The Technical Infrastructure of Data Centres

Best Practice Document

Produced by the CESNET-led Working Group
on Network Monitoring
(CBPD 121)

Author: Martin Pustka

November 2012

Document No:	GN3-NA3-T4-CBPD121
Version / date:	11.11.2012
Original language :	Czech
Original title:	Technická infrastruktura datových center
Original version / date:	1.0 of 11.11.2012
Contact:	Martin.Pustka@vsb.cz

CESNET is responsible for the contents of this document. The document was developed by a CESNET-led working group on network monitoring as part of a joint-venture project within the higher-education sector in the Czech Republic.

The research leading to the update and revision of this report has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 238875, relating to the project 'Multi-Gigabit European Research and Education Network and Associated Services (GN3)'.

Table of Contents

Executive Summary

The purpose of a data centre is to provide operational, network, server, computing, and storage infrastructure for IT services, with sufficient, scalable capacity to operate these services, using converged network technology (Ethernet, Fibre Channel, Fibre Channel over Ethernet), virtualisation of servers, and shared physical infrastructures.

This document describes the various options that should be considered when designing or operating a data centre, in order to provide the optimal environment for upgrades and for the expansion of capacity.

Data centre technology is currently developing at a rapid pace and many old rules no longer apply. The differences between LAN and SAN networks and server systems are no longer so clear, and at the same time, new rules, boundaries, and administrative roles are emerging.

The administrator of a virtual server is different from the administrator of HW (hardware or physical servers). The administration of physical servers is significantly simpler than in the past, and migration of an installed system to other hardware devices requires only a few minutes.

These changes require a change in the IT personnel needed to administer the infrastructure. A server administrator no longer needs to administer physical systems, but rather, only needs to integrate modules within its systems so that they communicate with the virtual infrastructure.

The administrator of a computer or SAN network is now more involved in the administration of the converged infrastructure network, which is combined with the administration of the infrastructure of the physical server in many solutions . In this case, the server no longer acts as a server in the traditional sense. The server is a technical device with its own CPU, memory, and connection to the network, and thus, the configuration in its server profile determines its function within the entire infrastructure.

A disk storage administrator often connects individual disk storage devices to server systems. A network with a converged infrastructure can connect different disk storage devices within different virtual SAN networks.

The administrator of the virtual server infrastructure continues to play a major role, and is responsible for setting up an optimal virtual server environment.

The use of such modern technology can result in significantly increased technical efficiency and more effective use of physical resources.

Finally, this document provides an example, describing the setup of the infrastructure of the data centre at the VŠB-Technical University of Ostrava.

This document does not describe how to set up specialised, supercomputer data centres requiring additional technology, such as low-latency Infiniband networks, but these supercomputer centres ***do*** use a lot of the technology described below.

1 Basic Concepts and Definitions

1.1 Computer Network Infrastructure

A network includes the routers and switches, which realise the internal and Internet connectivity of the data centre.

1.2 Server Infrastructure

A server infrastructure includes computer server systems, represented by stand-alone server systems with their own network interface, disk capacity, and physical administration of the system. In integrated blade solutions, these systems are represented by server modules which may (but not necessarily) have local disk capacity and whose administration may be centralised.

1.3 Virtual Server Infrastructure

A virtual server infrastructure provides an environment to operate virtual server systems. There are several virtualisation platforms available on the market (for example, VMWARE, XEN, Hyper-V, and KVM), each with different capabilities and licence policies. A virtual infrastructure operates on physical server systems and uses shared physical infrastructure to operate multiple virtualised systems.

1.4 Disk Storage

The infrastructure of disk storage includes control nodes and disk policies. Disk storage customarily offers disk policies with different performance levels and technologies (for example, FC, SATA; and 7.2/15kRPM). Centralised disk storage should use available technologies that can increase its performance (for example, disk cache, and distributing data across multiple disks) and improve its functionality (for example, de-duplication, replication, and thin provisioning).

Disk storage capacity is made available to the server systems through a SAN or converged network.

1.5 SAN Network Infrastructure

This type of infrastructure connects disk storage to the server systems that use it. In modern data centres, this type of infrastructure is based on Fibre Channel protocols, or FCoE for more modern data centres. The purpose of these systems is to ensure uninterrupted and redundant connection between the disk infrastructure and the server infrastructures.

1.6 Converged Network Infrastructure

A converged network is one that is able to transfer LAN/IP and SAN network data within a single physical network. Using a single infrastructure saves space, requires less energy, and reduces the costs of acquisition and operation of the infrastructure.

2 Traditional Data Centre Technology

This section of the document offers a brief overview of the technology traditionally used in data centres. It is important to understand these technologies in the event that they need to be integrated into modern data centres.

Traditional technology is made up of server systems with one or more integrated network interface, most often using Gigabit Ethernet (GE) technology. Recently, significantly loaded systems began to use gigabit channel-bonding technology (IEEE 802.3ad), or possibly 10GE technology. This type of network connection is mostly used for IP connectivity, and allows the connection of disk storage infrastructure, for example, when using NFS, iSCSI or CIFS protocols.

