

UNIT-5

BLADE SERVER

Q.114 What are blades? What are different types of blades? What are the advantages of implementing blade systems versus rack systems?

Ans:

- Blades are a new form factor for computer technology, which includes components such as servers, storage, and communications interfaces in a prewired chassis with shared components such as power, cooling, and networking.
- In contrast to the traditional horizontal positioning within a rack, blades are typically installed vertically in a blade chassis, like books in a bookshelf.
- There are various types of blades as follows:

1. Server Blades and PC Blades

- When one is implementing blade technology, blade servers are generally the starting point.
- In addition to building servers on blades, a number of vendors have delivered PC blade products.
- In a PC blade implementation, end users operate with only a monitor, a keyboard, and a mouse and the PC runs on a central blade.
- PC blades have the ability to recentralize certain aspects of distributed computing.

2. Storage and Network Blades

- Blade technology also extends to other components of computing, including storage and networking.
- Blade servers require access to information, the choice is whether to incorporate storage on the blade server itself, utilize storage or networking protocols to communicate with standard storage devices outside of the blade environment.

➤ **Reasons for implementing blade systems versus rack mounted servers are as follows:**

- Space savings and efficiency— packing more computing power into a smaller area.
- Consolidation of servers to improve and centralize management as well as to improve utilization of computer assets.

- Simplification and reduction of complexity ease of deployment, and improved manageability and serviceability.
- Return on investment (ROI) and improved total cost of ownership (TCO) through increased hardware utilization and reduced operating expenses.

Q.115 Discuss the different themes that apply to adaptation of blades and virtualization technologies.

Ans:

1. Bigger, Better, Faster, Cheaper

- **Bigger** describes overall ability— generally increasing the computing power and memory or disk capacity.
- **Better** usually addresses increased usability, manageability, and reliability (lower mean time between failures [MTBF]).
- **Faster** relates to component speed— whether it's a processor, bus, channel, or network— and I/O speed.
- **Cheaper** is just that— less expensive but hopefully not cheaper quality.

2. Miniaturization

- The size reductions and increased density possible with blade technology come directly from the miniaturization of the components.
- Miniaturization has ability to put more memory and high-performance processors together has meant that software operating systems and programs can be more sophisticated.

3. Decentralization and Recentralization

- More computers left the management and control of the Management Information Systems (MIS) organization the decentralization of IT management began.
- Decentralization often gave users' departments the ability to respond more quickly to business needs.
- Thus began the cyclical process of decentralizing and recentralizing IT resources.
- As the management of IT systems, networks, and storage has become increasingly complex, there has been a shift toward recentralizing much of the operational management of IT resources.

Q.116 Discuss the eras of evolution of computing starting from mainframes to consolidation.

Ans:

Q.117 Explain the evolution of storage technologies.

Ans:

- The **ENIAC**, hailed as the first electronic digital computer and recognized as the grandfather of today's computers, could only store 20 10-digit decimal numbers in its localized buffers— there was no central memory.
- Before the advent of magnetic disk storage, programs and data were stored on punch cards, paper tape, and magnetic tape.
- Storage technologies have evolved (or have become extinct) over the past 50 years. As computers got bigger, better, faster, cheaper, and smaller, so did the storage technologies.
- Within a few years, in many enterprise data centers large amounts of storage began to be decoupled from computers and attached to networks called **storage area networks (SANs)**.
- **Virtualization** has also been part of the history of storage, almost from the beginning. IBM used virtualization in the VM operating system to present both virtual disk and tape devices to the guest operating systems running as virtual machines.
- Storage virtualization comes in three major flavors, categorized by where the virtualization takes place.
- **Device-based virtualization** is done within the storage arrays, **Network-based virtualization** is done within the network itself, **Host-based virtualization** is done by software that resides in the host computers.

Q.118 What is clustering? Explain the evolution of clustering.

Ans:

- To understand clustering, it is important to note the communications required between clustered nodes to provide both failover and job sharing.
- A cluster, in its broadest definition, is two or more computers on a network, working together on one or more computational applications.
- Therefore, it is possible for the network, rather than being a local area network (LAN), to span across a wide geographical area — for example, the Internet.
- Each computer in the cluster is a node with its own processors, operating system, memory, and system storage.
- Initially clusters were simply separate workstations on a network.

- However, with shrinking footprints and increased costs for space, often cluster nodes were rack-mounted, one on top of the other.
- Today blade servers enable clustering in even smaller footprints, as well as shared power and cooling.
- The first commodity clustering product in the open-systems market was **ARCnet**, developed and released by Datapoint in 1977.
- ARCnet wasn't successful in the market and clustering lay dormant until DEC released their VAXcluster product in the 1980s.
- The University of **Tennessee** released the first public version of **PVM** in 1991.
- The PVM software is free from netlib and has been compiled on everything from a laptop to a Cray. This eventually led to what is known today as **high-performance computing clusters (HPCC)**.

Q.119 Discuss the evolution of grid/utility computing.

