THE POSSIBILITIES ARE INFINITE

FUJITSU

WHITE PAPER

DYNAMIC SERVER INFRASTRUCTURE

Design Guide

Business requirements and the development of IT architectures

Business processes and IT infrastructure

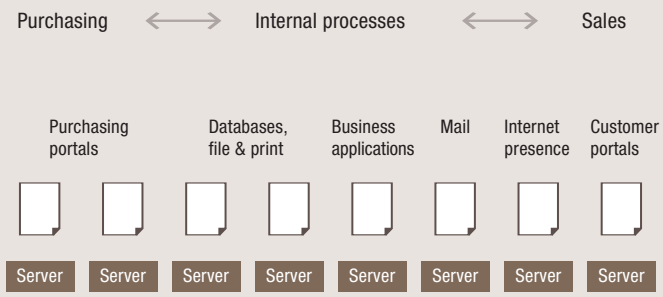

The future viability of an IT infrastructure is linked closely with the business process. A business process is typically supported by an array of applications. The purchasing department cooperates with its suppliers via a purchasing portal. Production planning transfers the purchased materials to its system. And orders from customers are processed by Sales via a customer portal. Each application is a link in a chain, with each link depending on the others.

From the point of view of the IT infrastructure, this means that several servers with different applications run alongside each other. Each and every one of them has to be available and deliver sufficient performance. If that is not the case, the business process is impacted directly – and can even be interrupted in the worst case scenario.

A dynamic server infrastructure addresses the needs of a modern application environment by means of a server pool concept. This pool is managed centrally. Software is not installed concretely on individual servers. Instead, the operating system and application software are provided on a central storage system. From there, they can be started on any server in the pool – on physical and virtual servers alike. Central management of the pool also includes high availability functions: If one server fails, the affected services are automatically started on another server in the pool.

A dynamic server infrastructure offers multiple benefits:

- Pooling servers makes it possible to greatly reduce excess capacities. System environments can be consolidated and utilization optimized.
- Entire business processes are protected by a single high availability concept. At the same time, less critical applications can be operated reliably at very little extra cost.
- The dynamic assignment of servers to applications enables rapid rollout of new applications. It also simplifies maintenance, reduces downtimes and permits rapid response to unplanned workloads.
- A pool can be built with a small number of servers and expanded gradually. Unlike with large system architectures, the up-front expenditure on initially unused capacities is lower.
- Last but not least, a dynamic server infrastructure is the ideal platform for increasingly important IT concepts such as SOAs (service oriented architectures) and cloud computing, which necessitate no fundamental changes to the architecture.

Dynamic server infrastructure

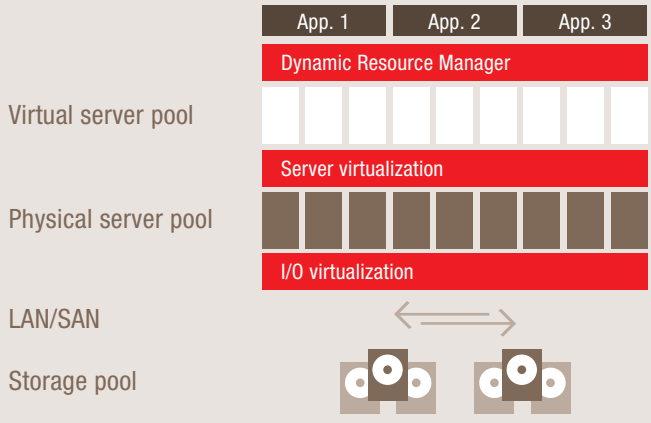

The technology behind a dynamic server infrastructure

A pool architecture is built on four pillars:

1. x86 server architecture

The first characteristic of a pool architecture is, of course, the fact that it consists of multiple servers. The best price/performance ratio can be achieved with x86 servers that run Windows or Linux. Just about all new applications are developed first on these platforms – and that is certainly true when it comes to SOAs and cloud computing.

Blade servers are ideal for this since they represent per se a pool architecture. Their design alone already produces a consolidation effect – in relation to I/O connections, power supply and cooling, for example. They are much easier to handle than tower or rack servers, and that pays off by reducing the administrative overhead and allowing more servers to be run by the same team.