Stand-alone server systems often have their own local disk storage infrastructure. In server systems, these storage devices are often integrated into a simple RAID, usually with a hardware controller. Software RAID solutions are more sensitive to interruptions and lowered the performance of the server system, so they are used less often.

Server systems are connected to the IP network using traditional L2 Ethernet switches and do not require redundancy against outages in switches, cabling or network cards. These often use IEEE 802.3ad technology and accessibility to the computer network can be improved by virtually bonding the physical switches, so that they appear as a single logical switch within the IEEE 802.3ad to the server system. This introduces such concepts as clustering, vPC (Virtual Port Channel), and VSS (Cisco implementation).

Disk storage is made up of disk arrays and their control modules. These disk systems then offer their storage capacity via the Fibre Channel protocol, or via application protocols implemented above the level of the IP protocol, usually using NFS, iSCSI, CIFS or SMB protocols.

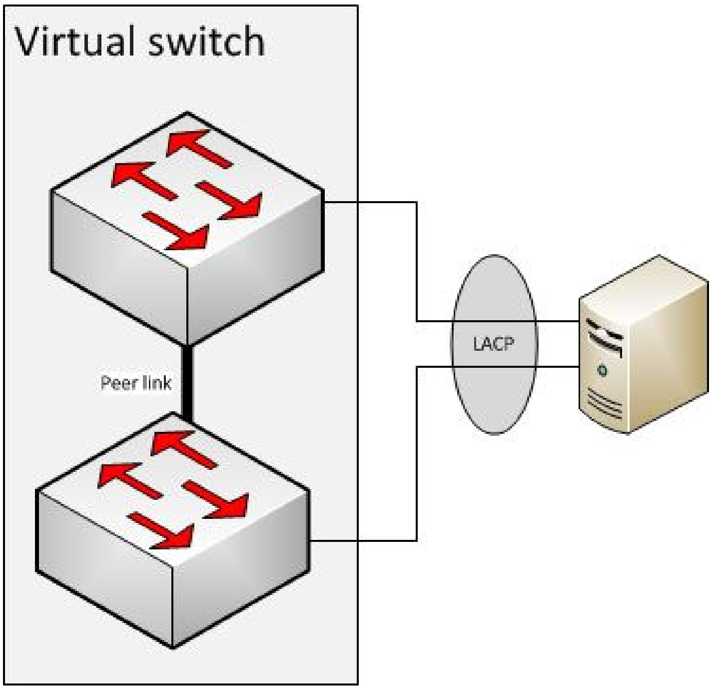

Figure 1: Virtual switches and connecting end systems

3 The Design of and Technical Options for Modern Data Centres

As a whole, the services offered by data centres remain the same: services that allow the infrastructure to operate and deliver centralised IT services.

The only difference is the actual infrastructure. The infrastructure of a modern data centre is often more efficient, offering greater and higher performance with consolidation options. This represents a more efficient use of hardware resources and often results in better performance at a lower cost.

The first innovation required is the implementation of a converged network. The aim is to set up a single, high-quality network infrastructure that is sufficiently accessible and redundant to integrate all of the protocols used in the data centre.

The second innovation is the new role of server systems. These are now perceived as technical devices that have their own processor performance, memory, and network interface. In contrast to existing systems, such a server usually does not have its own disk space, which is centralised, nor any local settings, such as a MAC address or BIOS settings. The integrated administration of such systems is based on server profiles applied to the "naked" server systems.

The third innovation is in the use of smart disk arrays, which is not really a revolutionary innovation because disk arrays have been used before. A smart disk array brings two advantages: it provides disk space by FC and other protocols, and also provides de-duplication, replication, compression and other services. Such advantages greatly improve the performance of the entire data centre. Modern arrays may also offer direct support for converged Ethernet.

3.1 Technical Capabilities of a Converged Network

A converged network greatly reduces the cable length required and the number of active components (LAN or SAN switches are no longer required), while offering easy scalability, redundancy, and performance.

As was explained earlier, all operating data passes through a single, converged network. This includes traffic of both IP/Ethernet and disk storage devices. While IP protocols are usually able to handle data loss easily, this is more difficult with FC protocols. This is why Ethernet traffic includes "lossless operations", which are given priority in order to maintain critical system resources. The FC or FCoE protocol is one such system.

Converged networks use 10GE technology, which may also be bonded. In comparison, a traditional SAN/FC network uses capacities of 1/2/4/8 Gbps, which cannot be bonded. In the future, capacity is expected to increase to 40GE, even to 100GE.

3.2 Technical Capabilities of Server Systems

As was stated earlier, servers are no longer perceived as separately administered entities, but rather as technical devices. With centralised administration, servers can be set up independent of technical devices. Instead of setting up a particular device, the *server profiles* to be applied to different physical servers are set up.

This offers immense flexibility when integrating server systems, handling outages and when dividing the capacity of the server system. Consider the required procedure following the failure of a component of a traditional server, such as its CPU. This usually requires the physical presence of a qualified administrator, and one of three scenarios.