Ans:

- The nomenclature grid computing brings to mind an electrical power grid, where a number of resources contribute to the pool of shared power to be accessed as needed.
- This is where the name utility computing originated. Although the promise of grid computing has not yet been realized, steady progress is being made.
- Grid computing is expected to see strong successes in financial services, oil exploration, medical research, security, entertainment, and mechanical engineering.
- Much like clusters, grid computing doesn't rely on high-powered computational engines; rather, it uses underutilized computational engines, located on the network, anywhere in the world.
- As blade servers meet the criteria of being computational engines located on a network, along with their related blade component technologies, they will certainly see a bright future in grid computing.

Q.120 Explain the evolution of Windows and UNIX server operating systems.

Ans:

➤ **Windows server operating systems**

- The evolution of PC and server operating systems experienced many parallel improvements.
- The growth from MS/DOS (PC/DOS) to Microsoft Windows in late 1985 added a graphical user interface, device-independent graphics, and the beginnings of a multitasking virtual environment.
- Throughout the evolution of Windows the operating-system software grew more stable, and the "blue screen of death" for which Windows was famous became less and less frequent.

- This evolution continues, with virtualization capabilities being added into the OS, first to Windows Server with Virtual Server and eventually to the Longhorn release with Windows Server Virtualization.

➤ **Unix server operating systems**

- In the open-systems market, UNIX also grew up.
- Unix Server operating system developed at Bell Labs in 1969 to run on multiple types of computers, UNIX has grown to become the basis for most high-end workstation environments, with its sophisticated OS features and a full featured GUI.
- Initially found only on mainframes, minis, and high-end microcomputers, several versions of UNIX were available for PCs by the end of the 1980s, but they ran so slowly that most people ran DOS and Windows. By the 1990s, PCs were powerful enough to run UNIX.
- In 1991 a young man named Linus Torvalds took on the task of developing a free academic version of UNIX that was compliant with the original UNIX.
- UNIX has evolved into virtualization as well, with the Xen open source hypervisor making its way into most UNIX versions going forward.

Q 121. Discuss the blade and Virtualization technology timeline.

Ans:

- Blades are a new form factor for computer technology, which packages ultrahigh density components including servers, storage, and communications interfaces in a prewired chassis with shared components such as power, cooling, and networking.
- In contrast to the traditional horizontal positioning within a rack (rack-mounted servers), blades are typically (though not always) installed vertically in a blade chassis, like books in a bookshelf.
- In addition to the high density, prewiring, and shared components, an important differentiator between blades and conventional servers is the incorporation of remote out-of-band manageability as an integral part of each blade “device.”
- This is fundamentally different from conventional servers (rack mount or stand-alone) where systems management has been designed as an add-on capability.
- Blade and virtualization technologies together provide critical building blocks for the next generation of enterprise data centers, addressing a two-fold challenge:
 - The first is to deliver on the ever-increasing need for more computing power per square foot under significant IT operations budget constraints.
 - The second challenge is the management of the geographic proliferation of operational centers, forcing the enterprise data center to operate and be managed as a single entity.

- Blade servers were officially introduced in 2001 by a small company called RLX Technologies as a compact, modularized form factor, with low-end servers, well suited for scale-out applications such as web serving.
- Scale out is used to describe increasing processing power by running an application across many servers, versus scale up, which describes adding more processors in the same server (symmetric multiprocessing or SMP).
- The major server vendors (IBM, HP, Dell, and Sun) entered the space soon after blade servers were introduced, with varying degrees of effort. As Intel and AMD delivered new-generation (smaller form factor) chips, these server vendors began offering products using chips from both companies.

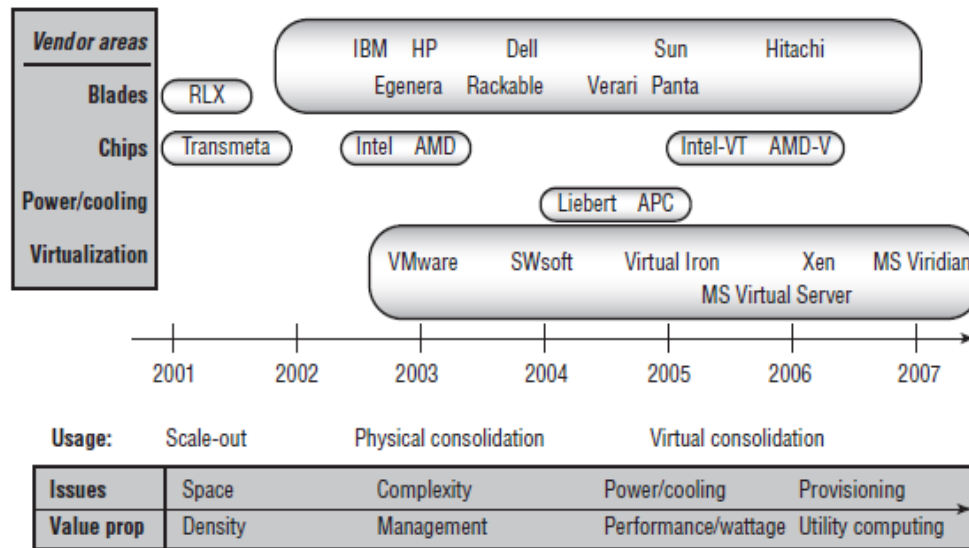


Figure 3.1: Blade and virtualization technology timeline

Q 122 Give an account of history of blade server systems.

