

Managing a Virtualization Project

In this chapter, we use the example of a large corporation that has decided to accelerate the penetration rate of virtualization within its information system. This company is already familiar with virtualization technology. Approximately 30% of its server environment has already been virtualized. It wants to move on to the second phase and virtualize on a large scale, including its servers that run critical applications. To ensure the project's success, a project team is set up.

Context

The IT system at this corporation consists of global servers with a high obsolescence rate for both hardware and software. With 362 servers (on x86 systems) spread out between two datacenters, more than 50% are older than four years and will reach the end of the manufacturer's support period at the end of the current year. The cost to replace the servers cannot be covered by IT budget allocations, which shrink every year. Certain archaic applications (10 years and older) are still used but for compatibility reasons cannot be migrated with the current hardware and operating systems. The company doesn't even have the documentation needed to reinstall some of these applications.

The IT department is under pressure. It must answer both to internal clients with new needs (40 new requests per year) and to corporate management, who demand higher service levels coupled with cost reductions. In a tight economic context, internal IT teams bear a significant workload because many agreements with service providers have not been renewed.

The datacenters are another source of great concern—they are reaching their limits in terms of floor space and energy consumption (only a few kilowatt in reserve). No new hardware can be added, which prevents any new project from being undertaken.

A *disaster recovery plan (DRP)* solution is in place, but it is not satisfactory; only 5% of the applications are secured on remote sites, and tests are sporadically performed.

Backups are also problematic; several solutions are in place, but there is no global management, and backup windows are often too narrow. Many restorations from tapes fail due to a lack of backup testing.

The company decides to rely heavily on the virtualization technologies best suited to face these issues. It decides to proceed with a large-scale deployment of virtualization for all applications, even the most critical.

Objectives

The IT department defines its objectives as follows:

- Cost reduction:
 - Reduction in the datacenter's power consumption
 - Infrastructure rationalization
 - Automation of certain tasks performed by administrators
- Securing the production site by implementing a DRP with a *recovery point objective (RPO)* of 0 and an *recovery time objective (RTO)* of less than four hours on the entire IT system.
- The IT system must be the company's growth engine. It must favor innovation, creativity, and new projects.

The IT department also identifies the following future projects that are outside the scope of the immediate objectives:

- Provisioning cloud-type services
- Providing end users and project teams with easy access to IT services
- Implementing chargeback tools
- Implementing a unified backup solution
- Implementing a third replication site
- Standardizing processes

Criteria to Choose a Solution

To evaluate the relevance of existing virtualization solutions in reaching its goals, the company looks at offerings from three top companies: Microsoft, Citrix, and VMware. After an in-depth examination of the different solutions, the company chooses VMware based on the following criteria:

- The product's maturity within large datacenters
- The upgradeable nature of the solution
- The possibility of attaining high levels of service through VMware's advanced features
- A single tool to administer and monitor the virtual platform and DRP; integration of Site Recovery Manager within vCenter
- A robust ecosystem, with numerous options (including migration tools, wide compatibility with backup tools, and reporting)
- The ability to implement a global backup policy
- Preexisting in-house VMware competencies

Project Stages

As shown in Figure 8.1, this virtualization project is planned in four stages: Planning, Design, Implementation, and Management.




Figure 8.1 Planning stages for virtualization project.

For any virtualization project, including critical production servers, the goal of the planning stage is a preliminary eligibility study. The planning stage includes two steps:

1. *Discovery*, which involves collecting all information on servers, storage, and production equipment to draft a production plan.
2. *Analysis*, which involves interpreting results from the discovery step and drafting optimized consolidation scenarios.

The Design stage involves proposing various target scenarios, taking into account the results of the eligibility study, and providing the technical specifications and performance metrics required for the target platform. This guarantees a level of performance at least equivalent to what it was before virtualization. (In fact, it is often much higher.) During this stage, the detailed reference architecture is drafted. This document serves in the drafting of the equipment specifications so that the various providers offer solutions that meet the needs as defined.

The Implementation stage defines the time phasing of migration for priorities defined by the IT department, such as end-of-server warranty contracts, urgency, or capacity. If the project's priority is to decrease energy consumption (kWh), groups of the most energy-intensive applications will be defined and migrated first. The implementation of the target platform will be accomplished by using *physical-to-virtual (P2V)* tools and in compliance with the detailed reference architecture.

The Management stage allows the application of best practices to add operational efficiency to the environment. Strict rules must be followed to maintain control. Tools complementing those offered by VMware can be used to manage day-to-day operations.

Planning

As mentioned in the preceding section, the Planning stage consists of discovery and analysis (see Figure 8.2).




Figure 8.2 At the Planning stage.

Discovery

This stage includes the *production plan*—the exact inventory of the servers and precise information about the use of the various elements over a defined period. Information collected is used to size the target architecture and determine which servers are not eligible for the virtualization project. Information collected is also used for the proper implementation of *Site Recovery Manager (SRM)*.

The project's scope includes x86 processor-based servers. Other types of servers, such as UNIX or mainframe, are excluded.

The study to gather information is realized through a collection server. The infrastructure’s physical servers send back all data related to production over a period that is representative of the activity. Usually, a period of 30 days is selected, excluding July and August (months that often are not representative of a typical production plan due to a slowing of business activity). The analysis tool chosen for this particular project is PlateSpin PowerRecon. The analysis stage can begin at the end of the 30-day period, covering the following technical angles listed in Table 8.1.

