

Protection of Business-Critical Applications in SUSE Linux Enterprise Environments Virtualized with VMware vSphere 4 and SAP® NetWeaver® as an Example

Version 1.1
Dresden, May 13, 2011



*CC Computersysteme und
Kommunikationstechnik GmbH*

Table of Contents

1	Introduction.....	5
2	Foundations for protecting business-critical applications	6
2.1	Infrastructure Reliability.....	6
2.2	Hardware Availability	6
2.3	Logical Fault Protection	7
2.4	Application Availability.....	7
2.4.1	Monitoring.....	7
2.4.2	Data Security	7
2.4.3	High Availability	7
2.4.4	Disaster Recovery	7
3	Increasing application availability in traditional UNIX environments	8
3.1	MC/ServiceGuard (HP-UX).....	8
3.1.1	Configuration types for the MC/ServiceGuard cluster	9
3.1.2	Quorum Server of the MC/ServiceGuard Cluster.....	9
3.2	Oracle Solaris Cluster.....	9
3.3	Veritas Cluster Server (VCS)	10
4	Virtualization and Partitioning on Non-x86 Systems	12
5	VMware vSphere Functions for Increasing Availability.....	13
5.1	Important basic functions of VMware technology.....	13
5.2	VMware vCenter Server Heartbeat	14
5.3	VMware High Availability	14
5.4	VMware vMotion	15
5.5	VMware Storage vMotion	17
5.6	VMware Fault Tolerance	18
5.7	VMware Site Recovery Manager	20
6	SUSE Linux Enterprise High Availability Extension	21
6.1	SUSE Linux Enterprise Server	21
6.2	SUSE Linux Enterprise Server for SAP Applications	21
6.3	SUSE Linux Enterprise High Availability Extension	22
6.4	SUSE Linux Enterprise High Availability Extension in a VMware vSphere Environment.....	24
7	SAP NetWeaver	26
7.1	Structure of SAP NetWeaver	26
7.2	Single Points of Failure (SPoF) of SAP NetWeaver	27

8	Reasons for Running Business-Critical Applications in a Virtualized Environment.....	28
8.1	Consolidation.....	28
8.2	High Availability.....	28
8.3	Disaster Recovery.....	28
8.4	Opportunities and Risks.....	29
9	Increasing Application Availability of SAP NetWeaver by Deploying SUSE Linux Enterprise High Availability Extension in a VMware vSphere Environment.....	31
9.1	Application Stack of SAP NetWeaver.....	31
9.1.1	Influence of VMware vSphere Functions on the Application Stack.....	32
9.1.2	Integration of Application Protection into the Application Stack.....	33
9.1.3	Deploying High Availability Central Services.....	33
9.2	Selection Criteria for Application Protection Solutions.....	34
9.2.1	Assignment of Functionality to the Application Stack Layers.....	34
9.2.2	Monitoring the Application.....	34
9.2.3	Non-Conflict Between Tools Used.....	35
9.2.4	Integration into the Management Environment.....	35
10	Design of a High Availability SAP NetWeaver Platform on the Basis of VMware vSphere 4.....	37
10.1	Design Aspects.....	37
10.2	Options for Increasing Availability of SAP NetWeaver.....	38
10.2.1	Using VMware High Availability.....	38
10.2.2	Using VMware Fault Tolerance.....	38
10.2.3	Using the SUSE Linux Enterprise High Availability Extension.....	38
10.2.4	Using the VMware Site Recovery Manager.....	39
10.3	Example Configurations.....	39
10.3.1	Example 1: VMware High Availability and SUSE Linux Enterprise High Availability Extension with NFS.....	39
10.3.2	Example 2: VMware High Availability and SUSE Linux Enterprise High Availability Extension with DRBD.....	42
10.3.3	Example 3: VMware High Availability and SUSE Linux Enterprise High Availability Extension with Shared Disk.....	43
10.3.4	Example 4: Disaster Recovery Scenario.....	45
10.3.5	Example 5: Deploying a High Availability NFS Service as a Central Resource.....	48
10.4	Running Business-Critical Applications in High Availability Clusters.....	49

11	Conclusion and Review.....	50
11.1	Consolidation of the IT Environment	50
11.2	Increasing Availability.....	51
11.3	Options for Disaster Recovery.....	51
11.4	Operating the Virtualized System Environment.....	52
12	Appendix.....	53
12.1	Trademarks.....	53
12.2	References.....	53

1 Introduction

This white paper presents technologies and solutions for increasing application availability for Linux systems in virtualized environments. An x86 environment is presented that meets high availability requirements based on the *VMware vSphere 4* virtualization platform and the *SUSE Linux Enterprise Server* operating system as the foundation for running *SAP NetWeaver*. Various options for increasing application availability are presented and evaluated, and recommendations are made for building a high availability *SAP NetWeaver* system platform. The white paper is intended primarily for system architects and IT managers.

SAP NetWeaver has for many years been the market-leading platform for supporting almost all value-adding processes in companies. These applications can therefore definitely be referred to as business-critical. Client/server computing and increasing use of open operating systems got their break in corporate IT with the increasing proliferation of SAP® R/3 – the predecessor of SAP ERP – since the beginning of the 1990's. In particular, UNIX systems from various manufacturers superseded the monolithic systems that served as the basis for SAP® R/2 and other comparable software solutions. This triggered the first great wave of consolidation, with users increasingly freeing themselves from the grip of proprietary systems. The systems became more flexible and cost-effective.

The ever increasing proliferation of x86-based Linux systems as a platform for SAP NetWeaver and the concurrent breakthrough in virtualization technologies have for some time been pushing the paradigm change initiated by client/server computing. This benefits users again, as the selection of hardware supplier is not made based on the specific UNIX distribution. With Linux as the operating system, the entire range of available x86 hardware can be used. The use of virtualization solutions ensures a more efficient utilization of existing resources and increased flexibility in terms of extensions or short-term reductions. Virtualization can reduce not just power consumption but also space usage and heat dissipation. The available capacity of IT is used more efficiently overall, reducing costs and CO₂ emissions.


CC Computersysteme und Kommunikationstechnik GmbH, based in Dresden, has many years of experience in the high-level design, practical design, and operation of high availability SAP environments under Linux. The company has active partnerships with all relevant hardware and software manufacturers and maintains a lively exchange of information and experience with the SAP LinuxLab in St. Leon-Rot, Germany.

2 Foundations for protecting business-critical applications

Dispensing with IT support for a company's business processes is inconceivable in this day and age. The integration of electronic data processing in all kinds of mission-critical application areas has become a standard that poses very high demands on IT operations. For the user, the focus is on:

- High performance
- Availability
- Stability
- Security

Achieving the lowest possible procurement and operational costs is another critical factor. There are a range of components that meet these requirements irrespective of the platform, operating system, and applications being used. Successful implementation requires a combination of all the components, which are detailed below:



2.1 Infrastructure Reliability

The reliability of the individual infrastructure components, starting from the quality of cabling through to fire-protection solutions for the data center, plays an important role in the availability of the overall system. It is the de facto foundation for Business Continuity. The physical design and other aspects such as uninterrupted power supplies (UPSs) are important in deploying a reliable infrastructure. Special turnkey data center units are also employed sometimes.

2.2 Hardware Availability

The active components, from the network switch to the server and storage hardware, are also critical to a successful Business Continuity strategy. Redundancies provided by RAID systems and dual power supplies or multiple network adapters for each server are common measures that significantly increase hardware availability.

2.3 Logical Fault Protection

Logical fault protection is equally important even if all the other aspects detailed here have been implemented successfully. Snapshots and Continuous Data Protection (CDP) that allow rewinding to a defined previous state help here as well. Shadow databases can also be a valuable aid.

2.4 Application Availability

2.4.1 Monitoring

Comprehensive monitoring of all important system parameters identifies possible pain points or anticipated bottlenecks at an early stage. Regardless of whether a file system is approaching its capacity limit due to continuous writes or a network adapter is generating packet losses, early detection allows appropriate preventive measures to be taken, thereby increasing availability.

2.4.2 Data Security

Data security and consistency always deserve the greatest attention. Generating snapshots on a regular basis is helpful in addition to using RAID storage systems and an appropriate backup concept. Mirroring to another storage system can provide additional security (see Disaster Recovery).

2.4.3 High Availability

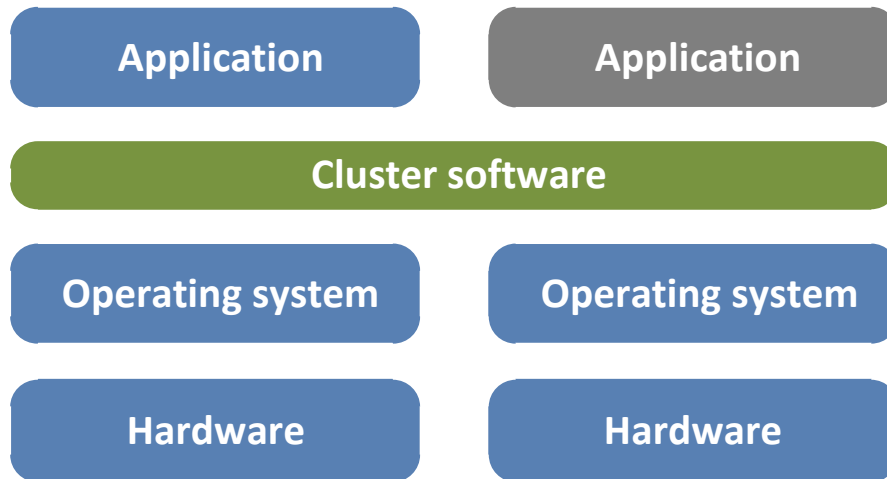
A cluster increases the availability of an application. The cluster nodes are monitored by a "heartbeat." Specific scripts are used to monitor the application's status in terms of application resources. In the event of a fault, the resources (and therefore the application) are switched to cluster node n+1. Clusters with two nodes are most common.

2.4.4 Disaster Recovery

In the event that the entire production facility fails, operations switch to a Disaster Recovery (DR) site. The data is mirrored to the DR site synchronously or asynchronously. There is no data loss in the case of synchronous mirroring. Mirroring can be storage-based or host-based. The switchover is performed manually; in a configuration with a quorum device, switchover can also be done automatically. Operations are switched back to the original production facility once the faults have been corrected and as soon as the data is again synchronized.

3 Increasing application availability in traditional UNIX environments

For many years, the application platforms for SAP R/3 were predominantly UNIX systems before x86 technology made its breakthrough. Stability, high performance, and overall availability of the application were major reasons for this. Cluster technologies that protect applications were an essential factor.



The following sections give an overview of various solutions for protecting applications in various UNIX derivatives. Further, this white paper will demonstrate that corresponding tools are available in virtualized environments based on x86 and that adequate availability can be achieved for applications.

3.1 MC/ServiceGuard (HP-UX)

The requirements for high availability systems are typically met by using cluster software. *MC/ServiceGuard* was developed by Hewlett-Packard (HP) as cluster software for UNIX systems (HP-UX and Linux). “MC” stands for “multiple computer.” With *MC/ServiceGuard*, “multiple” stands for a number between 2 and 16 on the HP-UX operating system and between 2 and 8 on Linux.

As usual these days, the cluster nodes can be easily installed at fibre channel/gigabit distance (10 km). If this is not far enough for the user in the event of a disaster, a larger distance can be achieved by using external storage systems utilizing synchronous or asynchronous data replication.

Reciprocal monitoring of nodes for availability is typically performed via a private network that should have a redundant design.

A serial connection can also be used for this heartbeat if only two nodes are connected together, though the connection should be secured using an additional LAN connection. In principle, a heartbeat can be implemented across multiple LAN connections, including together with user logons.

3.1.1 Configuration types for the MC/ServiceGuard cluster

The cluster's internal transfer function can be configured in three different ways:

- active-active
- active-standby
- rotating-standby

In the first configuration type, *active-active*, various services (in HP terminology: "HA packages") are made constantly available on all cluster nodes. If one node fails, a defined node takes over this service. In *active-standby* configurations, a service is provided exclusively on one computer.

After a fault, the service is started on a different computer. Both of these configurations are usual within symmetric cluster configuration and are also suitable for managing disaster scenarios.

In *rotating-standby*, the node with the lowest number of started services takes over the service from a failed node. The lower hardware requirements of this configuration are, however, achieved at the expense of possible performance losses in a hypothetical disaster case (failure of half of the computers when distributed to two locations).

3.1.2 Quorum Server of the MC/ServiceGuard Cluster

A *Quorum server* is a special feature of the *MC/ServiceGuard*. *Quorum devices* are defined to prevent attempts to start the application if communication has been accidentally and completely interrupted between two cluster members (split brain). This prevents multiple instances of the application from being available and prevents data inconsistencies from arising.

A transfer is only performed if the quorum device can be exclusively locked. The quorum server centrally provides such devices for up to 50 cluster environments, but for a maximum of 100 nodes.

3.2 Oracle Solaris Cluster

The *Oracle Solaris Cluster* is an extension of the Oracle Solaris operating system to include features for increasing the availability of the operating system and business-critical applications.

The agents integrated into the *Oracle Solaris Cluster* allow protection of a wide range of applications. These application-specific agents are able to start or stop the corresponding application or, in case a cluster node fails, to carry out failover of the application to another member.

The central components of the *Oracle Solaris Cluster* are:

- the heartbeat monitor
- the Cluster Membership Monitor
- the HA framework

Thanks to the fact that the central components of the *Oracle Solaris Cluster* are integrated at the kernel layer, it is possible to detect fault situations and failures without delay and initiate failover measures.

The status of all cluster servers is monitored via the *heartbeats*. If a server goes offline and thus loses its heartbeat, it is isolated from the remaining cluster servers, and failover of the applications to another server is started.

The *Cluster Membership Monitor* allows consistent integration of the server hardware into the entire cluster. It also coordinates the configuration of the cluster layer and the Resource Group Management.

The *HA framework* guarantees each member of the cluster a consistent view of the Cluster Configuration Database and the cluster status.

Protection of business-critical applications is ensured in the *Oracle Solaris Cluster* by a combination of a variety of necessary resources. These resources include, for example, global storage devices, file systems, and network devices, and ensure the functionality of all application modules. These resources are organized into resource groups, which allow all necessary resources to be started in the correct dependency and allows failover to a cluster server.

Failover File Services or Global File Services in the *Oracle Solaris Cluster* permit operation of failover clusters or parallel clusters.

