clock ticks to execute
  - An ADD instruction typically costs 1 tick to compute
- The speed at which the CPU operates is defined in GHz (billions of clock ticks per second)
  - A single core of a 2.4 GHz CPU can perform 2.4 billion additions in 1 second

Processors – Word Size

- Each CPU is designed to handle data in chunks, called words, with a specific size
  - The first CPUs had a word size of 4 bits
  - Today, most CPUs have a word size of 64 bits
- The word size is reflected in many aspects of a CPU's structure and operation:
  - The majority of the internal memory registers are the size of one word
  - The largest piece of data that can be transferred to and from the working memory in a single operation is a word
  - A 64-bit CPU can address 17,179,869,184 TB of memory (64-bit word)

Intel x86 Processors

- Intel CPUs became the de-facto standard for many computer architectures
  - The original PC used a 4.77 MHz 16-bit 8088 CPU
  - A few years later, Intel produced the 32-bit 80386 and the 80486 processors
- Since these names all ended with the number 86, the generic architecture was referred to as x86
- In 2017, the latest Intel x86 model is the 22-core E5-2699A Xeon Processor, running on 2.4 GHz

AMD x86 Processors

- Advanced Micro Devices, Inc. (AMD) is the second-largest global supplier of microprocessors based on the x86 architecture
  - In 1982, AMD signed a contract with Intel

- Relatively expensive
- High end disks
- Spinning disk platters with a rotational speed of 10,000 or 15,000 rpm
- Typically have 25% of the capacity of SATA or NL-SAS disks
- Uses the SCSI command set that includes error-recovery and error-reporting and more functionality than the SMART commands used by SATA disks

Mechanical Hard Disks

- Near-Line SAS (NL-SAS) disks
  - Have a SAS interface, but the mechanics of SATA disks
  - Can be combined with faster SAS disks in one storage array

Solid State Drives (SSDs)

- Based on flash technology
- Connected using a standard SAS disk interface
- Consume less power, and therefore generate less heat
- No moving parts and no vibration
- Data can be accessed much faster than using mechanical disks

Disk Capacity - Kryder's Law

- The density of information on hard drives doubles every 13 months
- An average single disk drive in 2025 will hold more than 20,000 TB (20 PB) of data

- Reasons:
  - High clock speeds make connections on the circuit board work as a radio antenna
  - A frequency of 3 GHz means a wavelength of 10 cm. When signals travel for more than a few cm on a circuit board, the signal gets out of phase with the clock
  - The CPU can heat up tremendously at certain spots, which could lead to a meltdown

CPU: Multi-core CPUs

- A multi-core processor is a CPU with multiple separate cores
  - The equivalent of getting multiple processors in one package
- The cores in a multi-core CPU run at a lower frequency
  - Reduce power consumption
  - Reduce heat (no hot spots)
- Trend is to have CPUs with tens or even hundreds of cores

CPU: Hyper-threading

- Certain Intel CPUs contain a propriety technology called hyper-threading
  - For example the Core i3/i5/i7 and Xeon CPUs
- Hyper-threading makes a single processor core virtually work as a multi-core processor
- Hyper-threading can provide some increase in system performance by keeping the processor pipelines busier

Virtualization performance

- Consolidating multiple virtual machines on one physical machine increases CPU usage and reduces CPU idle time
  - This is a primary driver for the use of virtualization
- The physical machine needs to handle th disk and network I/O of *all* running virtual machines
  - This can easily lead to an I/O performance

anyway), creating instruction overlap

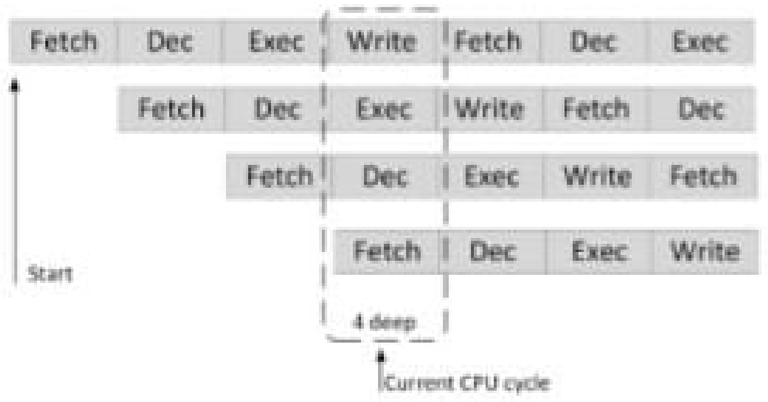

CPU: Prefetching and branch prediction

- Prefetching:
  - When an instruction is fetched from main memory, also the next instructions are fetched and stored in cache
  - When the CPU needs the next instruction it is already available in cache
- Unfortunately, most programs contain jumps (also known as branches), resulting in cache misses
  - The next instruction is not the next instruction in memory

CPU: Prefetching and branch prediction

- The cache system tries to predict the outcome of branch instructions before they are executed by the CPU (called branch prediction)
  - In practice more than 80% of the processor instructions are delivered to the CPU from cache memory using prefetching and branch prediction

CPU: Superscalar CPUs

- A superscalar CPU can process more than one instruction per clock tick
- This is done by simultaneously dispatching multiple instructions to redundant functional units on the processor

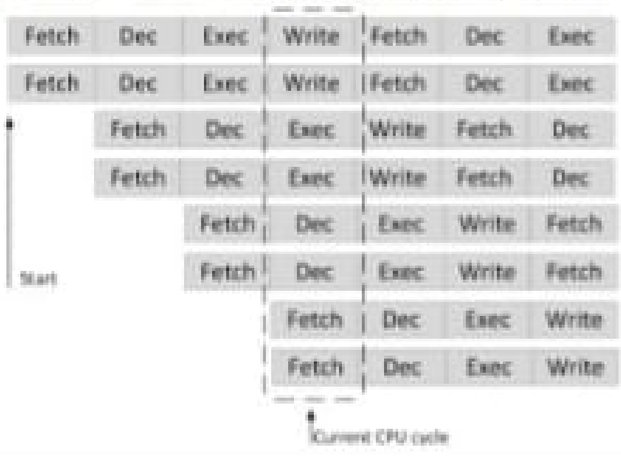

CPU: Multi-core CPUs

start and stop systems or applications on demand
- Rapid elasticity – A cloud is able to quickly scale-up and scale-down resources, leading to elasticity of resources
- Resource pooling – Instead of providing each application with a fixed amount of processing power and storage, cloud computing provides applications with resources from a shared pool
- Measured service – The actual resource usage is measured and billed. There are no capital expenses, only operational expenses
- Broad network access – Capabilities are available over the network and accessed through standard mechanisms