Table 8.1 Planning Objectives During Discovery Phase

Audit Theme	Goals	Notes
OS type	Verify the compatibility of the OS with VMware vSphere 5’s compatibility matrix.	Unsupported operating systems are not eligible, because this means there will usually be no tools allowing P2V or manufacturer support for the OS.
CPU use rate	Allows the determination of the average use rate of the current server park’s calculation power. Useful for the sizing of the target architecture.	Servers with intensive CPU use are not usually good candidates for virtualization, but if the goal is to simplify the <i>DRP</i> , even a processor-intensive server may be put on vSphere (1 VM for one ESXi host).
CPU usage pattern	Allows figuring whether the CPU activity can be predictable (for specific hours of days of the week) or whether it is random.	Allows the determination of which servers are “at risk.” One determining factor to validate a physical machine’s virtualization eligibility.
Memory use determination	Allows the determination of the memory’s average use rate.	Memory is one of the first elements to take into account when sizing a virtual architecture.
Storage volume sizing	Defines required storage capacities and performance.	The choice of the destination storage architecture is extremely important. Storage is the most critical component of a virtualization environment, because if storage-space and performance (in IOPS) sizing is not defined correctly, the project could fail.
Network connections	Defines prerequisites to ensure network streams in optimum conditions.	The nature of servers (supported applications) and the report scheduling (batch, backups) should be analyzed in detail to determine network activity.
Profiling the application	Allows the determination of the criticality and performance of each application. The switch to virtual should be done only when service levels and performance are at least equivalent.	Special attention should be given to critical applications for their protection and during migration phases.

NOTE

During this period, it is imperative that the data collection operations and the team in charge do not interfere with current production, respecting the company's business reality.

Data collection follows a logical process. The activity of CPU, memory, network, and disk components is collected at all times. Every five minutes, a maximum value is retrieved. After one hour, an average is generated from these maximum values. Each hour is calculated in this manner to establish a profile over 24 hours. In this way, the average value over one month is found for each server. Each server's profile can be calculated, and a graphic representation can be created, as shown in Figure 8.3, showing the activity from all machines. The collection allows the precise determination of the servers' activity and their production plan.



Figure 8.3 Representation of the datacenter's production plan.

In Figure 8.3, each color stacked in the histogram represents the activity of one machine within the datacenter. The collected activity accumulation represents the maximum load observed that the virtualized platform will need to offer when it is in place to ensure, at a minimum, a performance equivalent to pre-virtualization. Note that a large portion of activity occurs at night: backups, batches, database indexing, defragmentation, and so on. On some servers, the busiest time is at night.

As shown in Table 8.2, limit thresholds must be set to determine which servers are eligible for virtualization. These thresholds are based on VMware recommendations as well as our own experience. Every organization may set the limit thresholds differently.

Table 8.2 Limit Thresholds

	Set Thresholds
CPU	6 GHz
Memory	6 GB
Disk	Disk transfer 1700 IOPS
	Bandwidth 20 MBps
Network	20 MBps

Servers above these thresholds will be isolated and not eligible for this virtualization project.

Operating Systems

As shown in Table 8.3, the collection provides a precise description of the servers.

Table 8.3 Distribution of Existing Physical and Virtual Servers

Total Number of x86 Servers	Number of Physical Servers	Number of VMs	Number of Servers Connected to the SAN	Number of VMs Replicated on the SAN
362	242	120	50	35

362 physical and virtual servers are included in the scope of this project. Approximately 30% of the servers have already been virtualized using vSphere 4.1. The majority of servers use internal storage, although some use *storage-area network (SAN) Fibre Channel (FC)* shared storage (50 servers). Only 35 servers are secured on a remote site. The goal for this project is to virtualize the physical servers and to reuse the VMs that already exist. Figure 8.4 shows the operating system environment.

Windows is used in 95% of the environment (four different versions of Windows). The vast majority use Windows 2003, but a trend toward Windows 2008 Server can be observed.

NOTE

All operating systems are found in VMware’s hardware compatibility matrix.

CPU Data Collection

Of the physical servers, 50% are uniprocessor servers and 44% are biprocessor servers. Only 6% use quad processors.




Figure 8.4 Operating systems in use.

NOTE

Experience shows that the more sockets a server has, the more difficult it will be to virtualize because of the presence of an application requiring significant resources.

An in-depth view (see Figure 8.5) of the servers' CPU use shows the activity.

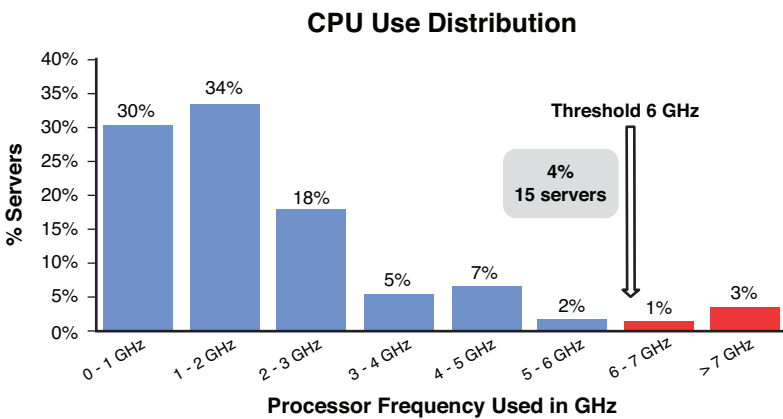


Figure 8.5 Server CPU activity.

As you can see in Figure 8.5, the majority (87%) of servers make reasonable use of the CPU (less than 4 GHz). There are 47 servers (13%) with high CPU activity (more than 4 GHz). 4 percent (15 servers) use the CPU above the threshold set at 6 GHz.

NOTE

These figures are global averages and do not rule out excessive use on a particular server. The average is calculated over periods representative of the activity.

Memory Data Collection

When the memory installed in all servers is added up, you can see that 64% of all memory is used; 36% is free (see Figure 8.6).




Figure 8.6 Total server memory usage.

An in-depth view, displayed in Figure 8.7, shows how used memory is distributed.

As Figure 8.7 shows, 90% of servers use a capacity equal to or less than 8 GB of RAM; these can be virtualized. Three percent use a significant amount of memory—more than 8 GB. These 10 servers are mainly SQL and Oracle databases. They are above the set threshold and will not be eligible for this project.

Servers with a significant amount of memory used must be analyzed in more detail to confirm whether they can be virtualized. In vSphere 5, licensing is partially linked to the configured memory of the VMs, so this must be taken into account when large memory capacities are necessary.

