Cloud Networking

ITN 4252 Advanced Topics in Networking

Partially based on Cloud Computing and Cloud Networking by Adel Nadjaran Toosi, Introduction to Cloud Computing by Indranil Gupta, and resources from AWS.

The Cloud ~~Hype~~ Reality

- Forrester’s Public Cloud Market Outlook (2022-2026):
  - The public cloud market is expected to surpass **$1 trillion by 2026**.
- Fortune Business Insights:
  - The cloud computing market is projected to grow from $677.95 billion in 2023 to **$2,432.87 billion by 2030**.
- Gartner:
  - Over **75% of governments** will operate more than half of their workloads using hyperscale cloud service providers by 2025.
- Market Trends:
  - Global uncertainty, data privacy concerns, and potential government overreach are driving greater demand for **sovereign clouds**.

Many many clouds

- Amazon Web Services
- Microsoft Azure
- Google Cloud Platform
- Oracle Cloud
- Alibaba Cloud, Dell EMC, Gigaspaces, Salesforce, DataStax, MongoDB, VMWare, Cloudera
- And many more…

But what exactly is a cloud?

What is a Cloud?

- It’s a cluster!
- It’s a supercomputer!
- It’s a datastore!
- It’s Superman!

- None of the above.
- All of the above.

“It was much nicer before people started storing all their personal information in the cloud.”

- Cloud = Lots of storage + compute cycles nearby

What is cloud computing?

- A technology model providing quick and easy access to shared, configurable system resources.
- Resources can be swiftly configured with minimal management, often via the internet.
- Companies can hire external cloud providers to manage their IT needs.
- Avoids substantial initial costs of setting up in-house IT infrastructure.

What is a Cloud?

- A single-site cloud (aka “Datacenter”) consists of
  - Compute nodes (grouped into racks)
  - Switches, connecting the racks
  - A network topology, e.g., hierarchical
  - Storage (backend) nodes connected to the network
  - Front-end for submitting jobs and receiving client requests
  - Software Services
- A geographically distributed cloud consists of
  - Multiple such sites
  - Each site perhaps with a different structure and services

