

Cloud Computing, Edge Computing, Big Data & Artificial Intelligence

Prepared by Dr. Mohamed Sleem

Foundations of FinTech
FNT611

Lecture 2



Cloud Computing, Edge Computing, Big Data & Artificial Intelligence

Lecture 2

FNT611

Foundations of FinTech

Prepared by
Dr. Mohamed Sleem

Chapter 1 Discussion Questions

Write a brief essay explaining the difference between a narrow definition and a broad definition of the digital economy.

	Narrow Definition of Digital Economy	Broad Definition of Digital Economy
Association	Primarily with the internet economy	All digital activities
Focus	Economic activities directly related to digital technologies, especially the internet	All economic activities utilizing digital technology
Limitation	Activities strictly associated with the internet	No specific limitation; includes both direct and indirect usage of digital tech
Examples	E-commerce, online banking, digital advertising	E-commerce, digital manufacturing, digital health services, smart agriculture
Inclusion of indirect usage	No	Yes

Chapter 1 Discussion Questions

Explain the impact of the 4th Industrial Revolution on how we live and work, focusing on its key features.

- 4th IR introduces disruptive technologies affecting living and working.
- Enhances interconnectivity via IoT, improving device communication.
- Increases coordination and efficiency in different sectors.
- Advanced robotics and AI reduce the need for human intervention.
- Real-time data offers immediate insights for process optimization.
- Facilitates informed decision-making.
- Advanced UIs like VR and AR change how humans interact with tech.
- Alters experiences in professional and personal settings.

Chapter 1 Discussion Questions

Explain the role of Artificial Intelligence (AI) in the digital economy and how it influences business processes and user experiences.

- AI is pivotal in the digital economy as an economic catalyst.
- Automates processes and enhances user experiences.
- Crucial for swift data analysis of extensive datasets.
- Enables personalized services and products.
- Boosts consumer engagement and sales.

Objectives

- 1. Understand how cloud computing enables users to access storage, processing, and software applications without the need for local infrastructure.**
- 2. Recognize the importance of efficient resource utilization in cloud environments, focusing on scalability and measured service.**
- 3. Distinguish between different cloud deployment models and their suitability for various business and organizational needs.**
- 4. Evaluate and select the appropriate cloud service model (IaaS, PaaS, or SaaS) for an organization's specific requirements.**
- 5. Explain the concept of Edge Computing and its role in improving response times and conserving bandwidth by processing data closer to its source.**

Objectives


- 5. Demonstrate an understanding of the potential of big data in deriving valuable insights for informed decision-making.**
- 6. Describe the principles of Artificial Intelligence (AI) in problem-solving, pattern recognition, and decision-making that emulate human cognitive functions.**
- 7. Advocate for the responsible and ethical use of AI, being cognizant of its benefits and potential pitfalls.**
- 8. These learning outcomes set clear expectations for what the learners should know or be able to do after reading and understanding the chapter.**

Cloud Computing

What is cloud computing?

Cloud computing refers to the delivery of computing services over the Internet, including storage, processing power, and software applications.

- **It allows users to access resources and services on-demand, without the need for physical infrastructure or local servers.**



Characteristics of Cloud Computing

Key Characteristics of Cloud Computing

On-Demand Self-Service: Users can provision resources and services as needed, without requiring human interaction with service providers.


Broad Network Access: Services are accessible over the internet via standard protocols and devices.

Resource Pooling: Computing resources are pooled together to serve multiple users, allowing for efficient utilization and scalability.

Rapid Elasticity: Resources can be scaled up or down quickly to meet changing demands.

Measured Service: Cloud service usage is measured, monitored, and billed based on actual consumption.


Cloud Deployment Models



Cloud Deployment Models

Public Cloud


- **Services are provided over a public network and available to anyone who wants to use them.**
- **It is a cost-effective option for businesses and individuals looking for scalability and flexibility.**
- **Public cloud providers, such as AWS, Azure, and GCP, offer a wide range of services accessible to the general public.**



Cloud Deployment Models

Private Cloud


- Infrastructure is dedicated to a single organization and may be located on-premises or off-premises.
- Private cloud environments are designed to meet specific security, compliance, or performance requirements.
- They offer enhanced control, customization, and privacy but require significant upfront investment.