Ans:

➤ **Before They Were Called Blades (pre-1999)**

- Although the official introduction of blade servers did not come until 2001, early versions of a PC on a card actually existed several years prior to that. Before the Internet explosion and broadband access, the early bulletin board operators were beginning to morph into the first Internet service providers (ISPs).
- These companies needed lots of servers to support sessions for their dial-up customers, and standard tower or rack-mount servers just took up too much space. Many enterprises were also supporting remote workforces using remote access/control software, such as Carbon Copy or PC Anywhere, for employees or partners to access in-house network resources from a remote PC or terminal.
- Several companies saw some initial success in this market niche in the mid-1990s, delivering PC-on-a-card servers that resided in a common chassis. By 1997,

however, Citrix Systems had achieved dominance in the remote access/control marketplace, displacing the Carbon Copy dedicated-PC approach.

➤ **Innovators and Early Adopters (1999–2001)**

- The initial version of what is now a blade server began in 1999 with an idea for a new, compact server form factor, designed for Internet hosting and web serving, from a reseller and consultant named Chris Hipp. Borrowing on the notion of vertical blades in a chassis (used already in networking), his initial blade server concept was to share components such as power and switches and eliminate unnecessary heat, metal, cables, and any components that were not absolutely necessary for serving up web pages.
- As a result, Hipp, whom some call the father of blade servers, formed a company called RocketLogix with the idea of building these devices. According to Hipp, by 2000, RocketLogix had filed six U.S. patents, with a total of 187 claims related to dense blade computing. With the addition of an executive team of ex-Compaq server executives including Gary Stimac, who helped legitimize the blade concept and raise substantial venture capital (including money from IBM), the company became RLX Technologies, Inc.
- The first product, the RLX System 324, shipped in May 2001, with RLX Control Tower management software released shortly thereafter. The blade concept was well-received and recognized by the industry as a revolutionary step.

➤ **Impact of the Internet Bubble (2001–2002)**

- By late 2001, the major systems vendors had announced their entry into the market, and HP, Compaq, Dell, and IBM were all shipping blade products by the end of 2002, with Sun joining in 2003.
- The initial growth of the blade market was fueled by the rapid growth of Internet data centers during that time.
- The initial furor of the endless need for both computing and I/O scalability, compounded by cost pressure (both for physical resources as well as IT skilled resources) led to a variety of additional blade-related companies and solutions entering the market.

➤ **Niche Adopters (2003–2004)**

- Right as the blade market began to take off and numerous vendors entered the game, the Internet bubble burst. As a result, the blade start-up companies struggled, and many disappeared. Others started focusing on niche areas in order to stay in business.
- In 2004, RLX exited the hardware business, and in 2005 they sold the Control Tower software to HP. Dell, who had entered the market in 2002, did not refresh their product at all during this time. With the market being so soft, blades did not

represent a new opportunity for Dell, and served only to cannibalize their existing server market. (As the market began to come back, Dell re-entered in 2004.)

➤ **Mainstream Re-emergence (2004–2006)**

- As IT budgets began to loosen and IT organizations had money for projects again, blades found significant traction. In 2005, blade server revenue exceeded \$2 billion. This growth puts the adoption rate for blades on much the same growth curve as the adoption rate for rack servers.
- By this time, the market also had shaken out, with IBM and HP taking the lion's share of the market, with a combined total share of more than 70%. Dell and Sun remained as systems players, along with several smaller players who had reasonable success.

Q 123 How and when did virtualization originate? Give an account of history of virtualization.

Ans:

- Although many think of virtualization as new technology of the 21st century, the concept of virtual computing goes back to 1962, when virtual memory was first used in a mainframe computer.
- In the early days of computing, physical memory, or core, was extremely small (64K–256K). This meant that programmers had to be highly aware of program size and resources, even to be able to run a program, and had to write programs that could overlay themselves in memory.
- Virtual memory allowed the creation of a virtual partition or address space, completely managed by the operating system, stored on disk, and paged in and out of physical memory as needed. This concept is now used in many operating environments, including PCs, workstations, and servers, freeing programmers from the burden of memory management.
- In that same era, virtual machines with virtual disk and virtual tape were used, as in IBM's VM/370 operating system, allowing system administrators to divide a single physical computer into any number of virtual computers, complete with their own virtual disks and tapes. The operating system simulated those devices to each virtual machine and spread the resource requirements across the corresponding physical devices.
- All of these early capabilities are reappearing now in various forms throughout enterprise computing: virtual storage subsystems; virtual tape libraries; storage virtualization at the host, network and disk level; server clustering; and virtual server hypervisors are all part of today's landscape and are delivering the same type of significant improvements to technology as their predecessors did.
- The early history of virtualizing CPU processing dates back to the 1970s with IBM's VM/370 mainframe operating system (see Chapter 2, "Evolution of Computing Technology— Setting the Stage"). VM/370 and its derivatives have

run on all IBM mainframes since that time, providing the ability to initiate multiple instances of all IBM mainframe operating systems on a single mainframe (with the number of virtual machines limited only by processing capacity).