1. The server must be shut down (if the malfunction did not already cause this), the faulty component replaced, and the system restarted.
2. Worse cases require connecting the disks to another server (if a similar model is available), reconfiguring the LAN/SAN network, setting up BIOS, and so forth.
3. The worst-case scenario requires data recovery from backups stored on another HW, and a similar reconfiguration as in 2.

All these scenarios block services for a relatively long period of time and require the physical presence of qualified personnel.

However, with centralised administration, the server profile of the faulty server module can be applied to another module, which is then activated. All this can be achieved remotely, within a few minutes.

The server profile contains all of the server's necessary settings:

- a definition of the LAN network interfaces and MAC addresses;
- a definition of the HBA interface, WWPN and other settings needed for SAN networks;
- a boot order from the provided disk systems;
- the BIOS settings;
- the other settings relating to the equipment used.

All connections required within a centrally administered server are made automatically.

A frequent disadvantage of the integrated server modules of the blade system is their limited expansion using special PCBs or boards. For this reason, not all demands can be met by this integrated system. However, the market offers stand-alone servers that can be integrated into a centralised server administration.

One advantage of integrated server modules is the small number of redundant or test systems that are immediately available during outages. Another advantage is the fully remote administration of the server systems. In this case, the data centre administrators perform all local administration.

3.3 Disk Storage Technologies

Disk capacity is provided from the data centre server system through a SAN network. Disk arrays of different performance levels and capacities can be connected to this network and provided to the connected server systems in the SAN network. This centralised storage reduces administrative demands, while providing a high-performance, flexible storage space for the data centre.

These storage systems must support the FC/FCoE interface. Due to compatibility (or, more appropriately, incompatibility) issues, it is important to consider the compatibility of the equipment to be used by the data centre. This includes compatibility between the disk storage, the network switches, the host bus adapters, and the applications.

In addition to compatibility, disk storage that can fulfil all requirements should be selected. In anticipation of further technological advances, it is wise to combine different storage technologies within a single storage system – from inexpensive and slow SATA disks for non-critical applications, to very fast disks with high throughput (SAS, FC, or SSD).

Our choice of disk storage also takes the following criteria into consideration:

- synchronised replication;
- snapshots of individual, logical disks;
- prioritisation of operations over a single disk section;
- de-duplication;
- thin provisioning;
- FC, FCoE, iSCSI, NFS, and CIFS protocols;
- the option to specify the FS operating over the disk space;
- full redundancy (including power supply, controller, ports, and paths to disks).

3.4 Redundancy

One of the important features of a data centre's stability is its tolerance to malfunctions of one or more components. By dividing the individual sections into layers (as shown in Figure 2), it is fairly easy to achieve redundancy and expand within the individual sections. It is generally advisable to design a data centre so that it can tolerate component failures in all layers. Separation of the administration of layer components so that each layer can tolerate configuration errors should also be considered. This approach increases the tolerance, but makes higher configuration demands, since two components need configuration rather than one.

Generally, modular units with centralised administration provide good tolerance against technical errors, which can include hardware, cable, or port failures. However, such common administration is less tolerant to configuration or program errors, which can include administrative errors or faults in service software.

For these reasons, the administration is often divided among the various components in the layers.