Cloud Deployment Models

Hybrid Cloud


- **Combines public and private cloud environments, allowing for flexibility and data sharing between the two.**
- **Organizations can leverage the benefits of both public and private clouds, ensuring optimal resource allocation.**
- **Hybrid cloud deployments enable workload portability and seamless integration between different environments.**



Cloud Deployment Models

Community Cloud


- **Community cloud is a deployment model where infrastructure and services are shared among a specific community or group of organizations.**
- **It caters to the needs of a particular community, such as government agencies, educational institutions, or research organizations.**
- **Community cloud provides a cost-effective solution while addressing specific requirements and compliance standards of the community.**



Cloud Deployment Models

Infrastructure as a Service (IaaS)


- IaaS provides virtualized computing resources over the internet.
- Users have control over the operating systems, storage, and networking components.
- They can provision and manage virtual machines (VMs), storage, and networks according to their requirements.
- Examples of IaaS providers include AWS EC2, Azure Virtual Machines, and Google Compute Engine.



Service Models

Platform as a Service (PaaS)

- PaaS offers a platform for developing, testing, and deploying applications.
- Users can focus on application development without worrying about infrastructure management.
- PaaS providers manage the underlying infrastructure, including servers, storage, and networking.
- Developers can leverage pre-configured environments, development frameworks, and deployment tools.
- Examples of PaaS providers include Heroku, Google App Engine, and AWS Elastic Beanstalk.



Service Models

Software as a Service (SaaS)

- SaaS delivers software applications over the internet on a subscription basis.
- Users can access and use applications directly through a web browser or APIs.
- The provider hosts and manages the underlying infrastructure, application, and data.
- Users can typically customize certain aspects of the application to fit their needs.
- Examples of SaaS include Salesforce, Microsoft Office 365, and Google Workspace.



Benefits of Cloud


- **Cost Savings:** Pay for what you use, with no upfront infrastructure costs.
- **Scalability:** Easily scale resources up or down based on demand.
- **Flexibility:** Access resources and applications from anywhere with an internet connection.
- **Reliability:** Cloud providers typically offer high uptime and data redundancy.
- **Collaboration:** Enable seamless collaboration and data sharing among teams.

Common Cloud Computing Use Cases

- **Data Storage and Backup:** Store and back up large amounts of data securely.
- **Software Development and Testing:** Rapidly create and deploy applications in a scalable environment.
- **Web and Mobile Applications:** Host web and mobile applications in the cloud for global accessibility.
- **Big Data Analytics:** Process and analyze vast amounts of data using cloud resources.
- **Disaster Recovery:** Maintain data backups and recovery plans in the cloud for business continuity.


Cloud Architecture

- The design and structure of cloud computing environments, including the arrangement of components and the relationships between them.
- It involves various elements that work together to deliver cloud services and ensure reliability, scalability, and security.



Cloud Architecture

NIST Architecture of Cloud



Cont...

- **Cloud Consumer:** An individual or organization that engages in a business partnership with Cloud Providers and utilizes their services.
- **Cloud Provider:** A person, organization, or entity that is accountable for offering services to individuals or entities who are interested in them.
- **Cloud Auditor:** A party capable of conducting an impartial evaluation of cloud services, as well as assessing the performance, security, and operations of the cloud implementation.
- **Cloud Broker:** An entity responsible for managing the utilization, performance, and delivery of cloud services. Additionally, they act as mediators between Cloud Providers and Cloud Consumers, negotiating relationships between the two parties.
- **Cloud Carrier:** An intermediary that facilitates the connectivity and transportation of cloud services from Cloud Providers to Cloud Consumers.

Cloud Architecture Best Practices

Up-front Planning: Prioritize up-front planning to understand capacity needs and avoid unexpected production glitches. Continuously test performance during the architecture design phase.

Security First: Implement robust security measures to protect against unauthorized access. Safeguard all layers of the cloud infrastructure through data encryption, patch management, and strict security policies. Consider adopting zero-trust security models for enhanced protection in hybrid and multi-cloud environments.

Cloud Architecture Best Practices

Disaster Recovery Readiness: Automate recovery processes to minimize costly downtime and ensure swift recovery from service disruptions. Implement monitoring mechanisms to track capacity and utilize redundant networks for a highly available architecture.

