

IT Infrastructure Architecture

Infrastructure Building Blocks
and Concepts

Compute

Virtualization technologies

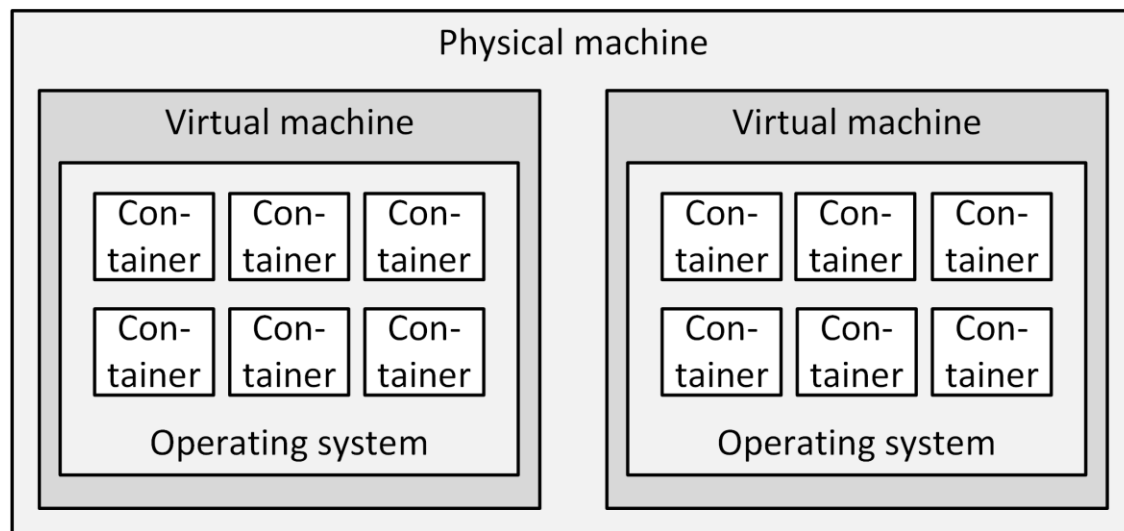
- Emulation
 - Can run programs on a computer, other than the one they were originally intended for
 - Run a mainframe operating system on a x86 server
- Logical Partitions (LPARs)
 - Hardware based
 - Used on mainframe and midrange systems

Virtualization technologies

- Hypervisors
 - Control the physical computer's hardware and provide virtual machines with all the services of a physical system
 - Virtual CPUs
 - BIOS
 - Virtual devices
 - Virtualized memory management
 - Three types:
 - Binary translation
 - Paravirtualization
 - Hardware assisted virtualization (most used on x86 servers)