3.3 Veritas Cluster Server (VCS)

In addition to the platform-specific cluster solutions MC/ServiceGuard and Oracle Solaris Cluster, the *Veritas Cluster Server* from Symantec is also a common solution in traditional UNIX environments for increasing the availability of business-critical applications. The main features of the *Veritas Cluster Server* include:

- intelligent failover rules for protected applications
- service groups and
- out-of-the-box support for many business-critical applications

A *Veritas Cluster Server* is defined as a network of systems that share the common cluster configuration and are connected via a common “interconnect network.” A cluster network can consist of up to 32 cluster servers that access the stored data of the applications on the shared storage devices in different ways.

Using the *Resource Agent*, the resources required for operation of the business-critical applications are grouped into one or more service groups. In addition to Application Resource Agents, the Veritas Cluster Server also includes agents for storage, file system, and network resources to ensure the basic requirements for operating the applications are met.

In the defined *service groups*, the resources are connected to and with each other based on their dependencies, which allows determination of a correct starting order for the resources. The Resource Agents manage the protected resources and applications. The resources on the cluster systems are started, stopped, monitored and, in case of failure, restarted on demand. If one or more resources of a service group cannot be restarted on the active cluster system, a failover is performed for this service group to another cluster system in the cluster network.

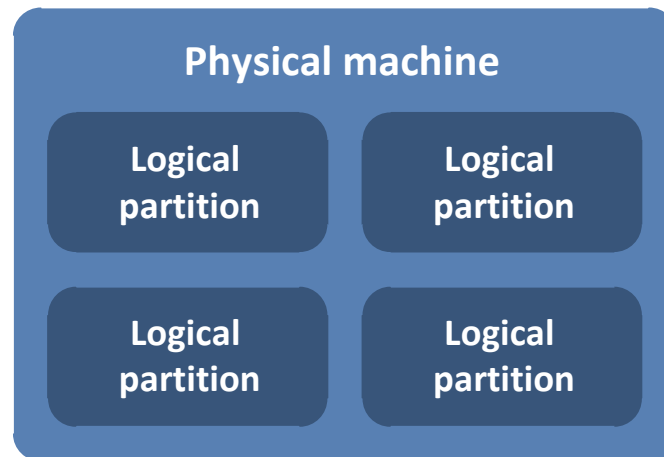
“N+1” and “N to N” cluster topologies are possible depending on the number of servers in the cluster network and the defined failover rules for the service group. In the “N+1” cluster configuration there is a free cluster system that functions as the failover target of the service groups. In an “N to N” cluster configuration, an application is by default active on all cluster systems. In case of failure, the failover rules decide based on the utilization and available capacity to which cluster system the failover of the faulty service group will be performed.

The cluster network can be controlled via the Veritas Cluster Server Management Console, which provides full view of all Veritas Cluster Servers. Using the *Veritas Cluster Server* thereby enables increased application availability in traditional UNIX environments.

4 Virtualization and Partitioning on Non-x86 Systems

The idea of distributing hardware resources as flexibly as possible to various systems comes from the mainframe field: one physical machine is divided (partitioned) into multiple units.

On *non-x86 systems*, hardware is often partitioned into *logical partitions* (LPAR, e.g. on IBM System p or z) or *logical domains* (e.g. on Sun Fire™ E10K).



In the process, resources are allocated to the LPARs/logical domains, for example:

- one CPU
- 2 GB main memory
- hard disk, etc.

In the LPAR/logical domain, an operating system (e.g. AIX) can now operate with the assigned resources. This allows one physical machine to have *multiple concurrent operating system installations* that have *direct access* to the *hardware assigned* to them. Other features include the *dynamic adjustment of assigned resources during operation* and moving an LPAR/logical domain from one physical machine to another.

Though the underlying technology for hardware partitioning and virtual machines (e.g. VMware) is completely different, several features are very comparable. A machine under high load can utilize the resources of a machine under low load, multiple operating systems can run concurrently on one set of hardware, and a machine can be moved to another set of physical hardware.

Beyond hardware partitioning, virtualized machines can also run on non-x86 systems such as the VM operating system on IBM's System z. This technology has capabilities and functions that are very comparable to those of the virtualization solution from VMware. In contrast to VMware virtualization technology, these solutions are bound to proprietary hardware.

5 VMware vSphere Functions for Increasing Availability

VMware, one of the leading manufacturers in the industry with many years of experience on the market, offers a wide range of virtualization products. VMware focuses on solutions whose usage in various areas pursues two main objectives: achieving *higher efficiency* with *increased flexibility*. Hardware agnosticism leads to greater flexibility and allows dynamic operation of the applications, in which the performance can be adjusted to the current demands.

The following sections present the relevant features and explain their main functionality. Depending on the licensing and hardware requirements, these solutions can also be used in combination.

5.1 Important basic functions of VMware technology

The *VMware vCenter Server* is the central configuration and management instance for every VMware infrastructure and also handles the monitoring of the individual virtual machines. In addition to manual configuration, automatically controlled actions can also be performed on the basis of measurement results. In the event the VMware vCenter Server fails, the resources can no longer be controlled centrally and automated processes such as load balancing by migrating individual virtualized systems can no longer occur. This can exert negative effects on the applications in the virtualized systems. *VMware vCenter Server Heartbeat* (see **section 5.2**) provides the ability to secure this instance against failure.

VMware resource pools (part of the VMware vCenter Server) can be used to form logical pools from CPU and memory resources and allocate them to specific virtualized servers, which can ensure a fixed level of resources for specific user groups. Because the resource pools are isolated from each other, changes made within one resource pool do not affect other, independent pools.

Using *VMware Dynamic Resource Scheduling* (VMware DRS), resource pools can be formed across multiple VMware hosts. These are referred to as Dynamic Resource Scheduling Clusters and also form the basis for VMware High Availability (see **section 5.3**).

Dynamic *load balancing between multiple VMware vSphere hosts* is possible within these cluster units. For this load balancing, current performance data is evaluated and special actions such as migrating individual virtualized systems to other hosts via VMware vMotion (see **section 5.4**) can be triggered according to definable policies.

The *queried performance data include* such factors as memory and processor utilization and network response times, but *no application-specific queries*, e.g. for the duration of specific SQL database accesses, can be defined.

VMware vStorage VMFS is a cluster file system that relies on shared storage systems to enable concurrent read/write access to the same file system by multiple instances of VMware ESX. This results in significantly easier allocation of resources to the virtualized machines and simplifies their management, as the entire operational status of a virtualized machine is efficiently stored at a central storage location.

5.2 VMware vCenter Server Heartbeat


VMware vCenter Server Heartbeat is used for the redundant design of the *VMware vCenter Server*, including the associated SQL database at a second standby system. In this process, the entire server, together with the applications running on it, is monitored and the data is transferred to a second system via host-based replication.

Factors such as performance attributes can be used to monitor the vCenter Servers for proper functioning. These factors are compared against target values and, in the event of corresponding discrepancies, lead to failover to a secondary system.

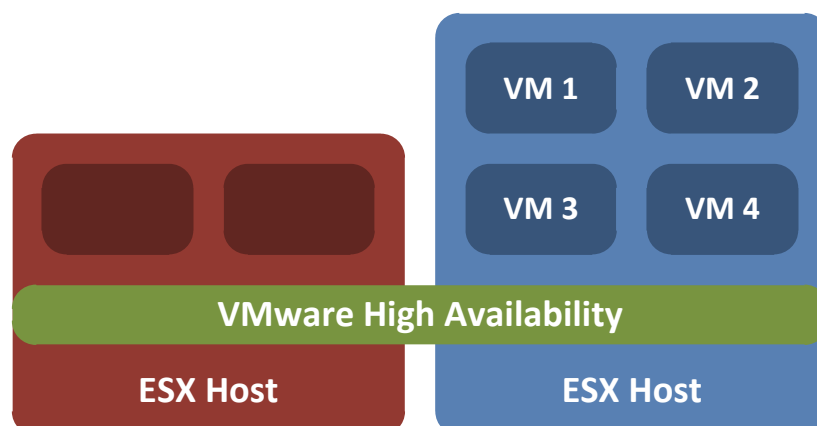
Communication with various core components of a system environment such as the global catalog on the domain controller, the primary DNS server, or the network gateway ensures that failure of network paths does not result in isolation of a VMware vCenter Server.

5.3 VMware High Availability

VMware High Availability (HA) is used to protect virtualized machines from hardware failures and faults in the host operating system.



To do this, clusters are formed consisting of multiple ESX hosts. The resources of these clusters form a system-wide pool. Status messages (heartbeats) are sent between all involved systems and are used to indicate that the machines within a cluster are functioning properly. If these status messages cease for a defined period of time, the virtualized machines on the system will be restarted on the remaining hosts in the cluster. Monitoring is performed by an agent that is part of the operating system.



In addition to monitoring at the host layer, the virtualized machines or individual applications can also be monitored by a heartbeat that is provided by the integration services (VMware Tools). If there is a disturbance in this communication, a reset can be triggered on the relevant virtualized machine. In addition to querying the status of the integration services, the I/O from the server to the virtual drives can also be used as a criterion for defining the system health.

Communication to individual applications (VMware Application Monitoring) must be implemented through a separately purchased SDK or explicitly supported by the particular application.

Because all recovery operations triggered by *VMware High Availability* appear as a cold start from the perspective of the guest operating systems, active memory contents and other information not buffered on hard disks or other storage media are lost.


The following conditions are necessary to use *VMware High Availability*:

- All hosts must have access to a shared storage system with the configuration drives and virtual drives for the virtual disks.
- All hosts in a cluster must be able to access the same network segments so that access to these systems by the virtualized machines is ensured.
- All hosts in a cluster must be able to communicate with each other via the management network.
- The integration services (VMware Tools) must be installed for monitoring individual virtualized machines.


When configuring a *VMware High Availability Cluster*, a definition is set as to whether resources should be reserved for a certain number of host failures and what priority the individual virtualized machines have in the case of failure.

5.4 VMware vMotion

The *VMware vMotion* feature is used for *nondisruptive migration of a virtualized machine between two virtualization hosts*. During this process, the requirements on both involved hosts are checked and, if the requirements are met, the necessary runtime information (e.g. current memory contents) is transferred. The target system then assumes complete operation of the virtualized machine.



This process is performed transparently from the perspective of the applications running on the system and *with guaranteed transaction integrity*. The configuration and the status of the network adapters belonging to the virtualized machine are also transferred.



Several requirements must, however, be met for these features to be used:

- Because no system data is transferred during a migration procedure with *VMware vMotion*, e.g. the virtual disks, these must be located in a storage location accessible by all hosts involved in the process (Shared Storage).
- The network communication paths of the virtualized machines must be ensured on the involved hosts and, additionally, a path for syncing the runtime information must be defined.
- Operation system-specific integration services (VMware Tools) are required on all virtualized systems. These services function as communication interfaces between the management operating system and the guest operating system.
- Within the virtualized machines, there may not be any other virtual devices connected that cannot be accessed from the target system. USB devices and RAW disk devices directly connected to virtualized systems are thus excluded too.
- The CPU architecture in the hosts must match or *Enhanced VMware vMotion Compatibility (EVC)* must be used.
- The virtual machine that will be migrated cannot be migrated concurrently in a Storage VMotion process.

Due to the requirement that shared storage be used for the virtual disks, it is not possible to migrate a virtualized machine between multiple data centers using *VMware vMotion*.

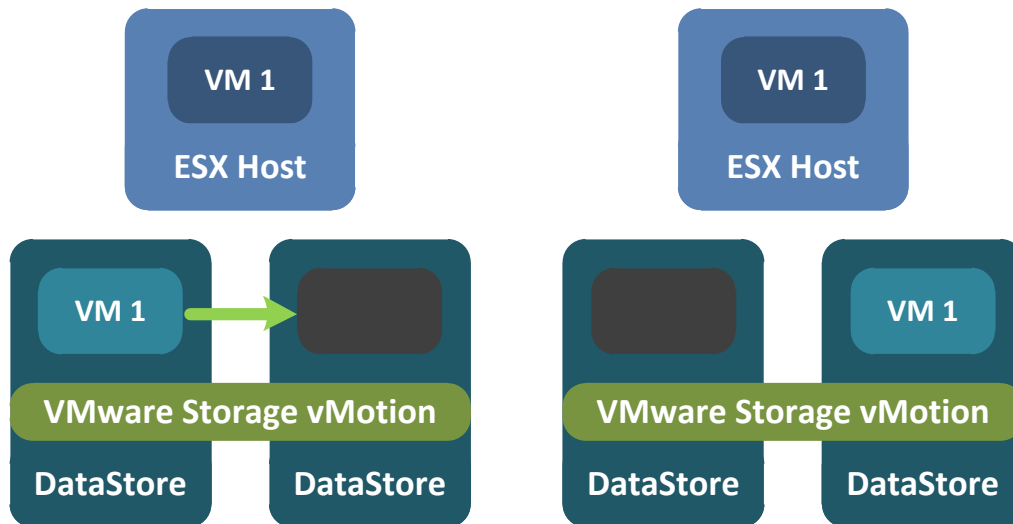
System resources from the host system are needed for the sync process during a migration. During the migration, communication does not occur between the integration services and the installed applications.

If the guest system of a virtualized machine is under load for an extended period of time with many I/O processes in the local RAM or the virtual disks, a vMotion process may be delayed or not executed. The time limit available for this can be configured.

Multiple concurrent *VMware vMotion* processes are only possible if the source and target host are within the same resource pool.

5.5 VMware Storage vMotion

VMware Storage vMotion is used for *nondisruptive migration between multiple datastores while virtualized machines are running*. This includes all associated data storage systems such as virtual disks or configuration files. When performing a migration with *Storage vMotion*, the virtual machine remains on a defined host. That is, the process should be considered separately from *VMware vMotion*.



The following requirements must be met to use *Storage vMotion*:

- Operation system-specific integration services (VMware Tools) are required on all virtualized systems. These services function as communication interfaces between the management operating system and the guest operating system.
- The virtualized machine cannot contain any snapshots.
- The host system must have direct access to the source and target datastores, which can be located on different storage systems.
- The use of virtual drives based on physical devices (RAW devices) is not supported.


In a *Storage vMotion* process, it is possible to separate the data drives from the configuration files and distribute them to different datastores. The type of virtual disks can also be changed to thin-provisioned or thick-provisioned.

During a *Storage vMotion* process, the *storage backend experiences heavier load due to disk I/O* and also has a *delayed response time for disk I/O processes in the guest operating system*. It is not possible to define bandwidth or performance constraints for Storage vMotion processes.