Timeline to the cloud

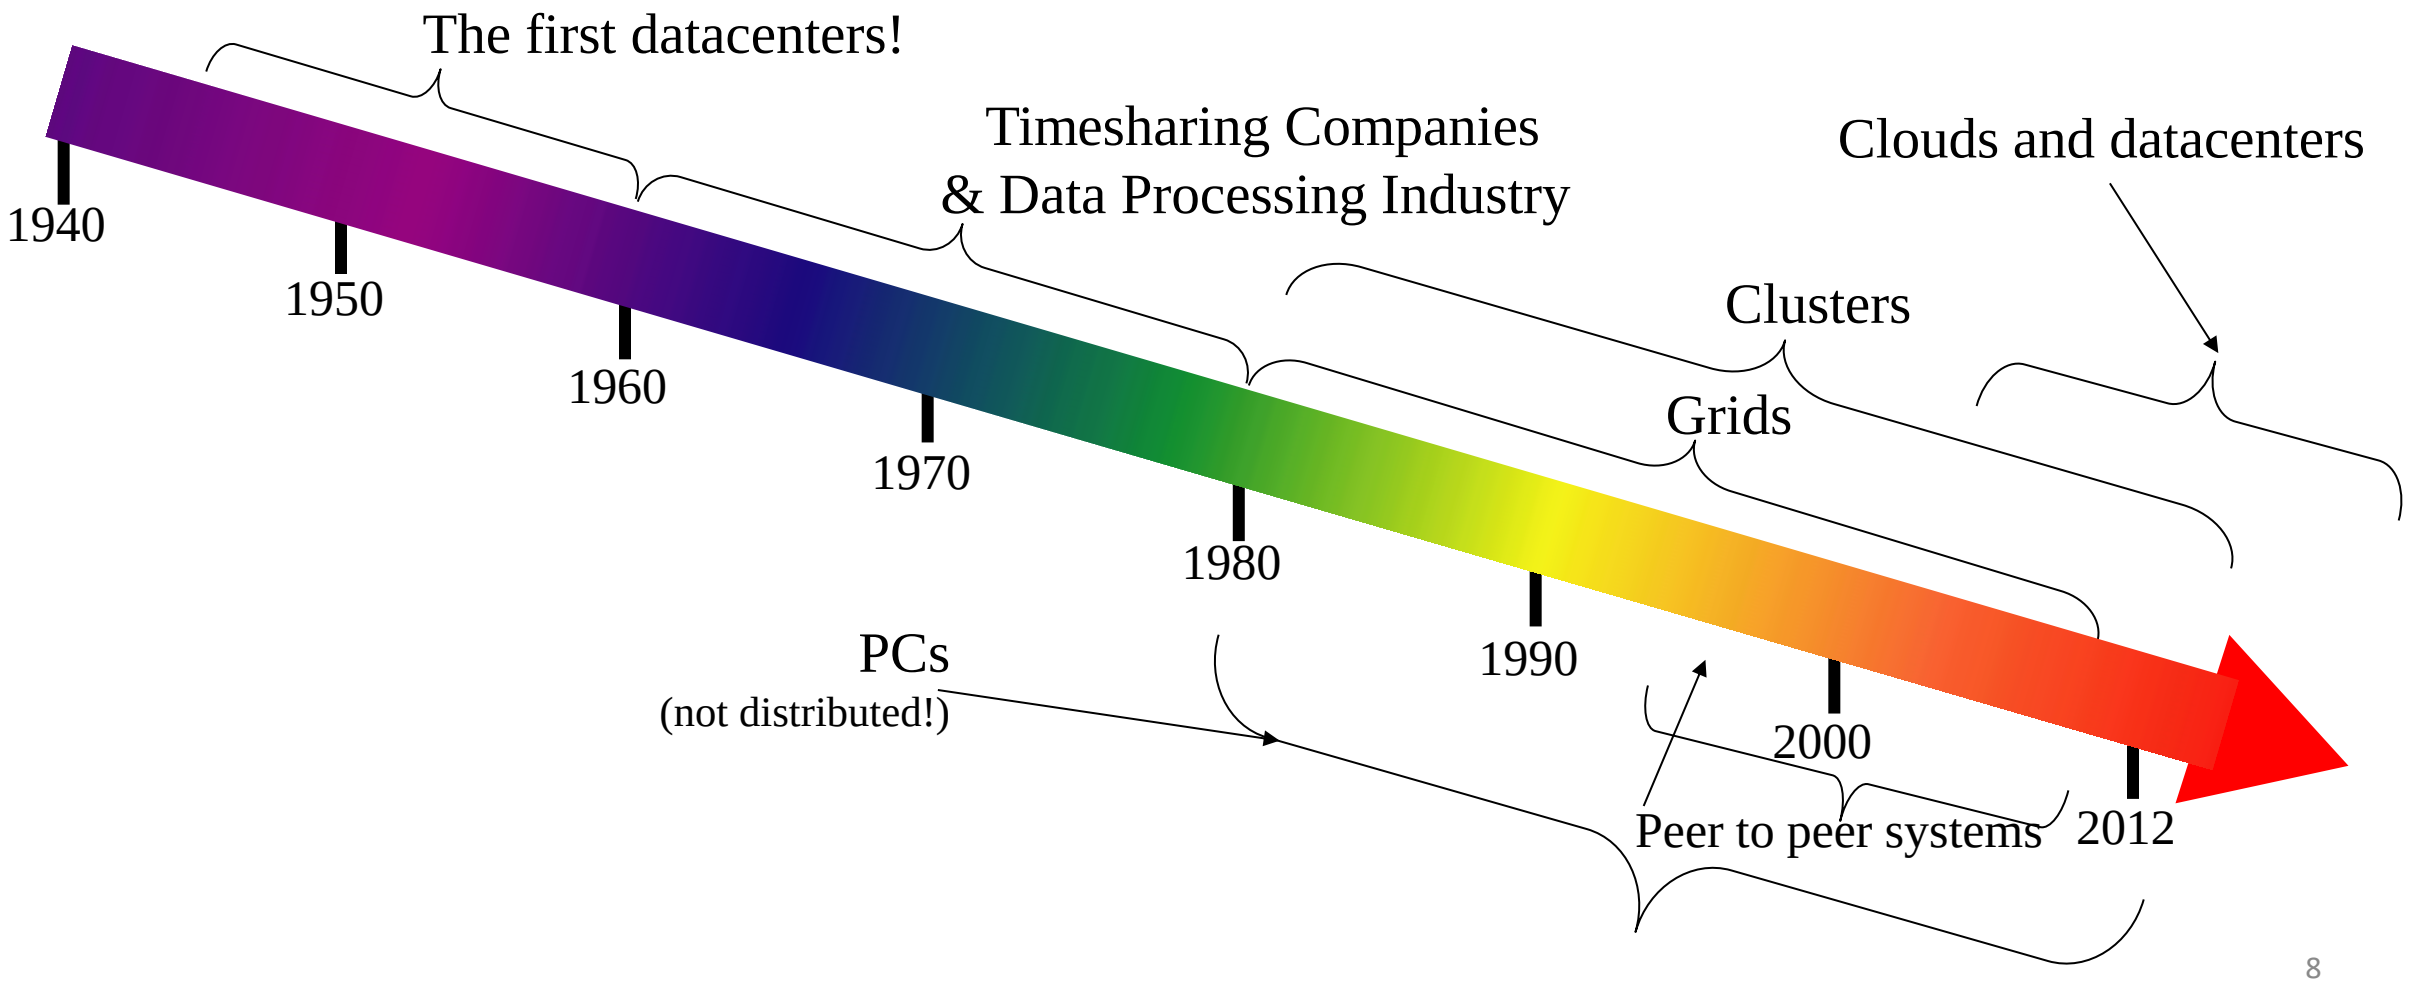

Why Cloud Computing?

Traditional On-Premises Resources

- Spend money upfront to purchase hardware
- Wait for server delivery
- Install servers in your physical data centre
- Make all necessary configurations

Cloud Compute Instances

- Use virtual servers to run applications in the cloud
- Provision and launch instances within minutes
- Stop using instances when workloads are finished
- Pay only for the compute time used, not when instances are stopped or terminated
- Save costs by paying only for the server capacity needed

Cloud Compute Services

- Provides secure, resizable compute capacity in the cloud as virtual instances.

