# Infrastructure Security Using Linux

Computer science / Cybersecurity

## Cloud Computing

9

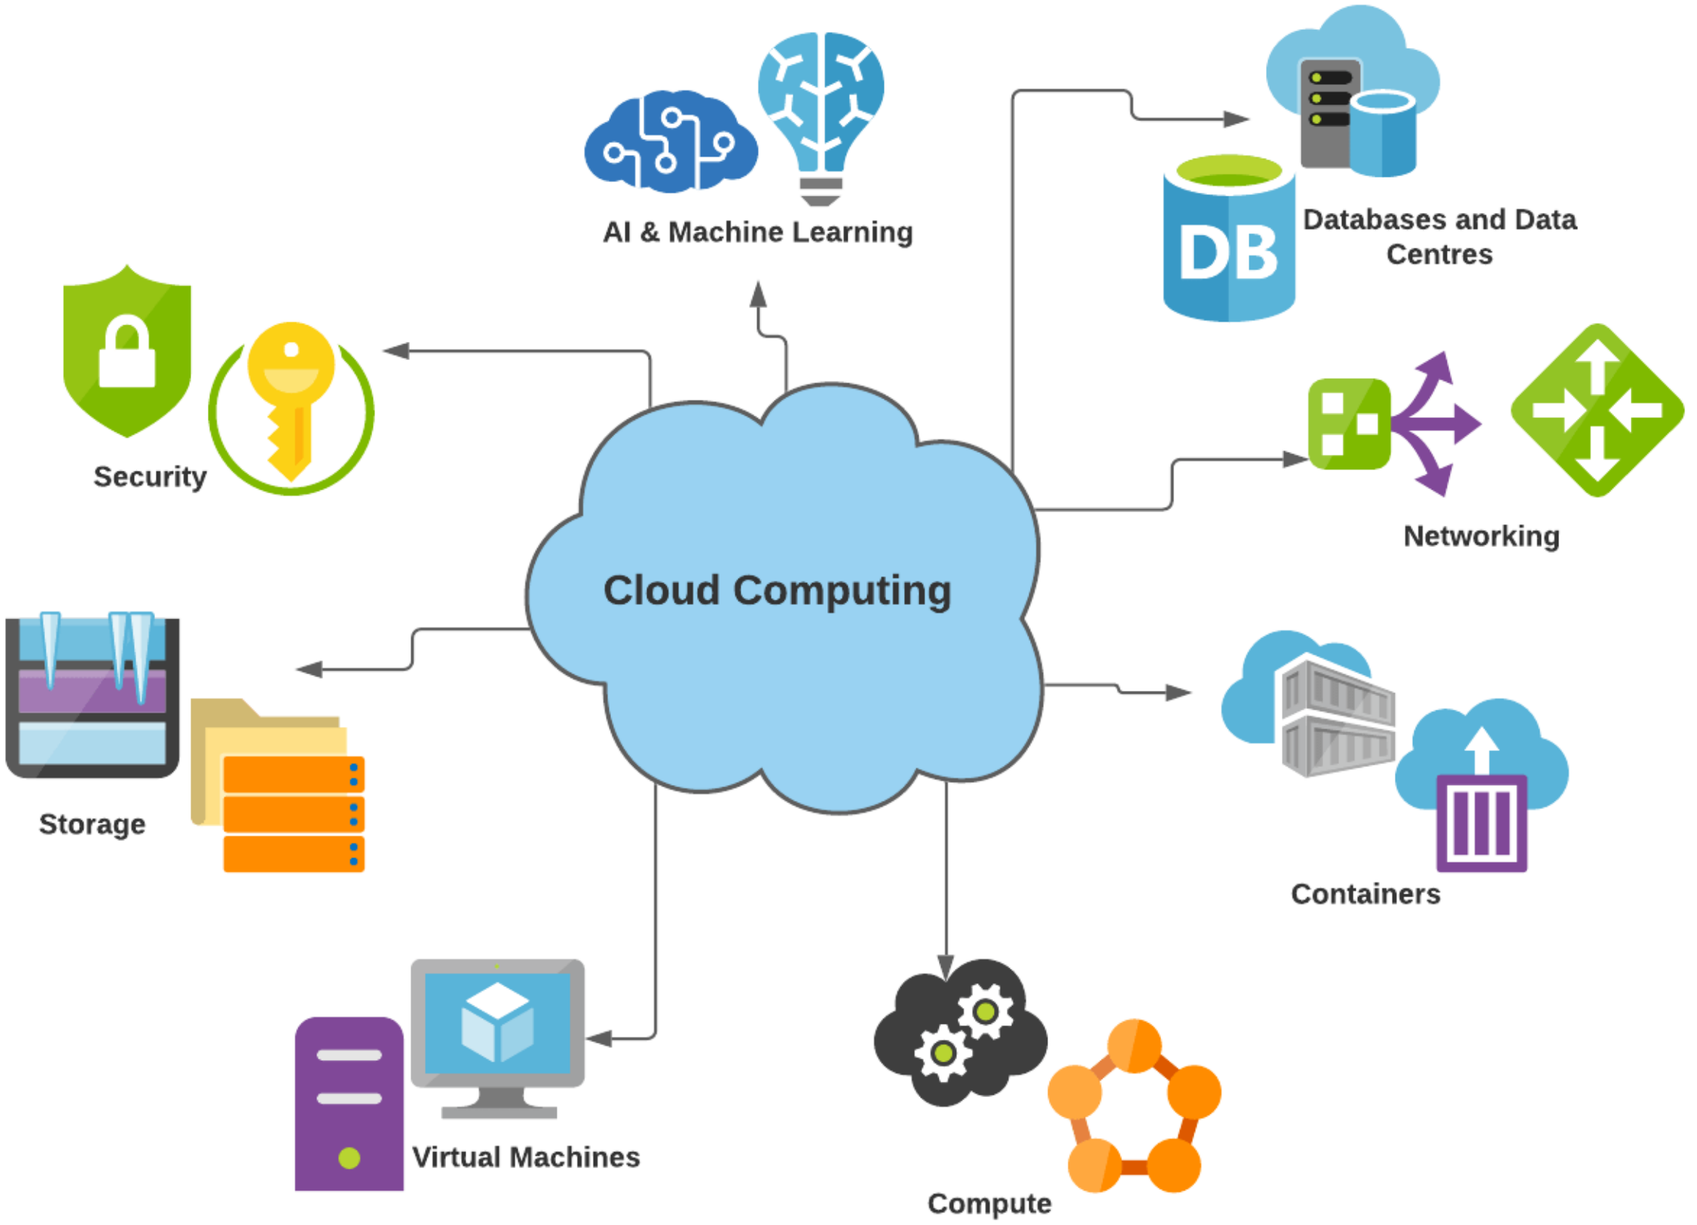
AI & Machine Learning
DB
Databases and Data Centres
Security
Networking
Cloud Computing
Storage
Containers
Virtual Machines
Compute