5.6 VMware Fault Tolerance

VMware Fault Tolerance (FT) is used to increase availability by minimizing the effects of a failure. The critical feature of this technology is that *switchover in case of failure can occur without disruption and without affecting the applications running on the system.*

The *VMware vLockstep* technology is employed to feed all input and calls within a virtualized machine to a second virtualized system. This unit is also called a *Fault Tolerance Cluster*. This second system functions as a clone and is automatically created and managed. This clone can *take over the tasks of the primary system at any time.*



After the clone system is activated, it immediately assumes the role of the primary system. Within a short period, a new clone system is automatically created. Because the clone systems are always located on a different host, this *also ensures against complete failure of a VMware host.*



If the requirements for using *VMware Fault Tolerance* are met, this feature can be activated and deactivated during operation, for example to secure virtual machines and the applications running on them during critical operational phases.

As with *VMware High Availability*, heartbeats are used to exchange status messages and verify the operational readiness of the cluster partner, though these are exchanged between the virtualized machines in milliseconds.

VMware Fault Tolerance can be used in combination with *VMware High Availability and Dynamic Resource Scheduling*. The hosts must support EVC in order to use DRS for *Fault Tolerance Clusters*.

The most important requirements and limitations for using *VMware Fault Tolerance* are:

- The VMs that will be protected must be members of a *VMware High Availability or Dynamic Resource Scheduling* cluster.
- All hosts must have FT-compatible processors with similar performance levels.
- The same ESX version must be used on all involved hosts.
- The virtualized machines must be on a datastore accessible by all involved hosts and have access to the same virtual networks.
- The virtual disks for the virtualized machines that will be protected must be preallocated.
- All members of the FT cluster can only have one virtual processor.
- The members of the FT cluster cannot contain any snapshots.
- FT cannot be used in conjunction with snapshots.
- In order to use FT, there must be an additional, dedicated gigabit network connection between the physical hosts.
- There must be at least two virtual networks within the FT cluster, with one for exchange of status information.
- FT-protected virtual machines cannot be migrated concurrently via *Storage vMotion*.
- Hot-plugging of virtual devices such as network devices is not possible within FT clusters.

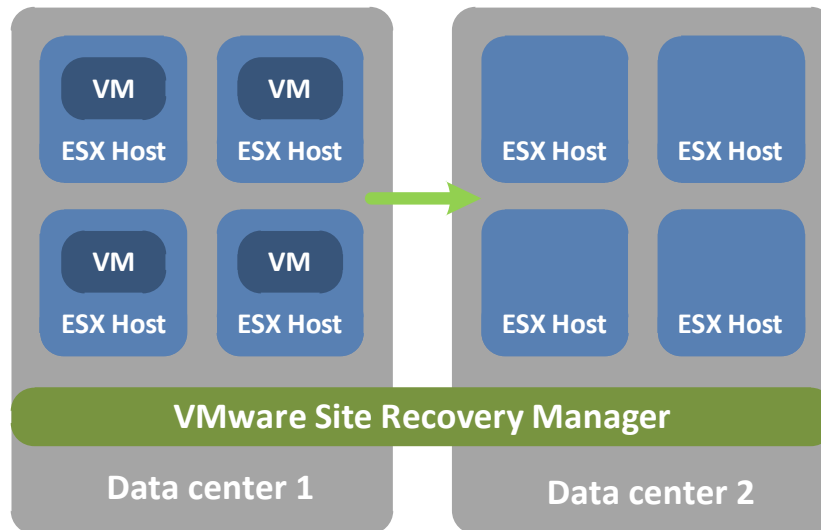
For the use of *VMware Fault Tolerance*, it must also be ensured that the VMs *occupy the appropriate resources*, although *only one instance can be used actively at a time*. The response times of the guest operating systems are extended due to the *vLockstep technology*, which can negatively impact the applications running on the systems.

In general, *VMware FT* is only suitable for applications with low to medium requirements for performance and memory bandwidth.

This solution *also offers no protection against operating system or program errors*, as no application monitoring is performed and changes to the operating system (e.g. updates) that require an interruption in operations cannot be made without downtime.

5.7 VMware Site Recovery Manager

The *VMware Site Recovery Manager (SRM)* is a framework for *automating Disaster Recovery (DR) operations*. It is designed to relocate the *operation of an entire data center to an alternative location* within a limited timeframe in the event of a failure. Various VMware and storage features are utilized for the implementation and are controlled according to a pre-defined policy. Due to the complexity of many virtualization environments, these policies consist of a variety of interdependent individual steps.



When setting up *VMware Site Recovery Manager*, mappings are defined to allocate resources between the primary data center and the disaster data center. These mappings contain network, resource, storage, and configuration settings for the virtualized machines. Resources can be grouped together to simplify administration, and individual policies can be defined for one or more of these groups. This makes it possible to, for example, implement partial switchovers.

The following requirements must be met in order to use *VMware Site Recovery Manager*:

- Each site must have a complete VMware vSphere environment, incl. hardware and associated licenses, including a vCenter management server, management clients, and a network and storage infrastructure, etc.
- The hardware resources on the DR site must be sized for the anticipated workload.
- The storage environment used must support array-based replication and have been specifically tested and approved for use with *VMware Site Recovery Manager*.
- *VMware Site Recovery Manager* uses a dedicated server application with its own database, which requires that corresponding resources be available on both environments.
- The virtualized machines must be given paths to virtual networks which correspond to the paths of the primary environment but do not necessarily have to be located in the same network segments.

All operations mapped in the VMware Site Recovery Manager must be manually planned and defined, which requires very precise knowledge of the system environment. Checks at the application layer are performed neither during the planning nor execution of the policies. These must be given separate and independent consideration.

6 SUSE Linux Enterprise High Availability Extension

6.1 SUSE Linux Enterprise Server

The *SUSE Linux Enterprise Server* by Novell is a *reliable, scalable, and secure operating system* for supporting business-critical processes in physical and virtualized environments. Companies can use it to operate an easy-to-administer, heterogeneous IT environment that ensures *interoperability with products from other manufacturers* such as Microsoft® and SAP and is also *more cost-effective* than using comparable UNIX systems.

The improved features of Version 11 offer a high level of reliability, availability, and easy maintenance. Existing systems can be used optimally and downtimes can be minimized. *SUSE Linux Enterprise Server* was *optimized for use in virtualized VMware environments* and provides excellent performance as a guest operating system *thanks to its optimized resource utilization*.

The modular design of the *SUSE Linux Enterprise Server* makes it possible to provide functions such as high availability clustering by installing *optional extensions*.

6.2 SUSE Linux Enterprise Server for SAP Applications

In early March 2011, Novell released *SUSE Linux Enterprise Server for SAP Applications*, which is based on *SUSE Linux Enterprise Server 11 Service Pack 1* and, in addition to *SUSE Linux Enterprise High Availability Extension*, also contains an integrated and automated installation workflow for the entire software stack, support for large in-memory databases, a special channel for SAP updates, and enhanced support for service pack overlaps. The solution was developed jointly by Novell and SAP and is aimed at companies that use SAP applications.

Even during installation, *SUSE Linux Enterprise Server for SAP Applications* offers various options (AutoYaST) for optimizing the operating system for subsequent use of SAP NetWeaver. These options relate, for example, to partitioning of the local disks in the system, including the installation of special packages (RPM) required for *SAP NetWeaver*.

There is also an option for manual installation, in which an XML configuration file can be created as a template for automated, custom installation of multiple systems.

Using *SUSE Linux Enterprise Server for SAP Applications* offers the following benefits:

- *Fast deployment*
An end-to-end installation framework to install and deploy SAP applications in hours, not days.
- *Optimized performance*
The product was validated and certified by both SAP and Novell to eliminate potential software incompatibilities.
- *Business continuity*
SUSE Linux Enterprise High Availability Extension is built into the product, providing a cost-effective cluster solution for both physical and virtual Linux deployments.
- *Integrated 24x7 SAP priority support*
Available from Novell or hardware partners.

- *Support for large in-memory workloads*
Optimizes paging behavior of Linux for large in-memory databases and applications such as SAP Business Warehouse Accelerator or SAP High Performance Analytic Appliance.
- *SAP Solution Manager*
Integration in SAP Solution Manager streamlines the resolution of support requests, reducing system complexity and lowering TCO.
- *Integrated support*
Seamless integration into the SAP Support structure through SAP Solution Manager provides you with support for application and operating system problems (through a priority support subscription).
- *Extended service pack support*
The standard 12-month support for SAP customers is extended to 18 months.
- *Separate, dedicated update channel for SAP*
Reduces operational risks by ensuring that all fixes and patches are approved by SAP and Novell.
- *SAP-specific JVM support and maintenance*
Ensures SAP-validated installation and updates via the dedicated SAP update channel.

The statements made in this document apply both to *SUSE Linux Enterprise Server High Availability Extension* and *SUSE Linux Enterprise Server for SAP Applications*.

6.3 SUSE Linux Enterprise High Availability Extension

The *SUSE Linux Enterprise High Availability Extension* makes it possible to increase application availability on *SUSE Linux Enterprise Server 11* by setting up a cluster.

In addition to the provisioned services, an application also contains the required additional resources such as IP addresses, file systems, volume managers, etc. File systems can be integrated in various ways, for example on shared disks with or without a cluster file system, as a network file system (NFS/SMB), RAW devices, or local file systems as replicated block devices. These resources are combined into a resource group that describes the application.

Among the components used are:

- OpenAIS (Application Interface Specification, AIS) as the cluster framework
- Pacemaker as the resource manager, with GUI for managing the cluster
- Resource agents for monitoring and controlling the applications and their required components
- Oracle Cluster File System OCFS2
- Cluster Logical Volume Manager, a special version of the Logical Volume Manager on Linux, which also works with cluster file systems
- Distributed Replicated Block Devices (DRBD) for data replication

By integrating these components, *SUSE Linux Enterprise High Availability Extension* provides support for scenarios that previously could only be implemented on traditional UNIX or mainframe environments.

The *SUSE Linux Enterprise High Availability Extension* includes, among other things, a resource agent for SAP instances. It is responsible for starting, stopping, and monitoring the services of an SAP instance. The following versions of SAP instances are supported:

- SAP WebAS ABAP Release 6.20 – 7.30
- SAP WebAS Java Release 6.40 – 7.30
- SAP WebAS ABAP+Java Add-In Release 6.20 – 7.30

Another resource agent is responsible for starting, stopping, and monitoring the database of an SAP system. This agent supports the following databases:

- Oracle 10gR2 and 11gR2
- DB2 UDB 9.x
- MaxDB

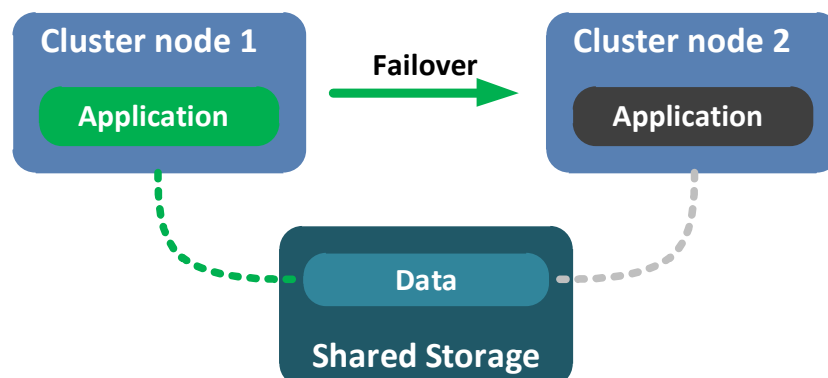
The *SUSE Linux Enterprise High Availability Extension* comes with even more resource agents such as for Apache, CTDB, DRBD, iSCSI, MySQL, libvirt, NFS, Postfix, Postgres, Squid, Tomcat, Xen, Oracle, DB2, Informix, WebSphere, and eDirectory.

The *SUSE Linux Enterprise High Availability Extension* can be deployed on both physical and virtualized machines. When using virtualized machines, it can also be combined with the virtualization solution's cluster functionality. In this case, it protects the applications while the virtualization solution's cluster functionality protects the virtualized machines.

In addition to increasing the availability of the applications, this cluster configuration also offers the following benefits:

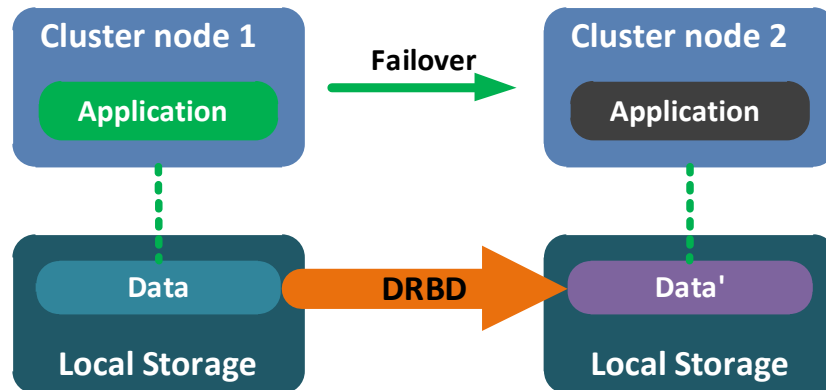
- Increased performance through load balancing
- Horizontal scalability
- Abstraction of physical hardware at the application layer
- High flexibility in the cluster design

Possible examples of these scenarios are described below with shared storage, with NFS and with DRBD as the I/O stack. In a simple usage case, an application is defined by a resource group and is active on a cluster node. A shared storage system is used in the cluster to provide the application data to all cluster nodes. The file systems are mounted to only one cluster node at all times.



The benefits of this scenario are a lower degree of complexity in the cluster design and an easy option for extending the cluster configuration to include other resource groups for other applications. The application data can also be provided to the cluster by an external, redundant NFS server. This does not change the functionality of the cluster and the application.

In another usage case, the application data is replicated between the cluster nodes synchronously or asynchronously to local disks or SAN LUNs using block replication (DRBD). The resource group describing the application contains a resource for controlling the DRBD devices. In the case of resource group failover, the DRBD resource switches the source and target of the replication.