Cloud deployment models

- A public cloud is accessible through the internet, and available to the general public
- A private cloud is operated solely for a single organization, whether managed internally or by a third-party, and hosted either on premises or external
- A community cloud is much like a private cloud, but shared with a community of organizations that have shared concerns
- In a hybrid cloud deployment, a service or application is provided by a combination of a public cloud, and a community cloud and/or a private cloud

Cloud service models

- Software-as-a-Service (SaaS) delivers full applications that need little or no configuration
  - Microsoft Office365, LinkedIn, Facebook, Twitter, and Salesforce.com
- Platform-as-a-Service (PaaS) delivers a scalable, high available, open programming platform that can be used by developers to build bespoke applications that run on the PaaS platform
  - Microsoft Azure Cloud Service and Google App Engine
- Infrastructure-as-a-Service (IaaS) delivers virtual machines, networking, and storage
  - Amazon Elastic Cloud (EC2 and S3) and Microsoft Azure IaaS

Cloud service models

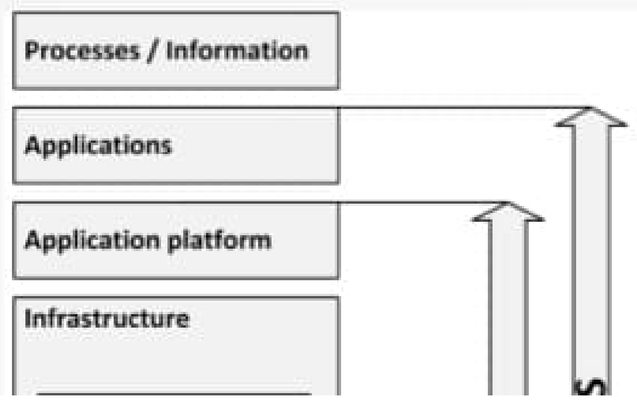

shifts in computing in recent years

- Definition NIST:

*Cloud computing is a model for enabling ubiquitous, convenient, **on-demand network access** to a **shared pool** of configurable computing resources that can be **rapidly provisioned** and released with **minimal management effort** or service provider interaction*

- Cloud computing is an outsourcing model, in which IT services are provided and paid based on actual on-demand use

Cloud characteristics

- On demand self-service – End uses can configure, deploy, start and stop systems or applications on demand
- Rapid elasticity – A cloud is able to quickly scale-up and scale-down resources, leading to elasticity of resources
- Resource pooling – Instead of providing each application with a fixed amount of processing power and storage, cloud computing provides applications with resources from a shared pool
- Measured service – The actual resource usage is measured and billed. There are no capital expenses, only operational expenses
- Broad network access – Capabilities are available over the network and accessed through standard mechanisms

Cloud deployment models

- A public cloud is accessible through the internet, and available to the general public
- A private cloud is operated solely for a single organization, whether managed internally or by a third-party, and hosted either on premises or external
- A community cloud is much like a private cloud, but shared with a community of organizations that have shared concerns
- In a hybrid cloud deployment, a service or application is provided by a combination of a public cloud, and a community cloud and/or a private cloud

Cloud service models

- Software-as-a-Service (SaaS) delivers full applications that need little or no configuration
  - Microsoft Office365, LinkedIn, Facebook, Twitter, an Salesforce.com
- Platform-as-a-Service (PaaS) delivers a scalab high available, open programming platform that can be used by developers to build bespoke

- Sufficient floor loading capacity

On-premises hosting

- Drawbacks:
  - On-premises datacenters don't scale well, as they are embedded in existing (office) buildings
  - The organization must have enough knowledge and staff available to manage the datacenter

Colocation

- A third party dedicated datacenter where racks, floor space, and network bandwidth can be rented
- Hosts and connects customer owned infrastructure components
- Provides:
  - Power
  - Cooling
  - Physical security
- Racks are empty – all infrastructure components must be provided and managed by the organization renting the colocation racks

Outsourcing

- Full infrastructure outsourcing is a subcontracting service in which some third-party purchases, deploys, hosts, and manages the infrastructure, and performs its lifecycle management
  - Managed using Service Level Agreements
  - Very rigid change management process
- Frees the organization from investing in hardware
  - Only operational cost
- The outsourcing organization must have:
  - A demand organization
  - A process to manage the outsourcing party

Enterprise infrastructure deployment

disk
- Installed at a third location

Voting and quorum disks

- The quorum acts as one vote in the voting system
- The quorum disk is always assigned to one (and only one) node at any time
- A faulty node releases it quorum assignment automatically
- The working node gets two votes: one from itself, the other from the quorum disk
- The faulty node will stop working, because it has only one vote

