

## Lesson 6.2: Core Banking and Legacy Systems – Quiz

Prof. Dr. Joerg Osterrieder

## Question 1

**Which function is NOT typically part of a core banking system?**

- A. Real-time ML-based fraud scoring on card transactions
- B. Interest accrual and fee calculation
- C. Account posting (debits and credits to customer balances)
- D. Regulatory reporting (general ledger extracts for supervisors)

## Question 1

**Which function is NOT typically part of a core banking system?**

- A. Real-time ML-based fraud scoring on card transactions
- B. Interest accrual and fee calculation
- C. Account posting (debits and credits to customer balances)
- D. Regulatory reporting (general ledger extracts for supervisors)

*[Answer hidden – compile with \solutionstrue to reveal]*

“Core banking” is the central ledger of record — account posting, interest, fees, GL extracts. Real-time fraud-scoring ML is usually a separate system (fraud-management platform or card-network service) that calls into core for balance checks, not part of the core itself.

## Question 2

**Which fundamental accounting principle does a general ledger enforce?**

- A. Mark-to-market: assets valued at current market price daily
- B. Single-entry bookkeeping: each transaction recorded once
- C. Cash-basis accounting: transactions recorded when cash changes hands
- D. Double-entry bookkeeping: every debit has a matching credit

## Question 2

**Which fundamental accounting principle does a general ledger enforce?**

- A. Mark-to-market: assets valued at current market price daily
- B. Single-entry bookkeeping: each transaction recorded once
- C. Cash-basis accounting: transactions recorded when cash changes hands
- D. Double-entry bookkeeping: every debit has a matching credit

*[Answer hidden – compile with \solutionstrue to reveal]*

The general ledger enforces double-entry bookkeeping, ensuring  $\text{Assets} = \text{Liabilities} + \text{Equity}$  at all times. Every financial transaction has offsetting debit and credit entries.

## Question 3

**Where does COBOL typically run in a modern ATM transaction flow?**

- A. On the downstream authorization host / core-banking mainframe
- B. On the ATM terminal itself (the user-facing hardware)
- C. COBOL is not used anywhere in the ATM flow
- D. On the card-network switch (Visa/Mastercard)

## Question 3

**Where does COBOL typically run in a modern ATM transaction flow?**

- A. On the downstream authorization host / core-banking mainframe
- B. On the ATM terminal itself (the user-facing hardware)
- C. COBOL is not used anywhere in the ATM flow
- D. On the card-network switch (Visa/Mastercard)

*[Answer hidden – compile with \solutionstrue to reveal]*

ATM terminals themselves usually run Windows or Linux. COBOL runs on the authorization host (typically IBM z/OS mainframe) that actually debits the account. Reuters (2017) estimated COBOL underpins about 43% of banking systems; the “95% of ATMs” folklore conflates terminal OS with back-end ledger.

## Question 4

**What is the primary risk of COBOL-based legacy systems?**

- A. COBOL cannot handle multi-currency transactions
- B. Mainframes are physically too large for modern data centers
- C. COBOL is too slow for modern transaction volumes
- D. The shrinking talent pool of COBOL programmers and undocumented business logic

## Question 4

**What is the primary risk of COBOL-based legacy systems?**

- A. COBOL cannot handle multi-currency transactions
- B. Mainframes are physically too large for modern data centers
- C. COBOL is too slow for modern transaction volumes
- D. The shrinking talent pool of COBOL programmers and undocumented business logic

*[Answer hidden – compile with \solutionstrue to reveal]*

The language itself performs well (99.999% uptime). The risk is the ageing talent pool (Micro Focus 2022 COBOL Market Survey: 47% of COBOL developers are 45+, 15% over 55) and decades of undocumented business logic embedded in the code.