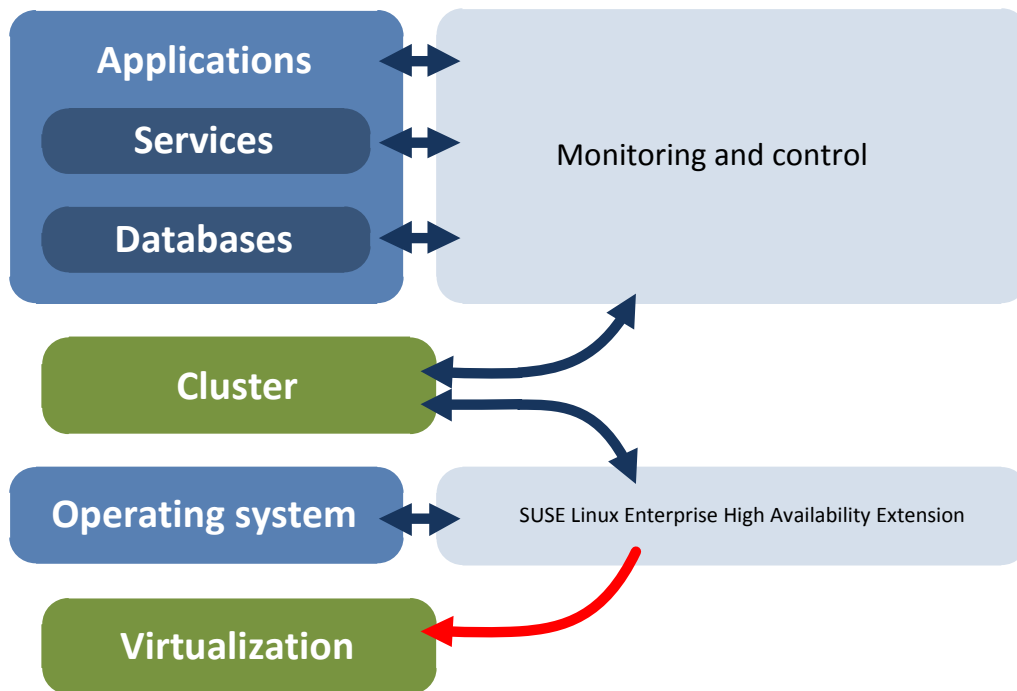


Comparatively low costs are a major benefit of this solution.

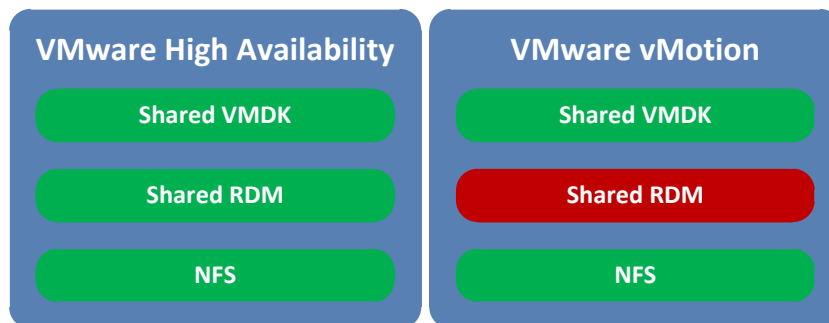
6.4 SUSE Linux Enterprise High Availability Extension in a VMware vSphere Environment

Currently, guest systems with the *SUSE Linux Enterprise High Availability Extension* can only be moved if the systems do not share any disk with other guest systems. *VMware vMotion* (see **section 5.4**) is currently released with virtualized systems whose shared data set is integrated via an NFS server. The use of shared, virtualized shared disks (VMDK) is described in VMware Knowledge Base article 1034165 (see link in References).

Note that *VMware vMotion* does not interrupt the *cluster communication layer* and therefore does not interrupt the cluster's heartbeat. When a virtualized machine is being moved, the most recently modified memory pages are also transferred.



Unlike VMware High Availability, VMware vMotion is subject to certain restrictions as shown in the figure below. All items marked in green are fully supported by VMware. The usage of shared VMDKs in conjunction with VMware vMotion will be supported in the future, but it is currently in the "Technical Review" status.



VMware High Availability Application Programming Interface

Version 4.1 of VMware vSphere provides an *Application Programming Interface* (API) that allows transmission of results from application monitoring to VMware High Availability and triggering of corresponding actions there. This interface will also enable the SUSE Linux Enterprise High Availability Extension to transmit the results from application monitoring to VMware High Availability and execute actions there. This extended functionality of VMware High Availability can be used to better integrate the SUSE Linux Enterprise High Availability Extension into the VMware environment.

7 SAP NetWeaver

SAP NetWeaver is a modern application server technology that provides the foundation for applications of the SAP Business Suite and ensures their seamless networking. The customer's business requirements determine which of the SAP Business Suite Applications based on SAP NetWeaver are deployed. This scope of services gives the functionality of the Suite's components great influence on the company's success.

SAP NetWeaver is a business-critical application, though it is not a single product but instead a combination of various interconnected SAP applications. The structure and importance of the individual components depends on the customer's requirements.

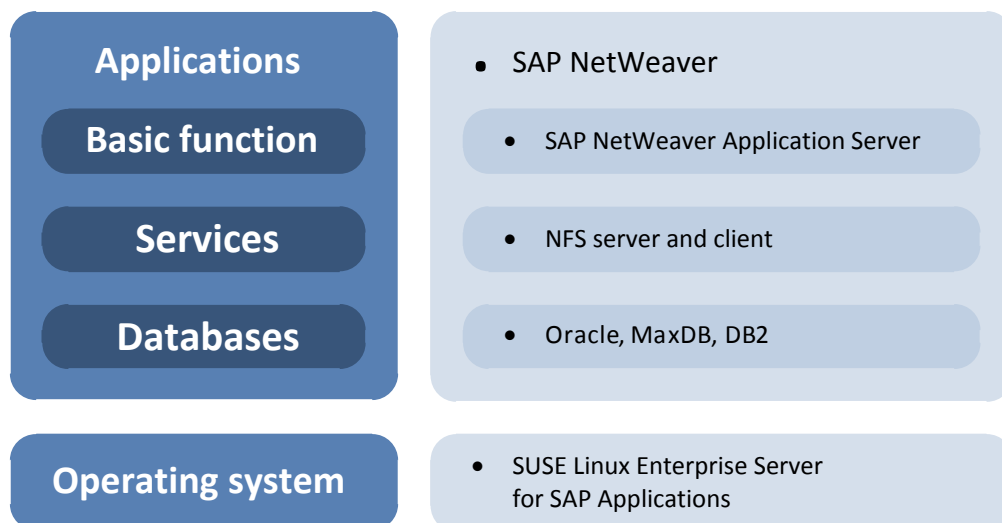
SAP NetWeaver provides a very versatile application server technology that lets companies coordinate their business requirements with the capabilities of IT, including for example the unification and consolidation of the corporate master data, the collection and analysis of information, and the rapid and targeted development of individual applications.

This white paper deals solely with the technical infrastructure component of SAP NetWeaver.

7.1 Structure of SAP NetWeaver

SAP offers an entire architecture as a scalable, fault-tolerant, and multi-tier solution that can be protected against failures thanks to the horizontal scalability of the NetWeaver application servers and cluster solutions. Cluster solutions such as SUSE Linux Enterprise High Availability Extension can be deployed both in physical and virtualized environments to protect components identified as *single points of failure*.

As shown in the figure, SAP NetWeaver can be structured into multiple functional units providing various tools for ensuring the required availability.



7.2 Single Points of Failure (SPoF) of SAP NetWeaver

An SAP architecture contains the following *single points of failure* (SPoF) that are non-redundant and need to be protected using suitable high availability tools:

- *Database*
The database is especially important in terms of performance and functionality. Every ABAP work process makes a private connection to the database at the start.
- *SAP Message Service*
The SAP Message Service is required to exchange information between SAP instances and regulate this exchange.
- *SAP Enqueue Service*
The Enqueue Service manages locking of business objects at the SAP transaction level.



The following SAP components can be defined based on the *Message and Enqueue Services*:

- *Central Instance (CI)*
The central instance is the first installed instance of the application server and includes up to NetWeaver Release 7.0 Message and Enqueue Services and in addition to SAP work processes enabling execution of batch and background workloads. The CI is a possible SPoF if it includes the described Message and Enqueue Services.
- *Central Services*
In newer versions of SAP NetWeaver, the Message and Enqueue Services have been separated from the CI and are operated as standalone services. Separate Central Services exist both for ABAP and for Java-based NetWeaver application servers. For ABAP environments, they are referred to as ABAP Central Services (ASCS).
- *Replicated Enqueue*
The Replicated Enqueue Server is run on a different host system (preferably on a virtualized machine) and contains a replica of the lock table in a shared memory segment. If a standalone Enqueue Server fails, it must be restarted by a high availability solution (cluster software) on the server on which the Replicated Enqueue Server is running. The restarted Enqueue server connects to the shared memory segment of the replicated Enqueue server and makes the existing locks available again after the move to the new host. The Replicated Enqueue Server can then be restarted on another host to provide continued redundancy.

The isolation of the Message and Enqueue services from the central instance on NetWeaver Release 7.0 and higher simplifies setup of a high availability solution for this part of the SAP architecture. This also makes the central instance start faster because it is stripped of these functions, which reduces the duration of the interruption.

8 Reasons for Running Business-Critical Applications in a Virtualized Environment

8.1 Consolidation

With the *VMware vSphere* virtualization solution, existing systems can be consolidated easily and hardware can be utilized more effectively, which reduces operating costs.

VMware vSphere combines the resources of physical servers, storage systems, and networks into resource pools that maximize efficiency and utilization, thereby establishing the foundation for creating a high availability and flexible IT infrastructure. The resource pool model offers a global management scheme including automatic optimization of physical resources. The need for logical resources can be assigned to different virtualized systems according to requirements.

8.2 High Availability

The VMware vSphere functions described in **section 5** counteract possible hardware failures in the virtualized infrastructure. These functions increase the availability of the entire environment primarily by protecting against:

- *Hardware failures at the server and storage system layer* by restarting and/or moving virtualized systems to a functional host system
- *Operating system errors* by monitoring the virtualized systems in conjunction with an automatic restart on failure

8.3 Disaster Recovery

Traditional *disaster recovery solutions* are usually expensive and high-maintenance. They require expensive hardware duplication at dedicated recovery sites and many complex and time-consuming manual procedures. The error rate is often very high on account of dependencies and inadequately tested failover processes.

In a virtualized infrastructure the disaster recovery strategy can be greatly simplified and a highly reliable switchover of services to the recovery site can be ensured. Production applications can be quickly, reliably, and cost-efficiently restored in the entire x86 infrastructure without having to restrict the recovery options to individual selected business-critical applications. Virtualized VMware machines that are independent of physical hardware make the requirements for the recovery site much more manageable, as the servers can be saved as images. This allows much more efficient migration of application, data, and configuration information to new machines, resulting in significantly shorter recovery time objectives (RTOs) compared to traditional physical failover solutions. In short, disaster recovery scenarios for multiple systems can be covered more cost-efficiently and reliably in a virtualized infrastructure.

8.4 Opportunities and Risks

The successful deployment of server virtualization technologies was made possible due to the tremendous increase in the performance and stability of the hardware used. The capability of running more than one system on the same hardware also provides the option of distributing services across multiple systems, thereby resolving unnecessary dependencies without requiring the purchase and operation of additional hardware.

Much simpler disaster recovery scenarios are possible due to hardware standardization in the virtualized systems. This applies both to simple recovery of a system after a possible hardware failure and to moving or failover of one or more virtualized systems to another host. It makes almost no difference whether this host or the pool of ESX Servers is located at the same site or a remote data center.

It must, however, also be noted that the required virtualization layer can represent an additional source of errors and may influence the availability of business-critical applications despite the sophisticated technology and a stringent quality control process followed when writing the software.

Using server virtualization increases the potential for faults resulting from:

- greater complexity of the application stack
- new potential causes of errors in the software of the virtualization layer itself and
- increased risk of failure due to special adjustments required in the device drivers; Their function is divided into two layers: they are now available in multiple instances to the virtualized systems and in the hardware-based layer they pass the collected requirements to the hardware.

New types of errors can also occur when using server virtualization:

- A host system failure affects far more applications and services simultaneously than in a traditional server environment and
- If tools that compensate for this effect are missing, using individual virtualization snapshots (for example) puts “all your eggs in one basket.”

Before technical implementation, the tools used in a virtualized environment, their functionality, and data protection need to be examined. For example, migrating a virtualized system from one host to any other host is not always suitable for increasing application availability.

When moving virtualized systems to other hardware in a pool, various aspects, especially regarding data consistency, must be observed:

- The suspend/resume technology requires a disk-based snapshot while the virtualized machine is running in order to make exactly the same consistent data available to the system again after it is moved.

If this process for establishing data consistency is initiated by an administrator or the system or via an interface while the system is running, the system and its applications and services will be available again with consistent data after moving to another host system. This mechanism is also utilized by VMware Consolidated Backup (VCB).

An unscheduled interruption occurs in the system's operation, it is too late to create a snapshot and the data for the affected applications may be inconsistent.

- The virtualization layer itself does not contain any mechanisms for restarting an application after an application error or for migrating the guest system within the resource pool due to an application error.
- Migrating a virtualized system to another host system is fundamentally not appropriate for resolving software or resource errors. This also applies to VMware Fault Tolerance.

The data associated with an application is of course an important component of every application. Server virtualization does not have a role in protecting this data against technical faults or logical faults in particular. Hence the need to take appropriate data protection measures in the context of the overall structure in question.

In terms of protecting against a disaster (failure of an entire unit, e.g. data center, entire virtualization environment, etc.), permanent replication of data to a second location is the secure course of action. This can be done through storage-based or host-based replication.

Another important cause of errors can be destruction of the data by logical faults, e.g. application errors, operator errors, or accidental deletion. Continuous Data Protection (CDP) technologies in conjunction with the relevant applications offer a wide range of possible solutions.

Another important component of a virtualization environment are systems for controlling and administering the entire environment, e.g. the VMware vCenter Server. Their basic functionality is designed so that a server failure does not lead to failure of the virtualization environment. Restrictions in the administration are possible. If features or modules are integrated into the VMware vCenter Server that directly affect the functionality of the virtualization environment, the system must be considered a critical application and subjected to application protection.