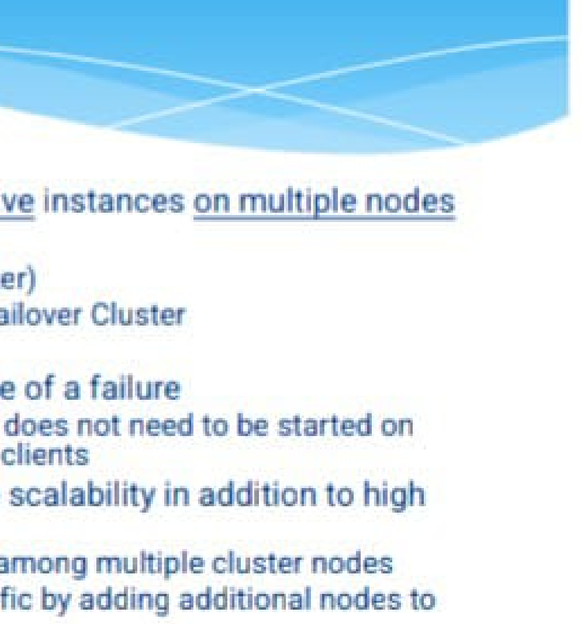

Cluster-aware applications

- Cluster-aware applications run active instances on multiple nodes
- Examples:
  - Oracle RAC (Real Application Cluster)
  - Microsoft SQL Server Always On Failover Cluster
  - Microsoft Exchange Server
- Enhances switch-over times in case of a failure
  - In case of a failure, the application does not need to be started on another node before it can service clients
- Cluster-aware applications provide scalability in addition to high availability
  - Client requests can be distributed among multiple cluster nodes
  - Handle increased demand and traffic by adding additional nodes to the cluster

Operating system performance

Operating system performance

- Windows was no real operating system
- In 1990, Microsoft Windows 3.0 was the first successful Windows version
- In late 1995, Microsoft released Windows 95, positioned as the new operating system for desktops
  - Windows 95 introduced the "start" button
- Windows targeted at workstations include:
  - Windows XP
  - Windows Vista
  - Windows 7, 8 and 10

Windows for servers

- In 1992, Windows NT was the first version of Windows designed to run on servers
  - A real operating system, not running on top of MS-DOS
- Windows NT 4
  - Included the Windows 95 style GUI
  - Companies started the switch from Novell servers to Windows NT 4
  - Some UNIX systems were being replaced by Windows NT 4 systems
- Windows 2000 introduced an implementation of LDAP directory services, called Active Directory
- The server operating systems were named after the year of release: Windows server 2003, 2008, 2012, and 2016

Windows – Stability

- While NetWare and UNIX would run for at least a year without crashing, it was not uncommon that a Windows server crashed once a day
- Causes:
  - The need for backwards compatibility of Windows
    - Every version of Windows needed to be able to run all already developed software without recompilation
  - Windows runs on all kinds of hardware
    - As opposed to UNIX or Apple systems, which are designed for specific hardware
    - The quality of third-party drivers was not always guaranteed

Windows – Security

- Windows security was weak
  - Windows was based on MS-DOS – a single user / single tasking operating system
    - Multi-user features and concurrently running multiple applications was built in later
  - Most Windows applications were not designed with multi user usage in mind
    - Applications had to run with the highest possible user

Operating System building blocks

- Drivers are small applications that connect specific hardware devices to the kernel
  - Printers
  - Network cards
  - Keyboard and mouse
  - Video screens
- Utilities are applications that are considered part of the operating system
  - User interfaces
  - Logging tools
  - Editors
  - System update processes

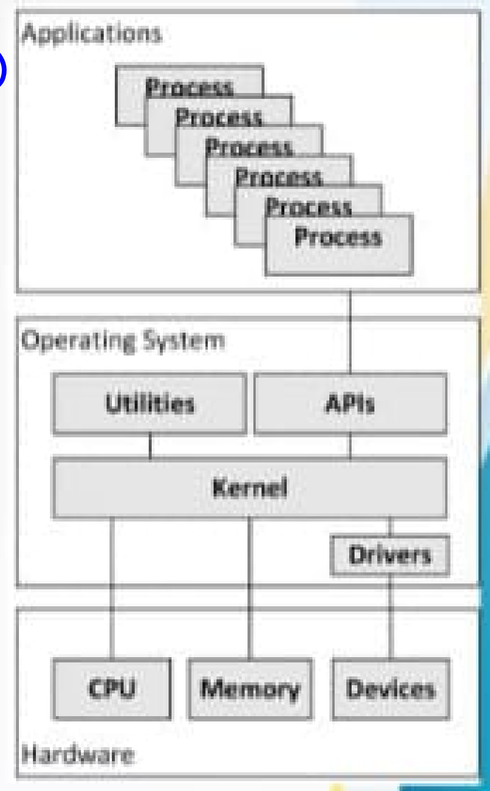

Operating System building blocks

- Applications consist of one or more processes that communicate with the operating system using system calls that are invoked through Application Programming Interfaces (APIs)

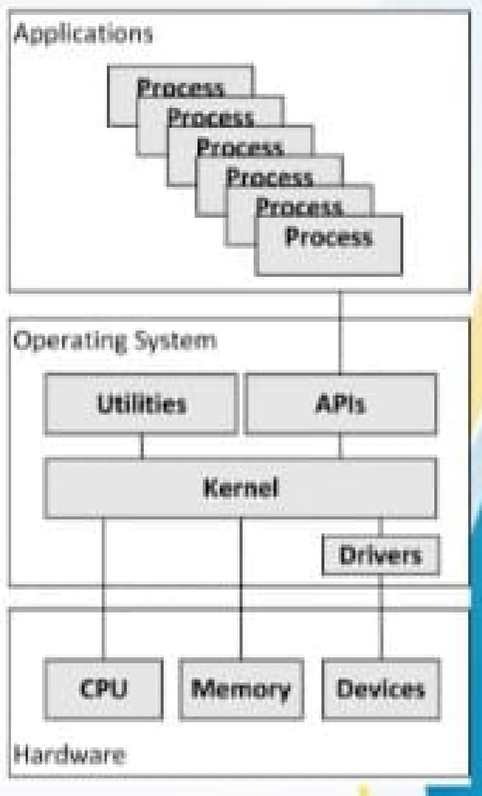

Process Scheduling

- Operating systems create the illusion of multiple running processes in parallel by scheduling each process to run only during a short time frame
  - This principle is also known as preemptive multitasking
  - Periodically, the operating system decides if a running process is to be suspended in favor of another process, or if the running process can keep on running for a while
- Process scheduling is fairly complex
  - Must be well-balanced
  - Switching processes introduces some overhead
  - The scheduling algorithm guarantees each process gets its fair share of CPU time
  - Because operating systems have evolved over decade scheduling algorithms are very sophisticated