- VM was used extensively within the IBM enterprise customer base for testing of and conversion to new hardware systems, operating systems, and applications. To some degree, it also allowed users to migrate from older, smaller systems onto a newer mainframe and consolidate with other systems and applications.
- As users moved to PCs and client-server computing, however, there was no equivalent solution for virtualization in the x/86 product line. In 1999, VMware delivered the first virtualization product for x86 workstations and then extended it to servers, creating a new market for server virtualization. Although a number of other vendors delivered PC and server virtualization solutions, VMware dominated the space, expanding their offerings to include hypervisor solutions (see the discussion later in this chapter), host-based solutions, and workstation products.
- In 2003, VMware was acquired by storage systems giant EMC, and now operates as a wholly owned independent subsidiary. Since VMware's initial product launch, other virtualization (and related) offerings have come from companies such as SWsoft who first delivered Virtuozzo for Linux in 2001; Virtual Iron, founded in 2003; and PlateSpin Ltd. (migration from physical to virtual), also founded in 2003.

Q 124. Discuss the market adaption of blades and virtualization.

Ans:

- According to IDC, more than 500,000 blade units were shipped in 2005, representing 63% growth over 2004. Blade revenue generated in 2005 was \$2.1 billion, representing 84% growth over 2004.
- Figure 3.3 shows the growth in the server blade installed base in the United States and worldwide. The growth rate of blades worldwide has been (and will probably continue to be) greater than the growth rate within the United States.
- One likely reason is that IT infrastructure outside the United States still lags, causing space and power availability to be even bigger issues outside the U.S. (particularly in places like India and China).
- Since blades offer significant advantages over rack servers in regard to these issues, their traction outside the United States is significant.
- In the fourth quarter of 2005, blade servers (including x86, EPIC, and RISC blades) accounted for 4.6% of total server market revenue. Figure 3.4 shows the growth in blade server shipments relative to total server shipments from 2002–2005.


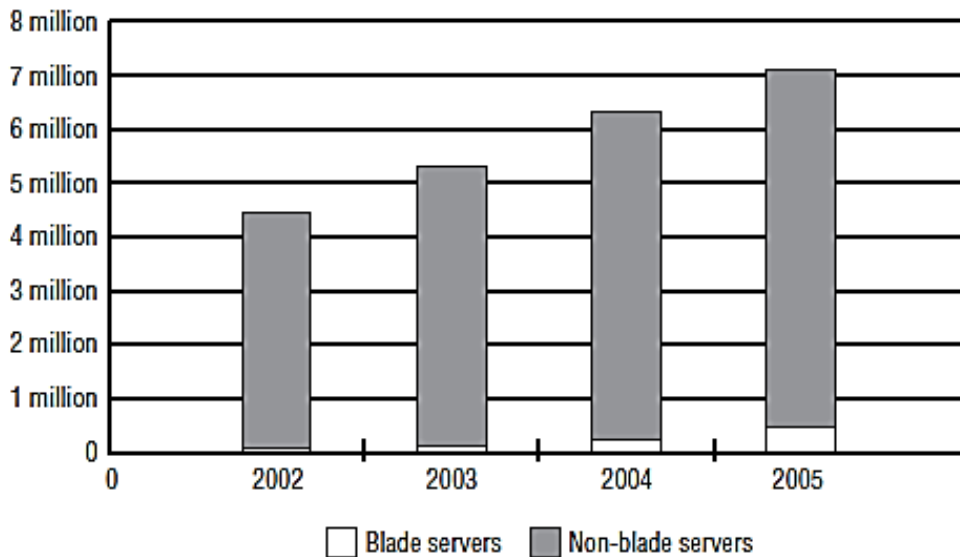


Figure 3.3: Server blade installed base from 2002 to 2005

Source: Worldwide and U.S. Blade Server 2005–2009 Forecast and 2004 Vendor Shares, Doc # 33631, July 2005



Q 125 Explain the usage of blade system

Ans:

➤ **Web Services**

- The initial vision that Chris Hipp of RLX held was based on high-density, low powered blades grouped together for the scale-out function of serving up web pages in a bursting-at-the-seams Internet market. This became an early sweet spot for blades— users needed lots of servers, geographically distributed, with lots of I/O (input/output) and low CPU requirements.