9 Increasing Application Availability of SAP NetWeaver by Deploying SUSE Linux Enterprise High Availability Extension in a VMware vSphere Environment

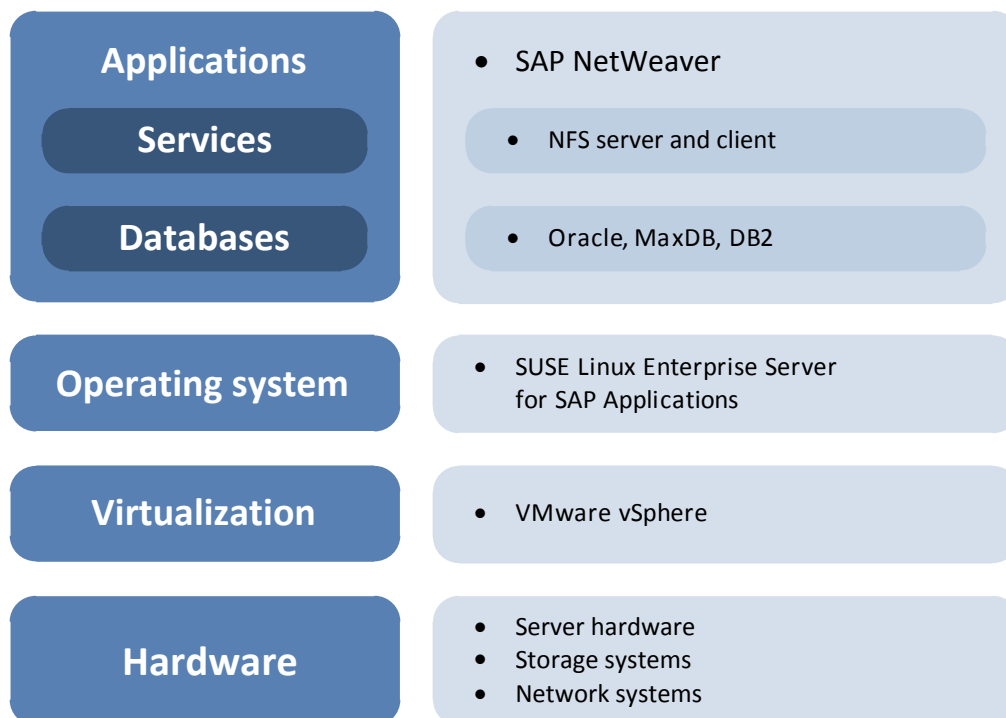
9.1 Application Stack of SAP NetWeaver

To ensure the availability of an application or service in accordance with a Service Level Agreement (SLA), all components that affect the availability of overall functionality must be examined. This relates to back-end processes on server systems, to the network in the broadest sense, and of course to the client used by the user. The statements in this white paper are limited to the server processes.

To structure the complexity of the requirements according to the required level of availability, it is helpful to assign layers for the application, the environment required for operation, and the necessary services. These layers are delineated by their functionality and are interchangeable according to function. The measures dealt with here for increasing availability or disaster recovery concepts relate to one or more layers. The totality of the measures taken yields the required availability of the application or service.

High availability and disaster recovery must be considered in conceptual terms when designing a system and cannot be revised later on.

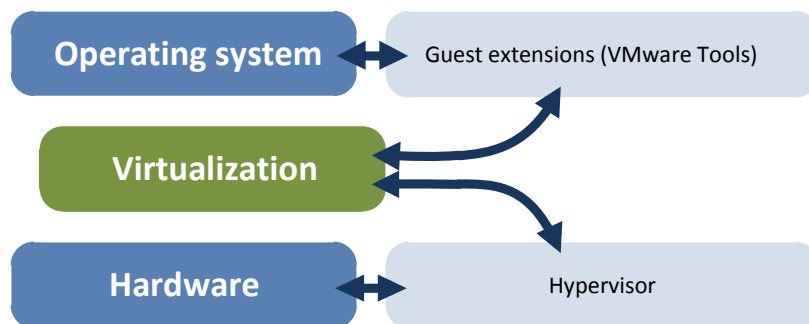
The totality of the layers required for the application operation is referred to as the *application stack*.



9.1.1 Influence of VMware vSphere Functions on the Application Stack

All functions of the VMware vSphere product range that are described in **section 5** represent solutions that significantly increase the availability of the underlying hardware (server and storage). The virtualization layer always acts as an interface between physical hardware and the guest operating system of the virtualized machine.

The hypervisor of the virtualization layer (*virtualization*) provides optimum access to the individual hardware resources (*hardware*). All other layers of the *application stack* that are *above the virtualization* are *denied direct access to shared hardware resources*. These layers thus also lose immediate dependency on these physical resources.



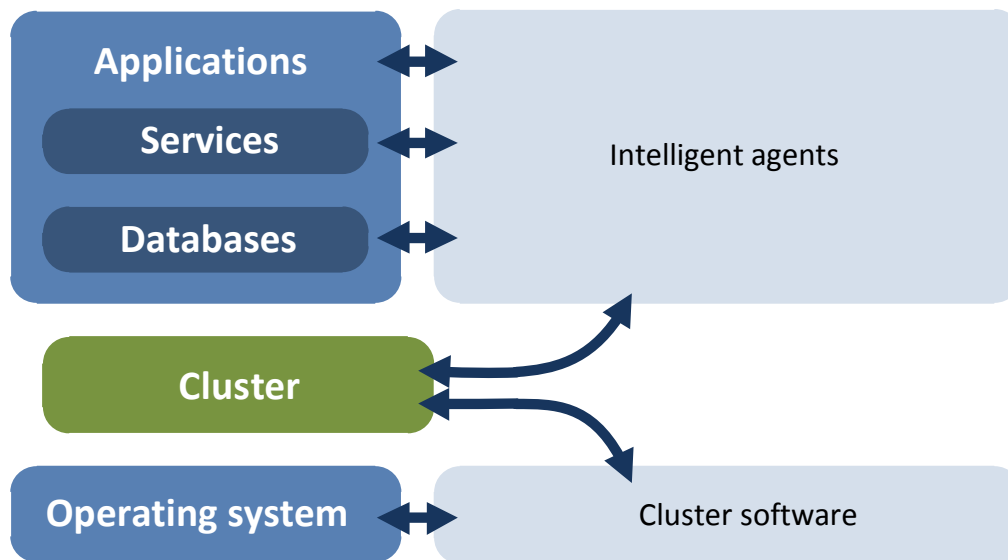
The interface between the *virtualization layer* and the *operating system layer* within a virtualized machine is provided by extensions to the guest systems (VMware Tools). Hardware abstraction allows these extensions to *efficiently use the components presented by the hypervisor*. The VMware Tools mainly include drivers for the emulated standard hardware (chipset, network), interfaces for communicating with the hypervisor, and general control functions for the operating system (pause, on and off, shutdown, and restart).

The *virtualization hypervisor* also has *no direct influence whatsoever on higher layers of the application stack*. Applications and services thus remain unaffected by the virtualization functions. The virtualization hypervisor can therefore *make no determination on the status of applications* other than performance data for the virtualized machine (e.g. processor utilization or memory usage) to respond automatically to any changes or faults.

Due to the functions of the virtualization described here, additional technologies are required to increase the availability of the higher layers of the application stack.

9.1.2 Integration of Application Protection into the Application Stack

The hardware layer and the operating system layer are secured by the tools for increasing system availability that are integrated in VMware. Products for ensuring application availability such as *high availability clusters* of course *cover protection and monitoring of the entire application stack*. These products can often be combined with data replication tools or already contain these products to also provide protection for data associated with the application. These products typically do not include functions for ensuring the availability of the virtualization layer.



If additional functions for protecting virtualized business-critical applications are used, they should take over continuous protection of the applications and services above the operating system layer. When designing the total solution, it can be assumed that both the hardware and the operating system will be adequately protected by the virtualization solution's functionality. Functional overlap that under different rules would affect a layer must be avoided.

The selected solution should therefore have deep knowledge of the structure, the function, and the testing capabilities of the application that will be protected. This tool should be modular so that it can handle the many possible combinations of database, services, and the application. Dependencies (hierarchies) are set up based on these resources. These dependencies reflect the structure of the application and at the same time map a policy for starting and stopping the application.

9.1.3 Deploying High Availability Central Services

When considering the availability of an SAP NetWeaver system, the availability of the infrastructure as well as the central services required for operation such as DNS or NFS services must also be considered.

Services such as DNS or directory services with conceptual designs allowing the setup of redundant systems do not require any additional measures.

If services such as NFS are provided centrally – sometimes for multiple SAP systems – then availability is especially important and appropriate measures need to be taken. In the case of NFS services, this can be done using appliances that meet the special availability requirements; alternately, the NFS service needs to be integrated into the landscape as a high availability service. A deployment example is provided later.

9.2 Selection Criteria for Application Protection Solutions

As already stated, a wide range of products and solutions can be used to achieve the planned availability for the business-critical application. Before deploying a solution or product for application protection, it is imperative that it be analyzed according to a wide range of criteria and a precise determination be made as to which features will be used.

When designing complex environments based on server virtualization, it is useful to define design rules so that, after installation, the required availability, stable operation, easy monitoring, support, disaster protection, and ultimately the desired efficiency of the solution can be achieved.

Among the factors that design rules should describe are:

- which service level will be ensured by which high availability technology
- which tools are permitted in which application layer and
- how functional overlaps and conflicts between tools will be resolved
- whether tools can act across layers
- how exchange of a product within a layer is supported (using proprietary tools)
- what requirements there are for operation of the environment

The following sections detail these various aspects in selecting tools for increasing application availability in virtualized environments.

9.2.1 Assignment of Functionality to the Application Stack Layers

Complex applications in particular typically consist of a combination of various program modules, libraries, services, and a database. As shown in the example of SAP NetWeaver, modules can be replaced with similar modules from a different manufacturer so that the application can be integrated into the existing infrastructure or special requirements for the system architecture can be met. A selection can be made between different operating systems and databases during installation of SAP NetWeaver.

The tools for application protection should also provide the flexibility required to leverage the various possible product combinations in accordance with the Service Level Agreement. This is handled by restricting or modulating the high availability tool to the particular layer in the application stack. This is the only way to ensure that replacing a functional component (e.g. database and operating system) does not interfere with the entire application design and the availability requirements that depend on it.

9.2.2 Monitoring the Application

The aim is to achieve the highest possible availability for the business-critical application. Detecting as many fault sources as possible in an application and the services it uses is necessary for documenting and then eliminating the sources of faults.

One of the most important components of any high availability solution is monitoring of the application being protected. Full monitoring requires comprehensive knowledge of the structure and functions of the application as well as the faults that may occur in them. Limiting consideration of these faults to hardware or operating system failure is not sufficient for a complete high availability solution. The depth of testing directly affects the achievable application availability.

It must also be noted that tools for ensuring availability are in principle not capable of detecting and resolving all faults:

- *Faults caused by lack of resources*
For example, there is no more disk space available for log files.
- *Failure of the entire environment, including primary and secondary hardware*
This is termed a disaster and special measures and rules apply.
- *The tool cannot make automated decisions, as data loss is expected*
After a system crash, the database integrity must be reestablished. Because this may be accompanied by data loss, the system should wait for operator input.

9.2.3 Non-Conflict Between Tools Used

When designing the overall solution, it is particularly important to ensure that the tools used work together harmoniously. This relates both to non-conflict between third-party tools and the features of the virtualization solution and also between the tools themselves. Beyond an advanced certification matrix covering the implementation and version of the virtualization layer, the variety of possible guest systems in their various implementations must also be considered.

A particularly important point is the absence of logical conflicts by the tools used. Various high availability solutions sometimes follow very different philosophies or stipulate different requirements. Especially serious problems may arise when, for example, tools to ensure high availability and those that ensure fault-tolerant operation of an application end up competing with each other.

If tools used to increase availability at different layers independently apply different automated measures, in the worst-case scenario this can lead to total failure of the application.

These potential conflicts can only be prevented by deploying a cleanly structured system with tools clearly assigned to the requirements for application and/or resource availability and the relevant layer.

9.2.4 Integration into the Management Environment

Although monitoring an application and its resources is an important element of every availability solution, the use of a general and comprehensive monitoring and management solution within the overall environment is still essential. There are many reasons why a management system for operating business-critical applications is necessary:

- *Support for end-to-end visibility*
Monitoring of infrastructure components (e.g. LAN and SAN) can be integrated into the management and monitoring system, which allows complete visibility over the environment and supports faster troubleshooting.
- *Preventive error detection*
The trend analysis capability allows early detection of possible errors resulting from resource exhaustion.
- *Resource monitoring*
Sources of errors that cannot be handled by high availability tools or fault tolerance solutions are visible in the monitoring system.

- *Restoration of redundancies*

High availability is based on the skillful use of as many redundancies as possible. If one of the resources fails, its functionality is assumed by the respective redundancy as inconspicuously as possible. The protection is typically against single-chance faults, so another failure of the resource active at the time can lead to total failure. The management system is responsible for identifying components that have failed and been automatically replaced so that full fault protection can be reestablished.

10 Design of a High Availability SAP NetWeaver Platform on the Basis of VMware vSphere 4

The previous sections have detailed a number of technical solution options. These differ in terms of technology, mode of operation, and the layer in the application stack that they affect. In conjunction with the Service Level Agreement, design guidelines need to be drawn up for the system that will be deployed. These guidelines serve as the basis for building and developing the solution.

10.1 Design Aspects

Various aspects are important when deciding on a combination of the possible solutions mentioned for increasing the availability:

- *Complexity of the solution:*
VMware High Availability and VMware Fault Tolerance are relatively easy to set up and configure. The effort required for cluster software can be significantly higher.
- *Scalability of the solution:*
VMware Fault Tolerance has a number of restrictions, for example limitation to one virtual CPU (vCPU), which greatly limits its usage in some scenarios.
- *Reduction in unplanned downtime:*
How quickly should SAP NetWeaver become available again after a hardware failure? Cluster software can reduce this time, as the time spent on booting the operating system in the guest system is eliminated.
- *Reduction in planned downtime:*
The guest system in a cluster can be patched by moving SAP NetWeaver to the other node and then updating the passive node. This makes it possible to keep the operating system up to date with minimal downtime.
- *Requirements from the Service Level Agreement (SLA):*
What SLAs does the customer want to implement? Should it be possible to continue running SAP NetWeaver when a data center fails (disaster)?

10.2 Options for Increasing Availability of SAP NetWeaver

There are various ways to increase the availability of central SAP services. With VMware High Availability, VMware Fault Tolerance and the VMware Site Recovery Manager, VMware offers basic protection options for the platform for running the central SAP components. In addition, cluster software can be deployed within the virtualized system.

