White paper
ServerView Resource Orchestrator Virtual Edition

Simplifying day-to-day server management operations in a mixed physical and virtual it environment

Content

Introduction

Coping with high complexity in managing consolidated server infrastructures

For many IT organizations blade server environment, in combination with virtualization technologies on various levels, is the IT infrastructure of choice to achieve considerable consolidation benefits. It also is the ideal basis for ensuring more flexible usage of server resources. However, managing the increasingly complex underlying technology stack presents problems that put enormous pressure on IT administration staff. The purpose of this white paper is to provide data center managers, IT architects and IT administrators with information about how ServerView Resource Orchestrator Virtual Edition (ROR VE) can reduce the complexity of day-to-day server management operations in mixed physical and virtual IT environments.

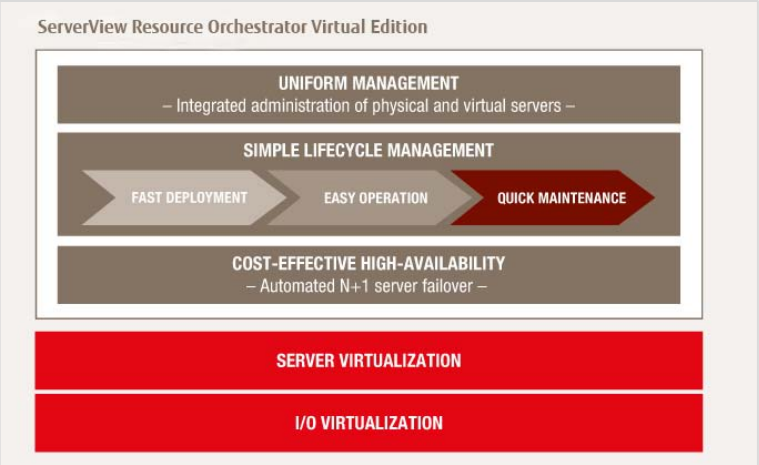

Fig 1: ROR VE key functional areas

How ServerView Resource Orchestrator Virtual Edition addresses operational server management challenges

Uniform and centralized management of physical and virtual servers

Mixed operation of physical and virtual servers will become a long-term reality in most data centers

Server virtualization is a key enabler for building dynamic IT infrastructures and many experts share the vision of a completely virtualized IT infrastructure. But despite the rapid growth rates of server virtualization products, the percentage of x86 server workloads running in virtual machines today is still low. This means that mixed operation of physical and virtual servers will be a long-term reality in most data centers. As a result, there will be a need for management tools that can handle the management of both physical and virtual environments. Fujitsu strongly believes that IT teams can work more efficiently if virtual server and physical server operations can be unified.

Coping with multi-hypervisor environments

The situation is even more critical for users who use products from more than one hypervisor vendor. IT organizations often start with one virtualization platform (e.g. VMware) and then expand their VM footprint using a second or even third type of hypervisor technology. For example, Microsoft offers an attractive product for server virtualization with its Hyper-V platform which is integrated in the operating system. Another example is Citrix, which is selected by customers because of its history and strong performance in desktop virtualization. These users generally implement the XenServer hypervisor in parallel to VMware.

Operating a multi-hypervisor environment is risky for IT organizations because they may get trapped in »siloed« virtualized pools. Taking advantage of each hypervisor's management tool would also mean accepting an even more complex management scenario. The approach here would be to use a standardized hypervisor- aware management tool for common basic daily operations and then to drop down to hypervisor–vendor management tools only for hypervisor-specific functionalities. This would prevent users from getting locked into specific hypervisor management tool sets. They would profit from the flexibility of a heterogeneous management tool, and basic management tasks could be performed independently of the underlying hypervisor technology.

Bringing together the management of physical and virtual server environments

ROR VE addresses the management challenges in a heterogeneous physical and virtual server environment by integrating the administration of physical and virtual servers as much as possible. Using a single console, ROR VE provides visualization of server resource states and common basic operations on physical and virtual servers. In addition to supporting physical server systems (Fujitsu PRIMERGY, PRIMEQUEST, SPARC Enterprise and selected third-party x86 servers), ROR VE also supports virtual servers from all major hypervisor vendors. For managing mixed hypervisor environments, ROR VE has a uniform interface for common administrative tasks like start, stop, reboot or live migration operations for all supported hypervisor products. Virtual resources and their relationships to physical systems are shown on a single console. In PRIMERGY blade server environments this includes the monitoring of physical and virtual network connections, which enables server administrators with limited network skills to easily confirm the network configuration. There is no longer any need to launch the management consoles of various hypervisor products for basic daily operations. However, for more advanced and complex management tasks, ROR VE also allows for integration with external management consoles. For details on specific hypervisor support, please see the support matrix in the ROR VE data sheet.

In addition to the traditional tree-based management console view, ROR VE features an innovative and simplified user interface especially designed for blade server management. Blade servers, along with their resource states, are realistically depicted as cabinet systems; ROR VE lets administrators manage a blade system remotely, but gives them the visual impression of standing directly in front of the server. Monitoring, startup, shutdown or reboot operations for physical and virtual servers can easily be performed directly using the management view, thus allowing even inexperienced staff to perform basic daily operations.