- As blades caught on in web services, many users liked the compact form factor, modularity, and cable less simplicity but had a difficult time paying more money for less computing power. They wanted higher-end processors for uses beyond web servers.
- **HPCC and Application Servers**
- A number of innovations in blade CPU configurations contributed to changes in blade usage. Blades were implemented with faster chips and went from single to multiple CPUs per blade. In addition, advances in chip technology (such as dual-core chips) deliver almost double the CPU power in the same footprint, with lower power and cooling requirements.
 - This increase in computing power allowed blades to move into more-CPU-intensive areas such as scientific research, mathematics, rendering, and HPCC. The addition of the second CPU, or coprocessor, allowed blades to evolve from being used only as low-end web servers to being used also as high-end application servers. Due to the additional I/O involved with application servers, the second CPU is needed for assisting in I/O processing.
- **Database Servers**
- When vendors started adding a second processor and high-speed interconnects, blade systems also became good candidates for parallel and clustered database applications such as Oracle Real Application Cluster (RAC) 10g. Oracle 10g was designed as a scale-out database and its parallel design allows it to run well on multiple processors at once, regardless of whether the processors are on the same blade (server).
 - The combination of multiple powerful processors and high-speed communications channels between blades creates an extremely high-performance environment for these applications. Standard databases, such as SQL or standard Oracle, although not designed for parallel operations, will also run well on these high-end blades.
- **Storage**
- Changes have also occurred in blade storage usage. Initially, when blade systems shipped, each blade contained some amount of disk storage. These first blade drives were all IDE (laptop-type) drives. The disks were very compact but had fairly high failure rates for hard-running servers.
 - The next step for most vendors was to move to SCSI drives, as they got smaller (2.5-inch drives). This increased costs, but also reliability. As the market matured and blade usage expanded, a buzz started around a movement toward storage blades, which IBM delivered.
- **Usage Evolution**

- As the market evolved, blade adoption began to grow in areas such as financial services, scientific research, retail, and health care. Egenera, which offered clustered single-system-image blade solutions, achieved tremendous success through porting of previously proprietary applications like those found on financial trading-room floors (from Sun or IBM systems) to their blade solution.
- Early on, the guideline for considering the move to blades was more than five servers doing homogeneous work in scale-out applications. The number five was based on the up-front costs of the blade chassis and the fact that unless the chassis is roughly half-full, it is not cost-effective.

Q 126 Enumerate the benefits of blades and virtualization.

Ans:

- With all this as a backdrop, it is now useful to look more deeply at the reasons introduced earlier for considering blades and virtualization as complementary technologies. Although some of these have not changed since the initial vision of blades and server virtualization in the early 2000s, additional benefits have emerged and/or become more significant as the technology has evolved.

Additionally, significant cost-reduction benefits can be achieved from implementing blades and virtualization, but it is not all about cost. Benefits of implementing blade technology include the following:

1. Space savings and efficiency, packing more computing power into a smaller area
2. Consolidation of servers to improve and centralize management as well as to improve utilization of computer assets
3. Simplification and reduction in complexity (modularity, ease of integration, prewired configurations, shared components), ease of deployment, and improved manageability and serviceability
4. ROI and improved TCO (increased hardware utilization, reduced OpEx)

➤ **Space**

- Packing more computing power into a smaller space was one of the main initial reasons for deploying blades. This reason is still significant for many data centers that are out of room.
- As IT organizations have responded to ever-increasing demands for additional processing power and storage capacity, many data centers have grown their IT hardware to the point of maxing out their floor space as well as their power and cooling capacity. In many cases, blade systems offer help in this area by providing a way to replace larger rack servers with the same or greater capacity in a smaller footprint.
- In fact, as blade vendors have implemented faster processors, two-way processors, and dual-core processors (with quad-core processors in 2007), the computing power per square foot has increased significantly over early blade systems. For

instance, 60 blade servers might now fit in the same physical space as 42 rack-mounted servers (these numbers vary from vendor to vendor).

➤ **Consolidation and Centralization**

- Consolidation has been viewed as a way to reduce costs by increasing server utilization, reducing the number of servers needed, and reducing the time required to manage those servers.
- Consolidating legacy servers and centralizing server functions onto blades using server virtualization solutions accomplish all of these objectives.
- Blades, in conjunction with virtualization software, provide an excellent platform for physical consolidation (centrally locating all the blade servers) and virtual consolidation (consolidating multiple server applications onto the same blade under VMware or other virtualization technology).

➤ **Simplification and Manageability**

- Although rack servers provided a good option for improving space utilization over towers, for most data centers they also introduced a nightmare in terms of cabling. The modular approach also makes adding/removing servers as easy as popping in a blade.
- If a blade fails, the LEDs indicate which blade it is, and it no longer requires a skilled tech to swap out the failing blade and replace it. The integration of the blade chassis management module allows for full remote-management operations, including power up/power down.

➤ **ROI and TCO**

Many of the reasons to implement blades and virtualization translate directly to money. These can be divided into ROI issues and reducing the TCO.