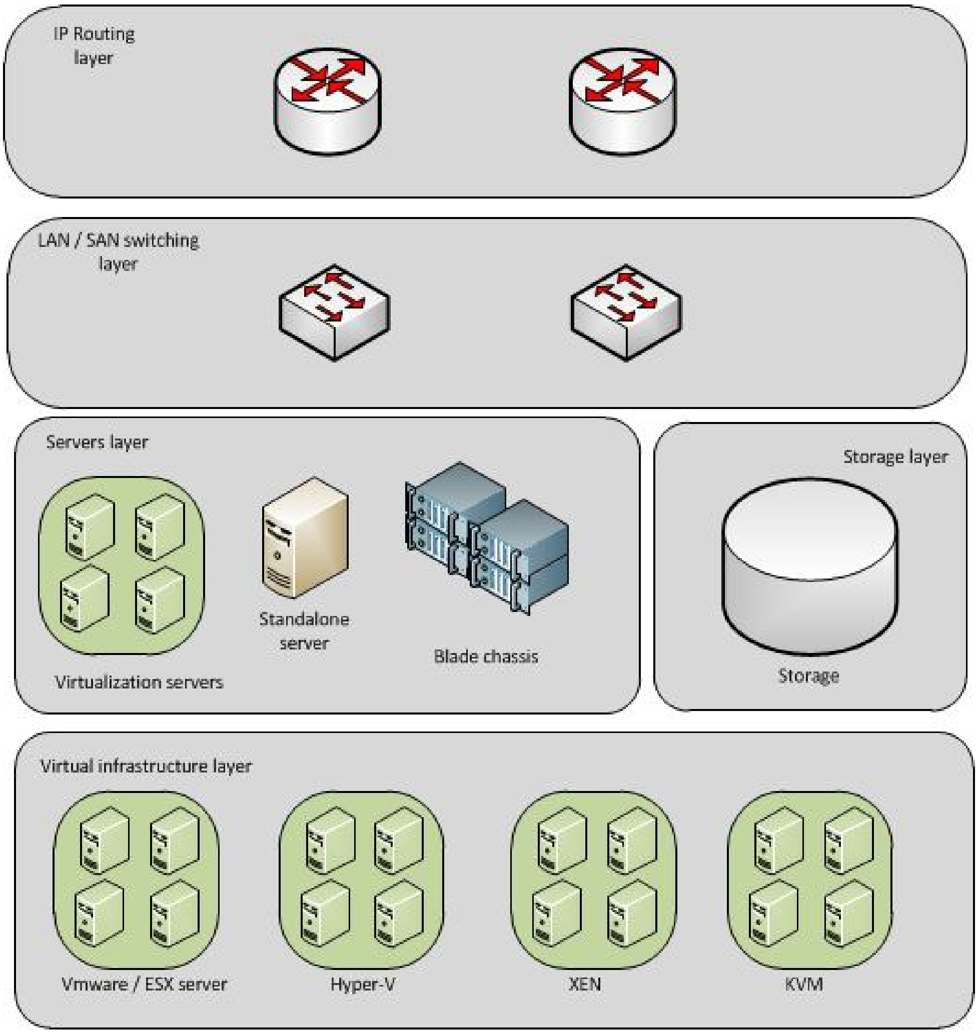

Figure 2: Data centre layers

3.5 Distributed Data Centre

When designing a data centre, consideration should be given to whether it should be broken up into two or more locations. Such separation can increase the entire system's tolerance against external factors, such as fires, flooding, or power outages. Two approaches are possible:

- a data centre separated by a distance of several hundred metres;
- a data centre separated by a great distance.

Distance and latency are the most significant limiting factors for the type of technology used to connect the data centre. These factors mostly limit SAN networks technologies. If the distance and latency are relatively minor, a distributed data centre with only slightly higher costs can be set up, in comparison with the costs of widely separated data centres.

If greater distances are chosen, infrastructure costs can be expected to double, at least. This is because of the required redundancy at each data centre and the connection costs. It is also possible that all services cannot be implemented in such a model (for example, the synchronised replication of disk arrays). In addition, the connection to the computer network can be significantly more complicated.

4 Integrating Traditional Technology into the New Infrastructure

When setting up a new infrastructure, old and operating equipment needs to be connected to the new system.

The problem is that older systems have separate LAN and SAN networks, and these older systems need to have access to devices running on the new infrastructure.

4.1 Older Server Systems

A 1GE interface can be connected to traditional access L2 switches that are connected to central LAN components. Another option is to connect these systems to L2 switches that are fully integrated into the new data centre infrastructure. These components, or their ports, only support a traditional LAN connection, and not FCoE.

The components can use the data centre's settings, and can connect to disk storage through an IP protocol (NFS, iSCSI, CIFS).

The server can then be connected to the SAN network using data centre switches, whose ports will be switched to FC mode. Another option is to connect to an older FC network, into which the disk capacity of new disk storages will be propagated.