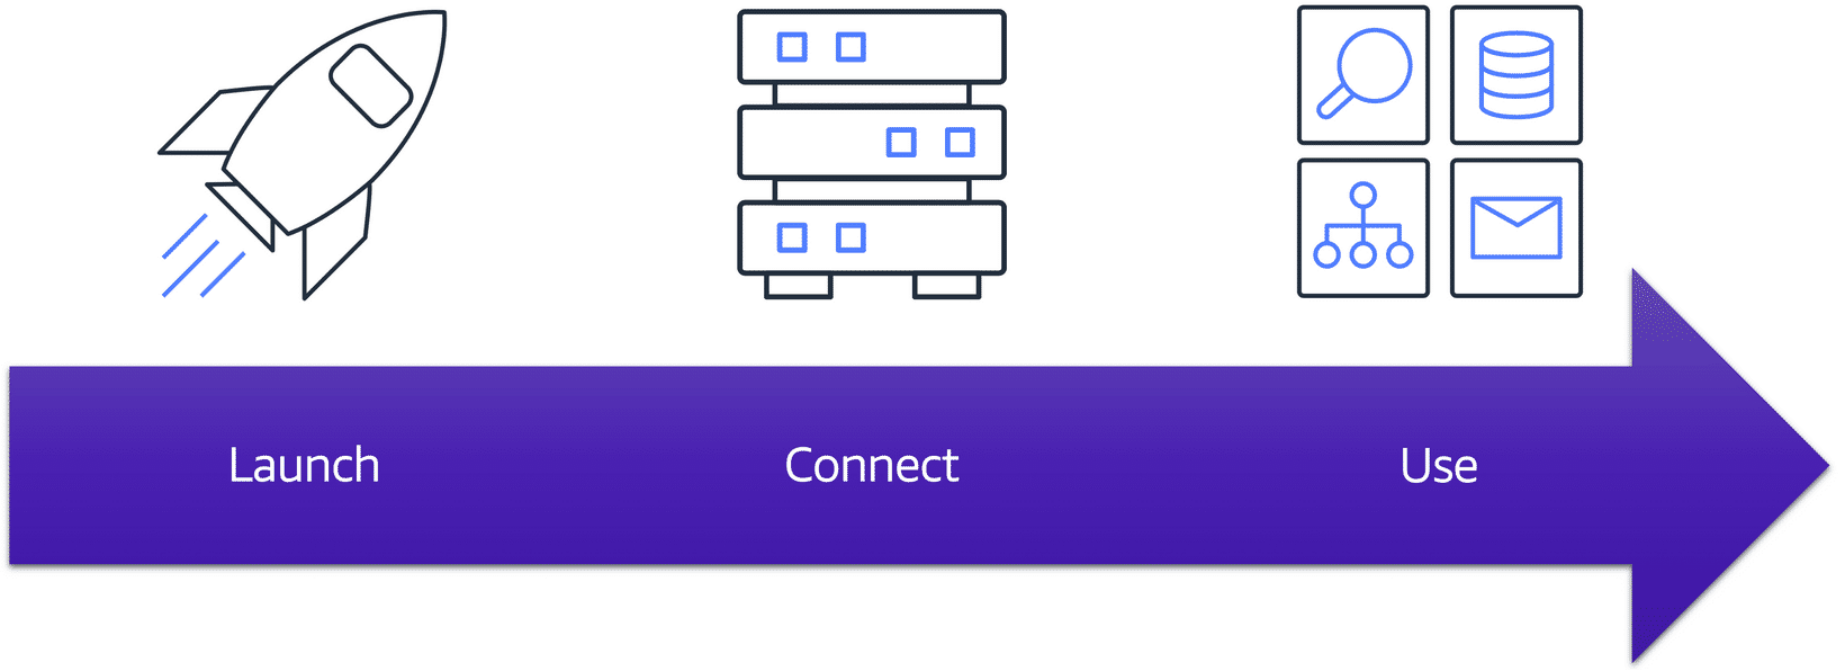

Cloud Deployment Models

- **Public**/Cloud-based deployment/Internet Clouds
  - Owned and operated by third-party service providers
  - Accessed via a web browser
  - Common uses: web-based email, online office applications, storage, test environments
- **Private**/On-premises deployment/Enterprise Clouds
  - Exclusively used by a single business or organisation
  - Provides more control over data, security, and quality of service
  - Can be located on-site or hosted by a third-party service provider
- **Hybrid** deployment/Mixed Clouds
  - Combines public and private clouds
  - Allows data and applications to be shared between them
  - Offers greater flexibility and more deployment options

Public/Cloud-based deployment

- Run all parts of the application in the cloud
- Deployment options:
  - Migrate existing applications
  - Design and build new applications
- Infrastructure options:
  - Low-level infrastructure managed by IT staff
  - Higher-level services reducing management, architecting, and scaling requirements
- Example:
  - An application with virtual servers, databases, and networking components fully based in the cloud

Private/On-premises deployment

- Resources deployed on-premises using virtualisation and resource management tools
- Increase resource utilisation with application management and virtualisation technologies
- Example:
  - Applications run on technology fully kept in your on-premises data centre

Hybrid deployment

- Cloud-based resources connected to on-premises infrastructure
- Integrate cloud-based resources with legacy IT applications
- Complies with regulations requiring certain records to be kept on premises
- Example:
  - Use cloud services for batch data processing and analytics
  - Keep legacy applications on premises while benefiting from cloud-based data and analytics services

Cloud Service Models

- Infrastructure as a Service (IaaS)
  - Choose virtual machines, operating system, memory, cores, and storage
  - Install and configure software as if it were a new server
  - No need to worry about server placement, air conditioning, or hardware maintenance
- Platform as a Service (PaaS)
  - Develop an app and submit the code to the cloud for deployment
  - No need to configure servers like Apache, Tomcat, Memcached, etc.
  - Infrastructure scales automatically if your app becomes popular
- Software as a Service (SaaS)
  - Use the application online
  - No need to buy a license, install, configure, or update the apps