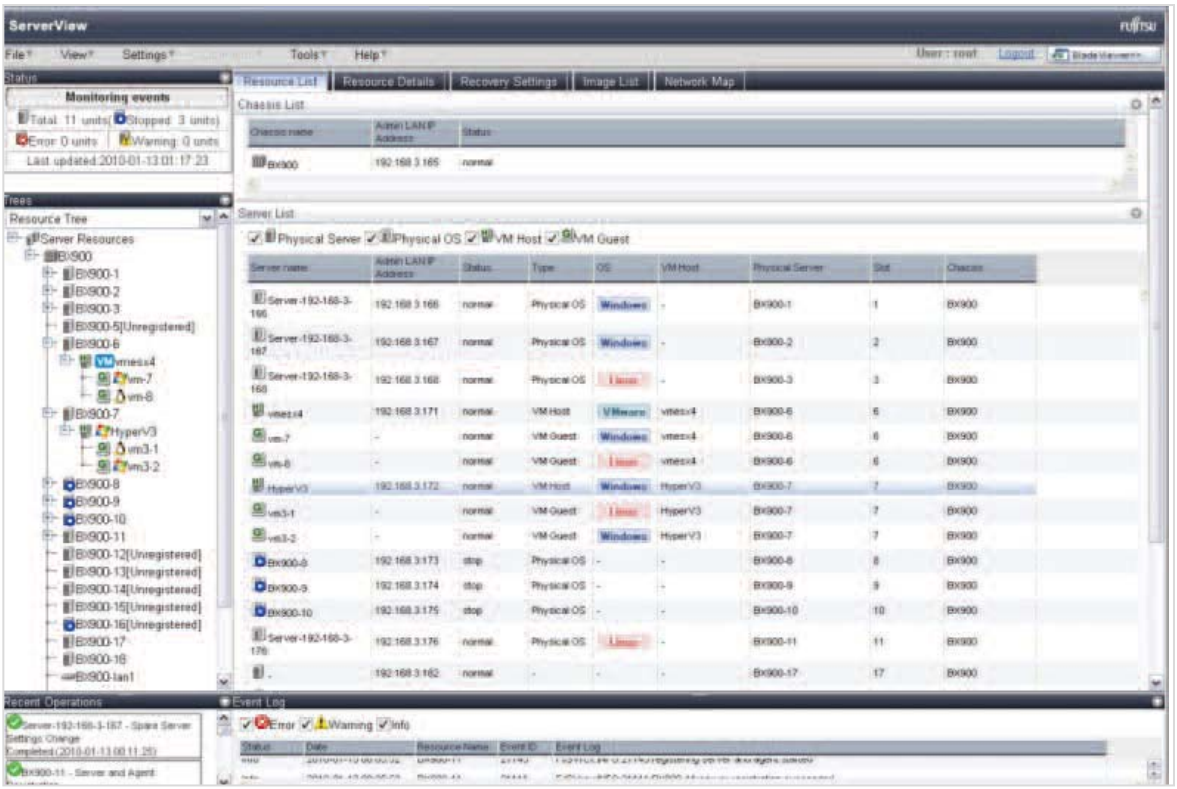

Fig 2: ROR VE advanced GUI - integrated management view on physical and virtual resources

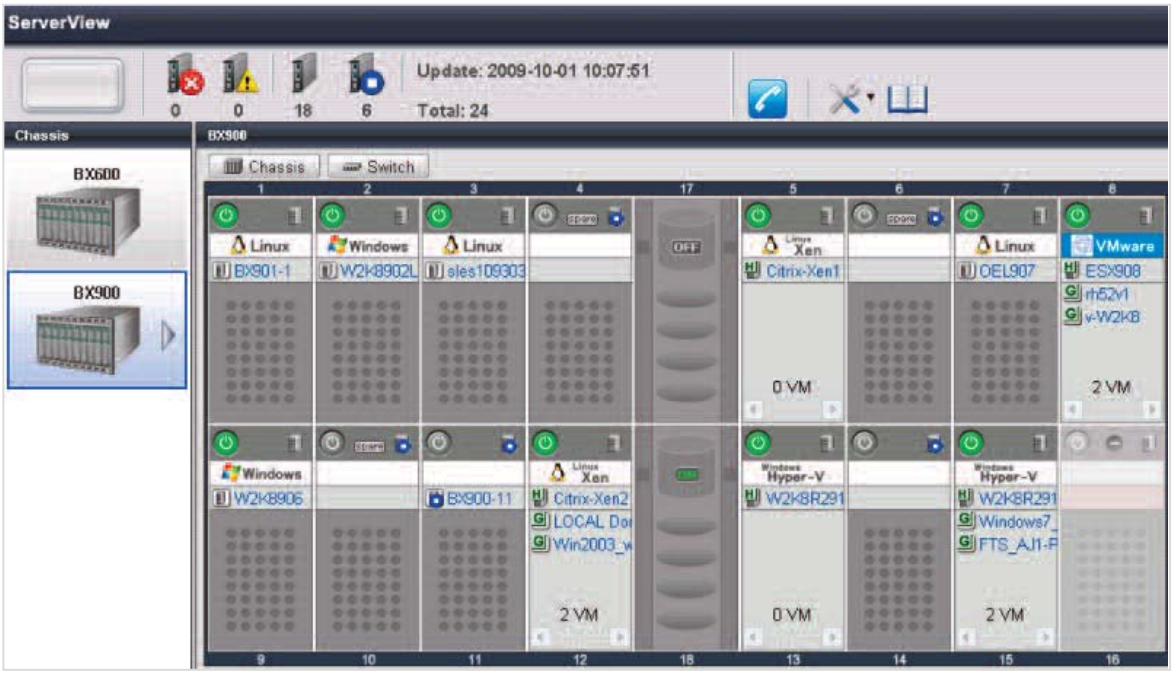

Fig 3: ROR VE BladeViewer GUI – optimized for easy daily blade server operation

Simple lifecycle management

IT departments are under constant pressure from business units that demand quick adaptation to changing business requirements. However, quick reaction is often not possible because many of the steps required to set up new servers or to make changes in the existing IT infrastructure still entail numerous manual administration tasks. Doing things manually, especially when managing a large number of resources, limits reaction times and wastes the valuable time of skilled administrators repeating the same tasks over and over again. It also poses a high risk of making mistakes. This is particularly true in light of the ever growing complexity of the technology stack admin teams have to manage today. The following section explains how ROR VE assists administrators in automating many of their daily management tasks. As a result, admin teams become more productive and more responsive. The same admin team can do more and react faster to changing business needs.

Automated server deployment via server cloning

ROR VE has an integrated server cloning feature that enables distribution of a cloned image (previously collected from the system disk of a reference server and stored on the admin server) to one or more other physical servers. When a cloned image is created, network-specific configuration settings such as host names and IP addresses are removed from the cloned image. The network-specific configuration is then dynamically re-configured on the servers to which the cloned image is distributed. Server cloning increases the productivity of administrators when adding new servers to a server farm. Compared to manual deployment processes, cloning significantly reduces the time required for installing the same operating system and software on multiple servers. As a result, administrators can now respond much faster to changing business requirements. In addition, software maintenance for multiple managed servers can also be performed more quickly and easily. Cloning servers allows users to install patches and software, or to modify installed software on a managed server before propagating those changes to other servers. After performing required maintenance tasks on one managed server, a cloned image can be collected from that server and deployed to other servers. This minimizes the time required for software maintenance of multiple managed servers while preventing operational mistakes.