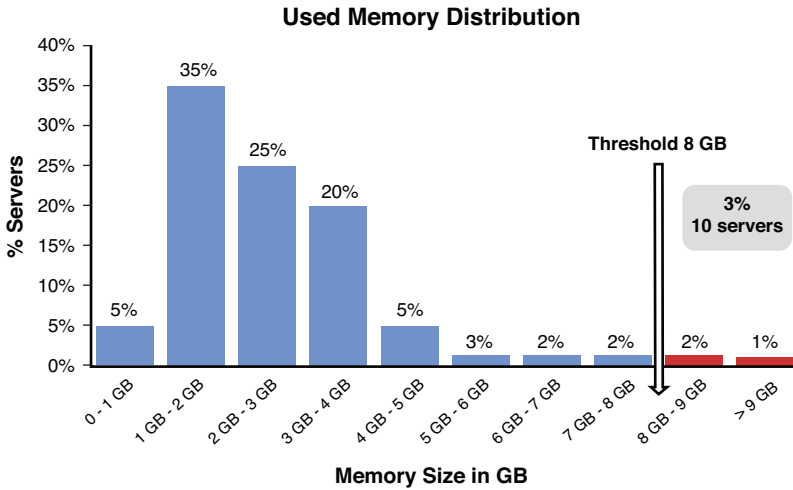


Figure 8.7 Detailed presentation of used memory distribution.

NOTE

A server with significant memory usage should not be excluded in all cases. Indeed, if the main goal is to set up a DRP simply, it can be interesting to keep such a server within the scope of the project.

Disk

There are 312 servers using local storage, and 50 are connected to the SAN. The total storage capacity (local and SAN) of the 362 servers is 25 TB, including 15 TB used, or 55%. This makes an average capacity per server of 40 GB used for 70 GB provisioned. Implementing an architecture with centralized storage will allow consolidation and improve the use rate of the storage.

NOTE

This information serves as a base to size storage. The size needed is one of the indicators used to choose the correct type of disk to use: vmdk, RDMp, or RDMv.

The graph shown in Figure 8.8 illustrates the number of servers per disk capacity used.

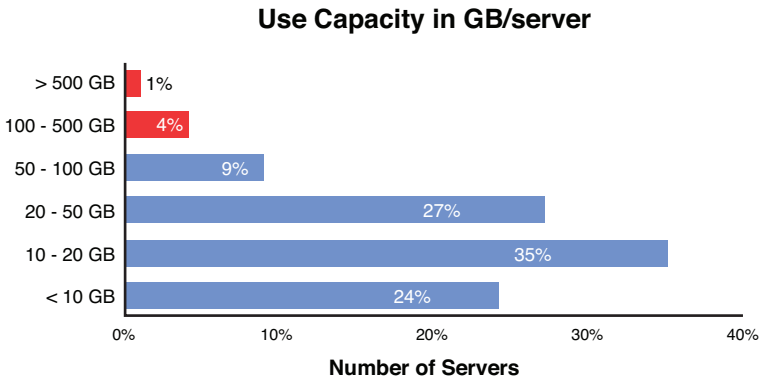


Figure 8.8 Number of servers per disk capacity.

Eighty-six percent of servers use a reasonable capacity, below 100 GB of disk, and these can be virtualized. Fourteen percent have significant volumes, including two above 1 TB. Large-volume servers require an in-depth study.

NOTE
Using *Raw Device Mapping (RDM)* mode disks is appropriate for large volumes.

Bandwidth is an important aspect of eligibility. Generally, a good candidate for virtualization requires an I/O disk bandwidth below 20 MBps. Fortunately, as shown in Figure 8.9, most servers fall below that limit.

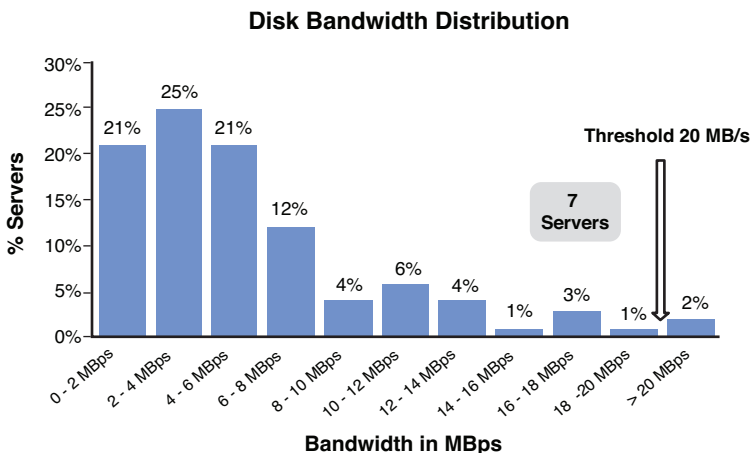


Figure 8.9 Two percent of servers (seven servers) require more bandwidth than 20 MBps.

An in-depth study is needed to identify values and periods. I/O activity, shown in Figure 8.10, is a critical element to monitor when servers hosting virtual machines share disks.

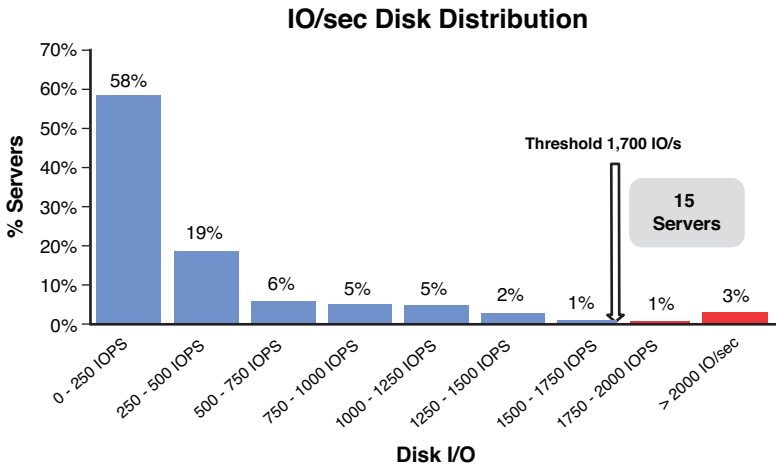


Figure 8.10 IOPS activity.

As you can see in Figure 8.10, the servers’ overall IOPS activity is reasonable. Fifteen servers have significant activity, greater than 1700 IOPS, and require in-depth study.