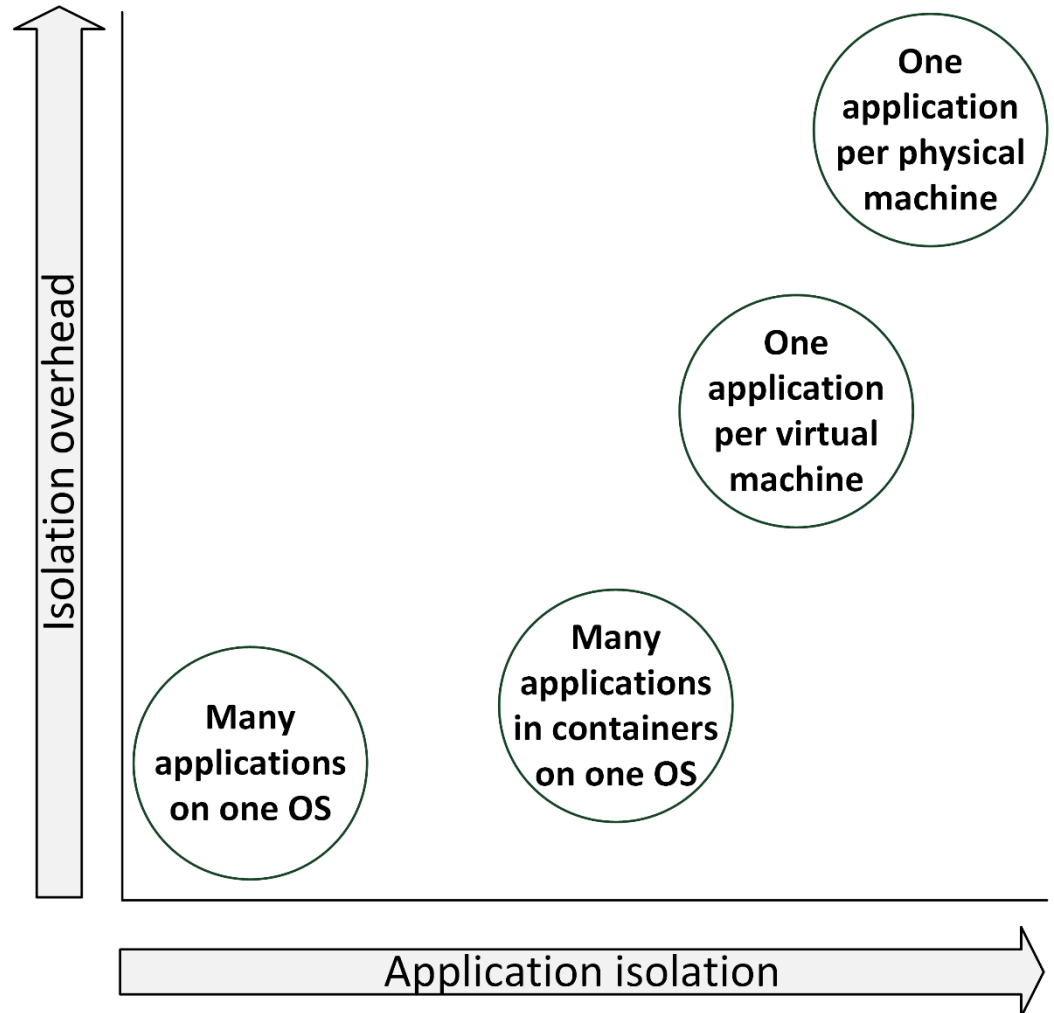
Container technology

- Container technology is a server virtualization method in which the kernel of an operating system provides multiple isolated user-space instances, instead of just one
- Containers look and feel like a real server from the point of view of its owners and users, but they share the same operating system kernel
- Containers are part of the Linux kernel since 2008



Container technology

- Containers are a balance between isolation and overhead of running isolated applications



Container technology

- Containers have a number of benefits:
 - Isolation
 - Applications or application components can be encapsulated in containers, each operating independently and isolated from each other
 - Portability
 - Since containers typically contain all components the application needs to function, including libraries and patches, containers can be run on any infrastructure that is capable of running containers using the same kernel version
 - Easy deployment
 - Containers allow developers to quickly deploy new software versions, as their containers can be moved from the development environment to the production environment unaltered

Container implementation

- Containers are based on 3 technologies that are all part of the Linux kernel:
 - Chroot (also known as a jail)
 - Changes the root directory for the current running process
 - Ensures processes cannot access files outside the designated directory tree
 - Namespaces
 - Allows complete isolation of an applications' view of the operating environment
 - Process trees, networking, user IDs and mounted file systems


Container implementation

– Cgroups

- Limits and isolates the resource usage of a collection of processes
 - PU, memory, disk I/O, network
- Linux Containers (LXC), introduced in 2008, is a combination of these
 - Docker is a popular implementation of a container ecosystem

Container orchestration

- Container orchestration abstracts the resources of a cluster of machines and provides services to containers
- A container orchestrator enables containers to be run anywhere on a cluster of machines
 - Schedules the containers to run on any machine that has resources available
 - Acts like a kernel for the combined resources of an entire datacenter



Mainframes

- A mainframe is a high-performance computer made for high-volume, I/O-intensive computing
 - Expensive
 - Mostly used for administrative processes
 - Optimized for handling high volumes of data
- IBM is the largest vendor – it has 90% market share
- The end of the mainframe is predicted for decades now, but mainframes are still widely used
- Today's mainframes are still large (the size of a few 19" racks), but they don't fill-up a room anymore