2. I/O virtualization

One requirement for the flexible assignment of applications to servers is a virtualization layer between the servers and the storage area network (SAN) as well as the local area network (LAN). This is because of the fact that, without I/O virtualization, any change in a server's allocation would necessitate many manual steps and close coordination between server, network and storage administration.

3. Server virtualization

The power of the latest server generation is sufficient for just about all workloads. Systems with a large main memory and many processor cores are typically utilized only to a slight extent. In order to improve the degree of utilization, servers are therefore being virtualized more and more through the installation of hypervisor products. As a result, a large number of virtual machines can run on one server, each of them acting like a complete, independent server. This not only raises the level of utilization, but also simplifies application provisioning, allows applications to move between servers without any downtime and enables low-cost, very simple implementation of high availability and disaster recovery solutions. The reason is that every virtual machine and its contents – i.e. the guest operating system and application – consists of just a few files that are easy to transfer.

4. Dynamic resource management

A wide range of resources have to be administered in a server pool – physical servers, virtual servers and their SAN and LAN addresses. In addition, storage capacity has to be allocated to every server. Central management is necessary to permit the dynamic administration of all these resources and thus the efficient operation of a server pool. It will probably never be possible to replace all the tools that control very specific elements in an overall architecture. However, it is worthwhile to make a distinction between highly specialized tasks that are carried out only a few times a year and daily tasks. For the latter, there is the central resource manager. It can be used to reduce the time and effort involved in having to coordinate every task with server, network and storage specialists. The organizational concept here is that storage and network administration supplies quotas with storage capacities and network addresses to server administration, which administers them with the resource manager. It is also very simple with this technology to implement automatic high availability for the entire server pool. Here a large number of production servers are protected with a handful of substitute servers. More highly developed resource managers should have such a function.

Fujitsu's dynamic server infrastructure

A dynamic server infrastructure is one of the key components of Fujitsu's Dynamic Infrastructures strategy. It comprises four main components:

PRIMERGY BX900

The PRIMERGY BX900 blade server system has been specifically developed to enable efficient operation of a server pool. A lot of computing power in a very small space, maximum main memory configuration and high I/O bandwidth means that all types of workloads – on both physical and virtual servers – can be centralized. Consolidation of network connections, integrated power management and central control of all physical components comprehensively optimize the pool architecture. The latest Fibre Channel and Gigabit Ethernet technology and a specially developed internal architecture for optimal network throughput eliminate any I/O bottlenecks.

In the past, high I/O throughput, CPU performance and main memory requirements were arguments against using blade servers in large server virtualization projects or for operating large databases and OLTP applications. The PRIMERGY BX900 is also ideally suited for these application scenarios: Blades with up to 32 processor cores and a main memory of up to 256 GB provide sufficient capacity even for extremely power-hungry applications. The PRIMERGY BX900 also boasts energy efficiency significantly. That dramatically reduces power consumption. Active power management and the highly developed CoolSafe™ cooling technology additionally improve the use of energy at data centers.

ServerView VIOM

ServerView VIOM helps ensure that servers can be deployed quickly for other purposes without incurring problems with network addresses and their administration. ServerView VIOM separates the network identity of a server from the hardware. It virtualizes the specific I/O parameters of a server and makes them available in hardware-independent server profiles. These profiles are stored centrally, so they can be assigned easily to a server blade. Operating systems and applications on a central storage system can thus be moved simply between servers without any change to their network addresses.

Hypervisor technology

Fujitsu supports hypervisor products from VMware, Microsoft, Citrix, RedHat and Novell in a dynamic server infrastructure end-to-end for the PRIMERGY BX900, ServerView VIOM and ServerView Resource Coordinator VE.

FUJITSU'S PLATFORM FOR A DYNAMIC SERVER INFRASTRUCTURE

- ☐ The **PRIMERGY BX900** blade server system was specifically developed for server pools.
- ☐ **ServerView Virtual I/O Manager (VIOM)** separates the network identity of a server from the hardware. It virtualizes the specific I/O parameters of a server and makes them available in hardware-independent server profiles.
- ☐ **Server virtualization** with VMware, Hyper-V and Xen: Hypervisor products from vendors such as VMware, Microsoft, Citrix, RedHat and Novell divide physical servers into logical virtual machines, allowing improved utilization of the servers in a pool.
- ☐ Dynamic resource management with **ServerView Resource Coordinator VE,** the central instance in a pool architecture that controls all physical and virtual resources.

ServerView Resource Coordinator Virtual Edition (VE)

A dynamic server infrastructure demands wide-ranging know-how. Administrators have to manage a mix of physical and virtual servers and master various virtualization technologies and multi-faceted combinations of operating system and hypervisors. They have to know how to deal with network addresses in assigning server resources, and they have to be familiar with starting operating systems over a SAN and allocating predefined storage capacity to a server. The task of ServerView Resource Coordinator VE is to avoid problems in such activities and simplify management as a whole. This powerful management system standardizes workflows in administering large pools of physical and virtual servers. It reduces the administrative overhead by automating processes, enabling the dynamic allocation of server resources and offering low-cost high availability.

The principle is extremely simple: Administrators define the physical and virtual resources they wish to assign to applications, and ServerView Resource Coordinator VE does the rest. They do not need to acquire in-depth knowledge of the new, often complex technologies involved – they just use them.

ServerView Resource Coordinator VE optimizes operation of server pools in three areas in particular:

Uniform management of physical and virtual server resources

ServerView Resource Coordinator VE consolidates and simplifies the administration of mixed environments and provides uniform control functions that can be used for physical and virtual servers. The management system supports blade servers, rack servers and tower servers.

Simple server lifecycle management

A server pool is not a static environment. New servers are added if workloads increase or new applications are rolled out. Servers are replaced and applications relocated from one server to another as resource requirements or workloads change.

To enable rapid allocation of servers, the operating system and applications are not installed on a concrete server. Instead, they are installed on a central storage system, from which the operating system and then the application is started on a server. This software centralization enables rapid assignment of servers. This is also true of the allocation of LAN and SAN addresses when one server is to assume another's role. In the case of blade servers, ServerView Resource Coordinator VE uses ServerView VIOM's very convenient I/O virtualization function. As a result, scenarios where days of installation work on a target system were required to relocate applications are a thing of the past. If an application is to run on multiple servers, the installation is cloned and can be run on the desired number of servers in a matter of minutes. Yet ServerView Resource Coordinator VE can do even more: It simplifies the replacement of servers for maintenance and rollout of new software versions, since the software images only have to be assigned to the servers. Tasks such as installation, rollout or replacement require just a fraction of the time and work compared with conventional environments. And because physical and virtual servers alike are supported, mixed environments are far more efficient to operate. Day-to-day work with physical and virtual servers using different technologies is carried out with a single tool – ServerView Resource Coordinator VE.

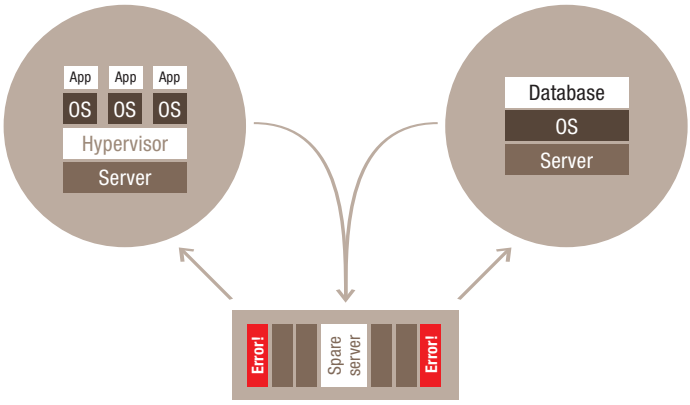

High availability for the entire server pool

If a server fails, ServerView Resource Coordinator VE automatically starts a spare system to prevent downtime. In this way, a large number of production servers in a pool can be protected by one or a few substitute servers (N:1 failover). This is a very low-cost alternative to a cluster, where a substitute is required for every production system. It should also not be forgotten that – due to the high costs and complexity of clusters – only a small portion of critical applications is generally protected. Such restrictions are no longer necessary. The spare servers can be jointly used in mixed, physical and virtual environments. And the high availability function complements existing high availability functions of the hypervisor vendors for their virtual machines and can be combined with them as desired.

How a dynamic server infrastructure supports the continuing development of IT

Rationalized IT infrastructure

If every application runs on a server to which it is explicitly assigned, the typical result is inefficient use of system resources and a high administrative overhead. The first step in optimization is therefore to relocate the applications from isolated servers to a pool architecture. Creation of the pool and use of server virtualization significantly increases resource utilization, reduces the number of required systems and makes management more efficient. This is particularly true of the PRIMERGY BX900 system, where handling of the physical servers is very simple – one of the reasons being that the I/O connections with all their cables are highly consolidated. And since all the physical and virtual servers are managed with just one tool – ServerView Resource Coordinator VE – no separate pools for VMware, Hyper-V, Xen, Linux or Windows have to be set up. There is one pool, enabling all options for rationalization to be leveraged. This also goes for high availability, since the N:1 failover of ServerView Resource Coordinator VE ensures that the number of standby systems required is sharply reduced.

To summarize: Improved utilization minimizes the number of servers required, high availability can be achieved at far lower cost, and the increased efficiency in system management enhances the IT team's productivity.

Dynamic IT infrastructure

Once the server pool is up and running, the resources can be used dynamically. An application with an increased workload can be relocated to a server with greater capacities, for example. Or an application can be started at the same time on several servers in order to offer sufficient performance. And since such changes are possible within minutes, unproductive less excess capacities are no longer needed. More sophisticated scenarios are also possible, for example moving server resources on the basis of schedules (for daytime and nighttime operation or for working days and the weekend) or in response to sudden changes in workload. These and other scenarios are very simple to handle thanks to a dynamic server infrastructure with PRIMERGY blade servers, ServerView VIOM and ServerView Resource Coordinator VE.

The primary benefits are: The degree of system utilization is further improved, resources can be made available to new services at short notice, and it is possible to respond quickly to peak loads.

Dynamic server infrastructure

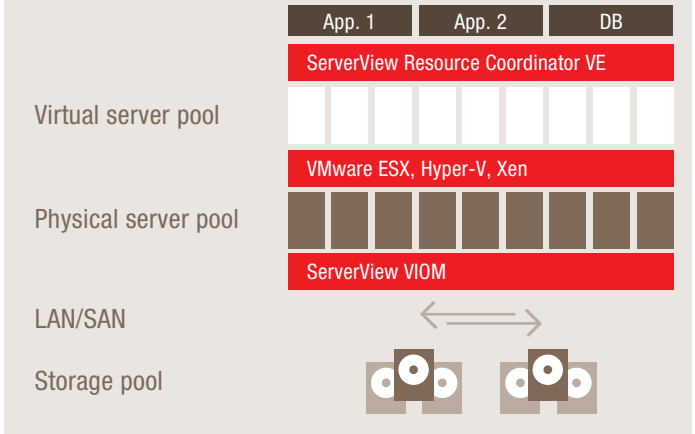

Summary

Service oriented IT infrastructure (SOI)

In coming years, SOAs will define the development of IT at the application level. The idea is to replace large, monolithic applications with their comprehensive functionality by a series of application modules with specific features. This is intended to speed up the implementation of changes in a business process at the application level – simply by adding or replacing application modules with specific features like building blocks. Each application module can be shared by several processes.

Of course, such a radical change to an application landscape also impacts the underlying IT infrastructure. Changes at the application level will increase and so it will become more and more difficult to determine the right server capacities. That is why a counterpart of a SOA is required – an SOI. A dynamic server infrastructure with its pool architecture is ideal for this. This is because an SOI can be built on the same components, without the need for any changes to the architecture. Above all, this is where the N:1 functionality pays off since, particularly in a SOA, the failure of an application module may have a direct impact on one or possibly several business processes.

The benefits: With an SOI, new application modules can be rolled out in a very short time and high availability ensured efficiently for business processes.

Pool architectures offer greater adaptability, higher efficiency and better quality of service.

Pool architectures are ideal for consolidation, dynamic IT and business processes based on SOA environments.

Fujitsu's dynamic server infrastructure offers a platform for the optimal operation of server pools.

Our dynamic server infrastructure consists of seamlessly integrated components: PRIMERGY servers, ServerView VIOM, hypervisor products from leading vendors and ServerView Resource Coordinator VE.

The crucial thing is the fact that we separate resource management cleanly from the underlying technologies and provide it as an end-to-end tool. That enables efficient and flexible operation, without the need to worry about the growing complexity of virtualization technologies.

ts.fujitsu.com

Published by

Fujitsu Technology Solutions

Mies-van-der-Rohe-Strasse 8, 80807 Munich, Germany

Printed in Germany
Order No.: 10722-3-0409-EN
Contact: ts.fujitsu.com/contact

For more information see ts.fujitsu.com