1. Reduced Infrastructure Cost

When considering an implementation that includes more than five or six servers, the overall infrastructure costs of blades versus other equivalent alternatives will be less. This is due in part to the sharing of devices (switches, power supplies, and so on) and in part to the reduction of cabling and connector costs.

2. Time to Return

A significant portion of the benefit that blades and virtualization bring is in the area of ease of implementation, deployment, and provisioning. This benefit can be measured by what is called time to return (on investment) or time to revenue. The modularity and prewiring of blade systems described previously work together to dramatically reduce the time need initially to deploy blade systems.

3. Reduced Operational Cost

All of the management benefits described here can be translated (directly or indirectly) into reduced OpEx. These include everything from initial deployment to adding blades to reconfiguring and re-provisioning.

4. Reduced Serviceability Cost

Finally, the simplicity and manageability of blade systems and virtualization solutions changes the level of expertise of the repair staff. Since everything except for a hardware swap can be done remotely, when there is a hardware failure on a blade, virtualization software can automatically restart the VM on another blade (high availability), keeping the applications running.

Q 127 What is datacenter in box? What are its features?

Ans:

- A newer trend in blade usage comes from the modularity and simplicity of blade server systems. Because blade systems have the ability to house a group of servers and switches with minimal cabling and strong remote-management capabilities, they are an excellent solution for remote offices/mini data centers.
- As a result, numerous solutions have come to market, following the “data center in a box” and “retail store in a box” models. These solutions package together a variety of servers, each performing different functions.
- For example, a chassis may contain one blade that serves as a web server, one or two blades performing as application servers running software such as BEA or SAP, one or more database server blades, and blades for middleware, along with the appropriate switch modules.
- A number of components may be included for various types of uses. (Examples might include firewalls, VPNs, and load balancers.) This “data center in a box” can be dropped into a remote location and managed remotely.

➤ **Firewall/VPN Blades**

- When a blade system is used in the Data Center in a Box scenario, security is often a key requirement. As with any remote server environment, there is a need to provide a firewall and possibly a VPN.
- These solutions are provided through security software operating on a standard blade server, from security vendors such as Check Point.

➤ **Load-Balancing Blades**

- As part of delivering scale-out applications and in remote data centers where multiple servers are running the same application, a load-balancing function is needed to manage the distribution of work across multiple servers.
- The load balancer acts as a virtual server, representing and responding to a single IP address and port, and distributing the requests to a selected physical server within the blade farm.

➤ **Ruggedized Portable Computing**

- Another successful extension of the Data Center in a Box solution is for situations requiring ruggedization for remote locations. Examples of these include blades systems for use on a ship, for military applications, for seismic simulations, and so on.

Q 128.What are the fundamental design objectives of blade server systems? Explain.**Ans:**

- The overall architecture of the blade server system was designed to minimize space requirements, simplify connections between computing components, and improve the manageability of the system components. The blade chassis leverages several principles from the chassis architectures, built for the telecommunications and networking industries, and applies them to servers.
- The most obvious of these principles is the placement of the blades into slots within the chassis (typically vertical, but not with all vendors). Another of these borrowed concepts is the sharing of common power supplies and fan assemblies at the rear of the chassis, which cool all of the blades. The chassis provides the interconnect between the blades and the switches, effectively eliminating much of the cabling mess created with rack servers.
- The fundamental design objectives for blade server systems include the following:
 - i. Aggregation of various components to be shared by the blades within the chassis, including power supplies, fans, I/O, and so on.
 - ii. Designed for manageability from the outset— introducing the idea of a chassis manager function, which delivers out-of-band management to all of the elements within the chassis. This enables one-to-many management of these components, either from the chassis manager itself or through integration with higher-level systems-management products.
 - iii. Redundancy of switches and chassis-management modules, going beyond redundancy of power supplies and fans.
 - iv. Shared modules (such as KVM and floppy/CD-ROM) primarily for use in diagnostics.

Q.129 State and explain the components of blade systems.**Ans:**

The blade system comprises the following components:

- i. Chassis (including shared power supplies, fans, and so on)
- ii. One or more blade servers
- iii. Redundant switch modules
- iv. Chassis manager module (also called the management module)
- v. Shared modules (primarily for diagnostics)
- vi. Keyboard-video-mouse (KVM) interface
- vii. Floppy and CD-ROM drive tray

➤ **Chassis**

- The chassis is the base component of blade systems and provides the following:
 - Housing for all the components of the blade system
 - Connectivity and plumbing for the various components via the midplane

- The chassis includes the power supplies and the fans that provide power and cooling to the various components of the blade system.




Figure: Blade chassis — back view

➤ Server Blades

- The blade server is the central computing element of the blade system. A blade server typically has either one or two CPUs (also called one-way or two-way blades). Four-way blades are also available but may take up two slots in the chassis, depending on the vendor.
- Each CPU processor can be single-core or multicore, adding more horsepower per socket. In addition to the processors, each blade has a variable amount of memory and a minimum of one network connection (Ethernet by default).