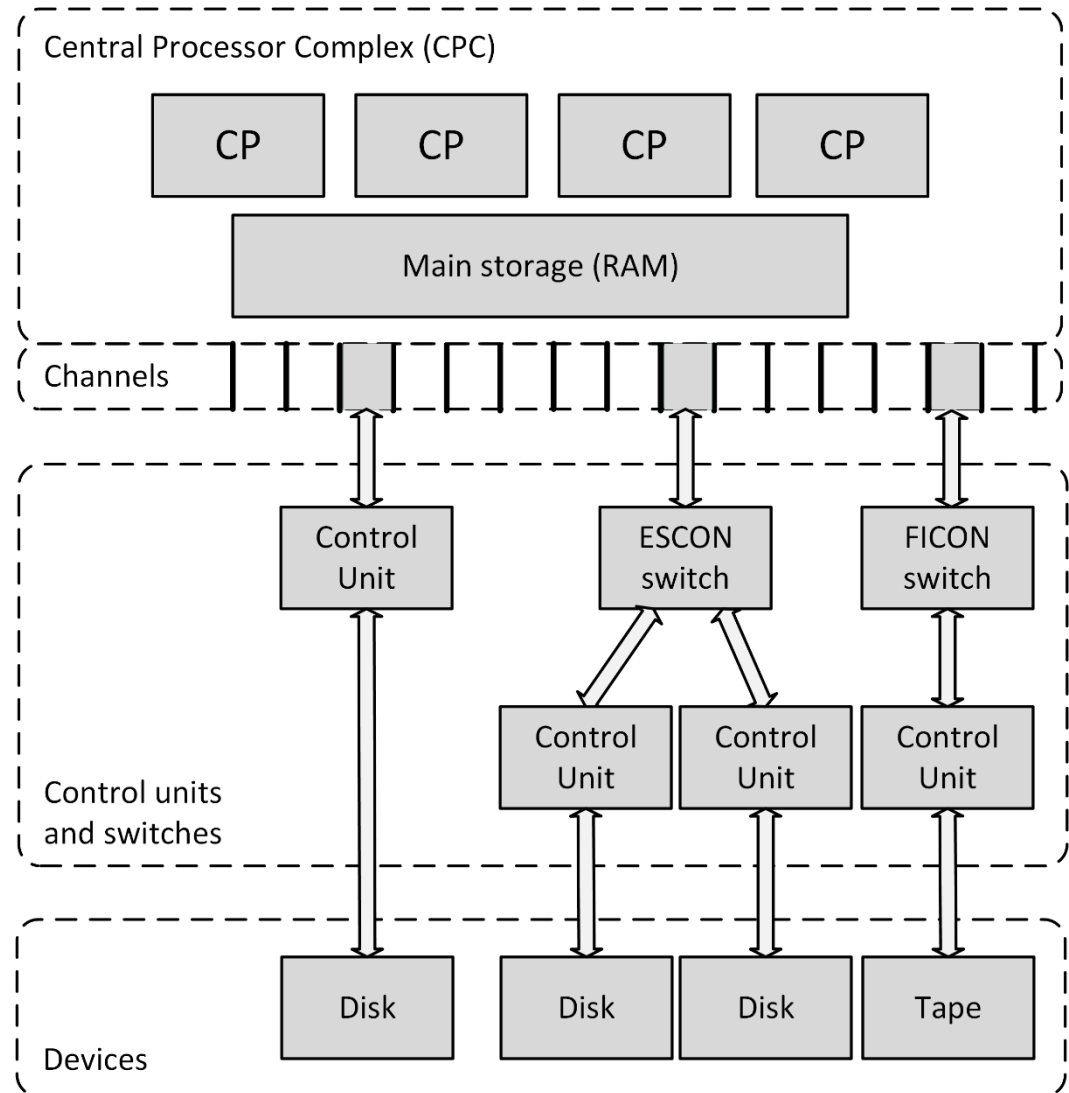
Mainframes

- Mainframes are highly reliable, typically running for years without downtime
- Much redundancy is built in
 - Hardware can be upgraded and repaired while the mainframe is operating without downtime
- The latest IBM z13 mainframe:
 - Introduced in 2015
 - Up to 10TB of memory
 - Up to 141 processors
 - Running at a 5GHz clock speed
 - Can run up to 8000 virtual machines simultaneously



Mainframe architecture

- A mainframe consists of:
 - Processing units (PUs)
 - Memory
 - I/O channels
 - Control units
 - Devices, all placed in racks (frames)



Mainframe architecture – PU, memory, and disks

- In the mainframe world, the term PU (Processing Unit) is used instead of CPU
 - A mainframe has multiple PUs, so there is no central processing unit
 - The total of all PUs in a mainframe is called a Central Processor Complex (CPC)
- Each book package in the CPC cage contains from four to eight memory cards
 - For example, a fully loaded z9 mainframe has four book packages that can provide up to a total of 512 GB of memory
- Disks in mainframes are called DASD (Direct Attached Storage Device)
 - Comparable to a SAN in a midrange or x86 environment

Mainframe architecture – Channels and control units

- A channel provides a data and control path between I/O devices and memory
- Today's largest mainframes have 1024 channels
- A control unit is similar to an expansion card in an x86 or midrange system
 - Contains logic to work with a particular type of I/O device, like a printer or a tape drive

Mainframe architecture – Channels and control units

- Channel types:
 - OSA
 - Connectivity to various industry standard networking technologies, including Ethernet
 - FICON
 - The most flexible channel technology, based on fiber-optic technology
 - With FICON, input/output devices can be located many kilometers from the mainframe to which they are attached
 - ESCON
 - An earlier type of fiber-optic technology
 - Almost as fast as FICON channels, but at a shorter distance

Mainframe virtualization

- Mainframes were designed for virtualization from the start
- Logical partitions (LPARs) are the default virtualization solution
- LPARs are equivalent to separate mainframes
- A common number of LPARs in use on a mainframe is less than ten
- The mainframe operating system running on each LPAR is designed to concurrently run a large number of applications and services, and can be connected to thousands of users at the same time
- Often one LPAR runs all production tasks while another runs the consolidated test environment