Cloud Service Models

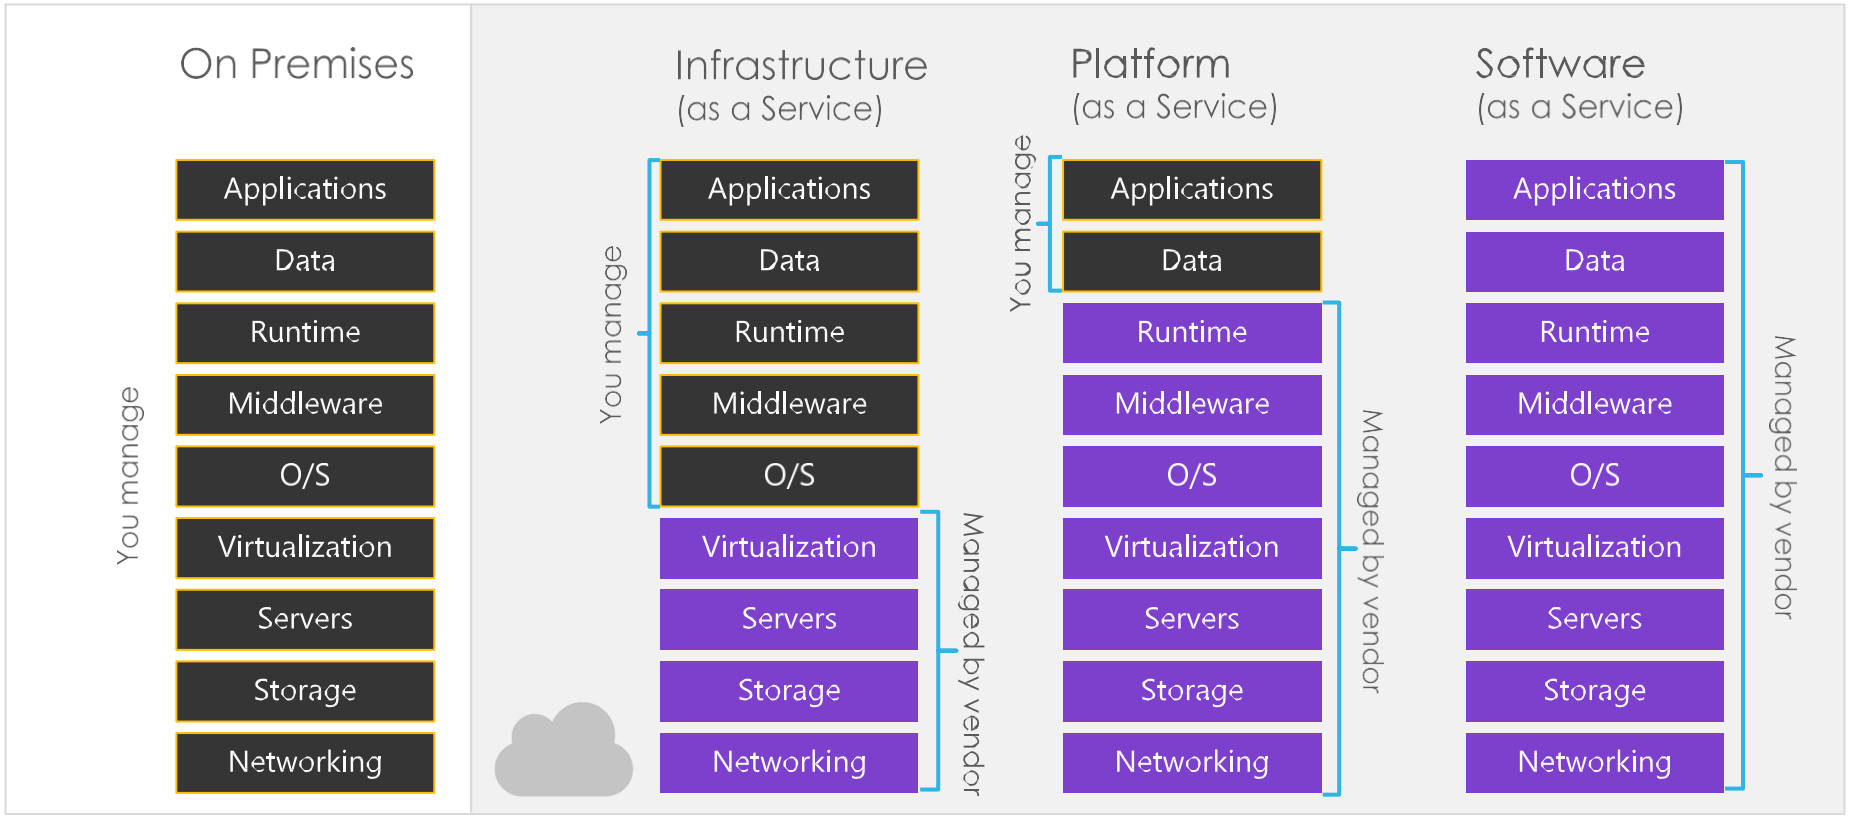

Four Features of Today's Clouds

- Massive Scale:
  - Ability to handle vast amounts of data and numerous simultaneous users
- On-Demand Access:
  - Pay-as-you-go, no upfront commitment, accessible to anyone
- Data-Intensive Nature:
  - From MBs to TBs, PBs, and XBs; daily logs, forensics, web data
- New Cloud Programming Paradigms:
  - MapReduce/Hadoop, NoSQL/Cassandra/MongoDB; high accessibility, ease of programmability, lots of open-source

Benefits of Cloud Computing

- Cost Efficiency:
  - Pay only for what you use, no upfront investment in data centres or servers
  - Lower costs due to economies of scale from providers
- Flexibility and Scalability:
  - Easily scale resources up or down based on demand
  - Quick access to new resources, enabling faster development and deployment
- Focus on Core Business:
  - Less time and money spent on managing infrastructure
  - More focus on applications and customers
- Global Reach:
  - Deploy applications globally with low latency
  - Serve customers worldwide efficiently

Global Infrastructure

- Cloud computing ensures efficient service to customers worldwide through a network of geographically distributed data centres.

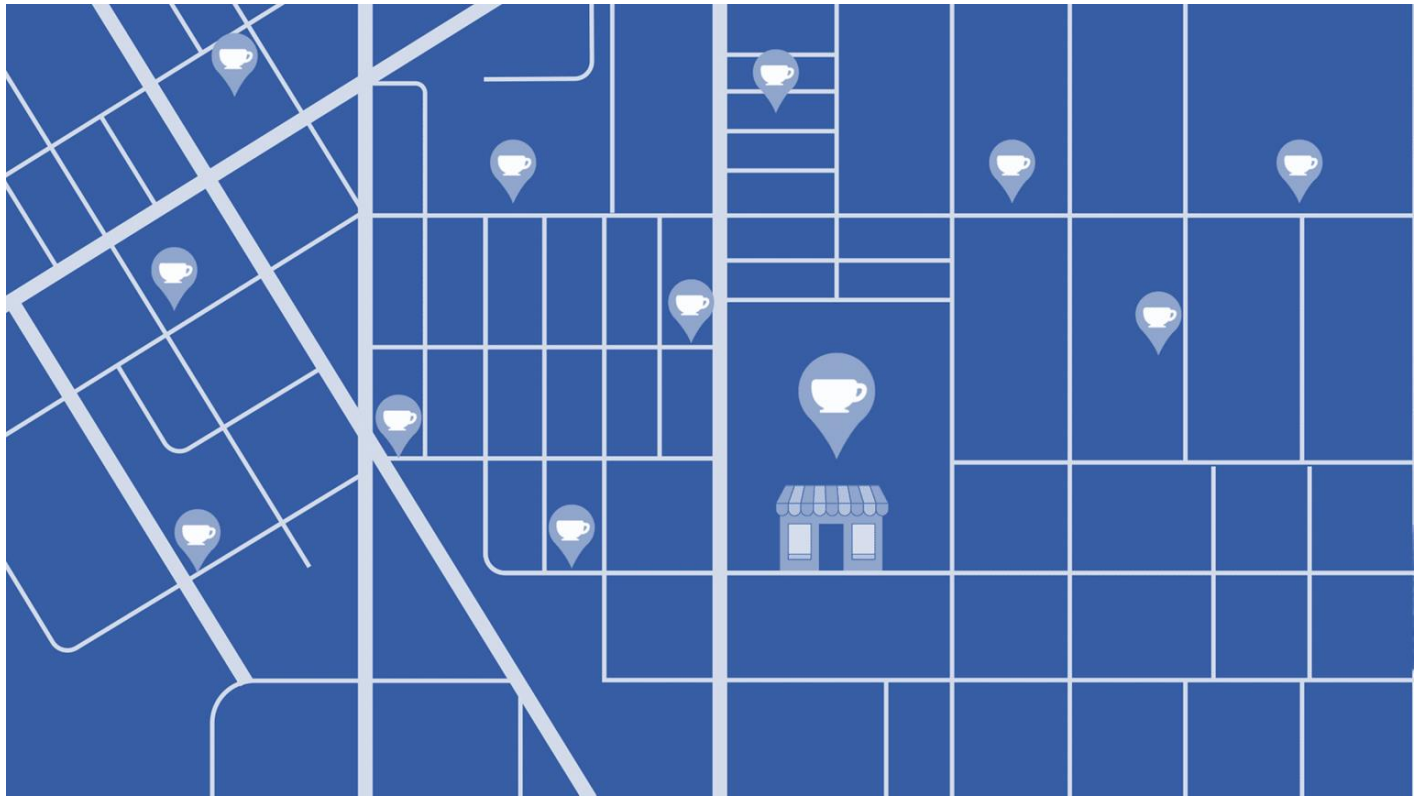

Cloud Networking

Cloud networking enables the global deployment of applications with low latency, ensuring efficient service to customers worldwide.

What is Cloud Networking?

- Connectivity to and between on-premises, edge, and cloud-based services (IaaS, PaaS, SaaS).
- Managed and maintained by a third-party cloud service provider.
- Includes virtual routers, firewalls, and network management software.
- Software-Defined Networking (SDN):
  - Centralises command and control in a master device.
- Network Functions Virtualization (NFV):
  - Virtualises physical networking devices and scales out across devices cost-effectively.
- SDN and NFV enable network cloudification.
  - SDN and NFV can be used separately
  - Combined use of SDN and NFV in cloud computing results in automated provisioning and centralised command and control.

Benefits of Cloud Networking

- Scalability:
  - Allows for quick and easy deployment and decommission of IT services.
- User Experience Assurance:
  - Provides real-time responsiveness to traffic demands.
- Programmable Network:
  - Uses software applications and programming interfaces to manage and control network traffic.
  - Offers improved automation, flexibility, and agility.
- Efficiency:
  - Streamlines and cost-efficient network management.
- Lowers operational costs
  - Eliminates high capital costs, maintenance costs, and regular hardware upgrades.
- Security:
  - Utilises cutting-edge infrastructure and highly secure physical network components due to economies of scale.

Virtual Private Cloud (VPC)

- A private network in a cloud environment, where you define your private IP range for resources.
- Can be public-facing (internet access) or private (no internet access).
- Public and private groupings of resources, defined by IP address ranges, subnets, in your VPC.
- Example:
  - Public Subnet: Public facing website for customers (internet access).
  - Private Subnet: Database server for customer data (no direct customer interaction).