10.2.1 Using VMware High Availability

VMware High Availability (see **section 5.3**) protects virtualized machines against failure of a VMware ESX host system. For this purpose, each server communicates with all other servers in the cluster using a heartbeat. If a server failure is identified, its virtualized machines are restarted on other hosts in the cluster. If the SAP database is running in one guest system and the central instance is running in another guest system and both guest systems are on the same host, then in the case of a failure both virtualized machines can be started on different host systems. However, *VMware High Availability* only starts the guest systems and does not check whether services within these virtualized machines start in full. If the services do restart in full, they are available again a short time after a hardware failure. Scripts are required in the guest system to start the SAP and database instances after a restart by *VMware High Availability*.

10.2.2 Using VMware Fault Tolerance

VMware Fault Tolerance (see **section 5.6**) takes a different approach, in which a virtualized machine is not run just once but instead concurrently on two different host systems. This means that a guest system has an exact copy that receives the same input and is always in exactly the same state as the original. If the original host system fails, the guest system can continue working on the second host with no downtime. No time is lost by restarting the services and the user does not perceive any failure. However, if an error occurs in the application, it will also be transferred to the copy.

There are several critical restrictions here. *VMware Fault Tolerance* only supports virtual machines with a maximum of one virtual CPU, which may be sufficient for small SAP systems, but in larger environments at least the virtual machine with the database requires multiple CPUs. Another solution is therefore needed in such environments.

10.2.3 Using the SUSE Linux Enterprise High Availability Extension

The availability of the SAP environment can also be ensured using the technology of *high availability clusters*. For SUSE Linux Enterprise Server, the SUSE Linux Enterprise High Availability Extension will be considered here. This product protects applications and services. It combines at least two servers into a cluster. The application is installed on all servers in the cluster and the user data is located on shared storage or is replicated between two storage systems. If the active server or the protected application fails, the next server in the cluster takes over this application. The cluster software can also check whether the SAP components are functioning fully within the virtualized system. If this is not the case, the services can be automatically restarted or operation can shift to another cluster node.

Deploying the SUSE Linux Enterprise High Availability Extension enhances the depth of testing for application protection. *VMware High Availability* can also be used to reestablish the original redundancy after a guest fails. However, when these two tools are combined, the trade-offs between the achievable benefits and the increasing complexity must be considered.

If using cluster software within virtualized machines, it must be ensured that no unwanted side effects arise between *VMware High Availability* and the *cluster solution*. The cluster tools included in the SUSE Linux Enterprise High Availability Extension installation ensure that the resources belonging to the application are started in the correct order and on the right node. Additional configuration is required in this scenario to ensure proper operation.

10.2.4 Using the VMware Site Recovery Manager

VMware Site Recovery Manager (see **section 5.7**) allows operation to be resumed in a disaster data center if an entire data center fails. In this case, all virtualized machines are then run in the disaster data center instead of the primary data center. Using VMware Site Recovery Manager, priorities can be set for moving virtualized systems to the secondary data center. This allows a reduction in the time before critical systems can resume operation and also lets systems with a lower service level be skipped at the start.

After a relocation affecting not just a few especially relevant systems, the entire SAP environment and all the components can be operated again. Because the entire VMware environment can be moved in this process, the previously mentioned measures (VMware High Availability, VMware Fault Tolerance, and cluster software within the virtualized machines) can be used again too. Continued protection against a hardware failure is therefore ensured in the disaster data center.

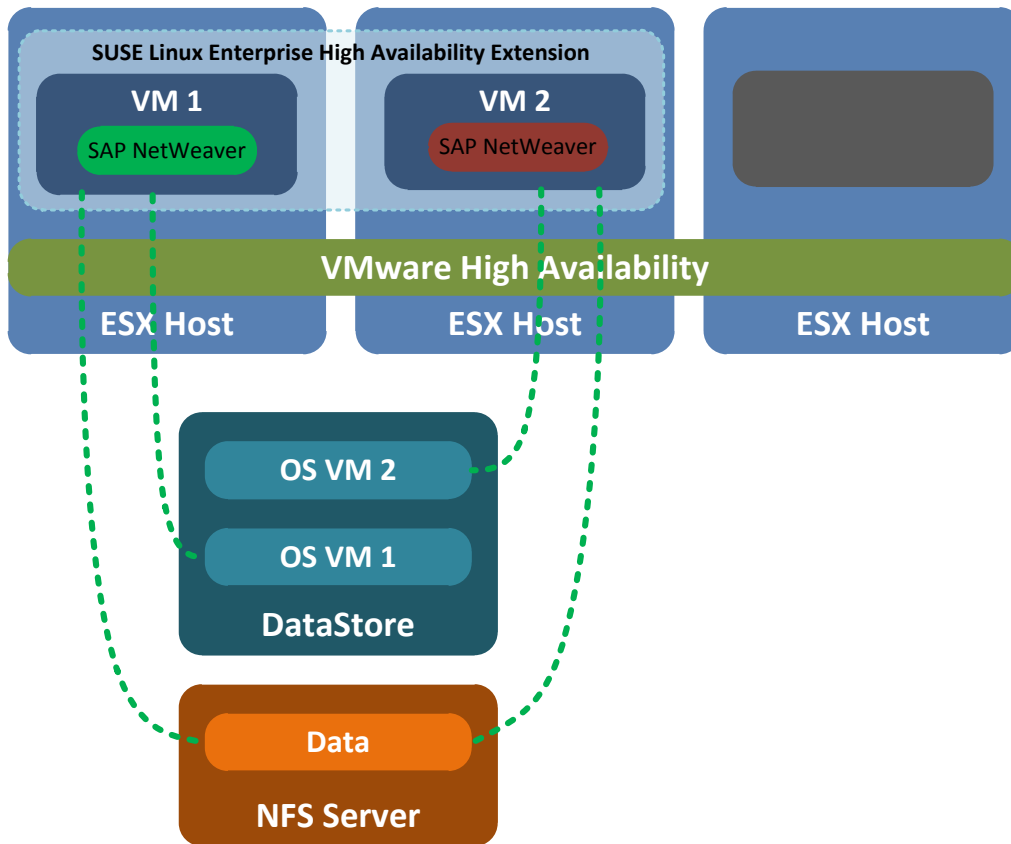
10.3 Example Configurations

The following configurations give abstract examples of the possible combinations for using the respective tools to increase the application availability of SAP NetWeaver. There are also other options for configuring the SUSE Linux Enterprise High Availability Extension (as described in **section 6.3**), but for the sake of clarity only four example solutions are presented below.

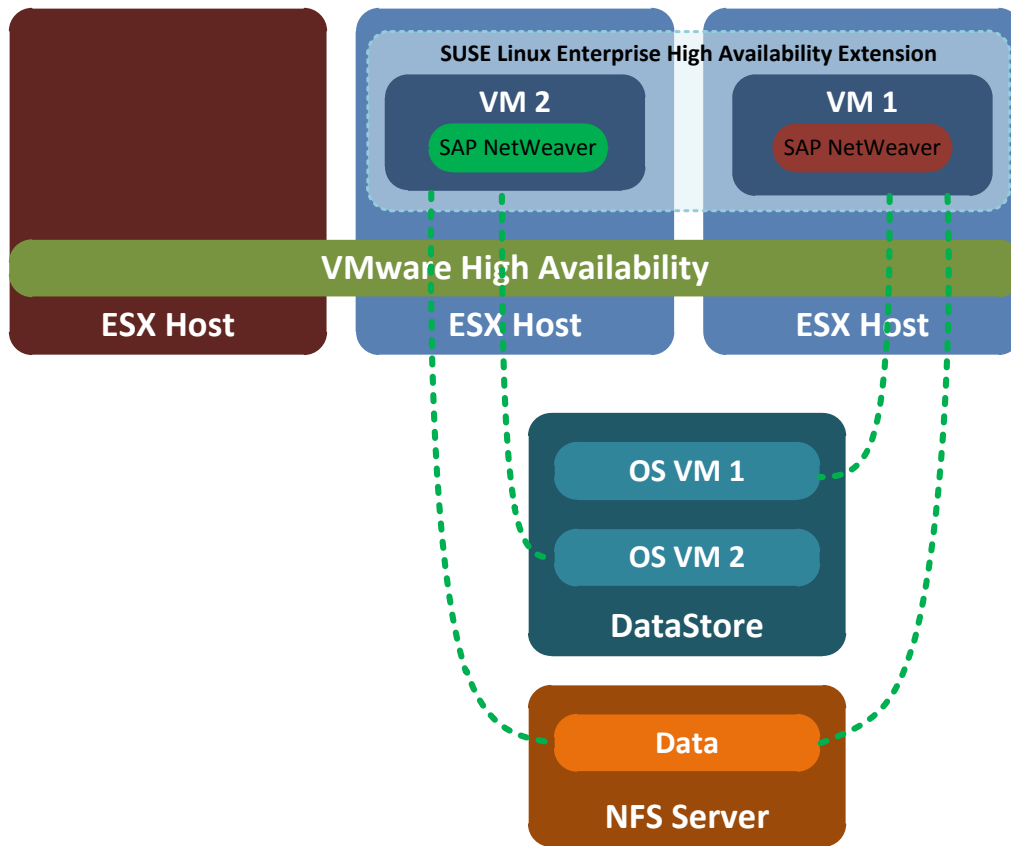
10.3.1 Example 1: VMware High Availability and SUSE Linux Enterprise High Availability Extension with NFS

In this first example, *SAP NetWeaver* is *protected against possible faults at the application layer by the SUSE Linux Enterprise High Availability Extension*. The SAP NetWeaver components described here (see **sections 7.1 and 7.2**) are monitored by the SUSE Linux Enterprise High Availability Extension. *In the event of a failure, they are moved to the second cluster node by restarting the affected services or a full failover of the entire application structure*. The shared application data area (for database and file systems) in this configuration is bound *to the cluster members by a central NFS server via NFS*.

The NFS server itself should have a redundant design to increase availability or be protected by a dedicated cluster. VMware vMotion can be used in this configuration. This allows, for example, the redundancy of guest systems to be reestablished after a host failure.



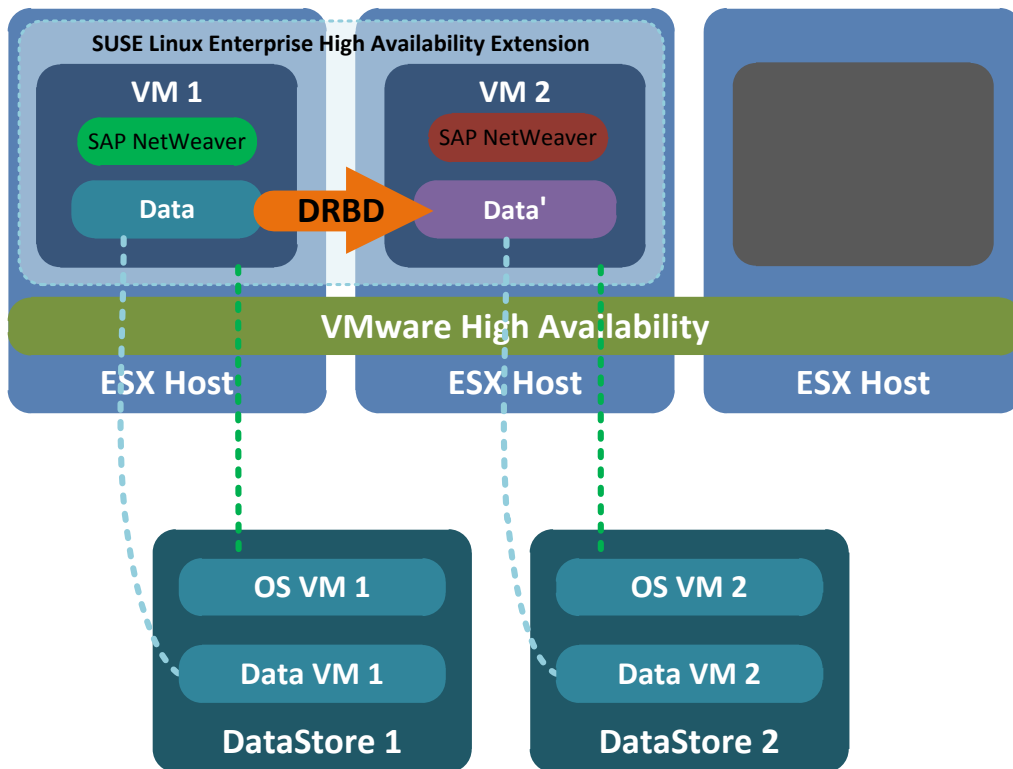
At the *hardware layer*, *VMware High Availability* protects against potential faults. If the first VMware ESX (where the active cluster node for virtualized machine VM1 is located) host fails in this configuration, VMware High Availability restarts virtualized machine VM1 on the third VMware ESX host *with a time delay*. During the time delay, SUSE Linux Enterprise High Availability Extension has already begun restarting SAP NetWeaver in virtualized machine VM2. The cluster framework prevents machine VM1 from attempting to start the application, although the same action is already being performed in virtualized machine VM2.



A *benefit of this configuration* lies in the support for VMware High Availability in conjunction with VMware vMotion. A running virtualized machine can thus be moved between two ESX host systems without influencing or impairing the operation of the cluster software and its monitored applications.