# Cloud Computing

- **Definition**: On-demand delivery of IT resources over the Internet with pay-as-you-go pricing.
- **How it Works**:
  - Eliminates the need to buy, own, and maintain physical infrastructure.
  - Offers services like computing power, storage, and databases.
- **Providers**:
  - Examples include ***AWS***, ***Microsoft Azure***, and ***Google Cloud***.

# The cloud in context

- **Benefits of Cloud Providers**
  - Offer advanced infrastructure that surpasses typical business capabilities.
  - Lower costs for running distributed computing services compared to on-premises data centers.
- **Programmatic Management**
  - Resources can be created and destroyed programmatically.
  - Allows for automation, rapid scaling, and bypasses slow procurement processes.
- **A New Mindset**
  - Automation replaces manual care:
    - **Cattle, not pets**: Servers are replaceable units, not individually managed assets.
  - This shift requires adapting to a model where hardware is abstracted and disposable.
- **Caveats and Planning**
  - The cloud is not an instant solution for cost reduction or performance improvement.
  - “Lift and shift” migrations need careful planning to avoid failure.

# Cloud platform choices

**Table 9.2: The most widely used cloud platforms**

| Provider | Notable qualities |
|---|---|
| Amazon Web Services | 900lb gorilla. Rapid innovation. Can be expensive. Complex. |
| DigitalOcean | Simple and reliable. Lovable API. Good for development. |
| Google Cloud Platform | Technically sophisticated and improving quickly. Emphasizes performance. Comprehensive big-data services. |
| IBM Softlayer | More like hosting than cloud. Has a global private network. |
| Microsoft Azure | A distant second in size. Has a history of outages. Possibly worth consideration for Microsoft shops. |
| OpenStack | Modular DIY open source platform for building private clouds. AWS-compatible APIs. |
| Rackspace | Public and private clouds running OpenStack. Offers managed services for AWS and Azure. Fanatical support. |
| VMware vCloud Air | Buzzword-laden service for public, private, and hybrid clouds. Uses VMware technology. Probably doomed. |

# Public, private, and hybrid clouds

- **Public Cloud**
  - Offered by third-party providers (e.g., ***AWS***, ***Google Cloud***, ***Microsoft Azure***).
  - Resources (servers, storage, etc.) delivered via the Internet.
  - All infrastructure is owned and managed by the provider.
  - Services are accessed and accounts are managed through a web browser.
- **Private Cloud**
  - Exclusively used by a single organization.
  - May be hosted on-premises or by a third-party service provider.
  - Operates on a private network, with dedicated services and infrastructure.