## Question 5

**Technical debt is best described as:**

- A. The implied cost of future rework from choosing quick solutions over better long-term approaches
- B. The cost of purchasing software licenses for legacy systems
- C. The financial debt a bank takes on to fund IT projects
- D. The depreciation of IT hardware over time

## Question 5

**Technical debt is best described as:**

- A. The implied cost of future rework from choosing quick solutions over better long-term approaches
- B. The cost of purchasing software licenses for legacy systems
- C. The financial debt a bank takes on to fund IT projects
- D. The depreciation of IT hardware over time

*[Answer hidden – compile with \solutionstrue to reveal]*

Technical debt is a metaphor coined by Ward Cunningham (1992). Like financial debt, you pay “interest” in the form of slower development, more bugs, and higher maintenance costs.

## Question 6

**What proportion of their IT budget do most traditional banks spend on maintenance (“run the bank”) rather than innovation?**

- A. 70–80%
- B. 40–50%
- C. 95–100%
- D. 20–30%

## Question 6

**What proportion of their IT budget do most traditional banks spend on maintenance (“run the bank”) rather than innovation?**

- A. 70–80%
- B. 40–50%
- C. 95–100%
- D. 20–30%

*[Answer hidden – compile with \solutionstrue to reveal]*

Most traditional banks spend 70–80% of IT budgets on maintaining existing systems, leaving only 20–30% for innovation. This is a direct consequence of accumulated technical debt.

## Question 7

**What happened during TSB Bank's core banking migration in 2018?**

- A. A big bang cutover that locked out 1.9 million customers and exposed some to other people's accounts
- B. A successful phased rollout over 2 years with no customer impact
- C. A strangler fig migration that was halted due to regulatory concerns
- D. A cloud migration that was reversed after data sovereignty concerns

## Question 7

**What happened during TSB Bank's core banking migration in 2018?**

- A. A big bang cutover that locked out 1.9 million customers and exposed some to other people's accounts
- B. A successful phased rollout over 2 years with no customer impact
- C. A strangler fig migration that was halted due to regulatory concerns
- D. A cloud migration that was reversed after data sovereignty concerns

*[Answer hidden – compile with \solutionstrue to reveal]*

TSB attempted a big bang migration from Lloyds' legacy platform to Sabadell's Proteo4UK. The failed cutover cost £330M, caused the CEO to resign, and 80,000 customers left the bank.

## Question 8

**The strangler fig pattern for software migration is named after:**

- A. A database migration technique invented by IBM
- B. A tropical vine that gradually grows around and replaces a host tree
- C. An economic theory about disruptive innovation
- D. A computer science algorithm for graph traversal

## Question 8

**The strangler fig pattern for software migration is named after:**

- A. A database migration technique invented by IBM
- B. A tropical vine that gradually grows around and replaces a host tree
- C. An economic theory about disruptive innovation
- D. A computer science algorithm for graph traversal

*[Answer hidden – compile with \solutionstrue to reveal]*

The strangler fig pattern, popularized by Martin Fowler (2004), is named after the strangler fig tree. The vine gradually grows around a host tree, eventually replacing it—just as new services gradually replace legacy components.

## Question 9

**Which migration strategy involves running old and new systems simultaneously, comparing outputs daily?**

- A. Canary deployment
- B. Big bang
- C. Parallel run
- D. Strangler fig

## Question 9

**Which migration strategy involves running old and new systems simultaneously, comparing outputs daily?**

- A. Canary deployment
- B. Big bang
- C. Parallel run
- D. Strangler fig

*[Answer hidden – compile with \solutionstrue to reveal]*

Parallel run operates both systems simultaneously and compares results. It offers maximum safety but at double the operating cost and with significant data synchronization challenges.