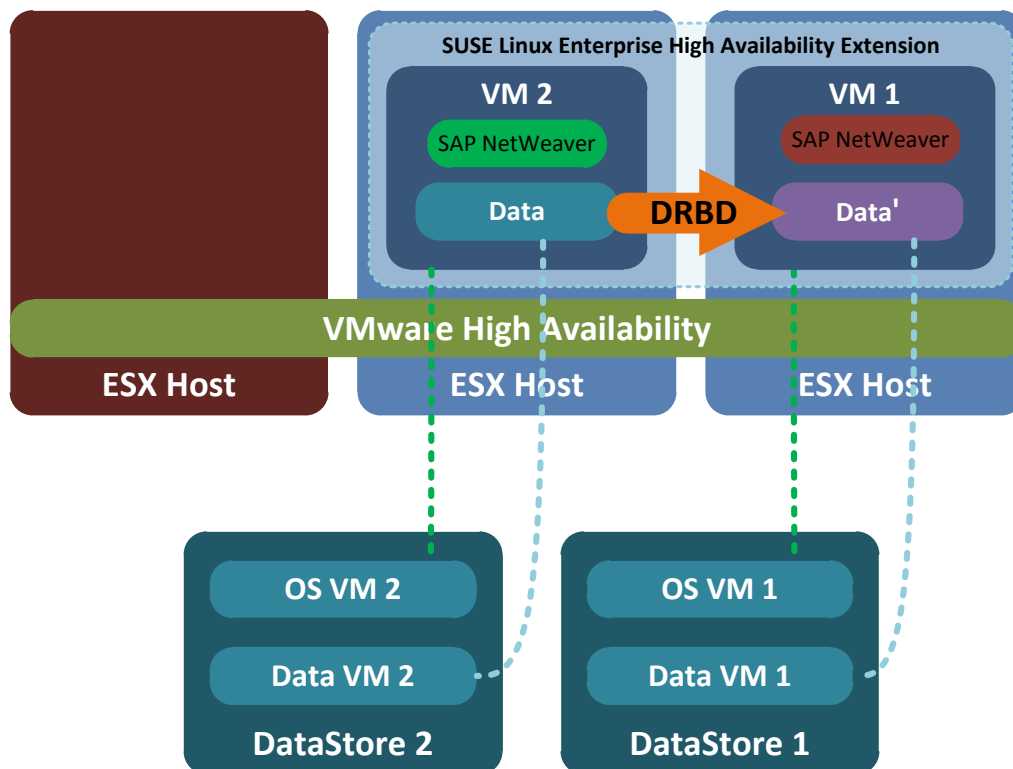
In this configuration, it must be ensured that sufficient network bandwidth is available. Otherwise, the data transfer over the network would negatively affect the data throughput and latency.

10.3.2 Example 2: VMware High Availability and SUSE Linux Enterprise High Availability Extension with DRBD



Protection at the application layer is, as in the previous example, ensured by using the *SUSE Linux Enterprise High Availability Extension*. It monitors the services required for SAP NetWeaver and in the event of failure it restarts them or moves them to the other cluster node. In this example configuration, host-based replication based on DRBD is used for the entire data area (databases, file systems). Each virtualized machine possesses dedicated, directly assigned virtual disks (VMDKs), one of which is used for the operating system and the second is used separately for the shared data area. The content of the data area is now replicated within the virtualized machine from the active node to the passive node using DRBD.

At the *hardware layer*, VMware High Availability protects against potential failures, just as in the first example. If the first VMware ESX host fails in this example, VMware High Availability restarts virtualized machine VM1 on the third VMware ESX host *with a time delay*.



The *benefit of this configuration* lies in the ability to use VMware vMotion to move the *active virtualized machine* in case of performance problems or for planned maintenance work on ESX host systems.

Minor performance penalties are to be expected due to host-based replication within the virtualized machines, as the data being transferred has to pass through multiple layers (datastore, ESX host, virtualized machine, and operating system) of the model.

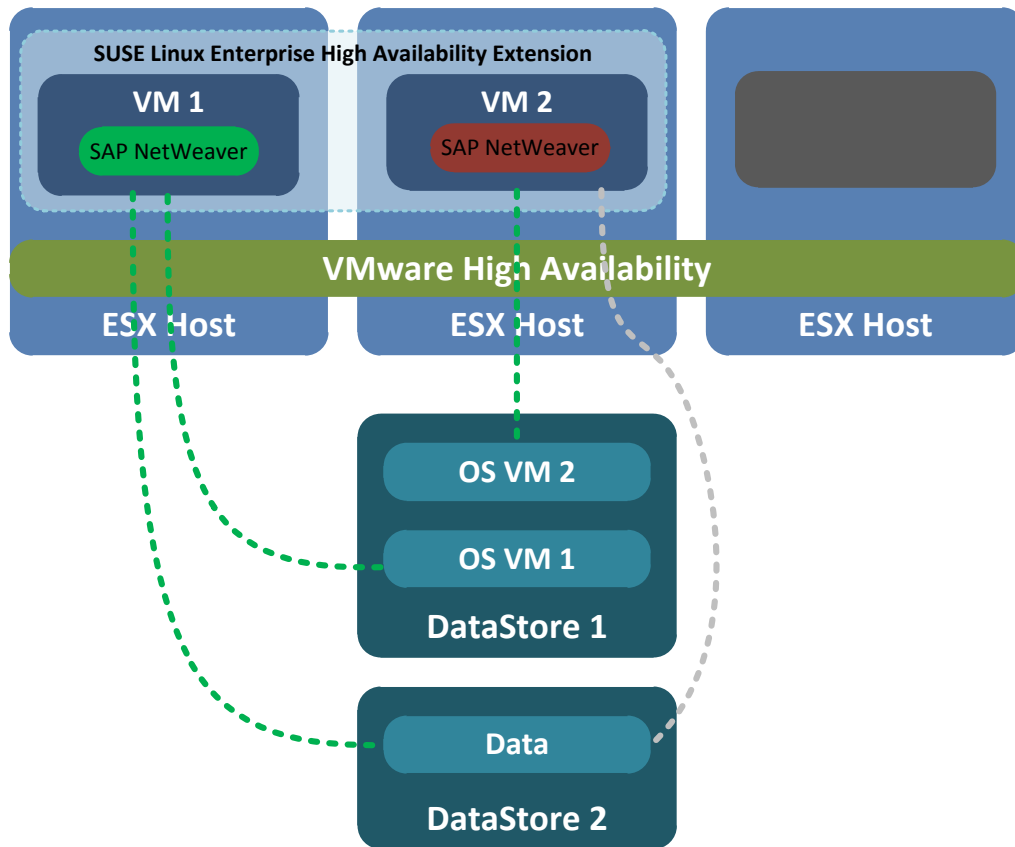
The *additional storage space* required due to host-based mirroring between both nodes using DRBD provides additional data redundancy when storing the mirrored volume on physically separate storage areas.

10.3.3 Example 3: VMware High Availability and SUSE Linux Enterprise High Availability Extension with Shared Disk

In the third example, no host-based mirroring is used for the SAP NetWeaver data areas. Instead, this *data area* (databases, file systems) is *made available* to both virtualized machines, i.e. both cluster members, for direct access as a *shared disk (VMDK)*. Since VMware vSphere 4, this function has been implemented with the *“multi-writer” option* described in VMware Knowledge Base article 1034165 (see References **12.2**; see also **section 6.4**). All virtualized machines in the cluster therefore have access to the shared disk at any time.

Exclusive access (for read and write operations) to this encapsulated VMDK disk must be ensured by the cluster software within the virtualized machines (I/O fencing).

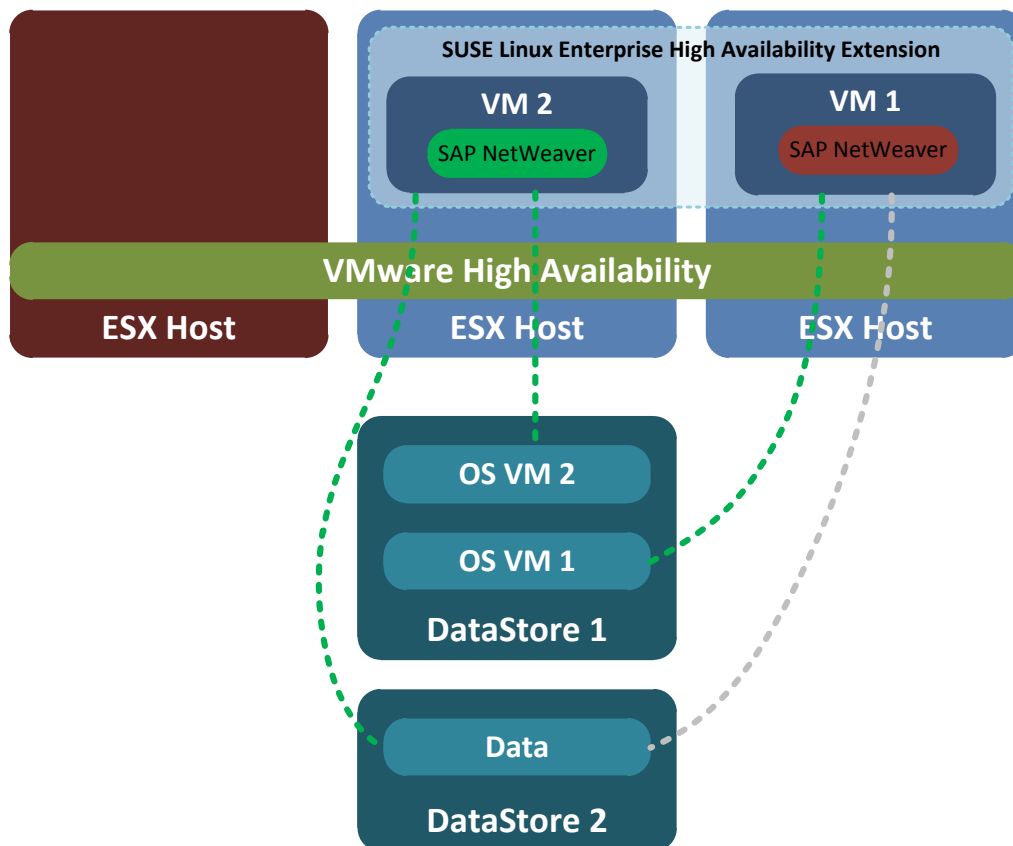
Alternately, a cluster file system (such as Oracle Cluster File System 2) can be used to enable competitive access by the involved cluster nodes. In this case, the entire data area is permanently bound to all cluster nodes and the integrity of the data must be ensured by the SUSE Linux Enterprise High Availability Extension.



As shown in the previous examples, *the SUSE Linux Enterprise High Availability Extension protects the SAP NetWeaver application and its services against potential failures.* The elimination of the DRBD resource reduces the load on the virtualized machines' virtual CPU.

The use of *VMware High Availability at the hardware layer accordingly offers protection against failure* of an ESX host system. With the introduction of the "multi-writer" option in VMware vSphere Version 4.0, VMware vMotion can also be used.

If the first VMware ESX host fails as in the previous examples, VMware High Availability restarts virtualized machine VM1 on the third VMware ESX host *with a time delay*. During this time delay, SUSE Linux Enterprise High Availability Extension has already begun restarting SAP NetWeaver in virtualized machine VM2 through a failover. When starting cluster node VM1, the SUSE Linux Enterprise High Availability Extension first checks which SAP NetWeaver resources have already been committed and starts the resources that are still uncommitted.

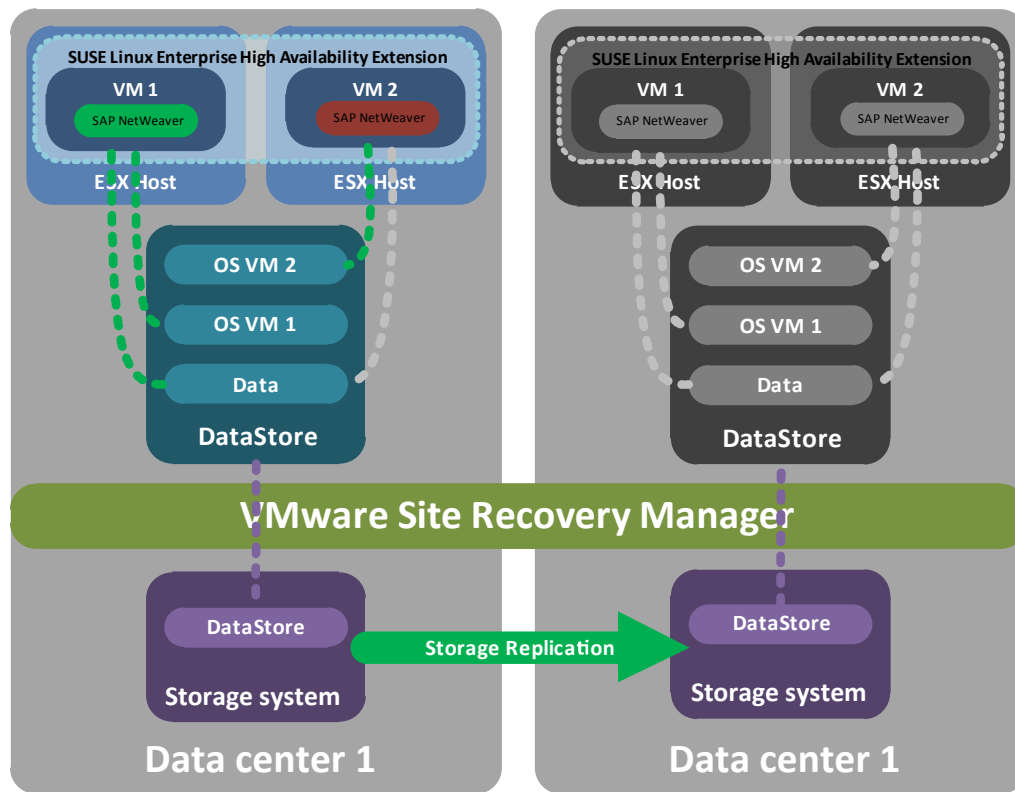


The *benefit of this configuration* lies in the relatively simple and flat structure and the flexible design in terms of performance of the virtual disks located on more than one datastore.