The peaks primarily take place at night. Activity is not necessarily a hindrance in transitioning to the virtual, but it comes at a price. Indeed, for the VM to offer the performance required, it must have its own dedicated RDM disks with a sufficient quantity of disk spindles to absorb the load.

If it is part of the target, dedicating portions of SAN for this VM is recommended. What would prevent virtualization would be caused by financial concerns, not technological reasons.

Network

As shown in Figure 8.11, 264 servers are using bandwidth of less than 20 MB/s. 10 servers use bandwidth greater than 20 MBps. Among these are servers for backups, antivirus software, and scanning applications. In a virtualized environment, 20 MBps is the limit for a machine’s network traffic. A Gigabit Ethernet adapter usually does not provide more than 70 MBps.

The peaks primarily take place at night and correspond to backup periods. A discussion involving the client must take place regarding backup issues of the target infrastructure. The network activity can be offloaded if a SAN architecture is set up.

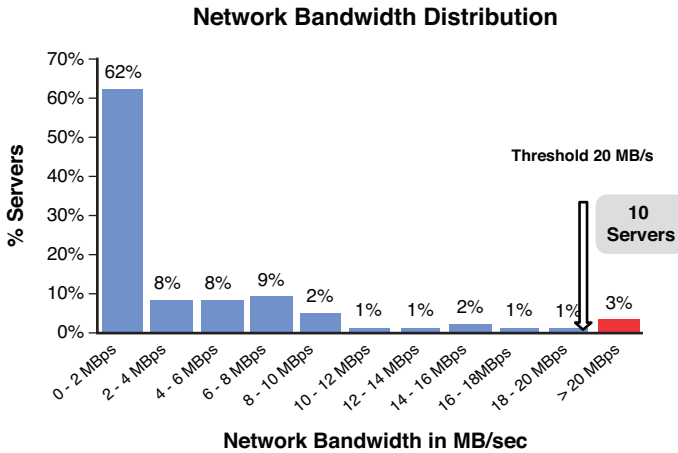


Figure 8.11 Network activity and bandwidth demands.

Applications

The collection provides a precise description of the application environment:

- SQL Server: 124 servers host SQL.
- Oracle Server: 50 servers.
- Approximately 100 existing VMs for tests, development, preproduction, and so on.
- The rest are related to infrastructure servers: *Active Directory (AD)*, *Domain Name System (DNS)*, web, files, printing, BlackBerry, business applications, and so on.

SQL Server: We have found 124 servers/VMs that host SQL. Certain servers have a significant queue length between 5 a.m. and 7 p.m. Performance issues can occur for servers with a queue length beyond 12. If they are retained for virtualization, it will be necessary to increase the number of disk spindles in building these servers' *logical unit number (LUN)*—one disk from the RAID group for every queue length unit above 12. This increase is important, or the contention of target disks could be displaced major concern, which would deteriorate the performance of other VMs.

Oracle Server: 50 servers run Oracle databases. The number of licenses required must be checked directly with Oracle. Some servers are very taxed in terms of CPU and disk I/O. Analyzing these servers will establish whether or not they are candidates for virtualization according to the defined thresholds.

NOTE

The queue length represents the I/O awaiting writing in the operating system. It is a waiting queue for I/O not yet acknowledged on disk. The cause can be an activity peak on the disk or a slowing down between the I/O manager and the file system (for example, an antivirus or a replication driver at the host level).

The queue length audit can identify the bottleneck in the queues between the storage and server. Currently, 14 servers show a queue length that justifies keeping them out of this virtualization project.

Analysis

The average CPU use rate on the company's entire server collection is 19%. This is a good rate compared with what we have seen with many other clients (whose use rate is usually less than 10% to 15%). The large proportion of uniprocessors explains this high CPU use rate. Yet 81% of resources remain unused. The CPU audit identified 24 servers with high activity, including 4 with random behaviors.

The overall server memory use rate is 74%. This is a good average.

In general, ideal candidates for virtualization are servers with predictable activity and relatively low resource requirements compared with what the server can provide.

Collection and analysis provide the average server setup in this company:

Average CPU activity between 1.5 GHz and 2 GHz, 2 GB of used memory, and storage space of 70 GB (40% of which is used), running Microsoft Windows 2003 Server

This average, provided for information only, serves as a reference.

To determine which physical servers will not be virtualized, *exclusion criteria* are based on the following:

- OS or hardware not found in VMware's hardware compatibility matrix (for example, fax cards, dongles, old operating systems)
- Intensive use of resources as compared with previously mentioned thresholds
- Service or performance levels below those of the physical environment (rare)
- Business reasons (for example, very sensitive applications)

Figure 8.12 shows the result of the eligibility inquiry for the servers in this project.

Eligibility Distribution

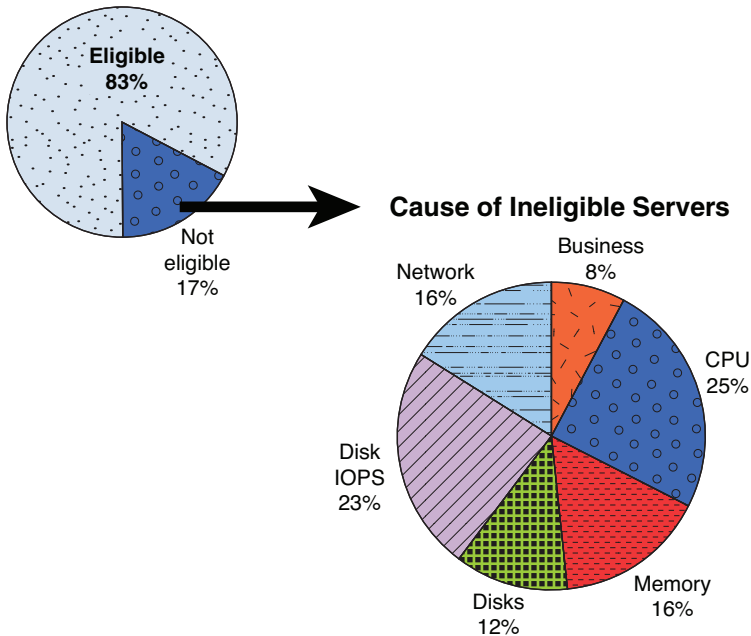


Figure 8.12 Various reasons leading to the ineligibility of servers for virtualization.