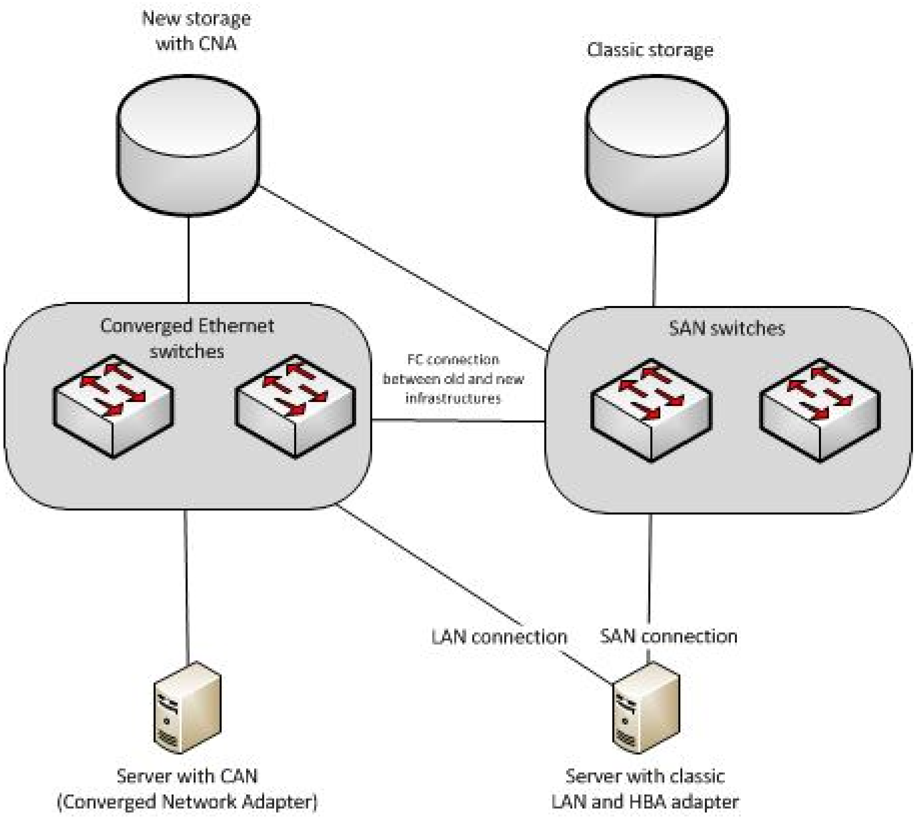

Figure 3: The principles of integrating old with new technology

4.2 SAN Networks and Older Disk Storage

It may happen that a system uses older disk storage technology that is still suitable for some purposes. Such equipment is usually comprised of a separate disk storage with an FC interface and a SAN network supporting only the FC protocol.

Some manufacturers support disk array virtualisation. This requires that equipment capable of connecting disk storage space using the FC protocol is connected to the system. With these virtualisation technologies we can still use old storage devices in a new data centre network. Usually, such virtualisation components are technologically more modern and are capable of providing this capacity with other protocols than FC. This means that they can be connected to newer, converged networks, and available through the iSCSI, NFS, and SMB protocols.

However, these connections require the arrays to be initialised, and this erases existing data.

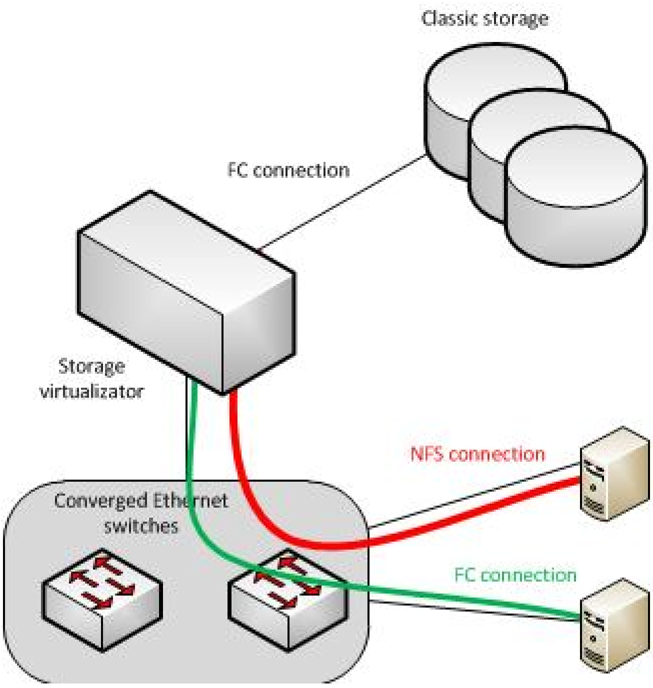

Figure 4: Virtualisation of disk arrays

5 The Integration of Virtual Server Infrastructures

Using modern technology and the capabilities of virtual servers, network functions or a virtual infrastructure can be integrated with the existing networks.

Because this is the latest technology, only Cisco Nexus 1000V is available for VMWARE and Hyper-V environments. This virtual switch has only one or two hypervisors, which control the functionality of *Virtual Ethernet Modules (VEM)*. Hypervisors backup each other and administer all VEM. Their CLI interface is identical to Cisco's traditional CLI.

A normal LAN connection (dedicated VLAN) is used for communication between VSMs and VEMs. VEM modules continue to work when VSM is not accessible, although their configuration cannot be changed. This means that VEM modules can still be used in a virtual infrastructure. VSM can also be virtualised, and the only consideration is whether one of the modules should be in the form of a physical server, set up outside of the virtual infrastructure.