# Public, private, and hybrid clouds

- **OpenStack**
  - Leading open-source platform for creating private clouds.
  - Backed by major companies like ***IBM***, ***Red Hat***, and ***Rackspace***.
- **Hybrid Cloud**
  - Combines public and private clouds.
  - Commonly used for:
    - Transitioning from local servers to public cloud.
    - Adding temporary capacity for peak loads.
    - Other organization-specific needs.
  - **Complexity Warning**: Managing two separate cloud environments can increase administrative challenges significantly.

# Cloud service fundamentals (main Categories)

- **IaaS (Infrastructure as a Service)**:
  - Provides raw compute, memory, network, and storage resources.
  - Delivered as virtual private servers (VPSs).
  - Users manage everything above the hardware layer, including OS, networking, and storage.
- **PaaS (Platform as a Service)**:
  - Developers supply their custom applications in a vendor-specified format.
  - The vendor runs the code on behalf of the user.
  - Users are responsible only for their code, while the vendor handles the OS and network.
- **SaaS (Software as a Service)**:
  - Vendor hosts and manages software.
  - Users pay a subscription fee for access.
  - Users manage neither the OS nor the application (e.g., hosted web apps like WordPress).


On-site
IaaS
PaaS
SaaS
Applications
Data
Runtime
Middleware
O/S
Virtualization
Servers
Storage
Networking
You manage
Service provider manages

# Regions and availability zones

- **Regions and Availability Zones**
  - **Region**:
    - A geographic location with cloud provider data centers.
    - Often named after the intended service territory (e.g., ***us-east-1*** in northern Virginia).
  - **Availability Zone**:
    - Collection of data centers within a region.
    - Connected by high-bandwidth, low-latency circuits for fast inter-zone communication.
    - Designed to be isolated from one another, ensuring that failure in one zone does not affect others.
- **High Availability and Multiregion Deployments**
  - Deploying across multiple zones and regions reduces the impact of localized failures.
  - Multiregion setups introduce complexity due to increased physical distance and latency.

# ** Virtual Private Servers (VPS)

- **Virtual Private Servers (VPS)**
  - Core cloud offering, also known as ***instances***.
  - Virtual machines running on the cloud provider’s hardware.
- **Instance Creation**
  - Based on ***images***, which are saved operating system states.
  - **Images** typically include:
    - A root filesystem.
    - A boot loader.
    - Optionally, additional disk volumes and custom settings.

# Networking

- **Virtual Networks**
  - Cloud providers offer virtual networks with customizable topologies.
  - Isolation options:
    - Keep systems separate from one another.
    - Keep systems hidden from the public internet.
- **Internet Accessibility**
  - Public addresses (e.g., ***Elastic IPs*** on AWS):
    - Leased from a provider's pool.
    - Directly accessible from the Internet.
  - Private RFC1918 addresses:
    - Restricted to your selected private network.
    - Not accessible from the Internet.

- **Accessing Private Systems**

  Use a ***jump server*** or ***bastion host*** open to the Internet.

  Connect through a ***VPN*** to your cloud network.

  Security advice: Minimize external-facing components.

# Storage

- The cloud vendors bill by the amount of data you store. They are highly motivated to give you as many ways as possible to ingest your data.
- Here are a few of the most important ways to store data in the cloud:
- **Object stores** contain collection of discrete objects (files, essentially) in a flat namespace. Object stores can accommodate a virtually unlimited amount of data with exceptionally high reliability but relatively slow performance. Examples include Amazon S3, Google Cloud Storage, and Azure Blob Storage.
- **Block storage** devices are virtualized hard disks that can be attached to instances. They are faster than object stores but are limited in size and are more expensive. Examples include Amazon EBS, Google Persistent Disk, and Azure Disk Storage.
- **Ephemeral storage** is local disk space on a VPS that is created from disk drives on the host server. Ephemeral storage is fast but is lost when the instance is terminated. It is useful for temporary files or caches. Examples include instance store on AWS and local SSDs on Google Cloud.

# Identity and authorization

- AWS is exceptionally strong in this area. Their service, called Identity and Access Management (IAM), defines not only users and groups but also roles for systems. A server can be assigned policies, for example, to allow its software to start and stop other servers, store and retrieve data in an object store, or interact with queues—all with automatic key rotation. IAM also has an API for key management to help you store secrets safely.
- Other cloud platforms have fewer authorization features. Unsurprisingly, Azure's service is based on Microsoft's Active Directory. It pairs well with sites that have an existing directory to integrate with. Google's access control service, also called IAM, is relatively coarse-grained and incomplete in comparison with Amazon's.

# Clouds: vps

- **Amazon Web Services**
  - By default, EC2 instances in VPC subnets do not have public IP addresses attached, rendering them accessible only from other systems within the same VPC.
  - Firewalls in EC2 are known as "security groups." If you don't specify a security group, AWS will assume the "default" group, which allows no access. To connect to the instance, adjust the security group to permit SSH from your IP address.
- **GCP**
  - **gcloud**(the CLI) initializes the instance with a public and private IP address. You can use the public IP with SSH, but gcloud has a helpful wrapper to simplify SSH logins.