Approximately 17% (62 servers) are considered “nonvirtualizable as they are.” They will not be considered for this project because decisions must be made quickly and these servers can become hindrances that put the entire project at risk. What matters at this stage is obtaining a good confidence level from users. Subsequently, virtualization possibilities can be explored. These could include the following:

- **Excessive disk and network bandwidth at certain times of night:** Is there a possibility of smoothing out the load?
- **Excessive CPU use rate between 10 a.m. and 7 p.m.:** Can the activity be spread over several VMs?
- **Excessive network bandwidth at 8 a.m.:** Can the charge be spread over a longer period to obtain a use rate below 20 MBps?
- **Significant queue length on several occasions:** To reduce the queue length, use more disk spindles in RAID groups.
- **A great deal of disk I/O over most of the night:** Is there a possibility of decreasing the load? The backup policy must be revised.

Design

In the Design phase, the focus is on the target architecture and sizing (see Figure 8.13).

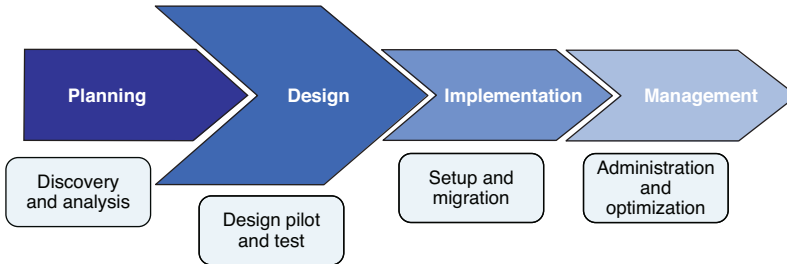


Figure 8.13 Design phase.

Target Architecture

The number of eligible physical servers that will be virtualized is 180 (242 servers minus 62 servers). With existing VMs, the target architecture will have a total of 300 VMs. To meet established goals, the target architecture is as shown in Figure 8.14.

The target architecture consists of two datacenters—Site 1 and Site 2—located 6 miles (10 km) apart. In each location, all ESXi host servers are in a cluster connected to a SAN FC, and the two storage arrays are synchronously replicated. The DRP is taken care of using SRM5.

NOTE

Implementing this type of solution leads to compliance with the Bâle III Agreements published December 16, 2010, requiring (of financial institutions, among others) up to two backup sites.

Sizing

The Discovery stage provides the precise capacities used, which is fundamental information for the best architecture design possible. VMware's SRM5 solution will be implemented to allow the integral recovery of one site on the other in case of a major incident. The architecture is spread over two sites, and all storage is replicated. Servers must be sized to be able to support the load if a site is switched over. The load must be able to handle 300 VMs in DRP. This capacity planning must also take into account upgradeability over three years, on a 30%-per-year basis.

Target Architecture

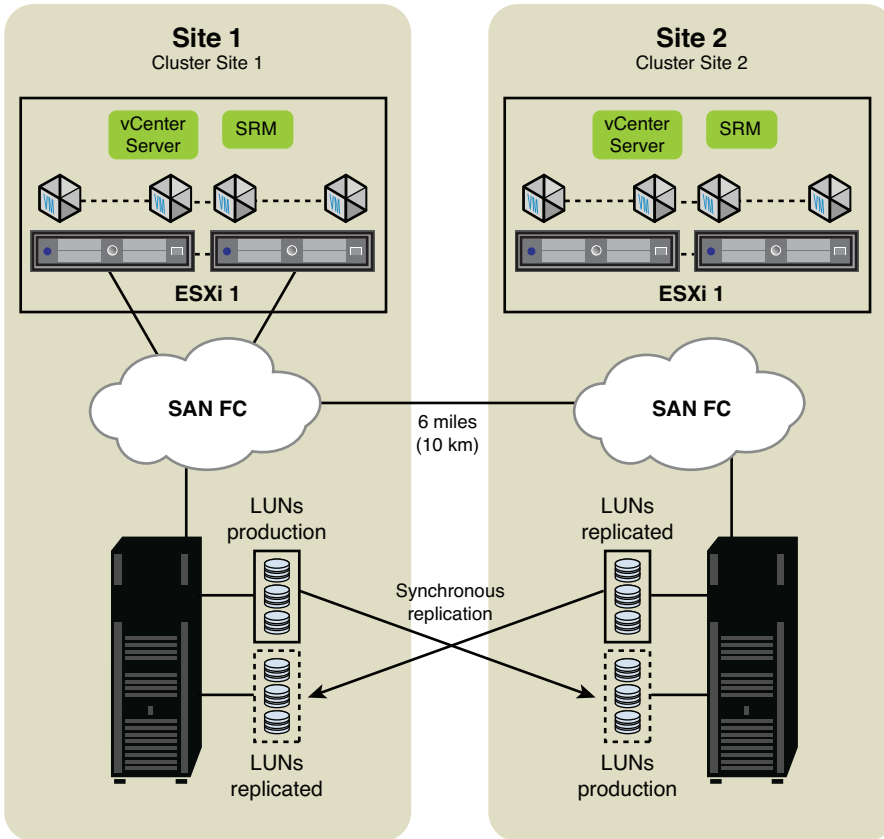


Figure 8.14 The target architecture.

ESXi host server: To provide the power needed in the most demanding scenario, it is necessary to install 16 Intel Westmere E5670 (2 processors, 6 cores, 2.93 GHz) servers (8 servers per site) with 96 GB of RAM per server.

Memory should be one of the first elements considered when sizing the target architecture. The average effective capacity of 2.3 GB per server represents a total memory of approximately 700 GB. Advanced overcommitment memory techniques can guarantee good conditions during activity peaks. They make the amount of memory reliable without the need to oversize the amount of memory.

TIP

A basic rule for sizing is to allot between 4 GB and 8 GB of RAM per physical core. For example, a server with a total of 12 cores should have between 48 GB and 96 GB of RAM.

When choosing a processor, it is preferable to choose a large number of cores.

TIP

Another basic rule for CPU sizing is to allot between 2 and 4 vCPUs per physical core. For example, a server with 12 cores allows 24 to 48 vCPUs to run. This rule applies only for Tier 2 and Tier 3 applications. For Tier 1 applications, allot one vCPU for one core.

