologies. In addition to resource management, I/O virtualization is covered, too.

By means of FlexFrame Infrastructure you reduce complexity of virtualization, increase speed of service delivery, improve service quality and you will lower your investment and administration efforts.

For IT operation, FlexFrame Infrastructure basically means management of resources instead of dealing with technology, therefore FlexFrame Infrastructure is the platform for building dynamic service-oriented IT infrastructures.

The first implementation of FlexFrame Infrastructure is PRIMERGY BladeFrame, a high integration of hardware and resource management in a sort of black box approach. Thereafter the operating principles of PRIMERGY BladeFrame were extended to the PRIMERGY blade server line.

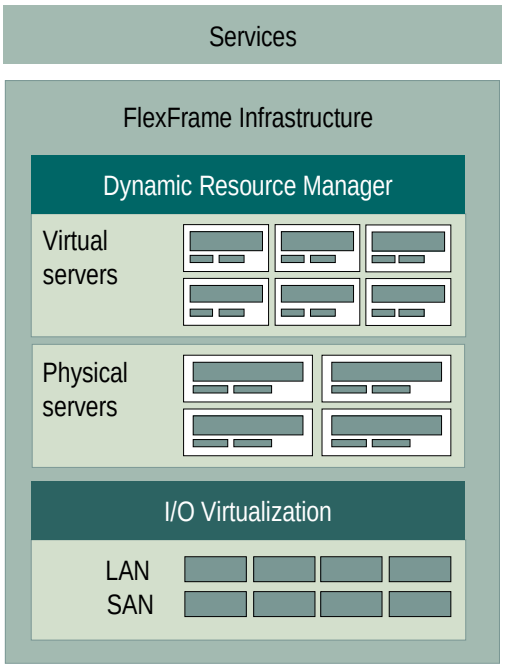

Resource management, the virtualization technologies used and I/O virtualization by FlexFrame Infrastructure require a more detailed presentation than is possible in the constraints of this White Paper.
A good point to start with is

http://www.fsc-mediaserver.com/ms/go.cfm?itemid=24990&jumpfactor=6

Summary

The one and only way to benefit from server virtualization does not exist. Individual use scenarios require different technologies. Virtual machines, in particular the hypervisor concept, are well accepted, but not optimum in every situation. One success factor will be a uniform and simple management of real and virtual resources in order to handle the complexity of virtualization easily, transparently and at affordable costs. FlexFrame Infrastructure is a platform that manages real and virtual resources. Thus a resource becomes a service. With FlexFrame Infrastructure we want to make our customer‘s transition to a service-oriented infrastructure as easy and smooth as possible.

Further information

http://www.fujitsu-siemens.com/solutions/it_infrastructure_solutions

http://partners.fujitsu-siemens.com/com/products/ddc/virtualization

Published by department: IP PM TOP

Günther Aust
Phone: ++49 8218044040
Fax: ++49 821 804 84040
guenther.aust@fujitsu-siemens.com
http://www.fujitsu-siemens.com/
Extranet: http://extranet.fujitsu-siemens.com/