Operating System Building Blocks

- The kernel is the heart of an operating system
  - Starts and stops programs
  - Manages the file system
  - Performs low level tasks that most programs need
    - The kernel schedules access to hardware to avoid conflicts if two programs try to access the same resource or device simultaneously

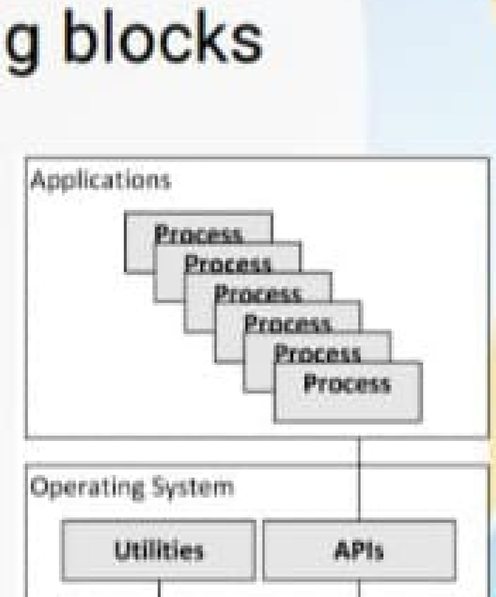

Operating System building blocks

- Drivers are small applications that connect specific hardware devices to the kernel
  - Printers
  - Network cards
  - Keyboard and mouse
  - Video screens
- Utilities are applications that are considered part of the operating system
  - User interfaces
  - Logging tools
  - Editors
  - System update processes

Operating System building blocks

- Applications consist of one or more processes that communicate with the operating system using system calls that are invoked through Application Programming Interfaces (APIs)

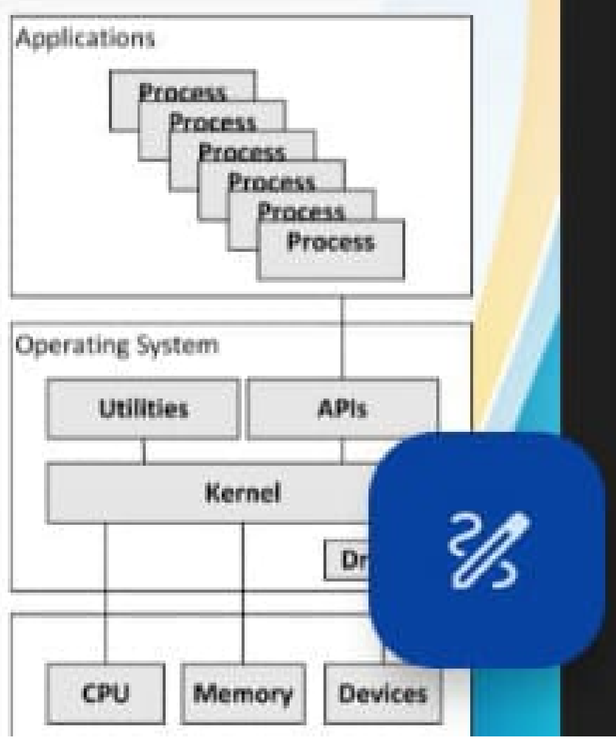

- Focused on management issues
- COBIT is a framework created by ISACA

Systems management processes - ITIL

- ITIL is the most used approach to implementing systems management processes
- Full life cycle of IT management
  - IT organization
  - 37 processes
  - Most implemented processes:
    - Incident management
    - Change management
    - Problem management
    - Release management
    - Capacity management

DevOps for infrastructure

- DevOps is a contraction of "developer" and "system operator“
- DevOps is typically used by teams developing and running functional software
- DevOps teams consist of:
  - Developers
  - Testers
  - Application systems managers
- Each team is responsible for developing and running one or more business applications or services

DevOps for infrastructure

- DevOps can also be used to develop and run an infrastructure platform
- In an infrastructure DevOps team:
  - Infrastructure developers design, test, and

– Change management
– Configuration management

Systems management processes - COBIT

- COBIT provides a structure for setting up:
  - IT Governance
  - IT organization
  - IT architecture
- Describes thirty-four IT processes
  - Management objectives
  - Associated measures
  - Performance indicators
  - Maturity levels
- Focused on management issues
- COBIT is a framework created by ISACA

Systems management processes - ITIL

- ITIL is the most used approach to implementing systems management processes
- Full life cycle of IT management
  - IT organization
  - 37 processes
  - Most implemented processes:
    - Incident management
    - Change management
    - Problem management
    - Release management
    - Capacity management

DevOps for infrastructure

- DevOps is a contraction of "developer" and "system operator“
- DevOps is typically used by teams developing and running functional softwa
- DevOps teams consist of:
  - Developers

Test stages

- Acceptance test
  - The final check of the delivered infrastructure
  - Consists mainly of verifying that all tests are performed and that defects found in previous tests are either solved or accepted
  - Leads to a discharge of the project for the delivery of the infrastructure

Go live scenarios

- Big Bang
  - At a set time, the existing system is switched off and the new system is immediately put in production, possibly after a short data migration run.
- Disadvantage:
  - It may be impossible to roll back to the old system after the system is live for some time
  - Downtime can occur when something goes wrong during the switchover

Go live scenarios

- Parallel changeover
  - Both the new and the existing system run simultaneously for some time (typically weeks)
  - Allows for testing the new system on both functionality and non-functional attributes
  - Ensuring it works with live production data before switching off the existing system
  - Switching back is possible at any time, minimizing risk
- Disadvantage:
  - The cost of maintaining both systems
  - Possible extra work to keep both systems in sync
  - Many system designs don't allow running two systems in parallel, for instance, if the system has many data interfaces with other systems

- After assembling the infrastructure, it should be tested
- Each test type has a predefined scope:
  - Functional tests
    - Ensure the infrastructure delivers the required functionality
  - Performance tests
    - Load, stress, and endurance tests
    - Prove the infrastructure has enough resources to run applications with the required performance
  - Security tests
    - Penetration tests and vulnerability scans
    - Prove security controls are in place and are functioning as designed