## Question 10

**Which of these is a banking-specific challenge for cloud migration?**

- A. Cloud-native applications cannot achieve regulatory compliance
- B. Cloud providers do not offer sufficient compute capacity
- C. Cloud platforms cannot support relational databases
- D. Data sovereignty regulations require customer data to stay in specific jurisdictions

## Question 10

**Which of these is a banking-specific challenge for cloud migration?**

- A. Cloud-native applications cannot achieve regulatory compliance
- B. Cloud providers do not offer sufficient compute capacity
- C. Cloud platforms cannot support relational databases
- D. Data sovereignty regulations require customer data to stay in specific jurisdictions

*[Answer hidden – compile with \solutionstrue to reveal]*

Data sovereignty (GDPR, FINMA, and similar regulations) requires that customer data remains within specific geographic boundaries. Banks must select cloud regions carefully and may need multi-cloud or hybrid strategies.

## Question 11

**Approximately what percentage of banking workloads were in public cloud as of 2024?**

- A. Over 80%
- B. Less than 5%
- C. About 50%
- D. About 15%

## Question 11

**Approximately what percentage of banking workloads were in public cloud as of 2024?**

- A. Over 80%
- B. Less than 5%
- C. About 50%
- D. About 15%

*[Answer hidden – compile with \solutionstrue to reveal]*

Only approximately 15% of banking workloads are in public cloud. Hybrid cloud dominates: sensitive data and core processing remain on-premise while analytics and development environments move to cloud.

## Question 12

**In a microservice architecture, how do individual services typically communicate?**

- A. By sharing memory within a single process
- B. Via APIs (synchronous) or message queues (asynchronous)
- C. Through a shared database that all services read from and write to
- D. Through direct file system access on a shared server

## Question 12

**In a microservice architecture, how do individual services typically communicate?**

- A. By sharing memory within a single process
- B. Via APIs (synchronous) or message queues (asynchronous)
- C. Through a shared database that all services read from and write to
- D. Through direct file system access on a shared server

*[Answer hidden – compile with \solutionstrue to reveal]*

Microservices communicate via APIs (typically REST or gRPC for synchronous calls) or message queues/event streams (Kafka, RabbitMQ for asynchronous communication). Each service owns its own data store.

**What is a key advantage of monolithic architecture over microservices?**

- A. Superior horizontal scalability
- B. Better fault isolation between modules
- C. Simpler to develop, debug, and deploy initially
- D. Independent deployment of components

**What is a key advantage of monolithic architecture over microservices?**

- A. Superior horizontal scalability
- B. Better fault isolation between modules
- C. Simpler to develop, debug, and deploy initially
- D. Independent deployment of components

*[Answer hidden – compile with \solutionstrue to reveal]*

Monolithic applications are simpler to develop, debug (single process), and deploy initially. Microservices introduce distributed-system complexity (network failures, eventual consistency, service discovery) that is only justified at scale.

**What does “API-first” design mean?**

- A. Only external partners can access the system through APIs
- B. The API contract is designed before the implementation, and the API is treated as the product
- C. APIs are prioritized over the user interface in the development schedule
- D. APIs are built after the application is complete, as an afterthought

**What does “API-first” design mean?**

- A. Only external partners can access the system through APIs
- B. The API contract is designed before the implementation, and the API is treated as the product
- C. APIs are prioritized over the user interface in the development schedule
- D. APIs are built after the application is complete, as an afterthought

*[Answer hidden – compile with \solutionstrue to reveal]*

API-first means designing the API contract (using specifications like OpenAPI/Swagger) before building the implementation. This enables parallel development and ensures consistent, well-documented interfaces.

**Apache Kafka is primarily used for:**

- A. Batch file transfer between banking systems
- B. Relational database management
- C. Distributed event streaming and real-time data processing
- D. Static web page hosting

## Question 15

**Apache Kafka is primarily used for:**

- A. Batch file transfer between banking systems
- B. Relational database management
- C. Distributed event streaming and real-time data processing
- D. Static web page hosting

*[Answer hidden – compile with \solutionstrue to reveal]*

Apache Kafka is a distributed event streaming platform used for building real-time data pipelines. In banking, it enables real-time fraud detection, instant balance updates, and regulatory event reporting.