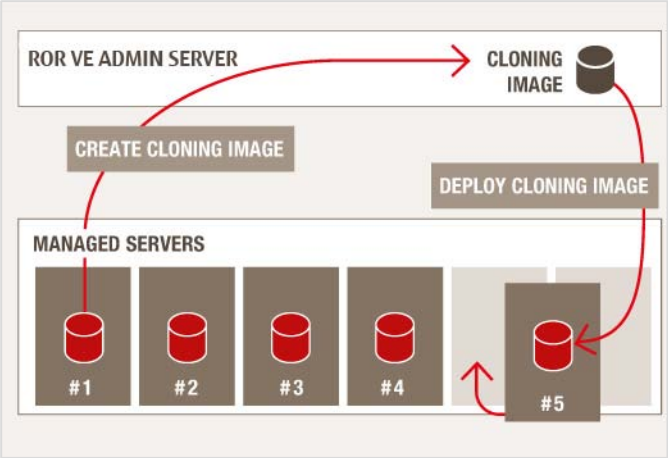

Fig 4: Integrated cloning engine

Automated initial setup or modification of a ROR VE configuration using the system pre-configuration tool

To assist administrators in creating or modifying a ROR VE configuration, ROR VE has a »System Pre-configuration« tool which supports the automated setup of multiple resources in a batch operation by simply importing a system definition file containing all the required settings and information of a ROR VE configuration. The system definition file (csv-format) can be produced either from an Excel design sheet or exported from a preconfigured system. Administrators thus have a powerful tool that allows them to easily design or modify an ROR VE configuration entirely offline at their desks. This reduces the setup time and minimizes on-site implementation work. In addition, for easier and secure system design, the Excel design sheet has improved usability features such as color-coded display of required fields and input validation, which helps prevent configuration errors at the earliest possible stage.

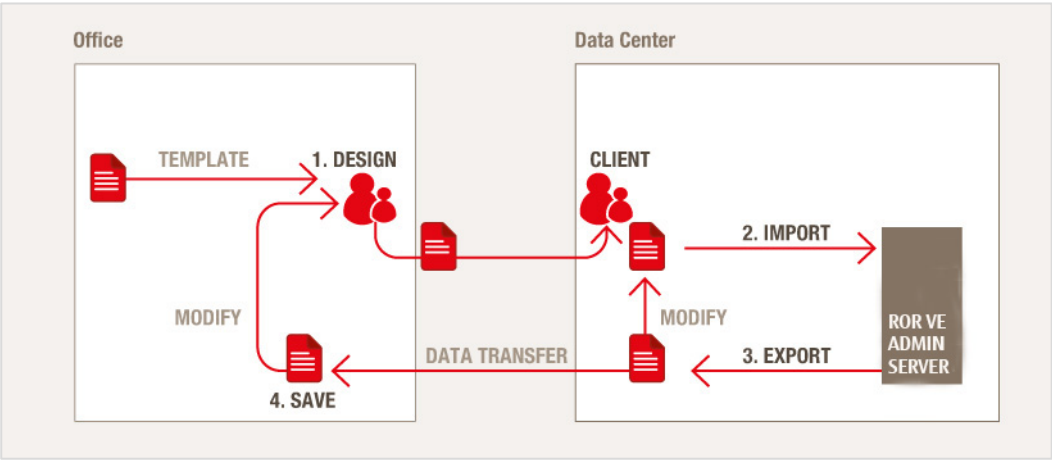

Fig 5: Setup and modification process of a ROR VE configuration

Resource monitoring

ROR VE can centrally monitor the configuration and status of servers or other managed resources directly from the console. This enables the identification of resources experiencing problems and thus reduces the time spent on system maintenance. Moreover, ROR VE can easily launch external management software to precisely locate faulty parts within a managed resource.

Monitoring is based on the following three components:

- Resources – The resource tree displays blade chassis, servers, LAN switches, power supplies (PDU or USV), physical OSes, VM hosts and their VM guests. If a hardware problem occurs on a server, the affected guest operating systems can easily be detected.
- Events – ROR VE displays events such as hardware failures, server switchover operations triggered by hardware failures and the results of every performed operation.
- Recent Operations – ROR VE displays the progress status of the various operations performed on resources.

Easy network monitoring with the integrated network map

For PRIMERGY blade servers, ROR VE provides a network map function which helps visualize and correlate physical and virtual networks. It also helps server administrators verify the network configuration. The comprehensive information can also be used as a basis for communication between server and network administrators to support them in coordinating their work.

The network map displays the following information:

- Network configuration of physical and virtual servers (including virtual switches and VM guests)
- Statuses of network links between all resources
- VLAN configuration affecting each physical and virtual server

Two different maps are available within the network map. A comprehensive map displays chassis, servers and their connections (network links) with adjacent LAN switches. A local map shows more detailed information about a selected resource (chassis or server). Up to two chassis can be displayed at a time.

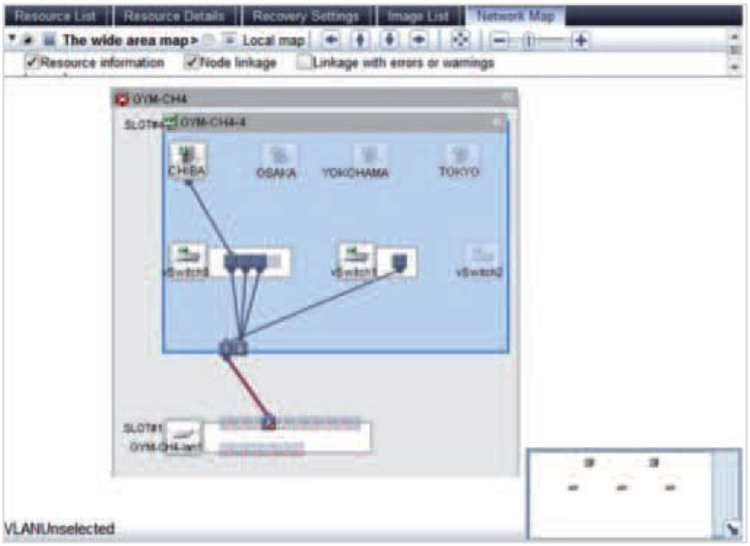

Fig 6: Integrated network map – helps server administrators to verify the network configuration

Monitoring power consumption

With the power monitoring feature, administrators can track the power consumption of their server infrastructure. The power consumption data can be displayed in a data graph (line or plot) or exported to a file in csv-format for further offline analysis.