Testing the infrastructure

- Availability of the infrastructure can be tested by:
  - Physical actions
    - Pulling cables from infrastructure components
    - Unexpectedly rebooting machines
  - A failover test, a fallback test and a disaster recovery test
  - Testing backup and recovery processes and disaster recovery plans

Test stages

- System integration test
  - The combination of components is tested including their interfaces
  - A system integration test checks both functional and non-functional requirements
- Fallback test
  - Fallback from the main datacenter to the secondary datacenter is checked on a technical level
- Migration test
  - Ensures applications are installed without errors and data from previous systems can be migrated to the new system as designed

Test stages

- Acceptance test

Bidding and tendering

Ntp ka b pocha tha ny stand for

- Offer
  - Suppliers provide the answers to the RFP
  - Includes an initial offer
- Terms and conditions negotiations
  - The purchase department starts negotiations with the suppliers that provided the best response to
- BAFO - Best and final offer
  - [illegible] a final price and SoW
- Award
  - Based on the BAFO, the purchase department awards the supplier with the deal

Ordering and delivery

- Ordering
  - Typically done by the purchasing department
  - Because the delivery time is often weeks after the purchase order is placed, it makes sense to start ordering the goods as early as possible
- Delivery
  - Beware that the person that physically receives the goods, is not always the one formally accepting the delivery
  - Before signing for delivery, check the boxes for any damage and check for completeness of the delivery
- Warranty period
  - During the warranty period, defects will be fixed without additional cost

Renewal

- When purchased goods are used for some time, they might need renewal
  - Hardware is often used for five years before it is replaced
  - Software typically has makor releases every few years
  - Service contracts are also often agreed upon for a fixed number of years
- Often a renewal of the hardware, software licenses, or service contracts leads to a new purchase process
- Systems management should have a Life Cycle Management (LCM) process implemented to

- A set of physical network switches can be programmed as a virtual network:
  - Hierarchical
  - Complex
  - Secured
- A virtual network can easily be changed without touching the physical network components

How does SDN work?

- In SDN (like anything virtualized), the software is decoupled from the hardware.
- SDN moves the control plane that determines where to send traffic to software, and leaves the data plane that actually forwards the traffic in the hardware.
- This allows network administrators who use software-defined networking to program and control the entire network via a single pane of glass instead of on a device by device basis.

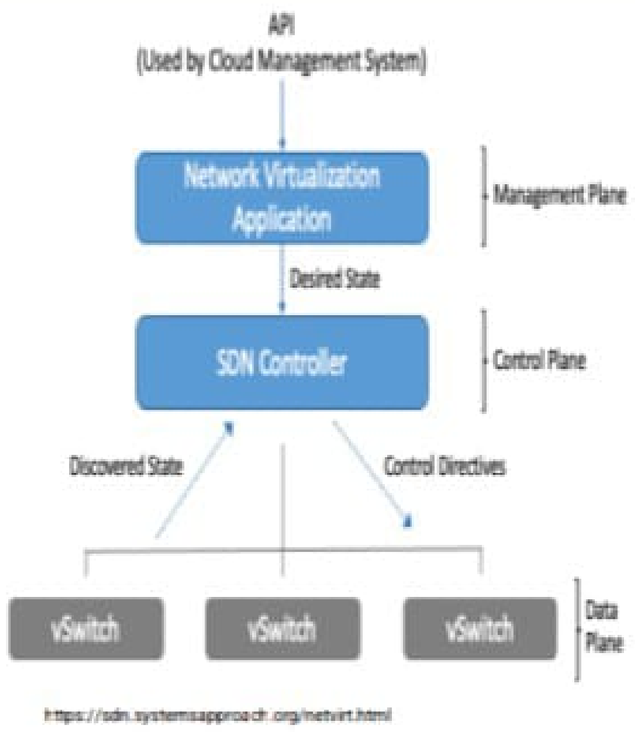

Software Defined Networking

- SDN can be controlled from a single management console
- Provides open APIs that can be used to manage the network using third party software
- In an SDN, the distributed data plane devices are only forwarding network packets based on ARP or routing rules that are preloaded into the devices by the SDN controller in the control plane
  - This allows the physical devices to be much simpler and more cost effective

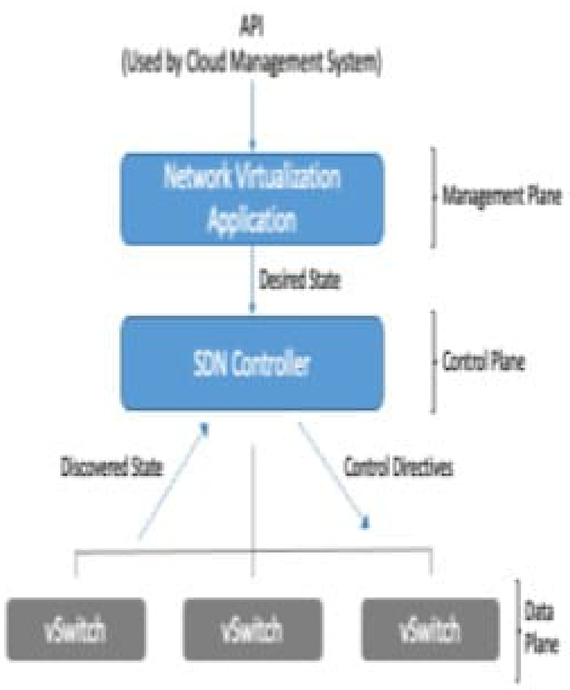

Software Defined Networking

- Control plane

– Renewal

Determine what is needed

- Before any purchase can be made, it must be crystal clear what is actually needed
- A Bill of Materials (BoM) is made that includes part numbers of all items
  - If allowed by the purch[illegible]g rules, work with suppliers to get the BoM
- Also a Statement of Work (SoW) is made
  - A SoW describes what the supplier will do, apart from delivering the goods
  - It must be clear from the start who does what
- The supplier can have specific requirements, like:
  - Is a loading dock available to deliver goods to the datacentre?
  - Is the elevator large enough to lift the equipment to the final destination?