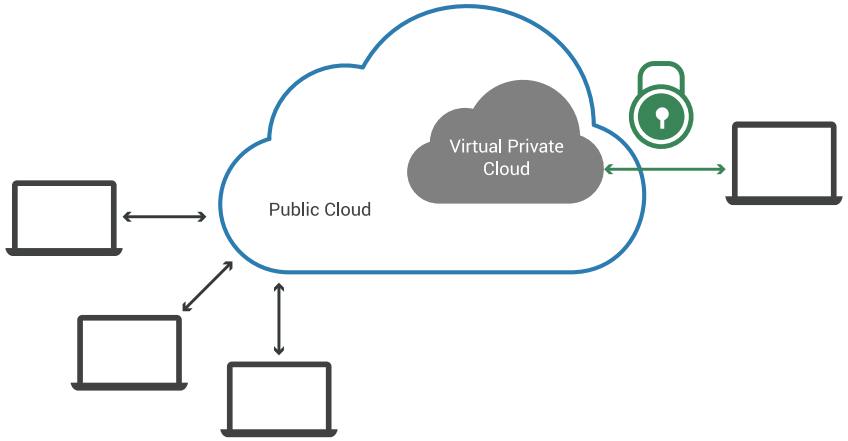

Subnets

- A subnet is a section of a VPC where you can group resources based on security or operational needs.
- Subnets can be public or private.
- **Public subnets** contain resources that need to be accessible by the public, such as an online store's website.
- **Private subnets** contain resources that should be accessible only through your private network, such as a database with customers' personal information and order histories.
- In a VPC, subnets can communicate with each other.
- For example, an application might involve instances in a public subnet communicating with databases in a private subnet.

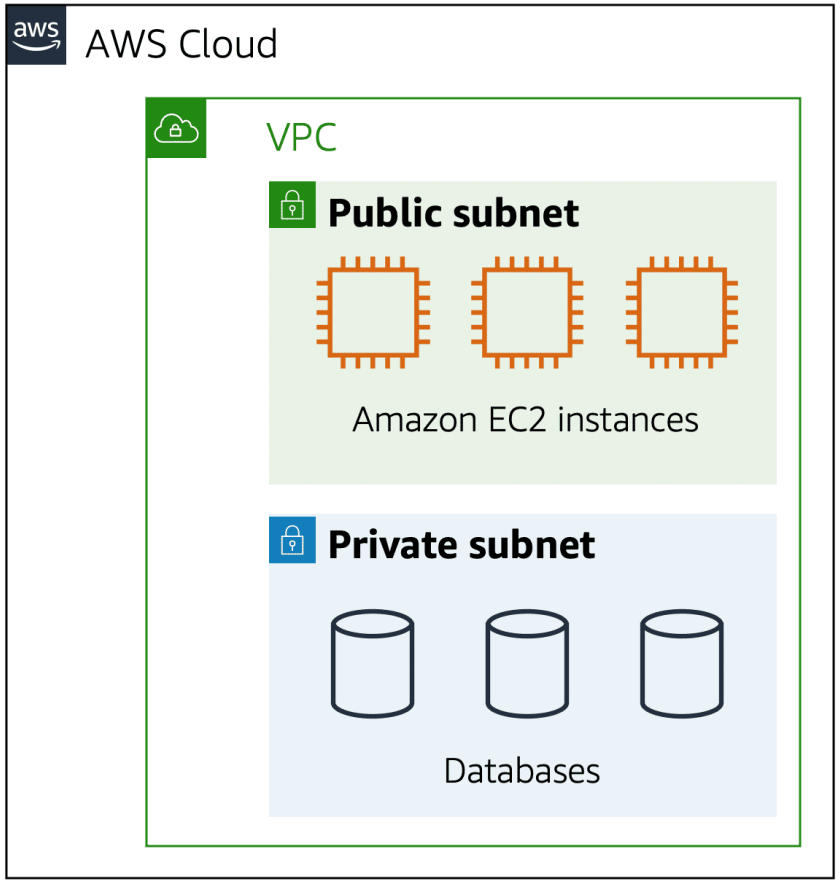

Virtual Private Cloud (VPC)

- Resources like instances and load balancers are placed inside the VPC and grouped into subnets.
- An internet gateway (IGW) allows public internet access to the resources.
- A virtual private gateway enables private network access and VPN connections.
- One VPC can have multiple gateways for different types of resources.

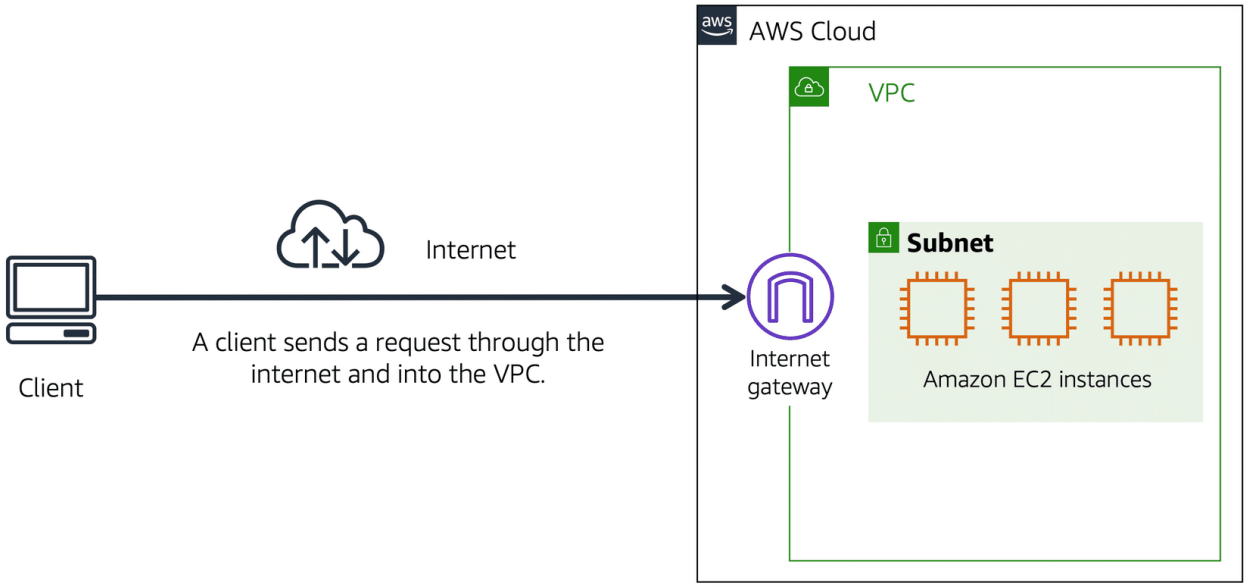

Virtual Private Gateway

- Use a virtual private gateway to access private resources in a VPC.
- A VPN connection encrypts your internet traffic from other requests.
- The virtual private gateway allows protected internet traffic to enter the VPC.
- It enables a VPN connection between your VPC and a private network, such as an on-premises data centre or internal corporate network.
- A virtual private gateway allows traffic into the VPC only if it is coming from an approved network.