Network card configuration is as follows:

- Two physical gigabit network cards for VMs
- Two gigabit network cards for administration
- Two Gigabit Ethernet cards for vMotion

TIP

One Gigabit Ethernet network card usually does not provide more than 70 MBps. A basic rule is to not allow more than 10 VMs per gigabit network card.

With regard to the SAN adapter, we recommend two physically distinct *host bus adapter (HBA)* cards with the addition of redundancy software (for example, EMC PowerPath VE) to take advantage of more advanced features than those offered by VMware (Round Robin). By spreading the I/O load over available paths, this type of software provides administrative comfort and leads to better I/O disk performance without any action from the administrator. If one or several paths are lost, it switches the flux to the remaining paths, automatically detecting HS links.

The required storage sizing is calculated based on the storage capacity actually used for the 300 eligible servers. The following rule is used:

Capacity to provide = (Effective capacity + 30% for upgradeability) + (Total RAM capacity) + (20% for snapshots) + 10% for safety

Therefore, the total volume is 12 TB + 3.6 TB + 1.3 TB + 3.3 TB + 2 TB = 22.2 TB.

The same volume must be reserved for the DRP, so double the preceding result (Total = 44.4) to arrive at the anticipated capacity (see Table 8.4).

Table 8.4 Anticipated Capacity

Site	Production	Replication	Total
Site A	12.2 TB	10 TB	22.2 TB
Site B	10 TB	12.2 TB	22.2 TB

Expect a 30% increase each year (see Table 8.5).

Table 8.5 Estimated Increase Over the Next Two Years

	2012	2013	2014
Number of ESXi5	12.2 TB	14 TB	16 TB
Number of VMs	300	350	400
Storage TB	22.2	29	38

The DRP involves 270 servers. If a site crashes, all VMs, no matter their environment or rank (there are three levels of criticality), must be active on the remaining site.

TIP

For VMFS datastores, we recommend 800 GB with a maximum of 32 vmdk per volume, with no more than 8 ESXi servers connected.

Figure 8.15 shows the configuration to implement.

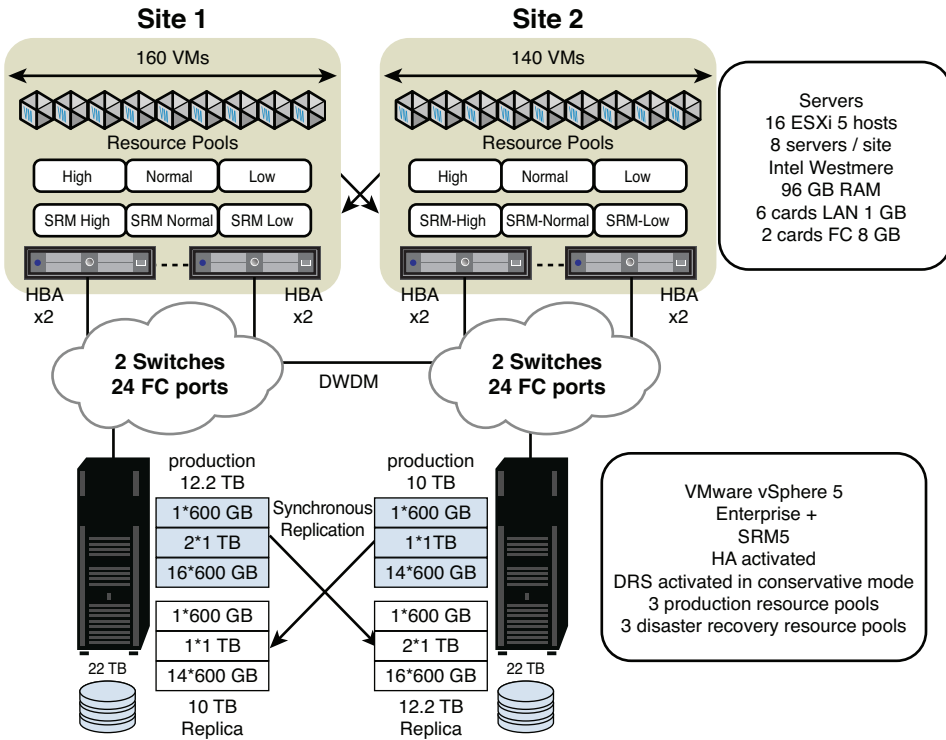


Figure 8.15 Configured target architecture.

Implementation

The previously defined elements are put into service in the Implementation stage, illustrated in Figure 8.16. This stage requires good communication between teams so that the schedule is respected and risk is managed. Prerequisites and deliverables must be rigorously monitored and re-evaluated, if necessary, to limit the zones at risk.

Installing the vSphere 5 Platform

Installation is done in the following order:

1. Physical installation of servers and the storage array in the racks
2. Firmware update of server components and BIOS
3. BIOS configuration

4. Installation of vSphere 5 onto servers
5. Basic configuration of the vSphere server
6. Installation of HBA cards redundancy software
7. Storage preparation, RAID creation, LUN creation, and so on
8. Zoning and masking of the new vSphere server
9. Configuration of local vSwitch and of its administration network
10. Configuration of VMkernel port for vMotion
11. Addition of vSphere server in vCenter
12. Application of host profile
13. Update through the vSphere server's Update Manager
14. Validation of the vSphere server's basic configuration
15. P2V migration

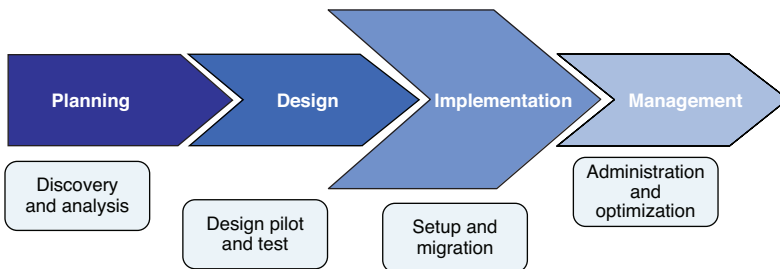


Figure 8.16 The Implementation phase.