Getting an offer

- In a large organization, the lead time for the internal procurement process can be several weeks, or even longer
  - Find a supplier
  - Handle contract issues and/or to get signatures from management to formally place the order
  - Procurement will try to get discounts
- It typically takes four to eight weeks for the supplier to deliver the goods

Choice of suppliers

- Most organizations use preferred suppliers for standard purchases
- Having a small number of preferred suppliers makes the purchase process easier
  - Contracts are already in place
  - Discounts can be negotiated because of larg[illegible] volume purchases
- Organizations often choose for a predefined

Ethics for IT infrastructure

Ethics f[redacted]ening the moral and societal i[redacted]f the [redacted] information technology systems. Here are some key ethical considerations for IT infrastructure:

- Privacy and Data Protection
- Transparency and Accountability
- Accessibility
- Security and Cybersecurity
- Sustainability and Environmental Impact
- Fairness and Bias Mitigation
- Intellectual Property
- Responsibility in AI and Automation
- Long-Term Planning and Scalability
- Human-Centered Design
- Ethical Decision-Making
- Compliance and Legal Adherence

Privacy and Data Protection

- Ensure that personal and sensitive data is collected, stored, and processed in a secure and transparent manner.
- Respect user privacy rights and comply with relevant data protection regulations such as GDPR (General Data Protection Regulation) or CCPA (California Consumer Privacy Act).
- Additional points:
  - Minimization of Data Collection
  - Informed Consent
  - Data Security
  - Data Retention and Deletion
  - Cross-Border Data Transfer
  - Data Breach Response
  - Data Protection Impact Assessments (DPIAs)
  - Data Governance
  - Data Transparency

Transparency and Accountability

- Transparency in IT infrastructure practices involves being open about policies, procedures, and data usage. It also includes taking responsibility for any errors or security breaches that occur and working to rectify them promptly.
- Key points to emphasize:
  - Clear Data Policies
  - Data Usage Notifications
  - Data Sharing and Third-Party Disclosures
  - Regular Transparency Reports

- DTAP:
  - Development
    - New software is developed or existing software is modified
  - Test
    - Software is tested by independent testers
  - Acceptance
    - Software is accepted by a delegation of the user population
  - Production
    - When all tests are successful, the software is deployed in the production environment
- DTAP environments are used in the software development process

DTAP environments

- DTAP suggests four environments, but often more environments are used
- A real-world DTAP environment could contain a:

Any three

  - Sandbox environment
    - Pre-development environment, where preliminary tests can be performed on new technology or solutions
  - Development environment
    - To develop new software and configurations
  - Test environment
    - Functionally test new software releases
  - User Acceptance test environment
    - Allow end users to functionally test new releases

DTAP environments

  - Non-Functional acceptance environment
    - Setup to be identical to the production environment to enable reliable performance, availability, and security testing
  - Hot Fix environment
    - Find fixes for production problems
    - Test fixes before they are deployed to production
  - Production environment
    - Runs the actual software for end users
  - A Systems Management environment is often used to manage the other environments

- Most infrastructure components generate log data
- Examples:
  - Network routers and switches
  - Operating systems
  - Applications
  - Databases
  - Firewalls
  - Intrusion detection systems
- Log data can be used to:
  - Correlate events
  - Identify sources of application issues
  - Identify trends to predict or even prevent unavailability
  - Find security vulnerabilities or security breaches

Capacity management

- Capacity Management ensures the timely availability of sufficient infrastructural capacity to process, transport, and store data now and in the future

Capacity management

- The following input is needed:
  - Monitoring of resources to detect trends
    - Reduced free disk capacity provides insight in when to purchase or free-up disk capacity
  - Business plans to anticipate on business changes that might have impact on the infrastructure
    - A marketing campaign during the summer time could justify temporary adding server capacity
  - Developments in technology
    - Upgrading servers when a higher capacity server blade becomes available

s
Manager
- BMC Patrol

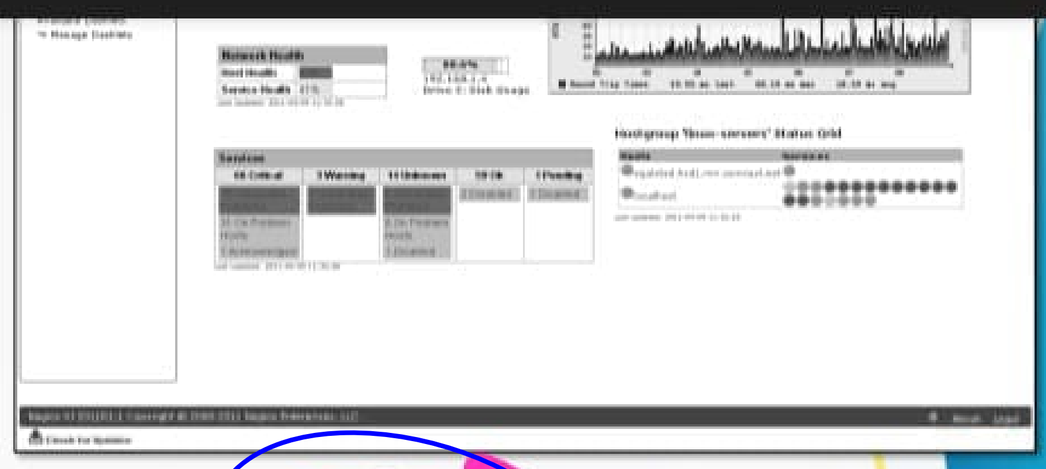

Management using SNMP

- Simple Network Management Protocol
- SNMP can be used to:
  - Remotely change or update configurations
  - Collect statistics and performance information
- Devices that support SNMP include:
  - Routers
  - Switches
  - Servers
  - Workstations
  - Printers

Management using SNMP

- SNMP uses a management/agent model
  - The agent runs on the monitored device
    - Has local knowledge of the system it resides on
    - Translates information to the SNMP protocol
  - A management server collects information from all agents
    - The Network Management System – NMS
    - Monitors and controls managed devices via the agents
    - Reading of values is done in regular polling intervals (like every 30 seconds)

Management using SNMP

- SNMP traps
  - A trap is an alarm that is sent to the NMS
  - A trap is sent immediately – no polling
- Security in SNMP is implemented using a