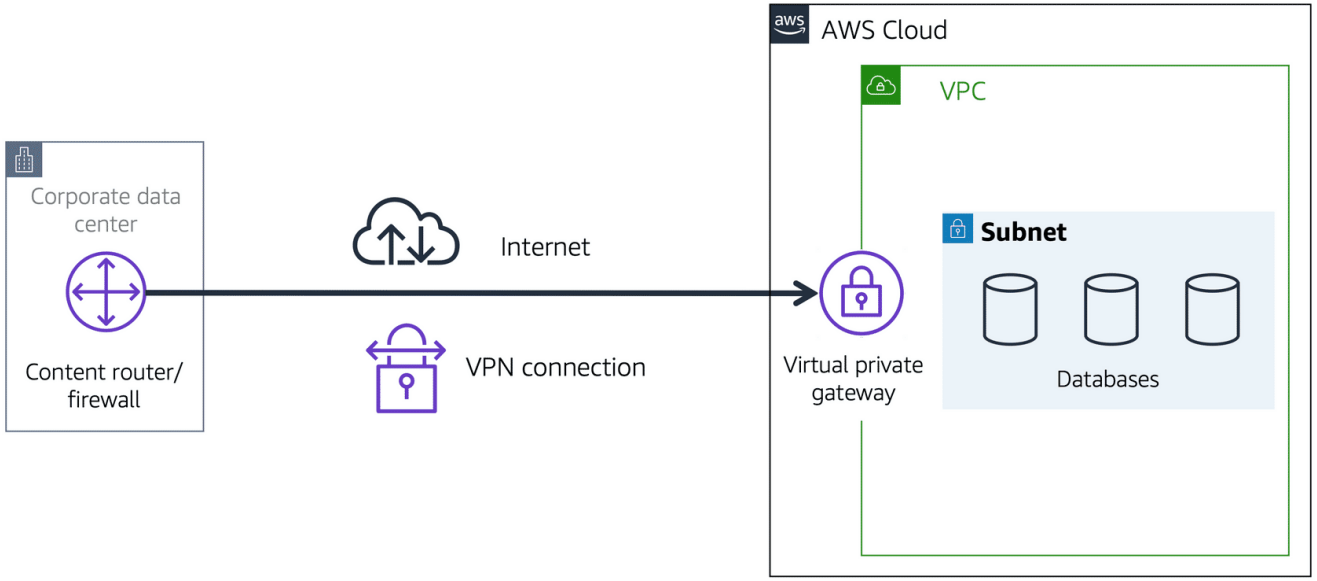

Network Traffic in a VPC *

- When a customer requests data from an application hosted in the cloud, the request is sent as a packet.
- It enters a VPC through an internet gateway.
- Before a packet can enter or exit a subnet, it checks for permissions.
- Permissions indicate who sent the packet and how it is trying to communicate with resources in a subnet.
- A network access control list (ACL) is the VPC component that checks packet permissions for subnets.

(in AWS)*

Network ACLs

- A network ACL is a virtual firewall that controls inbound and outbound traffic at the subnet level.
- Each cloud account includes a default network ACL.
- By default, the network ACL allows all inbound and outbound traffic, but you can modify it by adding your own rules.
- All network ACLs have an explicit deny rule to ensure that any packet not matching other rules is denied.
- For custom network ACLs, all inbound and outbound traffic is denied until you add rules to specify which traffic to allow.

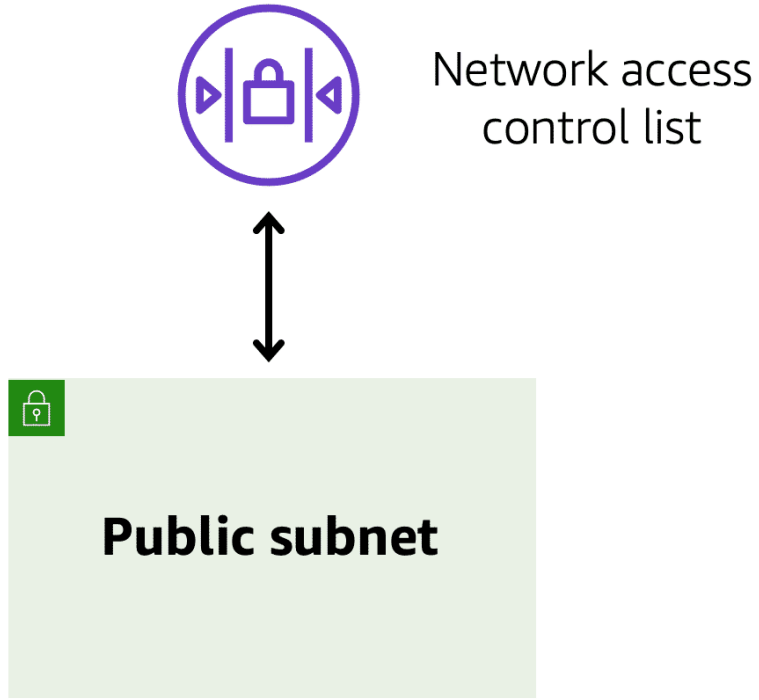

Stateless Packet Filtering

- Network ACLs perform stateless packet filtering, checking packets that cross the subnet border both inbound and outbound.
- They do not remember previous requests and check each packet against their list of rules.
- When a packet response returns to the subnet, the network ACL checks it again to determine whether to allow or deny it.
- After a packet enters a subnet, its permissions are evaluated for resources within the subnet.
- The VPC component that checks packet permissions for resources, such as instances, is a security group.

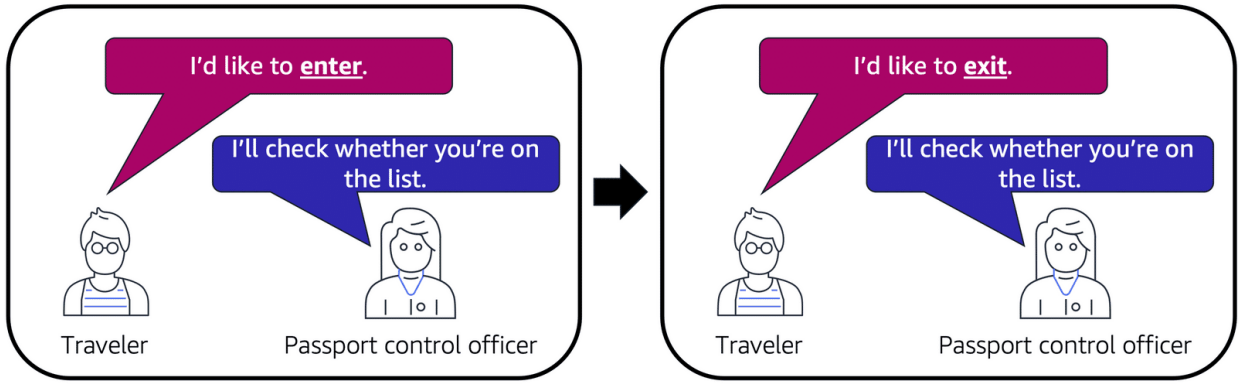

Security Groups

- A security group is a virtual firewall that controls inbound and outbound traffic **for an instance**.
- By default, a security group denies all inbound traffic and allows all outbound traffic.
- You can add custom rules to configure which traffic should be allowed; any other traffic is denied.
- Multiple instances within the same VPC can share the same security group or use different security groups for each instance.