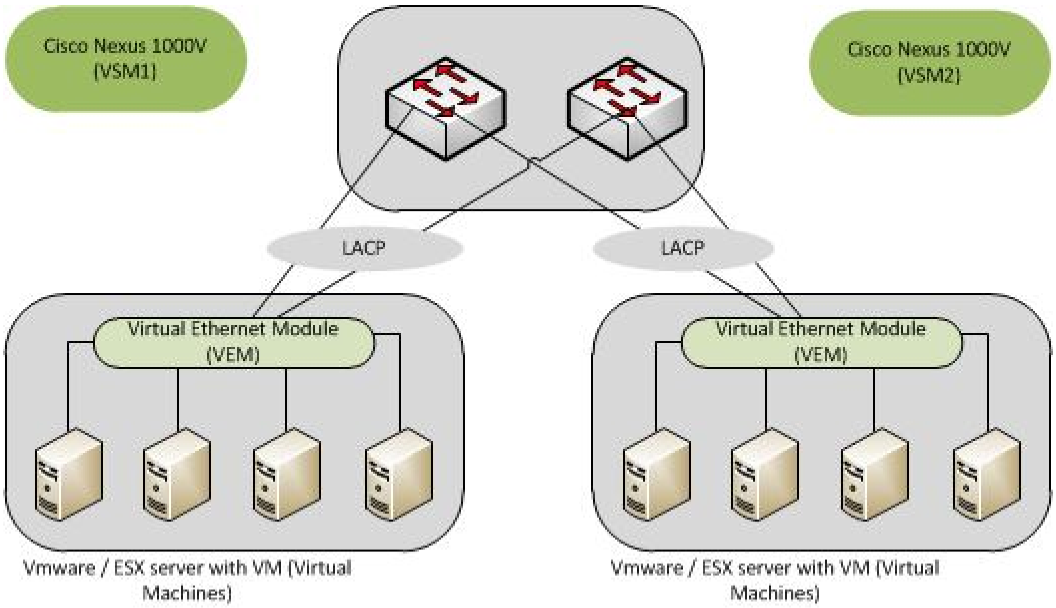

Figure 5: The topology to connect Cisco Nexus 1000V virtual switches

An example of a solution using a Nexus1000V module:

```
VSM1-NEXUS1000# sh module
Mod  Ports  Module-Type                       Model              Status
---  -----  --------------------------------  -----------------  ------------
1    0      Virtual Supervisor Module         Nexus1000V         ha-standby
2    0      Virtual Supervisor Module         Nexus1000V         active *
3    248    Virtual Ethernet Module           NA                 ok
4    248    Virtual Ethernet Module           NA                 ok
5    248    Virtual Ethernet Module           NA                 ok

Mod  Sw                  Hw
---  ------------------  ------------------------------------------------
1    4.2(1)SV1(5.1a)     0.0
2    4.2(1)SV1(5.1a)     0.0
3    4.2(1)SV1(5.1a)     VMware ESXi 5.0.0 Releasebuild-821926 (3.0)
4    4.2(1)SV1(5.1a)     VMware ESXi 5.0.0 Releasebuild-821926 (3.0)
5    4.2(1)SV1(5.1a)     VMware ESXi 5.0.0 Releasebuild-821926 (3.0)

Mod  MAC-Address(es)                         Serial-Num
---  --------------------------------------  ----------
1    00-19-07-6c-5a-a8 to 00-19-07-6c-62-a8  NA
2    00-19-07-6c-5a-a8 to 00-19-07-6c-62-a8  NA
3    02-00-0c-00-03-00 to 02-00-0c-00-03-80  NA
4    02-00-0c-00-04-00 to 02-00-0c-00-04-80  NA
5    02-00-0c-00-05-00 to 02-00-0c-00-05-80  NA

Mod  Server-IP        Server-UUID                           Server-Name
---  ---------------  ------------------------------------  --------------------
1    192.168.1.46   NA                                    NA
2    192.168.1.46   NA                                    NA
3    192.168.2.21   c61cfd34-472f-11e1-0000-0000000001bf  esx1
4    192.168.2.23   c61cfd34-472f-11e1-0000-0000000001af  esx3
5    192.168.2.22   c61cfd34-472f-11e1-0000-0000000001cf  esx2

VSM1-NEXUS1000# sh interface status up
...
Veth1          server1            up       21        auto    auto    --
Veth2          VMware VMkernel, v up       29        auto    auto    --
Veth3          server2            up       21        auto    auto    --
Veth4          VMware VMkernel, v up       29        auto    auto    --
```

Depending on the disk arrays used, consideration should be given to whether the infrastructure of the virtual server infrastructure should have access to disk space through NFS. The advantage is the ability to use *jumbo frames* or *thin provisioning,* which makes it possible to change the capacities assigned to server infrastructures without interrupting the operation.

Tests and real-time experience have shown that increases in performance and throughput are relatively small with jumbo frames. Furthermore, providing services in parts of a network where jumbo frames are not supported requires additional costs in IPv4 networks, because of the fragmentation of large packets. Therefore, careful consideration should be given to the suitability of this technology.

5.1 Server Infrastructure

Stand-alone servers or server modules that are a part of a server blade chassis can also be connected to a data centre.

These server systems can gain access to disk space on disk storage devices located in the SAN networks of the data centre. Locally operated disks are only installed in special cases. Usually, modular server systems cannot be expanded using other expansion cards, as is normal for stand-alone, physical servers.

In the data centre, stand-alone physical servers can be connected to the computer network with the following technologies:

- 1 Gbps Ethernet;
- 10 Gbps Ethernet (with FCoE support);
- 10 Gbps Ethernet (without FCoE support).

Server modules are the preferred method of operating server systems. These server modules can be administered via an administrative system that can administer the different modules or remote changes to HW server systems easily, while maintaining disk capacity (from disk array space connected to the SAN network), network connections, and basic server features (for example, MAC addresses).

However, it is advisable to design the entire infrastructure so that it is possible to connect other server systems or technologies from other manufacturers to it. In any case, the integrated administration of HW server modules will be limited to the capabilities of the hardware offered by each manufacturer, since, to date, no interface exists that can support multiple manufacturers.

If the design includes redundancy, the basic functionalities of data centre components allow them to substitute one another, so that the data centre's operation will not be limited if one of them fails. These components are redundantly connected to the rest of the network and their basic administration can be separated.