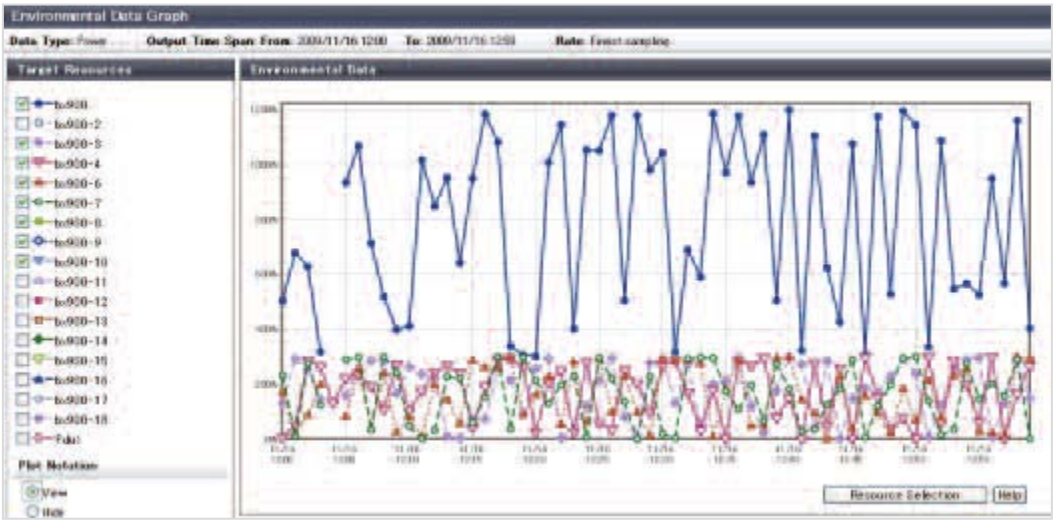

Fig 7: Tracking of power consumption trends in a PRIMERGY BX900 blade server environment

Streamlined server maintenance

Depending on the individual customer infrastructure, ROR VE offers multiple options for supporting maintenance scenarios.

- **VM guest live migration**

Thanks to the integration with hypervisor management software and VM hosts, ROR VE provides a function for triggering a live migration of VM guests between physical servers directly from the ROR VE console, or alternatively from a command line interface. ROR VE currently supports live migration the following x86 hypervisor platforms: VMware vSphere, Microsoft Hyper-V, Citrix XenServer, RedHat Xen/KVM and Oracle VM. Live migration supports moving virtual servers between physical servers without downtime. It helps reducing planned downtime for maintenance scenarios, and it is also a core technology for more advanced management functionalities like dynamic load balancing. For example, ROR VE uses VM live migration to trigger migration events in response to critical warning messages. This enables IT administrators to implement a pre-failure detection mechanism, which allows starting corrective actions before a real error occurs. In addition to keeping applications alive during maintenance work, VM live migration in combination with the remote power on/off capabilities of ROR VE can also be used to reduce overall power consumption. Temporarily relocating VM guests to a subset of physical servers makes it possible to shutdown any unused servers or even complete blade chassis in times of lower workload.

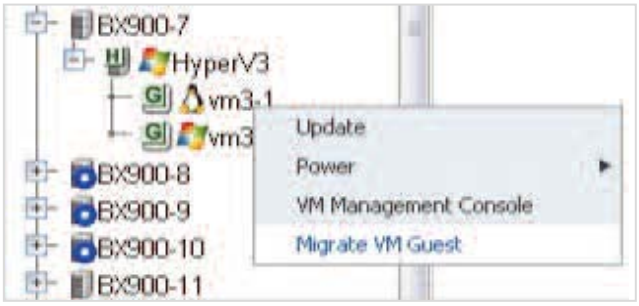

Fig 8: Trigger of VM live migrations directly from the ROR VE management console

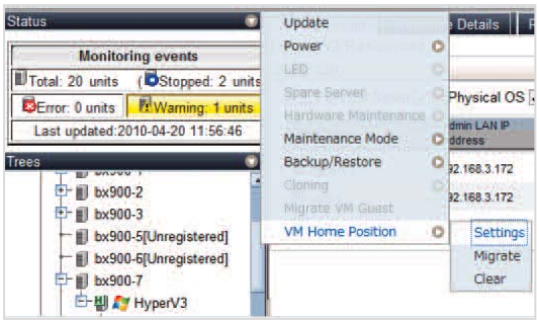

Fig 9: Setting the VM home position of all registered VMs

- **VM home positioning**

In all cases, where VM guests are temporarily migrated away from their original physical servers, ROR VE supports administrators to quickly and safely bring back VM guests to their original place. ROR VE stores the original position of each VM guest and at any time can migrate back all VM guests to their original place, at a push of a button.

- **Server switchover**

VM live migration is an excellent technology for operating a totally virtualized infrastructure. But as outlined before, operation of a mixed environment running virtualized and physical servers will be a long-term reality in most data centers. In such an environment VM live migration is only part of the whole scenario. For mixed environments ROR VE offers a much wider approach covering maintenance scenarios for physical servers running native operating systems as well as for virtual environments running hypervisor software. ROR VE is able to switchover complete VM hosts, including all running VMs, to a spare server. This facilitates server replacement scenarios in SAN-boot as well as local-boot configurations.

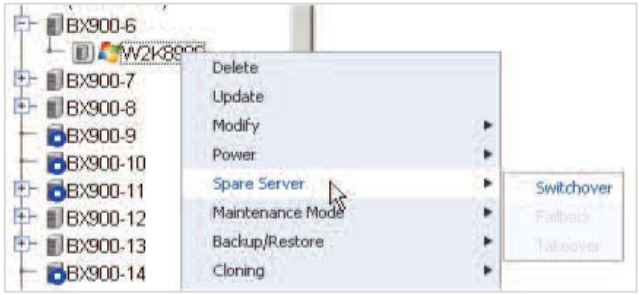

Fig 10: Initiate manual switch-over operation to a spare server

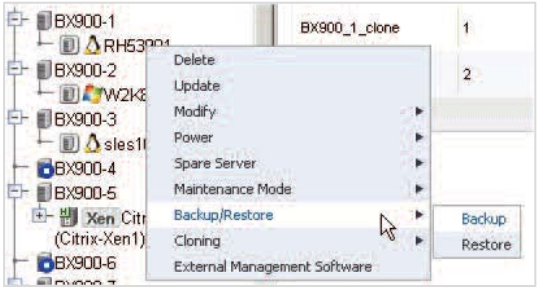

Fig 11: Backup of an OS image

- **Backup and restore of system images**

In local boot configurations, this function supports the backup and restoration of complete boot disk images including the operating system and possible further data stored on that disk. The images are centrally managed and stored on the disks of the admin servers. A system image backup should be created as a precautionary measure before performing software maintenance tasks such as installing patches and software, or modifying installed software. In addition, restoring a system image backup can be used to protect against server hardware failure in local boot scenarios.