Stateful Packet Filtering

- Security groups perform stateful packet filtering, remembering previous decisions made for incoming packets.
- When a request is sent from an instance to the internet, the security group remembers the request.
- The security group allows the response to proceed, regardless of inbound security group rules.
- Both network ACLs and security groups can have custom rules for traffic in your VPC.

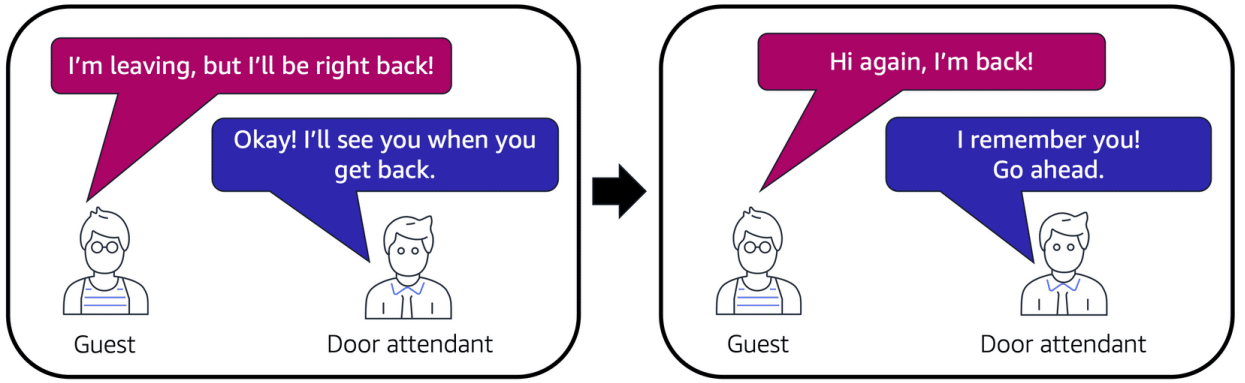

Network Traffic in a VPC *

- A packet travels over the internet from a client, to the internet gateway and into the VPC.
- Then the pack goes through the network access control list and accesses the public subnet, where two EC2 instances are located.

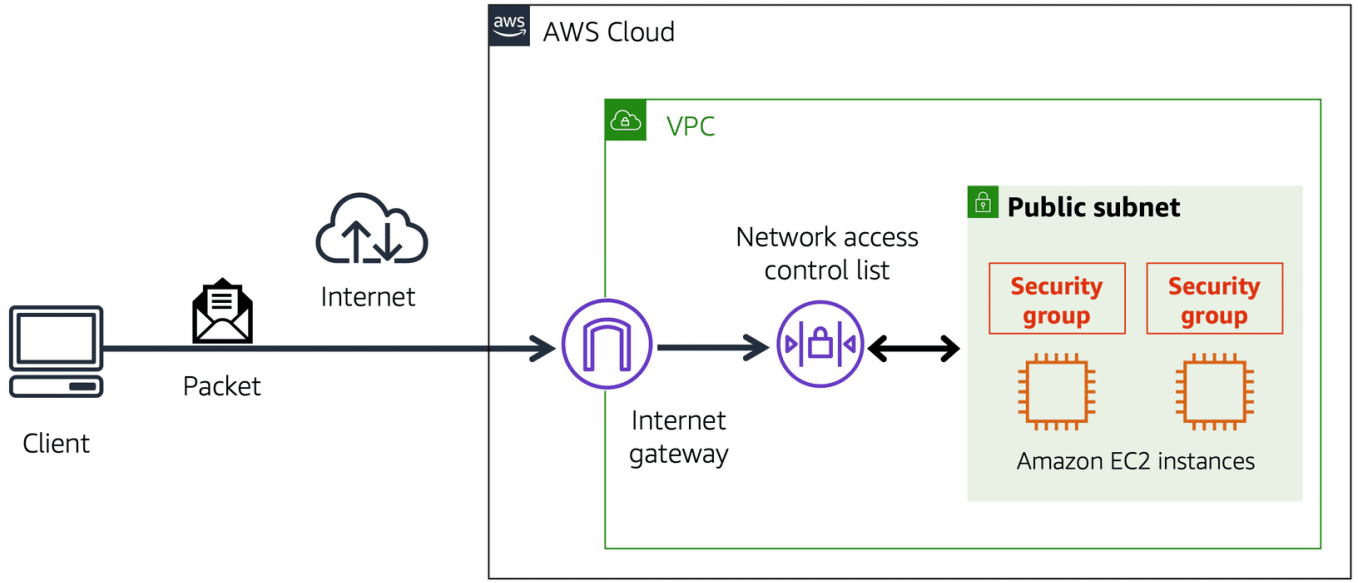

CDN with Cloud

- An application runs on multiple instances within an Auto Scaling group attached to a load balancer.
- A customer requests data from the application by accessing the website.
- DNS resolution identifies the corresponding IP address and sends this information back to the customer.
- The CDN connects to the load balancer, which sends the incoming packet to an instance.

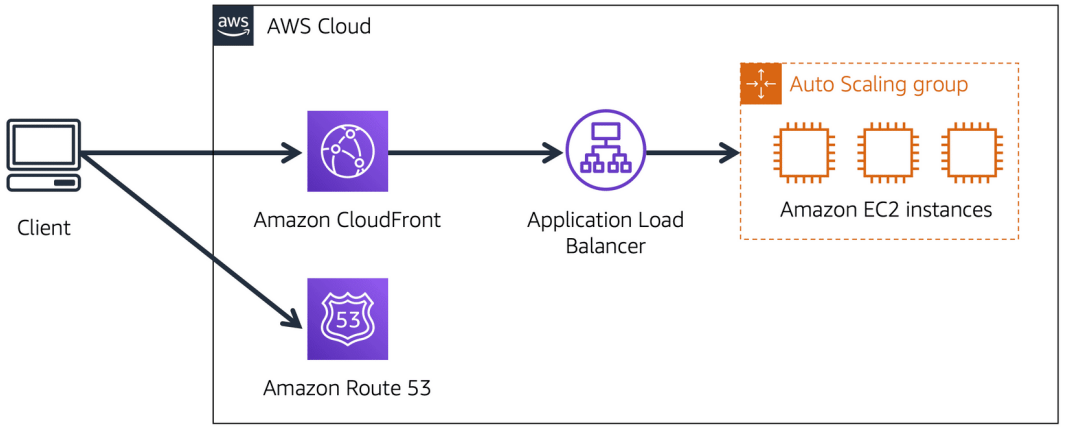

Multi-cloud Networking

- Multi-cloud networking involves using multiple cloud computing and storage services within a single network architecture.
- It enables easier access to and automated management of resources across multiple clouds and on-premises environments.
- Application and workload-awareness allows for understanding and managing the demands of specific applications and workloads.
- A SaaS-delivered control plane manages network traffic as a Software-as-a-Service.