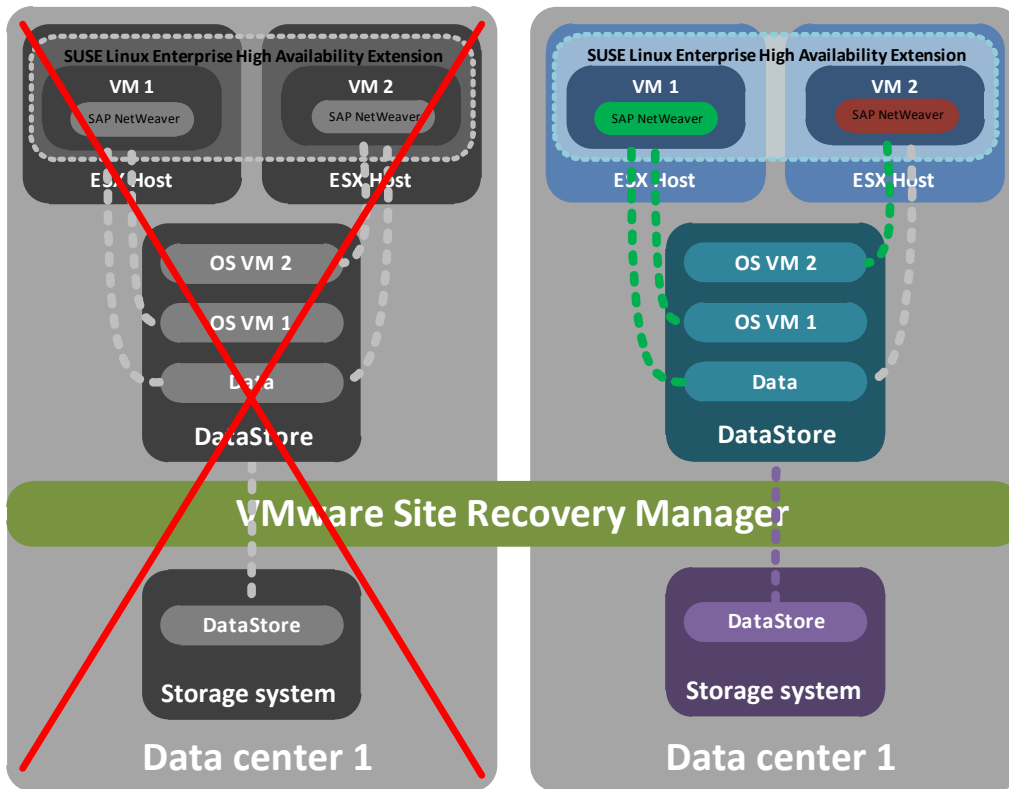
10.3.4 Example 4: Disaster Recovery Scenario

The fourth example deals with the *failure of an entire data center*. Protection against this scenario requires that the *complete infrastructure* of the primary data center be *redundantly available* in a second data center. In this configuration, *functions such as VMware High Availability and VMware vMotion are completely absent* between the data centers. *VMware Site Recovery Manager* (see **section 5.7**) is used to support switchover between the data centers. Dedicated mechanisms at the storage system layer handle replication of the complete datastores, rendering the process invisible to the entire VMware environment. Alternately, this replication can be replaced with host-based *DRBD replication*, a component of the *SUSE Linux Enterprise High Availability Extension*. The possible configurations are not detailed here.

As in the earlier examples, the SAP NetWeaver application is protected in the respective data centers by the *SUSE Linux Enterprise High Availability Extension*. In this configuration, the cluster software performs a failover of SAP NetWeaver upon failure of a service or an ESX host system.



Only when all ESX host systems or the storage at the primary site fail is the *entire VMware vSphere environment restarted in the secondary data center* with support of the VMware Site Recovery Manager. *Replicating data at the storage system layer ensures that a complete, consistent copy of the complete data set is available at the secondary site at all times.* This allows the VMware environment to start very quickly in the second data center.




When equipped accordingly, this configuration allows *protection of the entire VMware environment against disasters*. *Disaster training* held at regular intervals is an important *factor in the success* of this solution in particular. Mastering the mechanical aspects of the process is important, while eliminating differences in the configuration of both environments is as well.

Using fully redundant resources (dual server, network, and storage systems) makes it possible to run *full operations in the secondary data center* in accordance with the Service Level Agreement. In this sense, it is not necessary to switch operations back to the primary data center. *Complete switchover* of the environment to the secondary data center *also simplifies the disaster recovery process* because dependencies between systems are retained.

10.3.5 Example 5: Deploying a High Availability NFS Service as a Central Resource

Network File System Services (NFS services) are used to enable access for multiple clients to centrally stored data and configurations from the network. NFS services are widespread in Linux environments and are used often.

The *availability of this data storage* is correspondingly important. One way to achieve the desired availability is to *replicate the data between two systems* using *DRBD*. By also using *Pacemaker* as the *cluster resource manager*, both copies of the data are combined to form a cluster resource. The exported file systems of the NFS service *are also protected*. These are accessible via a virtual IP address.



A *high availability NFS service* consists of the following components:

- *DRBD resource for data replication* - the cluster manager switches the resource between the primary and secondary node as needed.
- *LVM volume group*, which is available on the currently active DRBD node.
- One or more *file systems* in logical volumes of the volume group that are mounted by the cluster manager on the currently active side of the cluster.
- One or more *virtual IP addresses* allocated by the cluster manager to the active node and enabling the NFS clients to connect to the NFS service. These IP addresses are therefore independent of the active physical node. A virtual NFS root export is required for NFS v4 but not for NFS v3.
- One or more *NFS exports* on which the data reside.

10.4 Running Business-Critical Applications in High Availability Clusters

In the previous sections, various technical options for increasing the availability of SAP NetWeaver as a business-critical application were presented. These examples were supplemented with a proposed solution for how SAP NetWeaver operations can be resumed quickly in a second data center after a disaster.

It should, however, also be noted that the *use of clusters* to increase the availability of business-critical applications must also *be given consideration at the management level*.

An important *requirement* for successfully deploying cluster technology is *well documented operating procedures* such as:

- *Instructions for management* that deal with the operational specifics of the application in the cluster.
- *Support in debugging and fault analysis* that takes the specific aspects of the cluster into account.
- *Notes on special issues* during installation of patches and updates in the cluster.
- *Instructions for adapting or extending* the cluster configuration.

Documents and training that provide support in case of faults or a disaster are especially important. These include:

- Regular failover and disaster *tests during maintenance windows* to gain experience with the technology and detect possible irregularities or errors.
- *Defining the switchback configuration* and the rules for when switchback to the primary system can occur.
- The *emergency manual* describing what constitutes a disaster and what organizational and technical measures need to be taken in such a case.

11 Conclusion and Review

VMware vSphere 4 is a mature, powerful server virtualization solution. It offers a wide range of features for operating the virtualized systems optimally. Various server and desktop operating systems are supported in a VMware environment. This white paper has exclusively addressed the operation of *SUSE Linux Enterprise Server 11*. The *SUSE Linux Enterprise Server 11* is complemented by the *SUSE Linux Enterprise High Availability Extension* for protection of business-critical applications. Options in addition to the features of the VMware solution have been presented for increasing application availability using the SUSE Linux Enterprise High Availability Extension.

Using *SAP NetWeaver* as an example, the *structure of a complex application* was examined and structured using a layer model. Various technologies ensuring the required availability were assigned to the layers. These are components of the VMware vSphere 4 platform, integrated error protection algorithms, and failover clusters based on the SUSE Linux Enterprise High Availability Extension.

After presenting the individual technologies and their purpose, design rules were created and the possible combination of the different technologies was discussed with respect to the availability requirements for protecting an SAP NetWeaver platform in a VMware vSphere 4 environment.

11.1 Consolidation of the IT Environment

A significant consideration in deploying a virtualization technology is the *consolidation of the existing IT environment*. Using the *VMware vSphere 4 platform*, a wide range of systems with various operating systems can be virtualized. The white paper presented that VMware vSphere 4 can provide a fully functional replacement for virtualization technologies such as LPAR in x86 systems. The following can thus be achieved in the first step:

- Reduction in the number of physical systems
- Standardization of the systems through virtualization
- Reduction in the hardware diversity of existing systems
- Establishment of a granular capability to extend the technical virtualization platform with x86 systems

This white paper shows that business-critical applications can be operated in accordance with availability specifications while running under virtualization. Using the sophisticated virtualization technology and the *performance of the x86 hardware platform, demanding applications such as SAP NetWeaver* can be run at high performance in a VMware vSphere 4 environment. Most UNIX applications are also available in Linux or Windows versions, which provides further possibilities for consolidation when using a VMware vSphere 4 platform.

- Reduction in the number of operating system platforms
- Virtualization of systems with business-critical applications
- Virtualization of systems with increased performance requirements

The goal of consolidation in a data center can be achieved in full using VMware vSphere 4.

11.2 Increasing Availability

To take appropriate measures to increase the availability of business-critical applications, their *structure first needs to be analyzed*. In complex applications such as SAP NetWeaver, a *large variety of components* act together to ensure the overall functionality.

This white paper mentioned the critical components for availability and presented options for protecting them. Examples of three high availability products and their utilization in different UNIX environments were described. In the following presentation of the technologies in the VMware vSphere 4 platform and the possibility of extended application protection using the SUSE Linux Enterprise High Availability Extension, it was shown that *availability requirements* as implemented in a UNIX environment can also be *achieved reliably in virtualized Linux systems*.

Combining the VMware vSphere 4 platform with the SUSE Linux Enterprise High Availability Extension yields the following benefits for protecting SAP NetWeaver:

- System protection based on VMware vSphere 4
- Straightforward support for system protection for a large number of guest systems
- Consolidation through common technology for Linux and Windows systems
- Support for Fault Tolerance for guest systems
- Technology for transparently provisioning system and storage resources (VMware vMotion, VMware Storage vMotion)
- Integration of the SUSE Linux Enterprise Server High Availability Extension in virtualized systems
- Resource agents of the SUSE Linux Enterprise Server High Availability Extension for extended application protection
- Protection of SAP NetWeaver components detected as SPoFs, by the SUSE Linux Enterprise Server High Availability Extension
- Various options for providing storage to virtualized systems with application protection (NAS, RAW device, VMDK)
- Using various capabilities of VMware vSphere 4 minimizes switchover times for protected applications after failure

The *combination of both products as described here* enables a *very adaptable high availability solution* for virtualized business-critical applications. The actual implementation always leads to a *custom solution* requiring good preparation.

11.3 Options for Disaster Recovery

Protecting one's IT environment against disaster always involves *balancing the likelihood of the event* and the *price and capabilities of a suitable solution*. The *ease of use in a disaster* is another factor that is always important when using a disaster recovery solution. *Data protection* plays an *especially important role* in preparations against a disaster.

Even though a distinction should be made between high availability and disaster recovery, they both utilize similar or identical technologies. Using *VMware vSphere 4* and the *SUSE Linux Enterprise High Availability Extension* in the design of a disaster recovery concept offers the following benefits:

- Support for inexpensive x86 hardware
- Capabilities of the VMware Site Recovery Manager for designing the switching behavior to the secondary data center
- Utilization of the capabilities of VMware vSphere 4 to reduce switchover times
- Integration of the SUSE Linux Enterprise High Availability Extension in the virtualized offers the same application availability after switchover as in the primary data center
- Ability to use block replication (DRBD) from the SUSE Linux Enterprise High Availability Extension to protect data in the event of a disaster

The VMware vSphere 4 platform and the SUSE Linux Enterprise High Availability Extension complement each other very well both for generating a copy of critical applications and running the applications in a secondary data center after a disaster. The ability to easily switch all protected applications with their tools to the secondary data center to ensure availability *means that both data centers are identical in terms of Service Level Agreement requirements*. If the hardware and infrastructure facilities are almost identical, there should be no need to quickly switch back to the primary data center. This avoids additional downtime.

For a *disaster recovery concept*, the VMware vSphere 4 platform together with the SUSE Linux Enterprise High Availability Extension offer small, scalable solutions for *data and application protection* as well as the capability *to continue operations* without limitations *at a reasonable price* in an equivalent secondary data center.

11.4 Operating the Virtualized System Environment

When *planning a virtualization solution*, *operation of the platform* needs to be considered. Operational concepts should contain *solutions for monitoring, failure protection, and emergencies*. It cannot be assumed that employees with the skills to spontaneously implement an emergency switchover will be available in all situations.

VMware vSphere 4 provides various options for supporting operation of a VMware environment:

- *VMware vCenter* as a central monitoring platform
- *VMware Site Recovery Manager* for failover management between data centers
- Interface for *integration of third-party applications* in VMware vCenter
- Interface for *integration of application protection tools* and VMware High Availability
- Support for *patching guest systems* through the VMware vCenter Update Manager
- Provision of *system management tools*

VMware vSphere 4 thus provides ideal conditions for operating a virtualization platform.

12 Appendix

12.1 Trademarks

- Hewlett-Packard® is a registered trademark of the Hewlett-Packard Company.
- HP-UX® is a registered trademark of the Hewlett-Packard Company.
- IBM®, System p®, System z® are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide
- Linux® is a registered trademark of Linus Torvalds.
- MC/ServiceGuard® is a registered trademark of Hewlett-Packard Company.
- Microsoft® is a U.S. registered trademark of Microsoft Corporation.
- NFS® is a registered trademark of Sun Microsystems, Inc.
- Novell® is a registered trademark of Novell, Inc.
- Oracle® and Java® are registered U.S. trademarks of Oracle Corporation, Redwood City, California.
- SAP® and SAP® NetWeaver are trademarks or registered trademarks of SAP AG in Germany and in many other countries.
- SUSE® is a registered trademark of SUSE LINUX AG, a Novell company.
- UNIX® is a registered trademark of the Open Group.
- VMware® is a trademark of VMware, Inc. in the United States and/or other jurisdictions.
- VERITAS® and VERITAS Cluster Server are trademarks or registered trademarks of VERITAS Software Corporation in the USA and/or other countries.
- All other marks and names mentioned herein may be trademarks of their respective companies.

12.2 References

Haas, Florian. 2010. *Highly available NFS storage with DRBD and Pacemaker* [PDF document]

Herschel, F., et al. 2010. *SAP on SUSE Linux Enterprise-Best Practice.* [PDF document] Massachusetts : Novell, Inc., 2010.

Novell. 2011. *SUSE Linux Enterprise Server for SAP Applications*
[URL]:<http://www.novell.com/products/sles-for-sap.html>

Schubert, Mike. 2010. *Einsatz des VMware Site Recovery Manager-Ein Erfahrungsbericht.*
[PowerPoint presentation] Dresden, Germany : interface systems GmbH, 2010.

SAP Note #1552925 *Linux: High Availability cluster solutions*

VMware. 2008. *Getting Started with VMware vCenter Site Recovery Manager.* [PDF document] Palo Alto : VMware, Inc., 2008. EN-000187-00.

VMware. 2010. *Reference Guide VMware vCenter Server Heartbeat 6.3.* [PDF document] Palo Alto : VMware, Inc., 2010. EN-000380-00.

VMware. 2009. *SAP Solutions on VMware® vSphere™: High Availability.* [PDF document] Palo Alto : VMware, Inc., 2009.

VMware. 2010. *SAP® Solutions on VMware® Business Continuanace.* [PDF document] Palo Alto : VMware, Inc, 2010.

VMware. 2010. *VMware® High Availability (VMware HA): Deployment Best Practices.*

[PDF document] Palo Alto : VMware, Inc., 2010.

VMware. 2009. *vSphere Availability Guide.* [PDF document] Palo Alto : VMware, Inc., 2009.

EN-000316-00.

VMware. 2011. *Knowledge Base Article 1034165* [URL]: <http://kb.vmware.com/kb/1034165>