➤ I/O Modules

- A blade chassis generally has two or three switch modules (also called switch blades or blade switches) to communicate between the blades and various external networks (both data networks and storage networks). These switch modules are typically redundant, as are the I/O options supported on the blade server.
- Two types of switch modules are available for blades: switch modules and pass-through modules. Switch modules are typically switches that provide aggregation of all I/O connectivity to the blades. In other words, each blade talks to the switch module, which provides standard network switching for all the blades in the chassis.
- Pass-through modules, on the other hand, provide one-to-one access from an external switch to the I/O connectivity on each server blade. In essence, this

provides a pass-through cabling function from each blade to an external switch. In order for the blades to communicate with each other or to an external network, at least one switch module or pass-through module must be present in the chassis.

➤ **Management Modules**

- Many of the blade chassis designs incorporate a management module that is responsible for controlling and managing the various components in the blade chassis, including the power supplies, the fans, the switches, and the blades.
- The management module is also responsible for powering up the blade system and its individual components, configuring the switches and the blades, and monitoring the operations of all the various components in the blade chassis. The management module can communicate directly with each blade and each switch in the chassis via an out-of-band, direct connection.

➤ **Other Shared Modules**

- The final components of a blade system include the media or device bay and the KVM module. These components were used initially in blade systems, primarily as shared diagnostics devices, connecting to one blade server at a time.
- In the early blade system chassis, the device bay was typically a floppy/CDROM device, used for boot purposes and/or diagnostics. With the advent of USB-based boot media support, KVM modules can be equipped with USB ports that provide another means to boot the blade server.
- Advances in KVM over IP now allow the KVM functions to be handled by the management module using standard IP protocol. Because of these types of advances, the trend is toward elimination of these components, allowing more space for other blade components.

Q 130 Explain server blades, I/O modules and management modules.

Ans:

➤ **Server Blades**

- The blade server is the central computing element of the blade system. A blade server typically has either one or two CPUs (also called one-way or two-way blades). Four-way blades are also available but may take up two slots in the chassis, depending on the vendor.
- Each CPU processor can be single-core or multicore, adding more horsepower per socket. In addition to the processors, each blade has a variable amount of memory and a minimum of one network connection (Ethernet by default).
- There also may be an I/O card (daughter or mezzanine card on the blade motherboard), which provides additional network connections, typically with redundant configurations (dual ports). The blade servers can communicate with

multiple network fabrics in addition to the management network. These include Ethernet, Fiber Channel, and InfiniBand.

- Depending on the vendor, each blade may be capable of local storage, configured with up to two hard-disk drives (HDDs). If HDDs are on the blade, they are placed on the front side of the blade in most designs, and are hot pluggable.
- Some vendors (such as Egenera) do not include drives on the blade at all (called diskless or anonymous blades). In this scenario, the blade server communicates to external storage on a NAS device or a SAN. Most vendors offer their blades with or without disk drives, leaving the choice to the user.

➤ **I/O Modules**

- A blade chassis generally has two or three switch modules (also called switch blades or blade switches) to communicate between the blades and various external networks (both data networks and storage networks). These switch modules are typically redundant, as are the I/O options supported on the blade server.
- Two types of switch modules are available for blades: switch modules and pass-through modules. Switch modules are typically switches that provide aggregation of all I/O connectivity to the blades. In other words, each blade talks to the switch module, which provides standard network switching for all the blades in the chassis.
- Pass-through modules, on the other hand, provide one-to-one access from an external switch to the I/O connectivity on each server blade. In essence, this provides a pass-through cabling function from each blade to an external switch.
- In order for the blades to communicate with each other or to an external network, at least one switch module or pass-through module must be present in the chassis. Supported switches for blade communications include Gigabit Ethernet, 10Gb Ethernet, Fiber Channel (2Gb and 4Gb), and InfiniBand.

➤ **Management Modules**

- Many of the blade chassis designs incorporate a management module that is responsible for controlling and managing the various components in the blade chassis, including the power supplies, the fans, the switches, and the blades.
- The management module is also responsible for powering up the blade system and its individual components, configuring the switches and the blades, and monitoring the operations of all the various components in the blade chassis.
- The management module can communicate directly with each blade and each switch in the chassis via an out-of-band, direct connection.