The use of 10GE technology is preferred, but often older systems with 1GE technology must be supported. The end systems can be connected to the computer network with the same capacity. Critical systems can be connected to multiple physical components to achieve a more redundant connection.

6 Scalability

One of the fundamental characteristics of a data centre is the ability to regularly expand its capacity without any, or only minimal, interruptions in its hosted systems and services. This can be achieved by migrating data and systems between different hardware, and this has to be capable of being done during operation.

These migration technologies are usually only compatible with equipment from a single manufacturer. However, with a layered architecture, the individual layers can be replaced or added with equipment from different manufacturers. Their replacement is possible over time, but it is not always easy to migrate between products from different manufacturers.

6.1 Computer Network

The scalability of a computer network depends primarily on the facility to expand its capacity and on the number of active component ports.

Capacity can be expanded by logically bonding physical channels (such as LACP technology) and by the division of operations. Such expansion is possible during operation and can be carried out without any interruptions.

The number of ports can be expanded by adding modules, or by replacing the components that cannot be expanded. With well-planned redundancy, it is possible to make a replacement without any interruption in operations. The switch must be made very carefully, because while making the switch, it is possible to have a short period of non-redundancy.

6.2 Server Infrastructure

With the increase in demand from hosted systems, the physical server must also be upgraded occasionally, or migrated entirely to a better-performing system server.

If a *blade chassis* is used, the physical server must be installed into the relevant chassis, and configured, or an existing server profile applied. Once the system is ready to be shut down, its migration to another physical server is a matter of five to ten minutes, depending on the start-up time.

6.3 Virtual Infrastructure

Ideally, host systems should be operated in a virtual infrastructure. This allows uninterrupted migration of virtual systems between physical virtualisation servers.

In this ideal situation, expansion of a virtual infrastructure's capacity only requires that additional server systems be connected to the virtual infrastructure.

Another option is to entirely replace the virtualisation servers with a higher-performing server. In this way, the virtualised systems can be migrated to another physical server, and the server profile applied to the higher-performance physical server, which can then be started and instantly used. This approach results in absolutely no down time.

6.4 Disk Storage

Expanding disk array capacity requires additional disk arrays or entire disk systems. This all depends on the disk limitations and the administrator's plan of how this newly added capacity will be integrated into the existing infrastructure.

Additional disk capacity can be connected fairly quickly and without interruptions in storage.

7 An Example of an Implementation at a University Data Centre

The technical infrastructure described in this section of the document was incorporated at VŠB - Technical University of Ostrava and includes the principles and approaches described above.

7.1 IP Network Layer

Network IP routing is handled by two Cisco Catalyst 6500 devices. The IP routing is handled by L3 switches that support both IPv4 and IPv6 protocols. Both components are administered independently.

7.2 LAN/SAN Layer

Two Cisco Nexus 5500 data centre switches make up the core switches of the data centre. These switches have then been expanded by satellite modules, Cisco Nexus 2000 FEX (Fabric Extender). The switches operate only on L2 layers and work with a converged Ethernet. Switches on this layer also support switching for the SAN network and connection to the older SAN network, where older storage fields are connected. Both components are administered independently.

7.3 Integrated Physical Server Infrastructure

The architecture of the centralised server is based on a Cisco UCS solution. This option was chosen for its support of converged Ethernet throughout the entire solution, its integrated and powerful administration of the entire environment, and its quick and easy scalability. Last, but not least, its price was very attractive.

7.4 Stand-alone Physical Servers

The primary task is to incorporate blade solutions to the data centre's server infrastructure. This is because there are some applications that require stand-alone server systems.

The integration of older and stand-alone servers into the data centre is underway and these are presently too cumbersome to migrate to a virtual or physical infrastructure. There are two primary reasons for this. The first is the licence limitations of the existing applications, which do not allow for a simple transfer to higher-performance, more modern servers.

The second reason is the need to maintain the operation of the physical server systems, which were acquired just before the setting up of the data centre. They are under service contracts and it would be too costly to interrupt their operation.

Independent servers are also connected directly into Catalyst 6500, while others are connected into the Cisco Nexus 2032 FEX module, equipped with 1GE ports that do not support converged Ethernet.

With the installation of special PCBs or boards, the stand-alone servers are operated outside of the Cisco UCS infrastructure. These servers are equipped with CNAs (Converged Network Adapters) or LAN/HBA adapters, giving them access to the data centre's disk infrastructure.

7.5 Disk Storage

NetApp 3240 was chosen as the primary platform for disk storage. This solution includes integrated disk arrays with different levels of performance and capacity. Besides providing disk capacity to the new converged network infrastructure using the FCoE protocol, this storage space is also accessible to the older SAN network, primarily for the purpose of migrating data from the older infrastructures.

Disk capacity is also provided to the IP network by NFS protocol. This uses a 10GE interface, whereby the disk arrays are connected directly into Cisco Nexus 5500 data centre switches.

The disk arrays also provide an interface for better integration into the virtual infrastructure, where communication between the disk storage spaces and the virtual infrastructure is handled more easily.