### What is CQRS (Command Query Responsibility Segregation)?

- A. A database backup strategy that creates read-only replicas
- B. A regulatory requirement for separating customer data from operational data
- C. A pattern that separates the data model for reading (queries) from the model for writing (commands)
- D. A cloud migration strategy that processes reads before writes

### What is CQRS (Command Query Responsibility Segregation)?

- A. A database backup strategy that creates read-only replicas
- B. A regulatory requirement for separating customer data from operational data
- C. A pattern that separates the data model for reading (queries) from the model for writing (commands)
- D. A cloud migration strategy that processes reads before writes

*[Answer hidden – compile with \solutionstrue to reveal]*

CQRS separates the read model (optimized for queries) from the write model (optimized for commands). This allows each side to scale independently and use different data stores optimized for their workload.

## Question 17

**What is the key difference between a data lake and a data warehouse?**

- A. A data warehouse stores raw data; a data lake stores processed data
- B. A data lake is always more expensive than a data warehouse
- C. A data lake can only store structured data; a data warehouse stores unstructured data
- D. A data lake uses schema-on-read; a data warehouse uses schema-on-write

## Question 17

**What is the key difference between a data lake and a data warehouse?**

- A. A data warehouse stores raw data; a data lake stores processed data
- B. A data lake is always more expensive than a data warehouse
- C. A data lake can only store structured data; a data warehouse stores unstructured data
- D. A data lake uses schema-on-read; a data warehouse uses schema-on-write

*[Answer hidden – compile with \solutionstrue to reveal]*

A data lake ingests raw data without requiring a predefined schema (schema-on-read). A data warehouse requires data to be transformed and structured before loading (schema-on-write). Lakes are more flexible but risk becoming “data swamps” without governance.

**In data mesh architecture, who owns the data?**

- A. The Chief Data Officer exclusively
- B. Domain teams who treat their data as a product
- C. A centralized data engineering team
- D. External cloud providers who host the data

## Question 18

**In data mesh architecture, who owns the data?**

- A. The Chief Data Officer exclusively
- B. Domain teams who treat their data as a product
- C. A centralized data engineering team
- D. External cloud providers who host the data

*[Answer hidden – compile with \solutionstrue to reveal]*

Data mesh decentralizes data ownership to domain teams (e.g., payments, lending, compliance). Each team owns and maintains its data as a product, with a federated governance model ensuring interoperability.

## Question 19

**What percentage of core banking migrations exceed their original timeline, according to industry data?**

- A. 20%
- B. 40%
- C. 70%
- D. 95%

## Question 19

**What percentage of core banking migrations exceed their original timeline, according to industry data?**

- A. 20%
- B. 40%
- C. 70%
- D. 95%

*[Answer hidden – compile with \solutionstrue to reveal]*

Approximately 70% of core banking migrations exceed their original timeline. Average duration for a large bank is 3–7 years. Success rates are higher with phased approaches (strangler fig) and early regulator engagement.

## Question 20

**A bank is evaluating cloud-native core banking vendors. Which factor is LEAST relevant to the decision?**

- A. Data sovereignty and geographic hosting options
- B. The programming language the vendor's engineers prefer personally
- C. Vendor's financial stability and long-term viability
- D. Regulatory approval for the target jurisdiction

## Question 20

**A bank is evaluating cloud-native core banking vendors. Which factor is LEAST relevant to the decision?**

- A. Data sovereignty and geographic hosting options
- B. The programming language the vendor's engineers prefer personally
- C. Vendor's financial stability and long-term viability
- D. Regulatory approval for the target jurisdiction

*[Answer hidden – compile with \solutionstrue to reveal]*

Vendor evaluation should focus on technology architecture, regulatory fit, commercial terms, vendor viability, and implementation track record. The personal language preference of the vendor's engineers is not a material evaluation criterion.