Maximize Performance: Optimize computing resources by continuously monitoring business demands and technological requirements. Efficiently manage and allocate resources to maintain optimal performance levels.

Cost Optimization: Leverage automated processes, managed service providers, and utilization tracking to reduce unnecessary cloud computing expenses. Regularly review and optimize resource allocation to ensure cost-effectiveness.

Conclusion

- **Cloud computing enables the delivery of computing services over the internet, eliminating the need for local infrastructure.**
- **Key characteristics of cloud computing include on-demand self-service, broad network access, resource pooling, rapid elasticity, and measured service.**
- **Deployment models include public, private, and hybrid clouds, offering flexibility and data-sharing options.**
- **Service models such as IaaS, PaaS, and SaaS provide virtualized computing resources, platforms for application development, and software delivery respectively.**
- **Cloud computing offers benefits like cost savings, scalability, flexibility, reliability, and enhanced collaboration.**
- **Cloud computing presents a transformative approach to computing, offering cost-effective solutions, scalability, and flexibility for organizations across various industries. By leveraging the benefits of cloud computing and aligning with the appropriate deployment and service models, organizations can drive innovation, improve efficiency, and adapt to evolving business needs.**


Remember

- **Understand your specific needs and requirements when selecting a cloud deployment and service model.**
- **Prioritize security measures such as data encryption and access controls to protect against unauthorized access.**
- **Plan for disaster recovery and high availability to ensure business continuity.**
- **Continuously monitor and optimize resource utilization for optimal performance and cost savings.**
- **Regularly assess and adapt your cloud architecture to evolving business needs and emerging technologies.**

Edge Computing



A distributed computing paradigm that brings computation and data storage closer to the sources of data.



This is done to improve response times and save bandwidth. Rather than sending vast amounts of data to a central data center or cloud for processing, the processing happens directly at the edge of the network, where the data is generated by devices, sensors, or users.

Components of Edge Computing:

- **Edge Devices:** These are the hardware devices at the edge of a network, which can include sensors, industrial machines, or other devices that produce or collect data.
- **Edge Gateways:** intermediaries between edge devices and the main data center or cloud. They process data from edge devices and determine what needs to be sent to the central system.
- **Edge Server:** Localized servers that can process data and even host local applications, reducing the need to send data back to a central data center.
- **Network Connectivity:** Ensures that the edge devices and servers can connect to the central system when needed. It's usually a mix of local networking (like LAN) and broader network connectivity (like WAN or cellular).
- **Edge Software:** Software platforms and applications specifically designed for edge computing environments. They often involve real-time processing capabilities.

Requirements for Edge Computing:

- **Low Latency:** Edge computing must provide quick data processing to offer real-time or near-real-time insights and actions.
- **Bandwidth Efficiency:** By processing data locally, edge computing reduces the need to continuously send large amounts of data over a network, saving bandwidth.
- **Security:** Given that data is processed at the edge, security protocols need to be robust. Edge devices and gateways can be more exposed to physical and cyber threats.
- **Reliability:** Systems must be robust and able to function even when connectivity to the main data center or cloud is compromised.
- **Scalability:** As more devices get connected, the edge computing infrastructure should be scalable to handle increasing data and processing needs.

Cloud Computing versus Edge Computing:

Feature	Cloud Computing	Edge Computing
Definition	Centralized computing resources hosted in data centers that deliver computing services over the internet.	Decentralized computing paradigm that processes data closer to data sources, such as IoT devices or local computing systems.
Data Processing Location	In remote data centers often far from the data source.	At or near the source of data, like sensors, IoT devices, or local appliances.
Latency	Typically higher due to data having to travel longer distances.	Lower, as data processing happens near or at the data source.
Bandwidth Consumption	Can be high since all data, whether processed or not, is sent to the cloud for processing.	Typically reduced as only necessary data or insights might be sent to a centralized system or cloud.
Use Cases	Big data analytics, large-scale applications, centralized storage, and backup solutions.	Real-time analytics, IoT applications, autonomous vehicles, smart cities, and augmented reality.
Scalability	Highly scalable as resources can be provisioned or de-provisioned on demand in large data centers.	Scalable but depends on the capabilities of local devices or the edge infrastructure.
Security Concerns	Data breaches at centralized data centers, data in transit vulnerabilities.	Potential vulnerabilities at many endpoints, concerns about the physical security of edge devices.
Cost	Costs associated with data transfer, storage, and processing in the cloud.	Can be reduced due to less data transfer to central locations, but there might be costs associated with edge infrastructure.

What is Big Data?



Big data is an all-encompassing term for any collection of data sets so large and complex that it becomes difficult to process using on-hand data management tools or traditional data processing applications.



Big data refers to extremely vast amounts of multi-structured data that typically has been cost prohibitive to store and analyze.



NOTE: however, big data is only referring to digital data, not the paper files stored in the basement at FBI headquarters, or piles of magnetic tapes in our data center.

Types of Big Data

In the simplest terms, Big Data can be broken down into two basic types, structured and unstructured data.

Structured – Predefined data type

- Spreadsheets and Oracle Relational database

Unstructured – is non pre-defined data model or is not organized in a pre-defined manner.

- Video, Audio, Images, Metadata, etc...

Semi-structured – Structured data embedded with some unstructured data

- Email, Text Messaging

A Big Data Platform must:

Analyze

Analyze a variety of different types of information

- This information by itself could be unrelated, but when paired with other information can illustrate a causation for an various events that the business can take advantage.

Analyze

Analyze information in motion

- Various data types will be streaming and contain a large amount of "bursts". Ad hoc analysis needs to be done on the streaming data to search for relevant events.

Cover

Cover extremely large volumes of data

- Due to the proliferation of devices in the network, how they are used, along with customer interactions on smartphones and the web, a cost efficient process to analyze the petabytes of information is required

A Big Data Platform must: (2)

Cover varying types of data sources

- ❑ Data can be streaming, batch, structured, unstructured, and semi-structured, depending on the information type, where it comes from and its primary use. Big Data must be able to accommodate all of these various types of data on a very large scale.

Analytics

- ❑ Big Data must provide the mechanisms to allow ad-hoc queries, data discovery and experimentation on the large data sets to effectively correlate various events and data types to get an understanding of the data that is useful and addresses business needs.

Five Characteristics of Big Data

Big Data is defined by five characteristics:

Volume:

Today's services generate massive data, demanding efficient engineering and parallel processing.

Tools like Hadoop, HBASE, and Cassandra are crucial for managing this scale.

Unique architectural approaches are needed for analysis, as traditional storage methods might not apply.

Velocity:

Big Data's value is in its timeliness, with users expecting real-time responses.

Traditional data warehousing struggles with delivering immediate results from streaming data.

Five Characteristics of Big Data (2)

Variety:	Big Data combines multiple data forms: structured (fixed fields), unstructured (text, multimedia), and semi-structured (mix of both, like medical records).
-----------------	---

Veracity:	Data quality varies greatly. One-third of business leaders doubt their data's trustworthiness, making trust establishment challenging as data sources increase.
------------------	---

Variability:	Data types and sources frequently change.
---------------------	---

Note:	Big Data involves massive datasets challenging for traditional tools. It requires advanced, parallel software on numerous servers, as common tools struggle with its volume and complexity.
--------------	---


Introduction to AI

Artificial Intelligence, commonly referred to as AI, pertains to the creation of computer systems or software that can perform tasks that typically require human intelligence.

These tasks include problem-solving, understanding natural language, recognizing patterns, making decisions, and interpreting visual cues, among others.

AI aims to mimic cognitive functions that humans associate with the human mind, such as "learning" and "reasoning".

The Functionalities and Applications of AI



Major Goals of AI



Simulate Human Intelligence: Creating systems that can perform tasks requiring human intelligence such as visual perception, speech recognition, decision-making, and language translation.



Reasoning: Developing systems that can emulate human reasoning, solving puzzles or making logical decisions, often at a much faster rate than humans.



Knowledge Representation: Storing and retrieving knowledge for problem-solving. This involves understanding human knowledge and representing it in a manner that's accessible and usable by machines.



Planning and Navigation: Machines that can plan tasks or navigate through various environments, like autonomous vehicles that can determine optimal routes.



Natural Language Processing (NLP): Understanding, interpreting, and generating human language in a way that is valuable. This could be in the form of chatbots, translation services, or sentiment analysis.

Artificial Intelligence:




FIGURE 5.1 Architecture of the AI machine.

Benefits of AI

Efficiency and Automation: AI can automate routine tasks, process large amounts of data at high speeds, and perform operations faster than humans, leading to increased efficiency.

Data Analysis: With its capability to process vast amounts of data, AI can identify patterns, make predictions, and offer insights that might be impossible or very time-consuming for humans.

Personalization: AI algorithms, especially in areas like online shopping or content streaming, tailor experiences to individual users based on their behaviors and preferences.

Cost Reduction: Over time, the use of AI can lead to significant cost savings, especially in fields where automation can reduce the need for manual labor.

Error Reduction: AI systems, once properly trained, can reduce human errors, especially in fields that require precision like medicine or aviation.

Availability: Unlike humans, AI systems can work 24/7 without breaks, fatigue, or emotional disturbances.

Complex Problem Solving: AI can tackle and solve complex problems by analyzing multiple variables at once, which might be challenging for humans.

Examples of A I Benefits

Kabbage:

- **Benefit:** Automated Lending Process
- **Description:** Kabbage offers an online platform that uses machine learning to analyze various data points, from e-commerce transactions to social media activity, to determine loan eligibility in a matter of minutes. This approach reduces the need for lengthy traditional credit assessments and streamlines the lending process.

Robinhood:

- **Benefit:** Democratized Stock Trading
- **Description:** Robinhood is a commission-free trading platform that uses AI to offer personalized financial news and alerts to its users. The platform analyzes users' trading behaviors and preferences to provide curated news and trading suggestions, enhancing user engagement.

Examples of AI Benefits (2)

Adyen:

- **Benefit:** Fraud Detection and Risk Management
- **Description:** Adyen, a global payment company, utilizes machine learning algorithms to analyze transactions in real-time and detect potential fraudulent activities. By analyzing patterns and behaviors, their system can identify and prevent fraudulent transactions, thereby reducing financial loss and ensuring trust in their platform.

Betterment:

- **Benefit:** Robo-Advisory for Personalized Financial Planning
- **Description:** Betterment, a pioneer in robo-advisory, uses AI algorithms to provide personalized financial advice and automated investing services. Based on users' financial goals and risk tolerance, Betterment's algorithms suggest optimal investment strategies, making sophisticated financial planning accessible to the masses.

Limitations of AI:



Lack of Emotional Intelligence: AI lacks human emotions and can't understand nuances in human feelings, making it unsuitable for tasks that require emotional understanding.



Ethical Concerns: AI decisions, especially when unexplainable (as in some deep learning methods), can raise ethical issues, including biases in decision-making processes.



Dependence on Data: AI's performance is heavily dependent on the quality and quantity of the data it's trained on. Poor or biased data can lead to incorrect or biased outcomes.



Cost and Complexity: Setting up an AI system, especially sophisticated ones, can be expensive and complex, requiring expertise and resources not available to all.

Some limitations of A I Machines

Lack human touch and feel

Lack attention to non-task surroundings

Can lead people to rely on A I machines too much

Can be programmed to create destruction

Can cause many people to lose their jobs

Can start to think by themselves, causing significant damage

- Hypothetically ... no evidence of that!

These limitations are diminishing over time

Chapter Summary: Key Takeaways

Cloud Computing: Delivers computing services via the internet.

Cloud Characteristics: Features on-demand service, broad network access, and pooled resources.

Deployment Models: Encompasses public, private, hybrid, and community clouds.

Service Models: Defined as IaaS (Infrastructure as a Service), PaaS (Platform as a Service), and SaaS (Software as a Service).

Edge Computing: Processes and stores data closer to its origin, enhancing speed and efficiency.

Big Data: Encompasses massive datasets that can be structured, semi-structured, or unstructured.

Artificial Intelligence: Replicates human-like thinking and problem-solving.

AI Considerations: While AI boasts numerous advantages, it's vital to recognize its limitations and address ethical implications.

Discussion

1. Define cloud computing and its significance in today's digital landscape.
2. Explain the concept of Edge Computing and its components. Why is it essential in the age of Big Data and IoT?
3. Provide an overview of Artificial Intelligence (AI) and its primary goals. How does AI differ from traditional computing?



Thank You