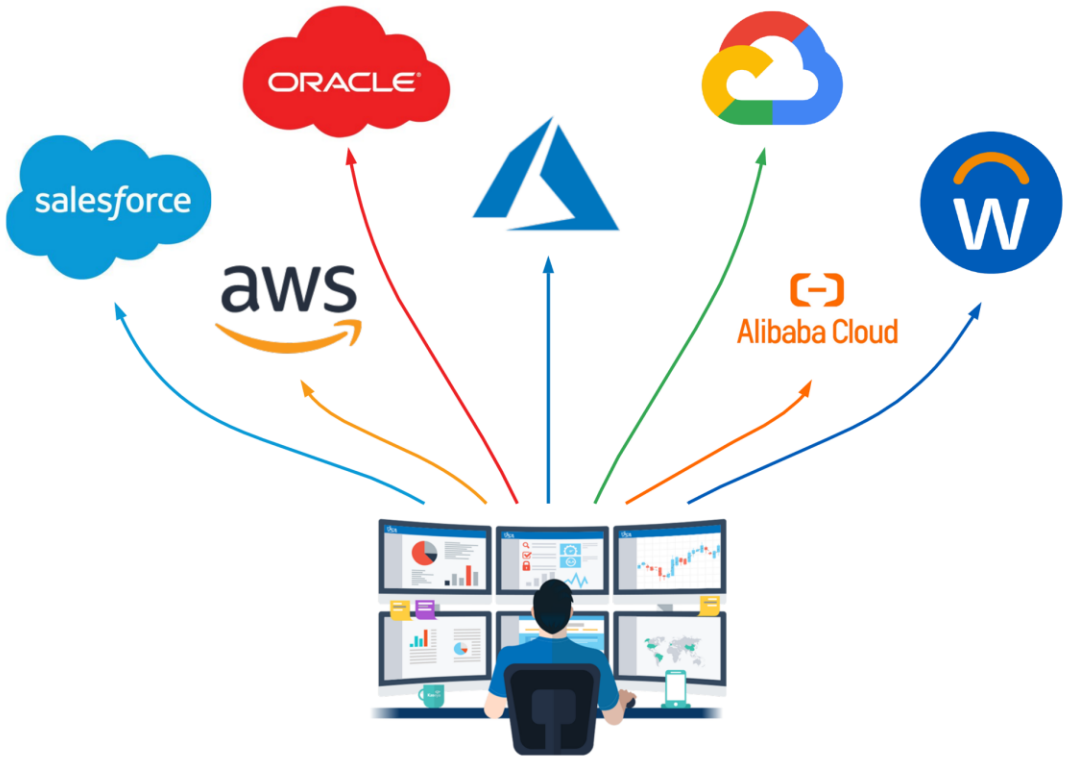

Multi-cloud Networking

- Integrations with public cloud providers ensure seamless operation with various public cloud services.
- Multi-cloud networking helps businesses avoid vendor lock-in, increase flexibility, and optimise costs and performance.
- Typical use cases include:
  - SD-WAN and SASE for optimised access to IaaS and SaaS services.
  - Multi-cloud SDN for consistent application-aware policy automation between on-premises and IaaS environments.
  - Hybrid application connectivity across an SD-WAN and between multiple public clouds/on-premises environments.

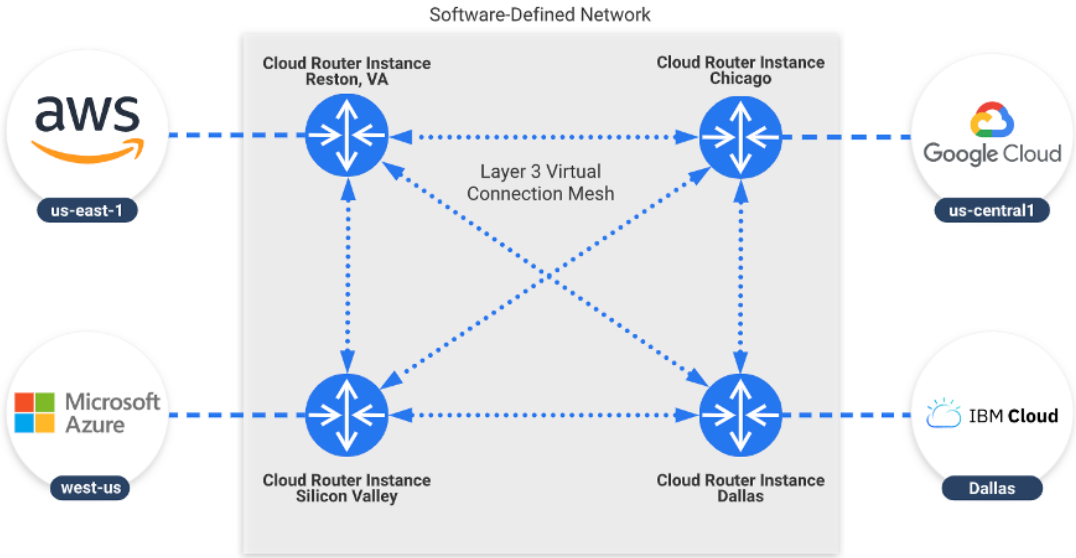

Hybrid cloud Networking

- Hybrid cloud networking is a subset of cloud and multi-cloud networking.
- It specifically relates to the connectivity between two clouds, such as:
  - On-premises private
  - Hosted private
  - Public
- It is also commonly used to describe:
  - Connectivity between an on-premises data centre or co-located facility and a public cloud.

Hybrid cloud Networking

- Hybrid cloud networking offers several benefits:
- Security:
  - Hybrid cloud infrastructures can provide enhanced security measures. Sensitive data can be kept on a private cloud while other data can be stored on a public cloud
- Cost-effectiveness:
  - With a hybrid cloud, organizations can optimize costs by using public cloud resources for less sensitive data and operations, while keeping critical operations and data in a private cloud or on-premises
- Flexibility and Control:
  - Hybrid cloud gives businesses the flexibility to choose where to host their applications and data based on regulatory, performance, or cost considerations

Additional resources

- https://aws.amazon.com/education/awseducate/
- https://cloud.google.com/learn/training/networking-security
- https://www.oracle.com/au/education/training/oracle-cloud-infrastructure/