Q 131 What are the CPU consideration for blades?**Ans:**

- After you have decided to use blades, sizing the CPUs is not much different than it is with traditional servers. Understanding your application requirements is the key to success for CPU sizing, as well as for memory and I/O connections.
- These become even more important if you are consolidating physical servers into multiple virtual machines on a blade.
- Requirements for hardware virtualization assists (such as with Xen and Windows Server Virtualization)
- Architectural requirements for special chips or chipsets (for example, some applications may require a particular chip that is available from only one vendor)
- Be clear when discussing CPUs, cores, and sockets. The total number of cores per blade is calculated by the number of sockets times the number of cores per socket. This can be especially important for software licensing, as some vendors may license by CPU or cores, and some by the number of physical sockets on the blade. Aim high on processor power, especially if you are consolidating through virtualization.
- Consider purchasing high-end blades with efficient power and cooling built in. Evaluate the cooling requirements of the processor-chip options you are considering (that is, evaluate comparable options from multiple vendors, including AMD and Intel).
- If you are using different classes of blades (for example, some low-end and some higher-end), you might consider grouping them by frame for manageability. However, if your high-end blades produce excessive heat output, consider mixing them within the frames to spread the heat, and managing them as a group of resources.
- If necessary for proper cooling, it's OK to leave some empty space in your chassis or racks. Even if you don't fully populate your chassis or racks due to power or cooling issues, blades still offer more computing power per footprint than traditional servers.
- If you are using or planning to use server virtualization, the tools for moving running virtual machines may allow movement only within matching processor groups. This may also be true for image deployment and server failover. Consistency in memory size may also be helpful.
- Keep this in mind as you make your blade-purchasing decisions. (This is one of the reasons that virtualization goes well with blades, because blades are typically bought as a group.)

Q.132 Discuss the memory and I/O considerations for blades.**Ans:****➤ Memory:**

- Many of the considerations listed for CPUs also apply to memory. Understanding your applications and their runtime characteristics is the best way to determine requirements when moving to blades. Databases may be CPU-heavy and require a lot of memory for efficient database record handling.
- In this case, the database back end should be run on a more powerful blade with lots of memory. If the middleware is I/O-intensive, don't stack all the middleware blades within the same chassis (spread them across chassis so you don't saturate the I/O connections).
- Aim high on memory as well as speed, especially when virtualizing.
- Understand how each application utilizes memory and CPU resources. If possible, use a performance-monitoring tool to baseline the current application performance parameters.
- If possible, use a test lab to size the hardware and to test the effects of higher-speed processors and/or more memory. Does adding more power make a difference, or is the bottleneck in I/O?
- Keep in mind that applications such as database engines or heavy crunchers like SAS will perform better the more resources they have (both higher processing speed and more memory).

➤ I/O:

- I/O issues need to be considered at various levels of blade selection and implementation. Chassis backplanes vary considerably in terms of types and speeds of interconnects and I/O-card options. There is also significant diversity in both the approach to blade I/O and in speeds of switch modules.
- In addition, some vendors offer pass-through options as well as switch modules, and others offer only pass-through.
- Evaluate all of your I/O requirements before deciding on a chassis. If you plan to run applications with low I/O requirements, there may be lower speed options that are more cost-effective.
- If, however, you will need high speed now or down the road, implementing a chassis with a high speed backplane/midplane will allow the blades in the chassis that do need the speed to leverage that backplane.
- Investigate your current environment for the number of NICs that are used per server. Ensure that the chassis is not limited to less than the number needed per blade.
- Consider the I/O speed needed per blade for your LAN connections. If you are running Ethernet, will GbE be enough, or will you need to upgrade to 10GbE within the next five years?

- After you have looked at individual blade requirements, remember to consider the possibility of saturation at the aggregated level (that is, the switch), especially if you are consolidating through virtualization.

Q.133.Explain the storage considerations for blades.

Ans:

- For I/O to storage devices, most blade chassis offer the option to use on-blade storage and/or to connect to a SAN or NAS. Diskless blades were introduced by Egenera in the very early days of blades and now are being implemented more frequently across all blade-vendor chassis.
- If you are connecting to a SAN, does your environment support boot from SAN, also known as PXE boot (Preboot Execution Environment)? This option is generally supported for all Fiber Channel connections and, more recently, for most iSCSI SANs as well (depending on which iSCSI initiator is used).
- Check with your SAN vendor to confirm any configuration requirements for SAN boot (for example, OS levels, firmware, and so on). • If you are connecting to a SAN and don't have a strong requirement for local disk storage on each blade, consider going diskless to simplify re-provisioning and management.
- Understand bandwidth requirements for your SAN connections. Fiber Channel HBAs are available for most blades in 2 GB and 4 GB options. If you are using iSCSI, would your performance benefit from a TCP Offload Engine (TOE) option on your IP NIC, if one is available?
- If your chassis supports switch modules from your SAN-switch vendor, it is important to verify that there are no special features configured on your external switch that are required but not available on their blade switch module.
- Understand your aggregated bandwidth requirements. Remember to verify not only that you have enough bandwidth at the blade level, but also at the aggregated level. If, for example, you are running 14 blades connected to an embedded switch module, will I/O become a bottleneck?
- This can become an even bigger issue if you are consolidating/virtualizing multiple virtual servers with high I/O requirements onto single blades. If this becomes a problem, it may be necessary to use the pass through option rather than switch modules for certain applications.
- If you require local HDD support, look for redundancy and hot-swap capability to help minimize disruption due to disk failures.
- If a particular server application requires a large number of direct connect HDDs (for example, 10), blades may not be the right choice.