Capacity management

- The following input is needed:
  - Monitoring of resources to detect trends
    - Reduced free disk capacity provides insight in when to purchase or free-up disk capacity
  - Business plans to anticipate on business changes that might have impact on the infrastructure
    - A marketing campaign during the summer time could justify temporary adding server capacity
  - Developments in technology
    - Upgrading servers when a higher capacity server blade becomes available

Deploying applications

DTAP environments

- DTAP:
  - Development
    - New software is developed or existing software is modified

RAID Penalty

- In RAID sets multiple disks are used to form one virtual disk (LUN)
- Writing data on multiple disks introduces some delay, known as the RAID penalty

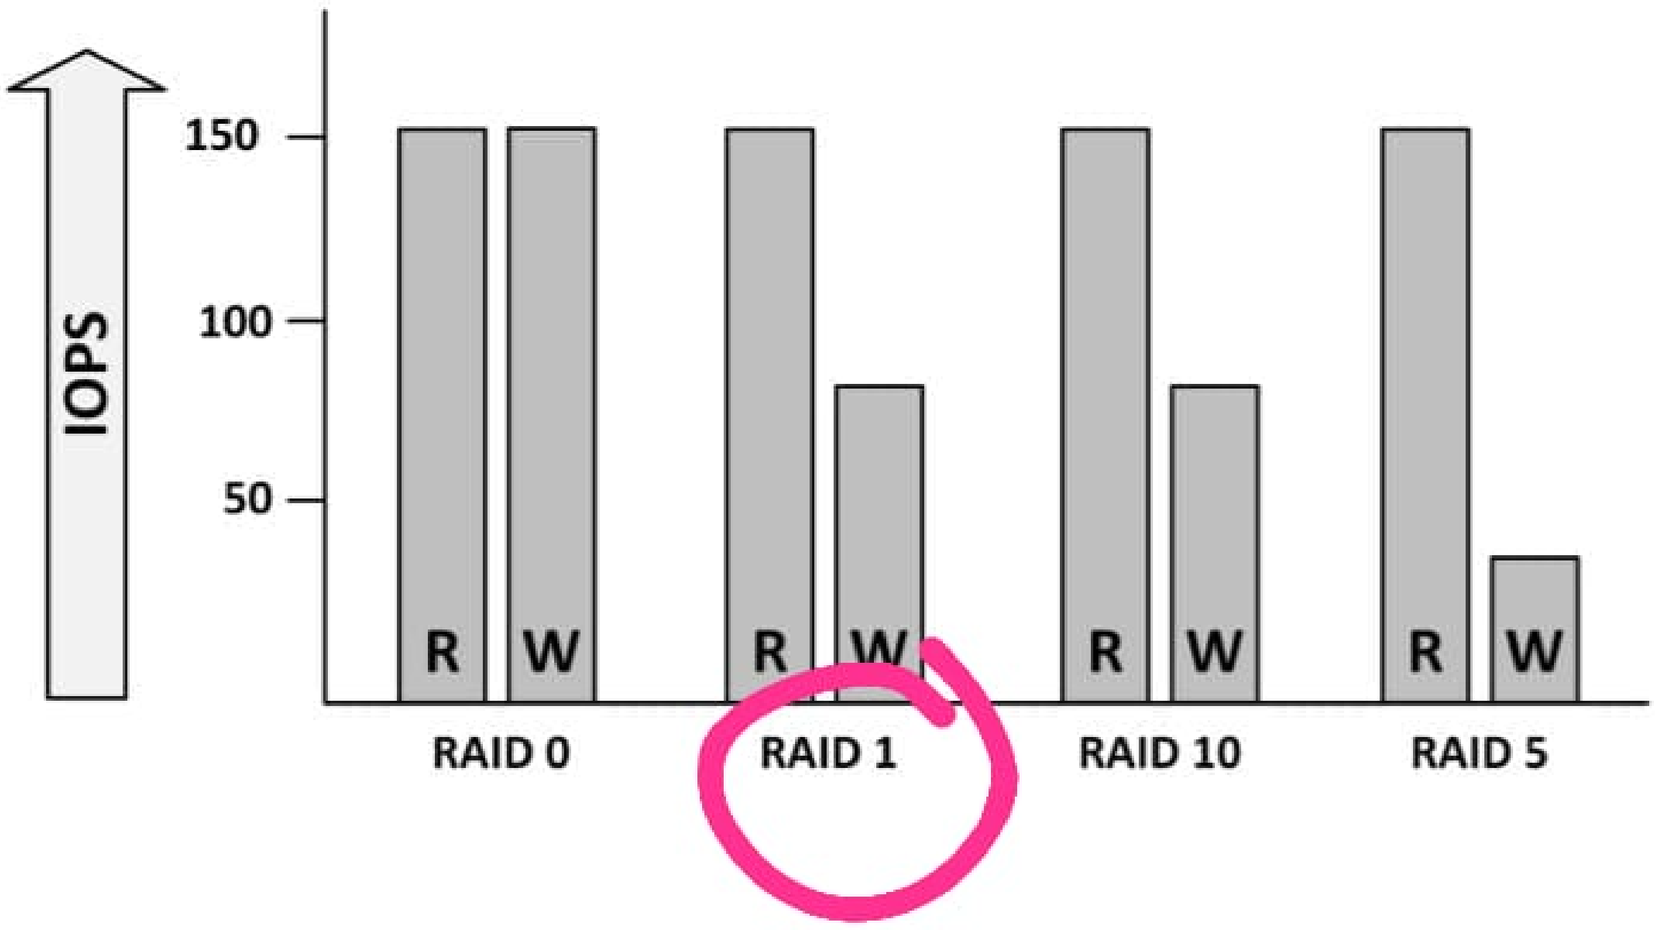

Interface Throughput

RAID

- Redundant Array of Independent Disks (RAID) solutions provide:
  - High availability of data
  - Improvements of performance
- RAID uses multiple redundant disks
- RAID can be implemented:
  - In the disk controller's hardware
  - As software running in a server's operating system