P2V Migration

The transformation of physical machines into VMs is organized based on the date at which support ends and on power consumption. The migration will take four weeks. PlateSpin Migrate is selected for the P2V migration because the potential period of time for service interruptions is very narrow. This tool is chosen instead of VMware Converter because it allows the differential synchronization of physical machines to ESXi servers by using planned tasks. Using the same tool, it is also possible to start a target VM and test it in isolation before its final switchover, thus ensuring the VMs indeed boot for critical applications.

Migration scenarios are drafted based on the following priority criteria:

- **Power consumption:** The goal is to rapidly free up energy resources used by the datacenters, using it to deploy all migration waves while ensuring the quality of services if a datacenter is lost. Therefore, the migration of the servers that consume the most power (in kWh) is considered a priority.
- **Hardware obsolescence:** Migrating older servers—some of which are nearing the end of the manufacturer’s warranty—lowers the risk of failure.
- **Migration dependencies:** Groups of physical machines must be migrated taking into account their interdependence.
- **New requirements**

Full synchronization is done the weekend before the migration. On the day of the migration, only blocks modified since the last synchronization point are synchronized and sent to the target VM.

Before performing P2V activities, it is necessary to prepare servers by shutting down antivirus software and resource-intensive applications. When the migration is done, hardware-related tools are no longer needed and must be uninstalled.

To decommission the critical physical machines, simply unplug the network cables for two weeks. (This is preferable to physically stopping the server to avoid hard drive problems.) If the virtualization migration goes wrong and the physical server has been decommissioned, it is possible to switch back through a *virtual-to-physical (V2P)* migration.

Implementation of SRM5

To implement SRM5, the following information must be known:

- Which virtual machines are part of the DRP.
- The topology of the VMs (particularly with regard to storage). Certain VMs can have RDM volumes that are invisible from the VM summary.
- What VM dependencies exist, so as to restart them in the proper order.

The boot sequence is based on criteria such as business needs, service-related VMs, infrastructure, and so on.

EXAMPLE

One boot sequence for the VMware environment is Active Directory, followed by DNS, then the vCenter database, followed by vCenter Server, and so on.

SRM requires thorough organization. VMs must find the same hierarchy of resources in the production site and the backup site. Network resources, resource pools, and VM organization files are remapped between both sites. Therefore, it is crucial to organize and name VMs properly so that you can find your way later.

Consistency Groups

Some storage array manufacturer replication solutions enable you to create consistency groups, which are essential.

Some VMs (for example, databases) can be stored in different datastores. Consistency groups for replication (called datastore groups) allow the logical management of the various LUNs as a single LUN, ensuring coherence between all LUNs.

EXAMPLE

If a VM has a vmdk virtual disk on Datastore A and a second vmdk on Datastore B, it is necessary to create a consistency group for storage replication. The consistency group calculates the interdependence of Datastore A and Datastore B. These datastores cannot be separated, and they form a coherent whole.

A VMFS datastore can be a unique LUN or can consist of several LUN extents. If the datastore consists of several extents, consistency groups guarantee that the replication state of the various LUNs used to make up a VMFS is coherent—all LUNs are in the same replication state.

On the two datacenters, a LUN of 2 GB must be introduced to all ESXi hosts within the target cluster; it is called a *placeholder*. It represents the configuration information for each protected VM (the source machine's vmx file) to ensure its business recovery on the target platform.

Mapping Resource Hierarchy

We strongly recommend properly identifying the resource pools between Site 1 and Site 2. For example, on the production site, three resource pools are created: High, Normal, and Low; on Site 2, there are SRM-High, SRM-Normal, and SRM-Low.

The prefix clearly identifies which pool hosts backup VMs and which hosts production VMs. This mapping of dedicated production resources keeps production and backup VMs separate. Indeed, when VM protection is activated, source VMs will be identified in the target environment. If production and backup VMs are not separated, they will be mixed together on the target site, making day-to-day administration more complex.

When the protection group is created, it is possible to add VMs into a datastore. In this case, the “protection group” detects that the VM is not considered as part of the business recovery plan. An email is sent, and a warning appears in the SRM administration console.

Management

After the architecture is set up, it must handle the load. Ensure that the CPU, memory, and I/O disk loads are not saturated, which would lower performance. The management phase, illustrated in Figure 8.17, is where you monitor and optimize the solution.

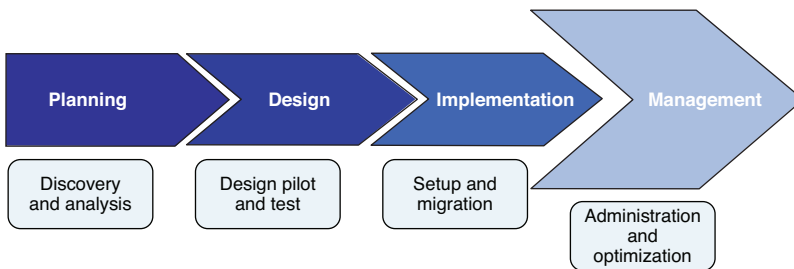


Figure 8.17 The Management phase.

The physical CPU load of ESXi host servers should not exceed 75%. If the CPU load is very intensive, use reservations for critical VMs to guarantee the minimum required. Critical VMs must have high share settings. Monitor the CPU ready time of the VMs’ vCPU; it must not exceed 2000 ms.

The DRS cluster is activated in normal mode. FT is activated for four very critical VMs, and HA for 25% of the servers.

Administrators must monitor the ballooning, memory swap in/swap out, and disk R/W latency, which must remain below 20 ms.

A simple action prioritizes resources in case of contentions. It is easy to create three resource pools (High, Normal, Low), which is also efficient for the distribution of CPU and RAM priority based on the VM’s criticality. vMotion and Storage vMotion simplify management and planned maintenance operations when acquiring new hardware. This reduces service interruptions during these phases, which would be delicate operations in a traditional physical environment.

Summary and Conclusion

This project is a success. The client is completely satisfied because both the schedule and budget were respected, and all goals set were reached. Cost reduction from energy savings is even above expectations.