Cost-effective high availability through auto-recovery of failed servers

In addition to using the server switchover functionality of ROR VE for planned downtime / maintenance scenarios, it can also be used to address unplanned downtime scenarios. By automatically triggering server switchover operations when a physical server fails, business applications can be resumed without any administrator intervention. Compared to manual recovery processes, server recovery time is reduced significantly, thus resulting in faster reaction to server failures. By leveraging this automated server switchover capability, ROR VE enables the implementation of a very cost-effective HA concept. The basic idea is to share one or a few spare servers between many production servers. This so-called N+1 high-availability concept can drastically reduce the number of required non-productive spare systems, thus leading to increased efficiency in server operations. Users see immediate tangible up-front cost reductions when planning to implement a HA solution. IT departments can now protect more servers without paying a premium for dedicated HA software. The following scenarios are examples of possible usages. The HA concept can add value in pure physical environments and in mixed physical and virtual environments. It can even increase service levels in environments already using cluster software.

Scenario 1: Providing HA for a pure physical server environment

This usage scenario addresses customers looking for more cost-effective solutions for implementing high-availability. ROR VE provides a less complex and less expensive alternative to traditional cluster solutions.

- ROR VE has no dependency on cluster-aware applications or specific OS types. Therefore, ROR VE can protect every application running in mixed Windows and Linux environments against hardware failures with one uniform HA solution.
- Cluster solutions still rely heavily on 1:1 server configurations – one spare server for every productive server. In a N:1 configuration with ROR VE only 1 spare server is needed for many productive servers.

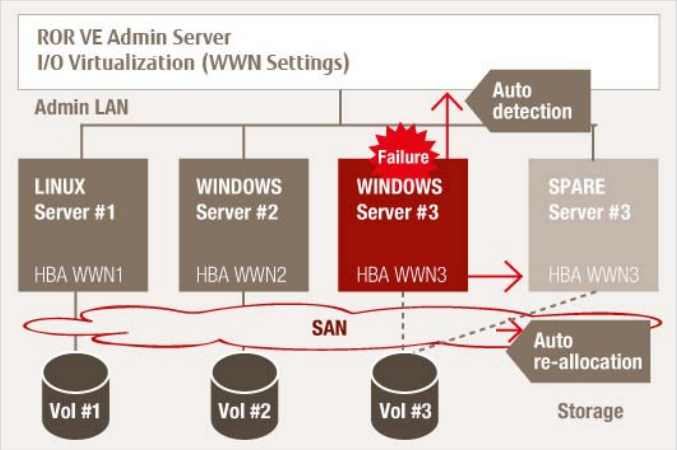

Fig 12: Automatic server failover including auto detection of server failure and auto re-allocation of storage volume

Scenario 2: Providing HA for a mixed physical and virtual server environment

This usage scenario addresses users operating a mixed physical and virtual server environment and who are looking for a HA solution with a uniform HA concept that covers both environments. Nearly all hypervisor vendors offer high-availability solutions for their own virtualized infrastructure. As long as server operations remain on a dedicated hypervisor platform, this approach would be sufficient. However, as outlined before, the factors that apply to server switchover for planned downtime scenarios are the same when it comes to unplanned downtime scenarios.

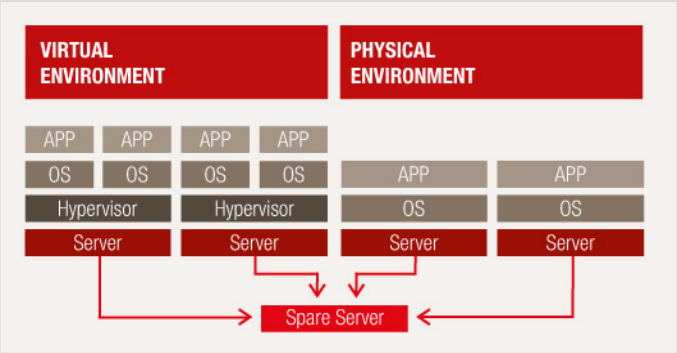

Fig 13: Sharing spare servers in mixed virtual and physical environments

A high-availability concept should operate on a broader scale, covering not only dedicated hypervisor islands, but also physical servers in combination with mixed hypervisor environments. The HA concept of ROR VE functions independently from the underlying hypervisor technology, thus offering a uniform HA solution that eliminates the need to manage several different high-availability solutions. Moreover, it reduces investments in spare server hardware because fewer spare servers are needed to protect a heterogeneous server environment. For restrictions on specific OS and hypervisor combinations please see the support matrix.

- **Sharing a spare server between physical OSes and VM guests**
  When using I/O virtualization technology, ROR VE supports the sharing of spare servers between servers with physical OSes and servers with VM hosts running VM guests by combining its own spare server functionality with the high-availability features available in each server virtualization product.

  For restrictions on specific OS and hypervisor combinations please see the support matrix.

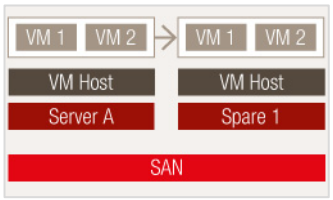

Fig 14: Case 1

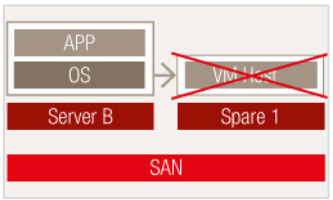

Fig 15: Case 2

  A server failure will trigger the following recovery operations:
  – Case 1:
  If the failed server was a VM host running VM guests, the high-availability feature provided with the server virtualization software will transfer the VM guests to the spare VM host server.
  – Case 2:
  If the failed server was running a physical OS, ROR VE will shut down the VM host on the spare server and switch the boot disk of the failed server over to the spare server.

- **Automatic recovery of a failed physical server in a virtual server farm**
  Another possible usage of ROR VE combined with the HA solutions from hypervisor vendors is ROR VEs ability to automatically replace a failed physical server in a virtual server farm. In the event of a physical server failure, the HA solutions of the hypervisor vendors ensure the availability of their virtual infrastructure, thus providing failover of single VMs to the remaining physical hosts within the virtual server farm. But repairing the failed physical system would still be a manual process. Depending on the readiness and availability of technical experts, this might take some time. And if there is not enough reserve capacity to cope with the additional load on the remaining VM hosts, it may not be possible to guarantee existing service levels. This is where ROR VE comes in. ROR VE can automatically replace a failed physical server in a virtual server farm. Restoring the original number of physical servers in the virtual server farm is much faster and results in higher service levels. Subsequent activation of load balancing tools from the hypervisor vendors (for example VMware DRS) would then take care of rebalancing the VMs across the restored server farm.

Underlying key technologies

SAN-boot technology

Booting a server's operating system from local disks has a long tradition and worked well in data centers for many years. But recent advances in SAN-boot technology offer a compelling alternative. SAN-boot technology allows any server on the network to boot from a shared Fibre Channel (FC) storage system located in the SAN. For example, separating the boot image from the server helps administrators in cases where servers need to be taken offline or replaced. They only need to connect a compatible spare server to the SAN and direct its BIOS to boot from the appropriate storage unit. In this way the process of connecting a spare server can be highly automated. As a result, this approach can dramatically decrease server downtime as there is no need to reinstall the OS and application software to make the spare server functional. ROR VE uses SAN-boot technology in combination with I/O virtualization technology to implement a very cost-effective high-availability solution. Moreover, many of today's advanced features of hyper-visor software like VM live migration also require SAN-boot technology.

Some of the key benefits of booting servers from SAN-based storage versus local disks include:

- Simplified storage administration – booting from a single FC storage unit on the SAN rather than individual servers reduces the management overhead of a distributed server environment and helps centralize storage management.
- Improved availability – booting from SAN leverages the advanced availability features of SAN storage systems including multi-path access, fault tolerance, battery backup, etc.
- Enhanced data protection – data protection features inherent in SANs (such as backup and restore, data migration, data replication and disk capacity expansion) can be utilized for the boot drive at no additional cost.

I/O virtualization

Each server I/O adapter comes with a pre-assigned default physical network address (WWN or MAC) used to identify servers on SAN or LAN networks. This makes it necessary to reconfigure network settings when performing server deployment, maintenance and recovery tasks. Reconfiguration of the network settings normally involves many manual steps like re-cabling, and it requires the assistance of network and storage administrators. Depending on the coordination and cooperation between the server admin and network admin teams, a network reconfiguration process can considerably slow down server admin tasks.

I/O virtualization technology puts an abstraction layer between the servers and the external networks. Instead of using default Media Access Control (MAC) addresses for Network Interface Controllers (NICs) and default World Wide Names (WWNs) for Host Bus Adapters (HBAs), virtual customer-defined addresses are used because they can be flexibly assigned to specific server blades. Although the hardware ships with default MAC addresses and WWNs, I/O virtualization resets these addresses prior to the boot sequence so that only the virtual addresses are used during operation.

Simply by changing server-side settings, I/O virtualization lets server administrators add, maintain or failover servers without affecting complex network settings, without re-cabling and without requiring the assistance of multiple administrators.

There are two ways to implement I/O virtualization within ROR VE:

- Using the HBA Address Rename Service (HARS)
- Using the ServerView Virtual I/O Manager (VIOM)

HBA Address Rename Service (HARS)

ROR VE offers an integrated basic functionality for the virtualization of SAN I/O addresses called HBA Address Rename Service. The service uses a set of virtual WWNs or Alternative World Wide Numbers (AWWNs) to overwrite the default WWN addresses on the host bus adapters in order to maintain the WWN settings in server replacement or failover scenarios. HARS offers SAN I/O address virtualization for PRIMERGY blade and rack servers as well as for third-party rack servers

ServerView Virtual I/O Manager (VIOM)

VIOM is an I/O virtualization technology specifically designed for blade server environments. VIOM uses a special server profile to store blade-specific information like I/O addresses, I/O connections or even boot parameters. Each server profile is stored in a central repository and serves as a pointer to a boot image, enabling administrators to run an application by simply assigning a profile to a server blade. As each profile can be moved not only within the same blade chassis, but also across hundreds of server blades residing in different chassis, VIOM paves the way for substantial improvements in scalability and flexibility. In contrast to HARS, VIOM can also virtualize MAC addresses used in LAN networks. This is helpful in applications using MAC addresses to generate license keys. Those applications will no longer run if they are moved to a blade with a different MAC address. Moreover, MAC address virtualization keeps IP configurations constant in environments

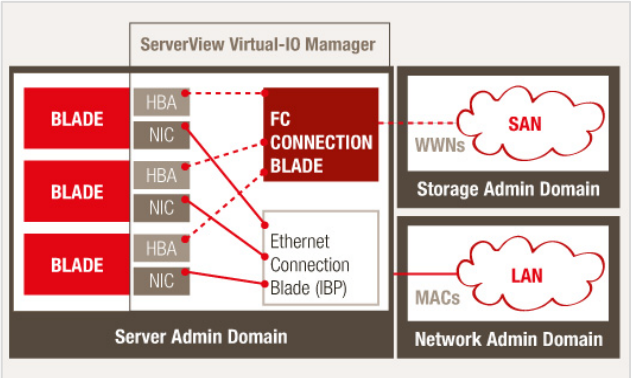

Fig 16: VIOM decouples server side network administration from external network administration

using DHCP servers, eliminating adaptation of DHCP server configurations. In combination with the Intelligent Blade Panel (IBP), a special I/O connection blade, it is not only possible to keep MAC addresses constant, but also to keep the network path to a specific uplink. This is particularly useful in configurations where many customers residing in separate physical or virtual networks should be operated within the same blade chassis. If a blade is moved or exchanged, IBP makes sure that the new blade will be connected to the correct customer network segment. The VIOM approach goes far beyond virtualization of WWN and MAC addresses of a server. In addition to WWN and MAC addresses, a VIOM server profile can contain other parameters like BIOS settings, also making it possible to keep those settings constant in server switchover or failover scenarios. The ultimate goal is to virtualize as many server parameters as possible to have a stateless blade, offering maximum flexibility for assigning server resources to applications. Beyond streamlining the management of I/O addresses, using VIOM within a blade server infrastructure has further advantages. For more detailed information about ServerView Virtual I/O Manager please see:
http://www.fujitsu.com/fts/viom

Server switchover

Server switchover is a feature that allows administrators to switchover applications from a primary server to a predefined spare server when the primary server fails or needs to be shut down for maintenance. This switchover is based on the so-called »cold-standby« method, in which the spare server is booted during switchover. Server switchover can be triggered manually by admin, or it can be triggered by the auto-recovery function, which can automatically switchover applications to a spare server when a hardware failure is detected on the primary server. Different switchover options are available depending on the boot configuration of each server. ROR VE supports servers booting from networked storage systems attached via Fibre Channel (SAN boot) as well as servers booting from local disks (local boot). For servers operating in SAN-boot environments, ROR VE uses I/O virtualization technology to switchover network addresses. For servers booting from local disks, ROR VE uses integrated operating system image management to restore a previous backup image on a spare server.

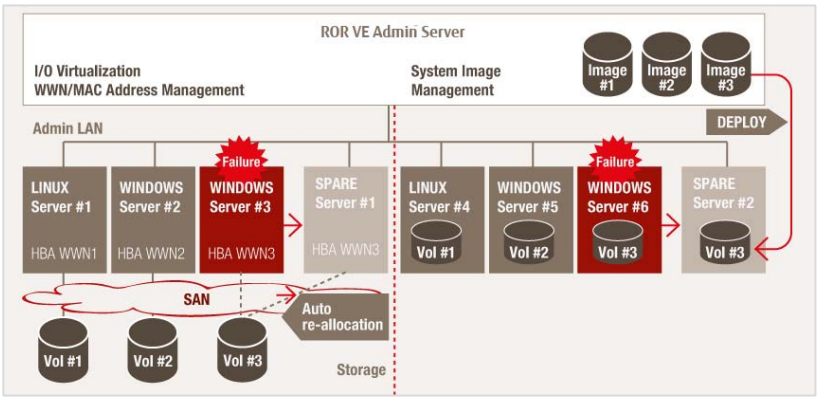

Fig 17: Server switchover in a SAN and local boot environment

Server switchover in SAN boot environments using I/O virtualization technology

Figure 18 shows the typical steps in a server switchover operation using the HBA rename service (HARS) functionality:

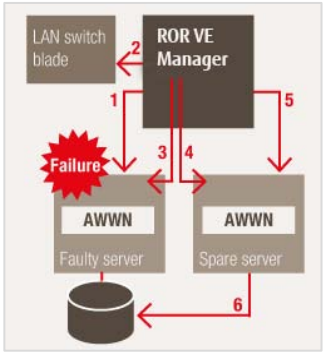

1. If the faulty server is still active (e.g. after manual switching), the ROR VE manager sends a shutdown request to it.
2. The ROR VE manager selects a spare server and changes the related VLAN setting(if required).
3. The ROR VE manager sends a request to delete the current Alternative World Wide Number (AWWN) setting on the faulty server.
4. The ROR VE manager sends a request to assign the same AWWN to the spare server.
5. The ROR VE manager powers up the spare server.
6. The spare server will boot from the original boot disk.

Fig 18: Server switchover using HARS

Figure 19 shows the typical steps in a server switchover operation using VIOM:

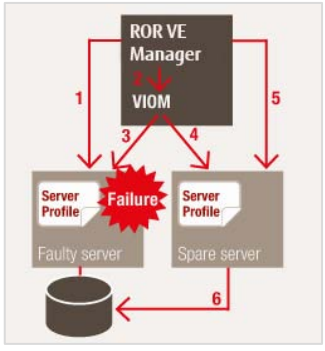

1. If the faulty server is still active (e.g. after manual switching), the ROR VE manager sends a shutdown request to it.
2. The ROR VE manager selects a spare server and sends a request to VIOM to initiate a change of the server profile.
3. VIOM un-assigns the server profile from the faulty server.
4. VIOM assigns the server profile to the spare server.
5. The ROR VE manager powers up the spare server.
6. The spare server will boot from the original boot disk.

Fig 19: Server switchover using VIOM

Server switchover in local boot environments using OS image backups

In a local boot environment, ROR VE restores a system image backup to a spare server which is then automatically booted. After switchover, the operating system and its applications will resume on the spare server from the state they were in at the last system image backup. Note that only the content of the first local disk (or boot disk) as seen by the BIOS of the managed server is subject to a backup or restore operation, including all partitions (Windows drives or Linux partitions) present on the boot disk. However, since additional disks (used as data disks), are not subject to backup and restore, their content cannot be made accessible to the spare server with a server switchover. When using more than one local disk, the backup and restore of such additional data disks should be performed using external backup software.

Figure 20 shows the typical steps in a server switchover operation using backup and restore of system images:

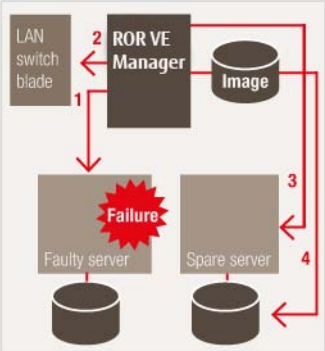

1. If the faulty server is still active (e.g. after manual switching), the ROR VE manager sends a shutdown request to it.
2. The ROR VE manager selects a spare server and changes the related VLAN setting (if required).
3. The ROR VE manager powers on the spare server, and sends a request to the spare server to initiate the image restoration.
4. The spare server starts restoration of a previously prepared system image.

After successful completion of the restore operation, the spare server is booted with the restored system image.

Fig 20: Server switchover using restore of a system image

Post switchover operations

Regardless of what action triggered a switchover, the user must decide whether to switch applications back to their original server (failback), or let the spare server take over those applications indefinitely (takeover). Choosing takeover will result in the spare server becoming the new active server.

Server switchover based on backup and restore takes approximately three minutes, plus the time required to restore the system image. Image restoration time depends on different factors such as disk space and network usage, but as an estimate, a disk of 73 GB will require 30 to 40 minutes (the transfer of the system images takes between 10 to 20 minutes, while system restarts and other configuration changes take another 20 minutes).

Server switchover based on I/O virtualization takes approximately five minutes, plus the time required to start up the operating system and services on the spare server. If a VM host was running on the spare server, the time required to shutdown the spare server must also be included.