RAID

- RAID can be implemented in several configurations, called RAID levels
- In practice, five RAID levels are implemented most often:
  - RAID 0 - Striping
  - RAID 1 - Mirroring
  - RAID 10 - Striping and Mirroring
  - RAID 5 - Striping with distributed parity
  - RAID 6 - Striping with distributed double parity

Layered Network Topology

- A network infrastructure should be built up in layers
  - Improve availability and performance
  - Provides scalability
  - Provides deterministic routing
  - Avoids unmanaged ad-hoc data streams
- Provides high availability
  - Because the layering provides multiple paths to any piece of equipment

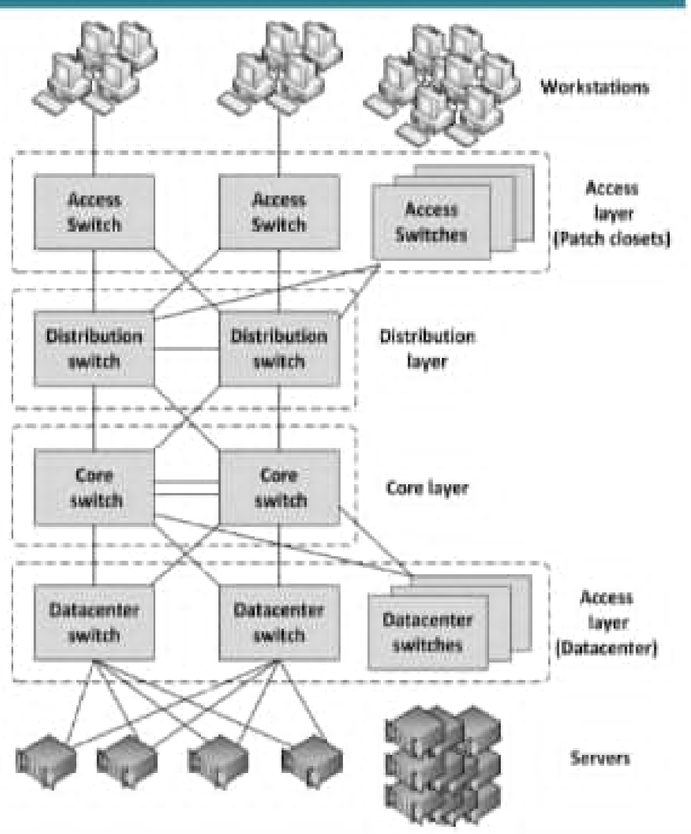

Layered Network Topology

- Core layer
  - This is the center of the network
- Distribution layer
  - An intermediate layer between the core layer in the datacenter and the access switches in the patch closets
  - Combines the access layer data and sends its combined data to one or two ports on the core switches
- Access layer
  - Connect workstations and servers to the distribution layer
  - For servers, located at the top of the individual server racks or in blade enclosures
  - For workstations, placed in patch closets in various parts of the building

Spine and Leaf Topology

- In a SDN, a simple physical network is used [illegible] can be programmed to act as a complex virtual network
- Such a network can be organized in a spine and leaf topology

Human-centered design (HCD)

Prioritizes the needs, experiences, and feedback of users and stakeholders.

- Key aspects of human-centered design:

– User Research
– User Personas
– Iterative Design
– Prototyping and Testing
– Accessibility and Inclusivity
– User-Friendly Interfaces
– User Engagement
– Empathy and Empowerment
– Human-Centered Metrics

Ethical Decision-making

- By fostering a culture that encourages employees and IT professionals to consider ethical implications and speak up about potential concerns, organizations can ensure that their technology aligns with ethical principles and values.

- Key Strategies:
– Ethics Training and Education
– Ethical Frameworks and Guidelines
– Whistleblower Protection

Compliance and Legal Adherence

- Compliance and legal adherence are critical aspects of IT infrastructure management, particularly concerning data protection, privacy, and IT security. Organizations are bound by various laws, regulations, and industry standards that dictate how they handle sensitive information, safeguard data, and ensure the security of their IT systems

2:16 0.77 KB/s 33%

IT602 FINAL TER...

Intellectual property

- Essential aspect of ethical IT infrastructure management. Intellectual property rights protect the creations and innovations of individuals and organizations, and it crucial to abide by the laws and regulations that govern these rights.
- Key points:
  - Software Licenses
  - Copyright Compliance
  - Open Source Software
  - Patent Respect
  - Trademark Protection
  - Intellectual Property Audit

AI and Automation

- Ensure that technologies are deployed ethically and do not have adverse effects on individuals, society, or employment.
  - Some key considerations:
    Ethical Design and Development
    Human-Centered Approach
    Privacy Protection

Long-term planning and Scalability

- Build a sustainable and efficient IT infrastructure that can adapt to future needs and growth.
- Key points to consider in long-term planning and scalability:
  - Scalability Assessment
  - Flexible Architecture
  - Cloud Solutions
  - Virtualization and Containerization
  - Automation and Orchestration
  - Future-Proof Technologies

- Accessible Forms
- Accessible Documentation
- Continuous Improvement

Security and Cybersecurity

- Fundamental aspects of ethical IT infrastructure management. Safeguarding data and infrastructure from cyber threats and unauthorized access is essential to protect user privacy, prevent data breaches, and maintain the trust of stakeholders.
- Key considerations for prioritizing security:
  - Risk Assessment
  - Data Encryption
  - Secure Authentication
  - Access Controls
  - Regular Security Audit
  - Patch Management
  - Incident Response Plan
  - Backup and Disaster Recovery

Sustainability and Environmental Impact

- Ethical IT infrastructure considers the environmental impact of technology. Organizations should adopt green IT practices, optimize energy consumption, and reduce electronic waste generation.
- Key practices to incorporate sustainability:
  - Energy Efficiency
  - Virtualization and Consolidation
  - Renewable Energy Sources
  - Green Data Centers
  - Energy Monitoring and Management
  - E-waste Management

Fairness and Bias Mitigation

- Critical ethical considerations in the development and implementation o
  infrastructure, especially when it involves automated systems and algor
- Key points to address fairness and bias: