

CAP Theorem

What is the CAP Theorem?

The CAP theorem states that a distributed system can only guarantee two out of the following three properties at any given time–

1. **Consistency (C)**– Every read receives the most recent write or an error.
2. **Availability (A)**– Every request receives a response, even if it is not the most recent data.
3. **Partition Tolerance (P)**– The system continues to function despite network partitions.

In essence, no distributed system can achieve all three properties simultaneously. This constraint forces system designers to make trade-offs based on their use case and requirements.

The Components of CAP Theorem

Consistency (C)

Consistency ensures that all nodes in a distributed system reflect the same data at any given time. For example–

- If one node updates a value, other nodes in the system should reflect that updated value immediately.

Example– Traditional relational databases like MySQL prioritize consistency.

Availability (A)

Availability ensures that the system always provides a response to user requests, regardless of the state of individual nodes. However, the response may not always reflect the most recent data.

Example– Systems like DNS ensure high availability even at the expense of providing slightly outdated data.

Partition Tolerance (P)

Partition tolerance ensures that the system continues to operate despite communication failures between nodes. Network partitions, caused by hardware failures or connectivity issues, are unavoidable in distributed systems.

Example– Modern distributed databases like Cassandra and MongoDB prioritize partition tolerance.

The Implications of CAP Theorem

The CAP theorem implies that designers must make trade-offs between consistency, availability, and partition tolerance—

- **CP (Consistency, Partition tolerance) Systems**— Focus on consistency and partition tolerance but sacrifice availability.
- **AP (Availability, Partition tolerance) Systems**— Focus on availability and partition tolerance but sacrifice consistency.
- **CA (Consistency, Availability) Systems**— Focus on consistency and availability but cannot handle network partitions.

In reality, partition tolerance is often a requirement for distributed systems, forcing a choice between consistency and availability.

Trade-offs in CAP Theorem

Choosing Between Consistency and Availability

The decision to prioritize consistency or availability depends on the use case—

- **Consistency-Centric Systems (CP)**—
 - Used in systems requiring strict correctness, such as financial transactions or inventory management.
 - **Trade-off**— May be unavailable during network partitions.
- **Availability-Centric Systems (AP)**—
 - Used in systems prioritizing uptime over strict correctness, such as social media feeds or caching layers.
 - **Trade-off**— May serve outdated data during network partitions.

Partition Tolerance as a Requirement

Partition tolerance is critical in distributed systems due to—

- Network failures, latency, or disconnections.
- Scalability requirements across geographically distributed nodes.

Since network partitions are inevitable, the CAP theorem forces designers to choose between consistency and availability during these events.

Consistency Models in CAP

Consistency models define how and when updates to data become visible in a distributed system. Key models include—

Strong Consistency

- Guarantees that all clients see the same data after an update.
- Ensures correctness but can increase latency.
- **Example**— Traditional databases using ACID transactions.

Eventual Consistency

- Ensures that all nodes converge to the same state eventually, though not immediately.
- Prioritizes availability and performance.
- **Example**— NoSQL databases like DynamoDB or Cassandra.

Weak Consistency

- Provides no guarantees about when data updates will be visible.
- **Example**— Systems with eventual synchronization during low-priority tasks.

Real-World Applications of CAP

Relational Databases (CP Systems)

Relational databases prioritize consistency and partition tolerance—

- **Use Cases**— Financial systems, e-commerce transactions.
- **Example**— MySQL, PostgreSQL.

NoSQL Databases (AP Systems)

NoSQL databases prioritize availability and partition tolerance—

- **Use Cases**— Content delivery, social media, IoT systems.
- **Example**— Cassandra, MongoDB.

Distributed Caches (CA Systems)

Distributed caching systems prioritize fast responses—

- **Use Cases**– Web page loading, search results.
- **Example**– Redis, Memcached.

Challenges of CAP Theorem in Modern Systems

While CAP theorem provides a useful framework, it has limitations–

- **Binary Choices**– In practice, systems often achieve a balance between consistency and availability rather than strict adherence to one.
- **Latency Issues**– Strong consistency models can introduce significant delays.
- **Evolving Architectures**– Modern systems often employ hybrid strategies to mitigate CAP constraints.