Three hundred VMs run on 16 ESXi hosts (8 servers per datacenter). The infrastructure in place provides high service levels with an RPO of 0. Failover tests show the RTO is less than four hours. IT teams are fully trained and perform two annual failover tests.

The system proved itself robust when work being performed by an electrician required a complete electrical interruption of the datacenter. Through an SRM5 planned migration, teams switched all production to Site 2. This operation—usually cumbersome and delicate in a traditional physical environment—was performed without encountering any problem and required a preparation period that had been reduced to a minimum.

New requests (about 50 in the queue) were all addressed quickly as soon as the infrastructure was in place.

The infrastructure rationalization is impressive. At the end of the project, as shown in Figure 8.18, only 78 servers remain instead of the 252 physical servers—a reduction of 69%.

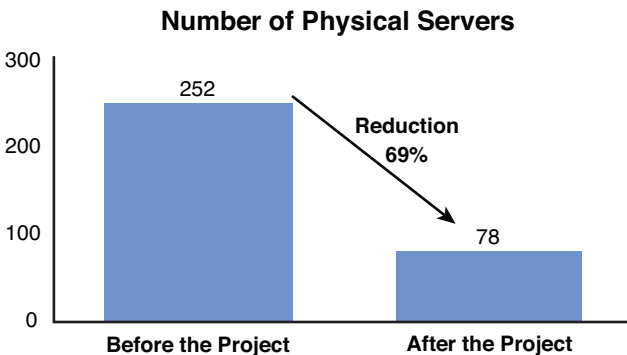


Figure 8.18 Physical server inventory before and after the project.

At completion of the project, a total of 252 existing physical servers (242 physical servers + 10 older-generation ESX servers hosting 120 VMs) were replaced by 78 physical servers (16 ESXi servers + 62 nonvirtualized physical servers).

As shown in Figure 8.19, power consumption after virtualization is 16 kWh; 48 kWh have been freed, resulting in a reduction of 75%.

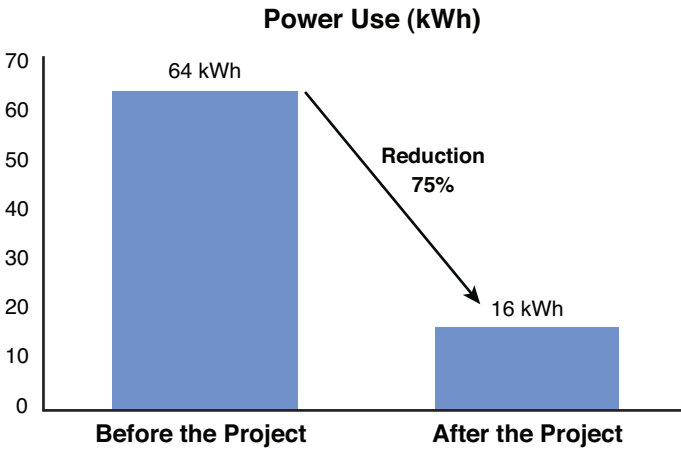


Figure 8.19 Power consumption before and after the project.

As shown in Figure 8.20, floorspace used was significantly decreased, from approximately 740 U (1 U = 1.75 in or 4.445 cm) to 240 U, a 68% reduction. The consolidation ratio is 1:18, a good average. Power use was decreased by 75%.

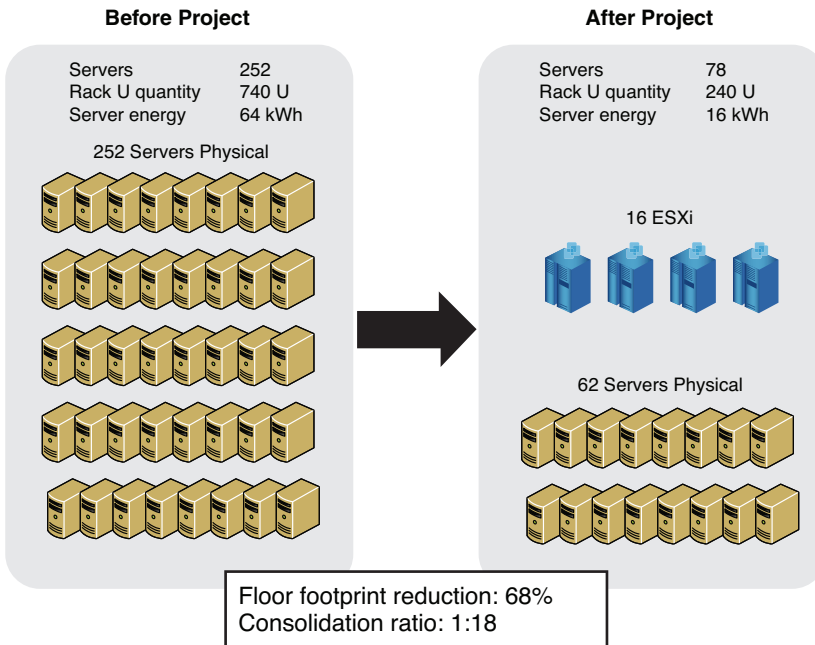


Figure 8.20 Floorspace used before and after the project.

This company now wants to automate a number of procedures and offer service catalogues for internal clients. Backup solutions offering de-duplication are being studied to replace the various existing solutions with a unified solution.

Through this example of a virtualization project, we can see the advantages this technology brings to the datacenter.

A Compelling Story

As you read about the project and challenges that the organization faced, you may have thought that it sounded familiar! In any economy, but especially a bad one, businesses look for any way to reduce costs, and virtualization using vSphere 5 is a great way to do that. In this project, you learned how to size a VM host and a VM correctly, and then converted some physical servers to VMs. To do so, you inventoried the needs of the 242 physical servers, made some decisions about what could best be virtualized, and performed the tasks. The thresholds set in the project were determined by the project team, but every organization may have different thresholds, and vSphere can virtualize most x86 servers. This project also implemented VMware SRM5, so the organization also can benefit from a comprehensive DRP implementation.

The savings were startling, reducing the number of physical servers by 69%, and reducing power by 75%. With green IT initiatives currently underway in so many industries, savings go beyond simply financial concerns. Virtualization saves money, time, the environment, and in some cases, your business!