Key solution components

Managed Servers

As »managed servers« running the production workloads or being used as spare servers, ROR VE supports Fujitsu PRIMERGY, PRIMEQUEST and SPARC Enterprise servers as well as selected third-party x86 servers. Depending on the operating system, those managed servers run ROR VE agents that are used to control and monitor server operations. For details on specific platform support please see the ROR VE support matrix and data sheet.

Storage

In general, there are two options for where the managed server's storage resides, either on a SAN or on local disk storage. SAN storage could be an ETERNUS array or any other certified SAN storage system. The other option is to use the local disks on the managed servers. Fujitsu recommends using SAN storage because ROR VE operation is optimized for use in SAN environments. In the special case of blade server environments in combination with ServerView VIOM, there is also the option of connecting to a storage system attached via iSCSI.

Admin server

ROR VE configurations need a dedicated admin server as the central management station. The admin server runs the ROR VE manager software.

I/O virtualization

When using VIOM for I/O virtualization, the VIOM manager has to be installed on the same server with the ROR VE manager. When using an optional HBA address rename service as a backup for WWN assignment in case of a ROR VE manager failure, this service has to be installed on a different system (server or PC).

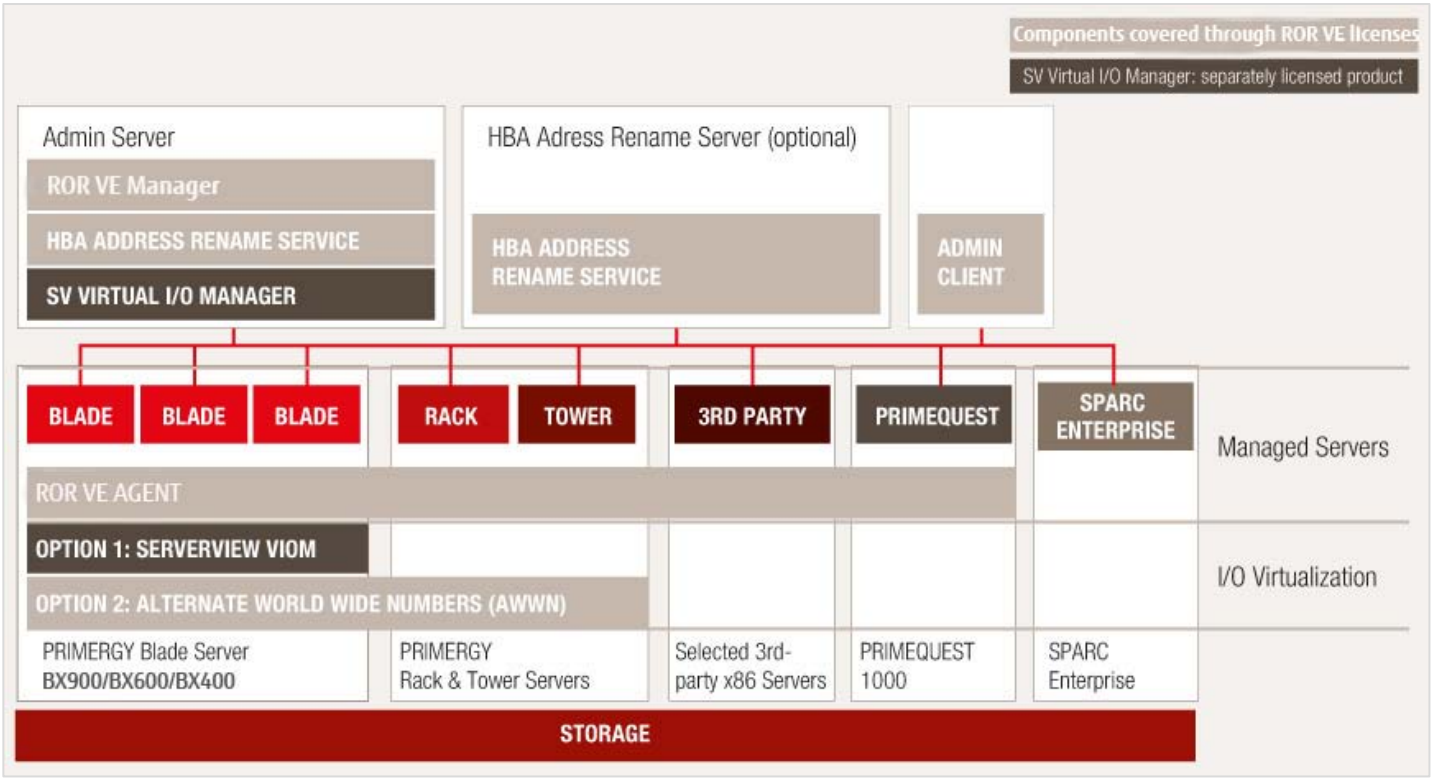

Summary

More and more enterprises are responding to the current economic environment with an increased investment focus on virtualization and automation technology that promises cost efficiencies. On the hardware level, IT organizations increasingly choose pooled blade server architectures as an ideal foundation for building a future proof IT infrastructure.

But despite the high expectations for more efficiency and agility generated by such a new technology stack, the top priority is still business continuity. Users are no longer willing to spend a fortune on high-availability solutions. They want to benefit from the design of a pooled server architecture enabling implementation of new HA solutions that are far more cost-effective than traditional HA approaches.

Fujitsu's BX900 blade server platform is ideal for building such a next-generation IT infrastructure. The design of the BX900 is especially optimized for use in virtualized environments supporting all major hypervisor vendors. I/O virtualization with ServerView VIOM further streamlines the management of I/O connections. And finally, ROR VE provides optimized management of the complete technology stack.

Through automation of server lifecycle operations covering both physical and virtual servers, ROR VE enables more efficient use of server resources. By consolidating portions of the physical and virtual server management tasks, ROR VE makes daily operations much simpler and more effective, providing answers to many IT managers looking for ways to manage the increasingly more complex underlying technology stack. At the same time, ROR VE offers a cost-effective way to implement high-availability for complete server environments, not only for selected mission-critical applications, but for all systems covering a complete business process.

More information

For more information about ROR VE and related technologies discussed in this white paper, please visit the following sites:
http://www.fujitsu.com/fts/ror-ve
http://www.fujitsu.com/fts/viom
http://www.fujitsu.com/fts/products/computing/servers/primergy/blades/index.html

Contact
Fujitsu Technology Solutions
Website: www.fujitsu.com/fts
2012-04-24 ww en