7.6 Connecting Older Technologies

The older server systems were connected to the central Cisco Catalyst 6500 switches or local LAN switches, but now, the preferred connection is to FEX modules with 1GE metallic ports.

The older disk arrays are connected to older SAN Brocade switches and preferably to Cisco Nexus 5500 switches. This older infrastructure is also connected by FC ports to the new infrastructure, making their capacity available to the new systems. Over time, the older disk arrays will be removed from the system gradually, including the old SAN infrastructure.

Up until now, the LAN switches were also used in other areas of the University computer system.

This infrastructure will continue to be operated for less critical applications but service support from the manufacturer will no longer be updated. Once a system's technical lifespan has been reached, it will be removed from operation.

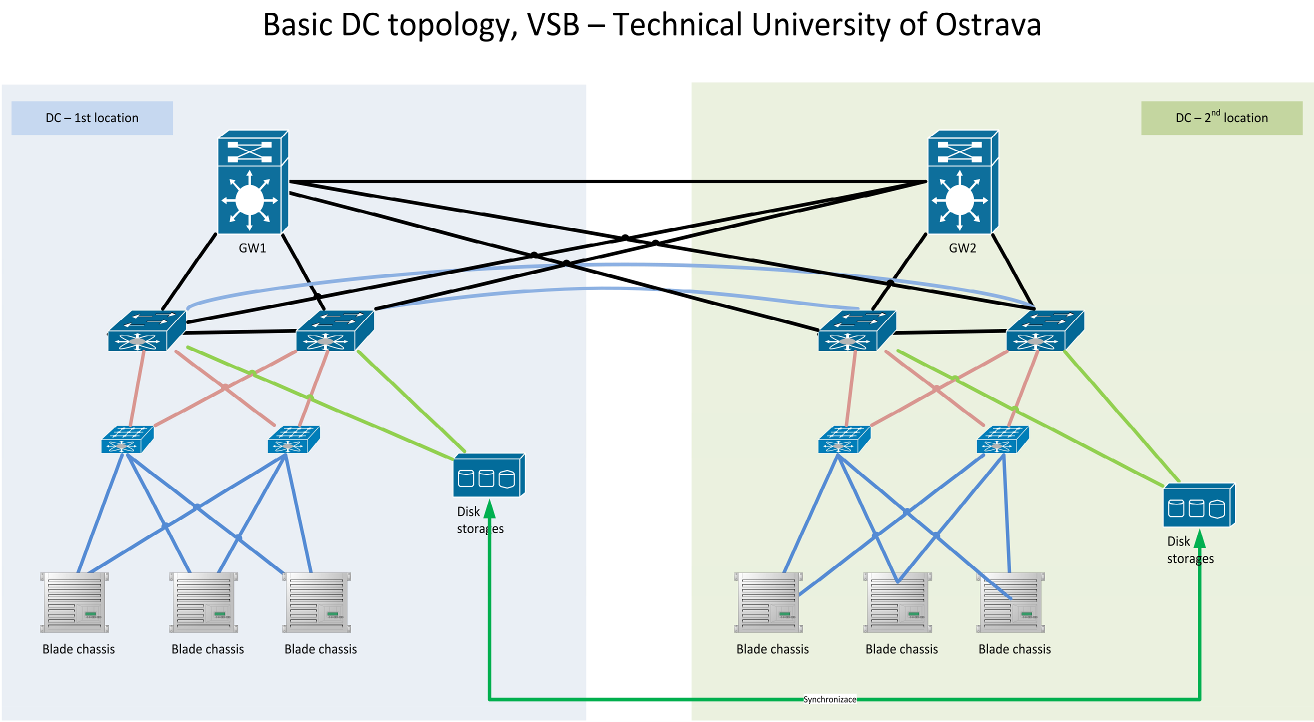

Figure 6: Basic DC topology

Glossary

ATA	Advanced Technology Attachment
BIOS	Basic Input Output System
CIFS	Common Internet File System
CLI	Command Line Interface
CPU	Central Processing Unit
DC	Data Centre
FCoE	Fibre Channel over Ethernet
FCP	Fibre Channel Protocol
FS	File System
Gbps	Gigabits per second
GE	Gigabit Ethernet
HBA	Host Bus Adaptor
IEEE	Institute of Electrical and Electronics Engineers
IP	Internet Protocol
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
iSCSI	Internet Small Computer System Interface
LACP	Link Aggregation Control Protocol
LAN	Local Area Network
MAC	Media Access Control address
NFS	Network File System Protocol
PCB	Printed Circuit Board
RAID	Redundant Array of Independent Disks
SAN	Storage Area Network
SAN/FC network	Fibre Channel Storage Area Network
SAS	Serial Attached SCSI
SATA	Serial ATA
SCSI	Small Computer System Interface
SMB	Server Message Block protocol
SSD	Solid State Drive
VEM	Virtual Ethernet Modules
VLAN	Virtual LAN
vPC	Virtual Port Channel
VSM	Virtual Switch Module
WWPN	World Wide Port Name